

Secure Deduplication of Encrypted Data without Additional Independent Servers

Jian Liu
Aalto University
jian.liu@aalto.fi

N. Asokan
Aalto University and
University of Helsinki
asokan@acm.org

Benny Pinkas
Bar Ilan University
benny@pinkas.net

Abstract

Encrypting data on client-side before uploading it to a cloud storage is essential for protecting users' privacy. However client-side encryption is at odds with the standard practice of deduplication. Reconciling client-side encryption with cross-user deduplication is an active research topic. We present the first secure cross-user deduplication scheme that supports client-side encryption *without requiring any additional independent servers*. Interestingly, the scheme is based on using a PAKE (password authenticated key exchange) protocol. We demonstrate that *our scheme provides better security guarantees than previous efforts*. We show both the effectiveness and the efficiency of our scheme, via simulations using realistic datasets and an implementation.

1. INTRODUCTION

Cloud storage is a service that enables people to store their data on a remote server. With a rapid growth in user base, cloud storage providers tend to save storage costs via *cross-user deduplication*: if two clients upload the same file, the storage server detects the duplication and stores only a single copy. Deduplication achieves high storage savings [21] and is adopted by many storage providers. It is also adopted widely in backup systems for enterprise workstations.

Clients who care about privacy prefer to have their data encrypted on the client-side using *semantically secure* encryption schemes. However, naïve application of encryption thwarts deduplication since identical files are uploaded as completely independent ciphertexts. Reconciling deduplication and encryption is an active research topic. The current solutions either use *convergent encryption* [13], which is susceptible to *offline brute-force attacks*, or require the aid of additional *independent servers* [4, 24, 26], which is a strong assumption that is very difficult to meet in commercial contexts. Furthermore, some schemes of the latter type are susceptible to *online brute-force attacks*.

Our Contributions. We present the first *single-server* scheme for secure cross-user deduplication with client-side encrypted data. Our scheme allows a client uploading an existing file to securely obtain the encryption key that was used by the client who has previously uploaded that file. The scheme builds upon a well-known cryptographic primitive known as *password authenticated key exchange* (PAKE) [7], which allows two parties to agree on a session key iff they share a short secret (“password”). PAKE is secure even if the passwords have low entropy. In our deduplication scheme, a PAKE-based protocol is used to compute identical keys for different copies of the same file. Specifically, a client up-

loading a file first sends a short hash of this file (10-20 bits long) to the server. The server identifies other clients whose files have the same short hash, and let them run a single round PAKE protocol (routed through the server) with the uploader using the (long, but possibly low entropy) hashes of their files as the “passwords”. At the end of the protocol, the uploader gets the key of another client iff their files are identical. Otherwise, it gets a random key.

Our scheme uses a *per-file rate limiting* strategy to prevent online brute-force attacks. Namely, clients protect themselves by limiting the number of PAKE instances they will participate in for each file. Compared with the commonly used *per-client rate limiting* (used in DupLESS [4]), which limits the number of queries allowed for each client (during a time interval), our scheme is significantly more resistant to online brute-force attacks by an adversary who has compromised multiple clients, or by the storage server. Per-client rate limiting is not fully effective against such attacks because the adversary can use different identities.

At a first glance, it seems that our scheme incurs a high communication and computation overhead because a client uploading a file is required to run PAKE many times due to the high collision rate of the short hash. In fact, the number of PAKE runs for an upload request is limited to a certain level by the rate limiting strategy. For a requested short hash, the server only checks a subset of existing files (in descending order of file popularity) that have the same short hash. This implies that our scheme may fail to find duplicates for some requests, and this will certainly reduce the deduplication effectiveness. Surprisingly, our simulations in Section 6 show that this negative effect is very small. The reason is that the file popularity distribution is far from being uniform, and popular files account for most of the benefit from deduplication. Our scheme can almost always find duplicates for these popular files.

To summarize, we make the following contributions:

- Presenting the **first single-server scheme for cross-user deduplication** that enables client-side semantically secure encryption (except that, as in all deduplication schemes, ciphertexts leak the equality of the underlying files), and proving its security in the **malicious model** (Section 5);
- Showing that our scheme has **better security guarantees** than previous work (Section 5.2). As far as we know, our proposal is the first scheme that can prevent online brute-force attacks by compromised clients or by the storage server, *without* the aid of an identity server;

- Demonstrating, via simulations with realistic datasets, that our scheme provides its privacy benefits while retaining a utility level (in terms of deduplication effectiveness) on **par with standard industry practice**. Our use of per-file rate limiting implies that an incoming file will *not* be checked against all existing files in storage. Notably, our scheme still achieves high deduplication effectiveness (Section 6);
- **Implementing** our scheme to show that it incurs **minimal overhead** (computation/communication) compared to traditional deduplication schemes (Section 7).

2. PRELIMINARIES

2.1 Deduplication

Deduplication strategies can be categorized according to the basic data units they handle. One strategy is *file-level deduplication* which eliminates redundant files [13]. The other is *block-level deduplication*, in which files are segmented into blocks and duplicate blocks are eliminated [25]. In this paper we use the term “file” to refer to both files and blocks.

Deduplication strategies can also be categorized according to the host where deduplication happens. In *server-side deduplication*, all files are uploaded to the storage server, which then deletes the duplicates. Clients are unaware of deduplication. This strategy saves storage but not bandwidth. In *client-side deduplication*, a client uploading a file first checks the existence of this file on the server (by sending a hash of the file). Duplicates are not uploaded. This strategy saves both storage and bandwidth, but allows a client to learn if a file already exists on the server.

The effectiveness of deduplication is usually expressed by the *deduplication ratio*, defined as the “the number of bytes input to a data deduplication process divided by the number of bytes output” [15]. In the case of a cloud storage service, this is the ratio between total size of files uploaded by all clients and the total storage space used by the service. The *deduplication percentage* (sometimes referred to as “space reduction percentage”) [15, 18] is $1 - \frac{1}{\text{deduplication ratio}}$. A *perfect deduplication* scheme will detect all duplicates.

The level of deduplication achievable depends on a number of factors. In common business settings, deduplication ratios in the range of 4:1 (75%) to 500:1 (99.8%) are typical. Wendt et al. suggest that a figure in the range 10:1 (90%) to 20:1 (95%) is a *realistic expectation* [27].

Although deduplication benefits storage providers (and hence, indirectly, their users), it also constitutes a privacy threat for users. For example, a cloud storage server that supports client-side, cross-user deduplication can be exploited as an oracle that answers “did anyone upload this file?”. An adversary can do so by uploading a file and observing whether deduplication takes place [18]. For a *predictable file* that has low entropy, the adversary can construct all possible files, upload them and observe which file causes deduplication. Harnik et al. [18] propose a randomized threshold approach to address such online brute-force attacks. Specifically, for each file F , the server keeps a random threshold t_F ($t_F \geq 2$) and a counter c_F that indicates the number of clients that have previously uploaded F . Client-side deduplication happens only if $c_F \geq t_F$. Otherwise, the server does a server-side deduplication.

In Section 8, we survey the state-of-the-art of deduplica-

tion with encrypted data.

2.2 Hash Collisions

A *hash function* $H: \Phi \rightarrow \{0, 1\}^n$ is a deterministic function that maps a binary string in Φ of arbitrary length to a binary string h of fixed length n . The term *cryptographic hash function* is often used to denote that the function has nice cryptographic properties such as computing outputs which look random, being one-way, and making it infeasible to find collisions. We model the hash function H as a random oracle¹.

A cryptographic hash function with a long output length is collision resistant, whereas a hash function with a small output length has many collisions. As we will see in Section 4 and 5, we will use a *short hash* to improve efficiency and privacy.

2.3 Additively Homomorphic Encryption

A public key encryption scheme is *additively homomorphic* if given two ciphertexts $c_1 = Enc(pk, m_1; r_1)$ and $c_2 = Enc(pk, m_2; r_2)$ (i.e., encryptions of the messages m_1, m_2 using the same public key pk but different random values r_1, r_2), it is possible to efficiently compute $Enc(pk, m_1 + m_2; r)$ with independent randomness r , even without knowledge of the corresponding private key. Examples of such schemes are Paillier’s encryption [22], or ElGamal encryption [16] where addition is done in the exponent.

We use $E()/D()$ to denote symmetric encryption/decryption, and use $Enc()/Dec()$ to denote additively homomorphic encryption/decryption. We abuse the notation and use $Enc(pk, m)$ to denote $Enc(pk, m; r)$ where r is chosen uniformly at random. In addition, we use \oplus and \ominus to denote homomorphic addition and subtraction respectively.

2.4 Password Authenticated Key Exchange

Password-based protocols are commonly used for user authentication. However, such protocols are vulnerable to offline brute-force attacks (also referred to as *dictionary attacks*) since users tend to choose passwords with relatively low entropy that are hence guessable. Bellare and Merritt [7] were the first to propose a password authenticated key exchange (PAKE) protocol, in which an adversary making a password guess cannot verify the guess without an online attempt to authenticate itself with that password. The protocol is based on using the password as a symmetric key to encrypt the messages of a standard key exchange protocol (e.g., Diffie-Hellman [12]), so that two parties with the same password successfully generate a common session key without revealing their passwords. If the passwords are different then neither party can learn anything about the key output by the other party (namely, cannot distinguish that key from a random key).

Following this seminal work, many protocols were proposed to improve PAKE in several aspects, e.g., achieving provably security [6, 8], weakening the assumption (i.e., working in standard model without random oracles) [17, 3], achieving a stronger proof model [11, 10] and improving the round efficiency [19, 6, 20].

¹Our deduplication scheme (see Section 5) will be based on a PAKE protocol (see Section 2.4) that uses the random oracle model. Therefore we choose this model for all our analyses.

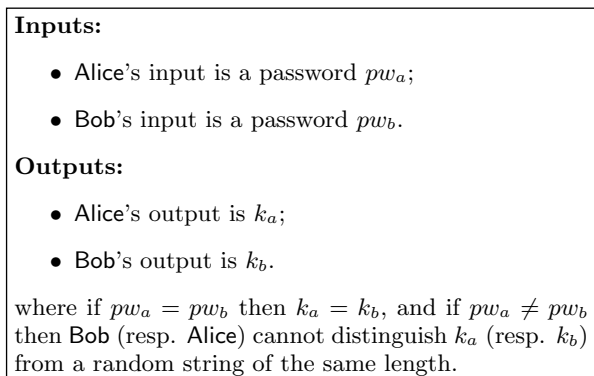


Figure 1: The ideal functionality \mathcal{F}_{pake} for password authenticated key exchange.

The ideal functionality of PAKE, \mathcal{F}_{pake} , is shown in Figure 1. We use it as a building block, and require the following properties in addition to the ideal functionality:

- *Implicit key exchange:* At the end of the protocol, neither party learns if the passwords matched or not. (In fact, many PAKE protocols were designed to be *explicit* so that parties can learn this information.)
- *Single round:* The protocol must be single-round so that it can be easily facilitated by the storage server.
- *Concurrent executions:* The protocol must allow multiple PAKE instances to run in parallel. There are two common security notions for such PAKE protocols. One stronger notion is “UC-secure PAKE” [11], which guarantees security for composition with arbitrary protocols, and with arbitrary, unknown and possibly correlated password distributions. The other notion is “concurrent PAKE”, defined by [6, 8], which is much more efficient than UC-secure PAKE. We therefore use a concurrent PAKE protocol in our work.

Our deduplication scheme (see Section 5) uses the SPAKE2 protocol of Abdalla and Pointcheval [1], which is described in Figure 2. This protocol is secure in the concurrent setting, in the random oracle model. It is very efficient, requiring each party to compute only three exponentiations, and send just a single group element to the other party. Theorem 5.1 in [1] states that this protocol is secure assuming that the computational Diffie-Hellman problem is hard in the group used by the protocol.

Same-Input-PAKE. In our deduplication protocol, one client will run multiple PAKE instances with other clients. The protocol must ensure that the client uses the *same input* in all these PAKE instances. We define this requirement in the functionality of *same-input-PAKE* described in Figure 3.

We list three possible methods for implementing the same-input-PAKE functionality: (1) The protocol can be based on the SPAKE2 protocol, by having Alice send a single first message X^* , thus using the same input in all these instances. This change is similar to the known transformation from Diffie-Hellman key exchange to ElGamal encryption. We do not prove or define security for this variant of PAKE, and leave it as a heuristic solution. (2) Alice can run independent SPAKE2 instances, with a different first message in each instance, and in addition prove in zero-knowledge that

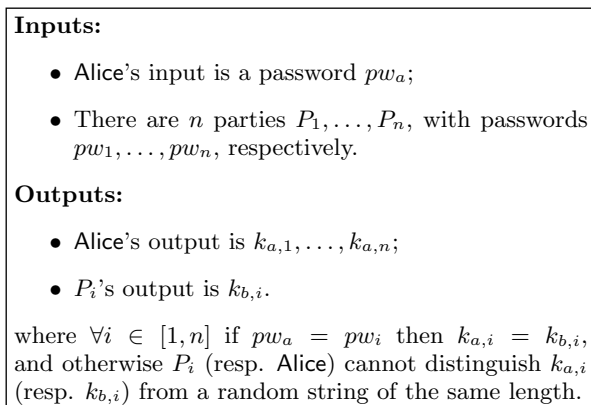


Figure 3: The ideal functionality $\mathcal{F}_{same-input-pake}$.

her inputs to all instances are identical. This proof, based on standard sigma protocols for Diffie-Hellman tuples and the Fiat-Shamir heuristic, requires only one additional exponentiation from Alice and two exponentiations from each other party. (3) The protocol can use generic protocols for non-interactive secure computation (NISC) [2]. These protocols are single round and secure against malicious adversaries. A variant called *multi-sender NISC* [2], has one party sending the first message of the protocol (committing to its input) and then has multiple other parties independently answering this message with messages encoding their input. The drawback of this approach in terms of performance is that the protocol requires an oblivious transfer for each input bit of Alice, and is therefore less efficient than protocols based on SPAKE2 or similar specific PAKE protocols.

3. PROBLEM STATEMENT

3.1 General Setting

The generic setting for cloud storage systems consists of a storage server (\mathcal{S}) and a set of clients (\mathcal{C} s) who store their files on \mathcal{S} . \mathcal{C} s never communicate directly, but exchange messages with \mathcal{S} , and \mathcal{S} processes the messages and/or forwards them as needed. Additional independent servers (\mathcal{I} Ss) can be introduced to assist deduplication [4, 26, 24]. But they are unrealistic in commercial settings² and can be bottlenecks for both security and performance. We do not require any \mathcal{I} Ss to take part in our scheme.

We assume that the parties communicate through secure channels, so that an adversary (\mathcal{A}) cannot eavesdrop and/or tamper with any channel.

We introduce new notations as needed. A summary of notations appears in Appendix A.

3.2 Ideal Model

We define the ideal functionality \mathcal{F}_{dedup} of deduplicating encrypted data in Figure 4. There are three types of participants: the storage server \mathcal{S} , the *uploader* \mathcal{C} attempting to upload a file and existing clients $\{\mathcal{C}_i\}$ who have already uploaded a file. A deduplication scheme for encrypted data is secure if no participant learns more information than is

²It is difficult to find business justification for an independent party to run an \mathcal{I} S solely for improving privacy in cloud storage services.

Public information: A finite cyclic group G of prime order p generated by an element g . Public elements $M_u \in G$ associated with user u . A hash function H modeled as a random oracle.

Secret information: User u has a password pw_u .

The protocol is run between Alice and Bob:

1. Each side performs the following computation:
 - Alice chooses $x \in_R Z_p$ and computes $X = g^x$. She defines $X^* = X \cdot (M_A)^{pw_A}$.
 - Bob chooses $y \in_R Z_p$ and computes $Y = g^y$. He defines $Y^* = Y \cdot (M_B)^{pw_B}$.
2. Alice sends X^* to Bob. Bob sends Y^* to Alice.
3. Each side computes the shared key:
 - Alice computes $K_A = (Y^*/(M_B)^{pw_B})^x$. She then computes her output as $SK_A = H(A, B, X^*, Y^*, pw_A, K_A)$.
 - Bob computes $K_B = (X^*/(M_A)^{pw_A})^y$. He then computes his output as $SK_B = H(A, B, X^*, Y^*, pw_B, K_B)$.

Figure 2: The SPAKE2 protocol of Abdalla and Pointcheval [1].

Inputs:

- The uploader \mathcal{C} has input F ;
- Each existing client \mathcal{C}_i has inputs F_i and k_{F_i} ;
- \mathcal{S} 's input is empty.

Outputs:

- \mathcal{C} gets an encryption key k_F for F . If F is identical to an existing file F_i then $k_F = k_{F_i}$. Otherwise k_F is random;
- Each \mathcal{C}_i 's output is empty;
- \mathcal{S} gets the ciphertext $E(k_F, F)$. If there is a ciphertext $E(k_{F_j}, F_j)$ that is equal to $E(k_F, F)$ in its storage, it learns j as well.

Figure 4: The ideal functionality \mathcal{F}_{dedup} of deduplicating encrypted data.

defined in the output of \mathcal{F}_{dedup} .

3.3 Design Goals

Threat Model. An adversary \mathcal{A} might compromise the uploader, the server, any subset of $\{\mathcal{C}_i\}$, or any collusion of these parties. The security of a single upload procedure can be captured by \mathcal{F}_{dedup} : requiring that the protocol implements the \mathcal{F}_{dedup} functionality according to the commonly used ideal-model/real-model security definitions. However, additional attacks are possible when considering the long-term operation of the system: a compromised *active* uploader might mount an online brute-force attack (as we described in Section 2.1); a compromised *passive* \mathcal{S} might mount an offline brute-force attack; a compromised *active* \mathcal{S} might mount an online brute-force attack by masquerading as \mathcal{C}_s , i.e., running the deduplication protocol for every “guess” and checking if deduplication occurs.

Security goals. We define the following security goals for our scheme:

- S1** Prevent online brute-force attacks by compromised active uploaders;
- S2** Prevent offline brute-force attacks by compromised pas-

sive \mathcal{S} ;

- S3** Prevent online brute-force attacks by compromised active \mathcal{S} (masquerading as multiple \mathcal{C}_s).

Functional goals. In addition, the protocol should also meet certain functional goals:

- F1** Maximize deduplication effectiveness (exceed realistic expectations, as discussed in Section 2.1);
- F2** Minimize computational and communication overhead (i.e., the computation/communication costs should be comparable to storage systems without deduplication).

4. OVERVIEW OF THE SOLUTION

Overview. We first motivate salient design decisions in our scheme before describing the details in Section 5.

When an uploader \mathcal{C} wants to upload a file F to \mathcal{S} , we need to address two problems: (a) determining if \mathcal{S} already has an encrypted copy of F in its storage and (b) if so, securely arranging to have the encryption key transferred to \mathcal{C} from some \mathcal{C}_i who uploaded the original encrypted copy of F .

In traditional client-side deduplication, when \mathcal{C} wants to upload F to \mathcal{S} , it first sends a cryptographic hash h of F to \mathcal{S} so that it can check the existence of F in \mathcal{S} 's storage. Naïvely adapting this approach to the case of encrypted storage is insecure since a compromised \mathcal{S} can easily mount an offline brute-force attack on h if F is predictable. Therefore, instead of h , we let \mathcal{C} send a short hash $sh = SH(F)$. Due to the high collision rate of $SH()$, \mathcal{S} cannot use sh to reliably guess the content of F offline.

Now, suppose that another \mathcal{C}_i previously uploaded $E(k_{F_i}, F_i)$ using k_{F_i} as the symmetric encryption key for F_i and that $sh = sh_i$. Our protocol needs to determine if this happened because $F = F_i$ and, in that case, arrange to have k_{F_i} securely transferred from \mathcal{C}_i to \mathcal{C} . We do this by having \mathcal{C} and \mathcal{C}_i engage in an oblivious key sharing protocol which allows \mathcal{C} to receive k_{F_i} iff $F = F_i$, and a random key otherwise. We say that \mathcal{C}_i plays the role of a *checker* in this protocol.

The oblivious key sharing protocol could be implemented using generic solutions for secure two-party computation, such as versions of Yao's protocol [29], which express the desired functionality as a boolean circuit. Protocols of this type have been demonstrated to be very efficient, even with

security against malicious adversaries. In our setting the circuit representation is actually quite compact, but the problem in using this approach is that the inputs of the parties are relatively long (say, 288 bits long, comprising of a full-length hash value and a key), and known protocols require an invocation of oblivious transfer, namely of public-key operations, for each input bit. There are known solutions for oblivious transfer extension, which use a preprocessing step to reduce the online computation time of oblivious transfer. However, in our setting the secure computation is run between two \mathcal{C} s that do not have any pre-existing relationship, and therefore preprocessing cannot be computed before the protocol is run.

Our solution for an efficient oblivious key sharing is having \mathcal{C} and \mathcal{C}_i run a PAKE protocol, using the hash values of their files, namely h and h_i , as their respective input “passwords”. The protocol results in \mathcal{C}_i getting k_i and \mathcal{C} getting k'_i , which are equal if $h = h_i$ and are independent otherwise. The next step of the protocol uses these keys to deliver a key k_F to \mathcal{C} , which is equal to k_{F_i} iff $k_i = k'_i$. \mathcal{C} uses this key to encrypt its file, and \mathcal{S} can deduplicate that file if the ciphertext is equal to the one uploaded by \mathcal{C}_i . Several additional issues need to be solved:

1. *How to prevent uploaders from learning about stored files?* Our protocol supports client-side deduplication, and as such informs \mathcal{C} whether deduplication takes place. In order to solve the problem, we use the randomized threshold strategy in [18] (see Section 5.1).
2. *How to prevent a compromised \mathcal{S} from mounting an online brute-force attack where it initiates many interactions with $\mathcal{C}/\mathcal{C}_i$ to identify F/F_i .* Each protocol interaction essentially enables a single guess about the content of the target file. \mathcal{C} s therefore use a per-file rate limiting strategy to prevent such attacks. Specifically, they set a bound on the maximum number of PAKE protocols they would service as a checker or an uploader for each file. (See Section 5.2.) Our simulations with realistic datasets in Section 6 show that this rate limiting does not affect the deduplication effectiveness.
3. *What if there aren't enough checkers?* If \mathcal{S} has a large number of clients, it is likely to find enough online checkers who have uploaded files with the required short hash. If there are not enough checkers, we can let the uploader run PAKE with the currently available checkers and with additional dummy checkers to hide the number of available checkers (See Section 5.3). Again, our experiments in Section 6 show that this does not affect the deduplication effectiveness (since the scheme is likely to find checkers for popular files).

Relaxing \mathcal{F}_{dedup} . The protocol we described implements the \mathcal{F}_{dedup} functionality of Figure 4 with the following relaxations: (1) \mathcal{S} learns a short hash of the uploaded file F (in our simulations we set the short hash to be 13 bits long). (2) \mathcal{F}_{dedup} is not applied between the uploader and all existing clients, but rather between the uploader and clients which have uploaded files with the same short hash as F .

We observe that in a large-scale system a short hash matches many files, and uploads of files with any specific short hash happen constantly. Therefore these relaxations leak limited information about the uploaded files. For example, since the

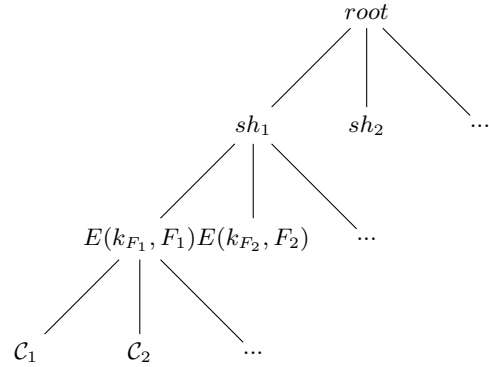


Figure 5: \mathcal{S} 's record structure.

short hash is random and short (and therefore matches many uploaded files), the rate with which a \mathcal{C} who uploaded a file is asked to participate in the protocol is rather independent of whether the same file is uploaded again.

5. DEDUPLICATION PROTOCOL

In this section we describe our deduplication protocol in detail. The data structure maintained by \mathcal{S} is shown in Figure 5. \mathcal{C} s who want to upload a file also upload the corresponding short hash sh . Since different files may have the same short hash, a short hash sh is associated with the list of different encrypted files whose plaintext maps to sh . \mathcal{S} also keeps track of clients ($\mathcal{C}_1, \mathcal{C}_2, \dots$) who have uploaded the same encrypted file.

Figure 6 shows a *basic* secure deduplication protocol. When \mathcal{C}_i stores $E(k_{F_i}, F_i)$ at \mathcal{S} , it also stores the short hash sh_i of F_i . (Step 0). When an uploader \mathcal{C} wishes to upload a file F it sends the short hash sh of this file to \mathcal{S} (Step 1). \mathcal{S} identifies the checkers $\{\mathcal{C}_i\}$ who have uploaded files with the same short hash (Step 2), and runs the following protocol with each of them.

Consider a specific \mathcal{C}_i who has uploaded F_i . \mathcal{C} runs a PAKE protocol with \mathcal{C}_i , where their inputs are the cryptographic hash values of F and F_i respectively, and their outputs are keys k'_i and k_i respectively (Step 3). The protocol must ensure that \mathcal{C} uses the same input (hash value) in the PAKE instances that it runs with all $\{\mathcal{C}_i\}$. Therefore we use a protocol implementing the same-input-PAKE functionality defined in Section 2.4. All the communication in the PAKE protocol is sent via \mathcal{S} , with no direct communication between clients.

At this point, a naive solution would be to just have \mathcal{C}_i send $E(k_i, k_{F_i})$ to \mathcal{C} . However, this would enable a subtle attack by the \mathcal{C} to identify whether F was previously uploaded.³ Therefore, the protocol continues as follows.

We assume that the keys output by the PAKE protocol are sufficiently long so that k_i can be divided into a left key and

³The problem with this approach is that \mathcal{C} learns keys k_i for multiple other clients, and should send \mathcal{S} information about each key, \mathcal{S} then tells \mathcal{C} which key index to use (and chooses a random index if no match is found). A corrupt \mathcal{C} might replace some keys with dummy values. If it is then told by \mathcal{S} to use an index of one of these keys then it knows that no match was found. The protocol must therefore send back to \mathcal{C} a key without specifying the index to which this key corresponds.

0. For each previously uploaded encrypted file $E(k_{F_i}, F_i)$, \mathcal{S} also stores the corresponding short hash sh_i .
1. Before uploading a file F , the uploader \mathcal{C} both calculates the cryptographic hash h and the short hash sh of F , and sends sh to \mathcal{S} .
2. \mathcal{S} finds the checkers $\{C_i\}$ who have uploaded files $\{F_i\}$ with the same short hash sh . Then it asks \mathcal{C} to run the same-input-PAKE protocol with $\{C_i\}$ ^a. \mathcal{C} 's input is h and C_i 's input is h_i .
3. After the invocation of the same-input-PAKE protocol, each C_i gets a session key k_i and \mathcal{C} gets a set of session keys $\{k'_i\}$ corresponding to the different C_i 's.^b
4. Each C_i splits k_i to $k_{iL}||k_{iR}$ and sends k_{iL} and $(k_{iR} + k_{F_i})$ to \mathcal{S} .
5. \mathcal{C} splits each k'_i to $k'_{iL}||k'_{iR}$, and sends \mathcal{S} its public key pk , $\{k'_{iL}\}$ and $\{Enc(pk, k'_{iR} + r)\}$ where r is chosen randomly.
6. After receiving these messages from $\{C_i\}$ and \mathcal{C} , \mathcal{S} checks if there is an index j such that $k_{jL} = k'_{jL}$.
 - (a) If so, \mathcal{S} uses the homomorphic properties of the encryption to compute $e = Enc(pk, k_{jR} + k_{F_j}) \ominus Enc(pk, k'_{jR} + r) = Enc(pk, k_{F_j} - r)$, and sends e back to \mathcal{C} ;
 - (b) Otherwise it sends $e = Enc(pk, r')$, where r' is chosen randomly by \mathcal{S} .
7. \mathcal{C} calculates $k_F = Dec(sk, e) + r$, and sends $E(k_F, F)$ to \mathcal{S} .
8. If \mathcal{S} already stores $E(k_F, F)$ it deletes $E(k_F, F)$ and allows \mathcal{C} to access the stored file $E(k_{F_j}, F_j)$. Otherwise, \mathcal{S} stores $E(k_F, F)$.

^aAll communication is run via \mathcal{S} . There is no direct interaction between \mathcal{C} and any C_i . \mathcal{C} 's input to the same-input-PAKE protocol was sent together with sh

^bNote that, with overwhelming probability, $k_i = k'_i$ iff $F_i = F$.

Figure 6: The deduplication protocol.

a right key, $k_i = k_{iL}||k_{iR}$, where k_{iL} is long enough so that the probability of two random instances of this key having the same value is small (namely, $|k_{iL}| \gg \log N$ where N is the number of clients participating in the protocol), and k_{iR} is sufficiently long to be used for symmetric encryption. The key k'_i is divided in the same way. Each C_i sends \mathcal{S} k_{iL} and $(k_{iR} + k_{F_i})$ where addition is done in a field that is also used for the additively homomorphic encryption. \mathcal{C} sends \mathcal{S} the set of pairs $\{k'_{iL}, Enc(pk, k'_{iR} + r)\}$, where r is chosen at random, the encryption is additively homomorphic and the private key is known to \mathcal{C} (Steps 4-5).

After receiving these messages from all C_i s, \mathcal{S} looks for a pair i for which $k_{iL} = k'_{iL}$. This equality happens iff $F_i = F$ (except with negligible probability). If \mathcal{S} finds such a pair it sends \mathcal{C} the value $e = Enc(pk, (k_{iR} + k_{F_i}) - (k'_{iR} + r)) = Enc(pk, k_{F_i} - r)$, computed using the homomorphic properties. Otherwise it sends $e = Enc(pk, r')$, where r' is chosen randomly by \mathcal{S} (Step 6). \mathcal{C} calculates $k_F = Dec(sk, e) + r$, and sends $E(k_F, F)$ to \mathcal{S} (Step 7). Note that if F was already uploaded to \mathcal{S} then $E(k_F, F)$ is equal to the previously stored encrypted version of the file.

THEOREM 1. *The deduplication protocol in Figure 6 implements \mathcal{F}_{dedup} with security against malicious adversaries, if the same-input-PAKE protocol is secure against malicious adversaries, the additively homomorphic encryption is semantically secure and the hash function is modeled as a random oracle.*

PROOF. We construct a simulator, and show that the execution of \mathcal{F}_{dedup} in the ideal model is computationally indistinguishable from the execution of the deduplication protocol in the real model. For the purpose of the proof we assume that the same-input-PAKE protocol is implemented

as an oracle to which the parties send their inputs and receive their outputs.

A corrupt uploader \mathcal{C} : We first assume that \mathcal{S} and the C_i s are honest and construct a simulator for the uploader \mathcal{C} . This simulator has access to \mathcal{F}_{dedup} , and needs to simulate \mathcal{C} 's view in the real execution.

The simulator observes the calls that \mathcal{C} makes to the hash function $H()$, which is modeled as a random oracle, and records the inputs and outputs as tuples $\{(F^j, h^j, sh^j)\}$. It then receives sh from \mathcal{C} , and observes the call that \mathcal{C} makes to the oracle of the same-input-PAKE protocol. If \mathcal{C} uses in that call a value h that appears in a tuple together with sh , the simulator records the file F appearing in this tuple. (Note that there is at most one such candidate, since h appears whp in only a single tuple.) Otherwise the simulator sets a random value to F . It then invokes \mathcal{F}_{dedup} with the input F and receives k_F . (Note that if F has been uploaded by any C_i , k_F is the key used to encrypt that file, otherwise k_F is random.)

The simulator records the output set $\{k'_i\}$ that \mathcal{C} receives from the same-input-PAKE. Then, when it receives a set of pairs $\{k'_{iL}, Enc(pk, x_i)\}$ from \mathcal{C} , it chooses a random index j , and checks if k'_{jL} is the same as its original value. If so, it sends $Enc(pk, k'_{jR} + k_F - x_j)$ back to \mathcal{C} ; otherwise it sends $Enc(pk, r')$ where r' is a random value.

A corrupt checker C_i : We prove security with relation to the relaxed functionality, where C_i also learns whether the uploaded file has the same short hash as F_i .

The simulator interacts with C_i in the real protocol and should extract C_i 's input to the functionality in the ideal model, namely (F_i, k_{F_i}) . We can assume that C_i has previously sent the encrypted file $E(k_{F_i}, F_i)$ to \mathcal{S} , and therefore the simulator need only find k_{F_i} and use it to compute F_i .

The simulator first observes whether sh matches the short hash of C_i . If that is not the case then the simulator provides a random k_{F_i} to the ideal functionality. Otherwise, the simulator observes C_i 's input h_i to the same-input-PAKE protocol, its output k_i , and the message (k_{iL}, x_i) that C_i sends to \mathcal{S} . If k_{iL} is different than the left part of k_i then the simulator provides to the functionality a random value k_{F_i} . Otherwise, it subtracts the right part of k_i , namely k_{iR} , from x_i , and provides the result as the input to the functionality.

A corrupt server \mathcal{S} : \mathcal{S} receives an input consisting of a set of encrypted files $E(k_{F_i}, F_i)$ associated with short hash values sh_i and C_i (we assume that the \mathcal{S} associates a single \mathcal{C} with each file). \mathcal{S} then receives from \mathcal{C} a short hash sh and first message for the same-input-PAKE protocol. Let $\{C_i\}$ be the set of clients who have uploaded files with the same short hash. C_i is marked as ‘‘OK’’ if it is selected by \mathcal{S} as a checker, i.e., \mathcal{S} has sent it a request to participate in the same-input-PAKE protocol, and forwarded its reply to the uploader (if \mathcal{S} is honest then all clients in $\{C_i\}$ are ‘‘OK’’). The simulator then pretends to be each selected C_i , sending \mathcal{S} random values for k_{iL} and $(k_{iR} + k_{F_i})$.

Next, the simulator invokes \mathcal{F}_{dedup} . It sets a bit $b = 1$ if it receives $E(k_F, F)$ and an index j from \mathcal{F}_{dedup} (in the case of a file match); otherwise, it sets $b = 0$. If $b = 1$ but client C_j is not ‘‘OK’’, the simulator changes b to 0. The simulator then pretends to be the uploader \mathcal{C} . If $b = 1$, it sets the message with respect to C_j to be $(k_{jL}, Enc(pk, x_j))$ (where k_{jL} is the same as before and x_j is random), and random values for the rest messages. Otherwise, it sets random values for all $(k'_{iL}, Enc(pk, k'_{iR} + r))$ pairs.

The simulator should now receive an encryption $Enc(y)$ from \mathcal{S} . If $b = 0$, it sends $E(k_F, F)$ (that was received from \mathcal{F}_{dedup}) to \mathcal{S} . If $b = 1$, it decrypts $Enc(y)$ and checks if $y = (k_{jR} + k_{F_j}) - x_j$. If so, it sends $E(k_F, F)$ and the index j to \mathcal{S} . Otherwise, it sends \mathcal{S} a random string of the same length with $E(k_F, F)$.

A collusion between a corrupt uploader and a corrupt server: The simulator in this case can invoke \mathcal{F}_{dedup} once, pretending to be both \mathcal{C} and \mathcal{S} , and providing \mathcal{C} 's input F (\mathcal{S} has no input to \mathcal{F}_{dedup}). It then receives the outputs of both parties, namely the key k_F , $E(k_F, F)$, and an index j (if there is a file match).

The simulation is similar to the case of a corrupted uploader, except that \mathcal{S} might choose only a subset of the checkers who have uploaded files with the same short hash. So the simulator begins as in the proof of a corrupt \mathcal{C} and extracts \mathcal{C} 's input F from the random oracle. Then it invokes \mathcal{F}_{dedup} with F . If a match was found, the simulator observes the operation of \mathcal{S} and checks if C_j is ‘‘OK’’ (as was defined in the proof of a corrupt \mathcal{S}). If so, it uses (F, k_F) as C_j 's input to the protocol, and random values as other checkers' inputs. Otherwise, it uses random values for all checkers' inputs.

A collusion between corrupt checkers and a corrupt server: The input of the adversary to the ideal functionality \mathcal{F}_{dedup} is the pairs (F_i, k_{F_i}) of the corrupt C_i s. Its output is the same as in the case of a corrupt \mathcal{S} .

The simulation is similar to the case of a corrupted \mathcal{S} . The main difference is that the corrupt C_i s send inputs to the same-input-PAKE functionality and receive outputs $k_i = k_{iL} || k_{iR}$. It can extract these values from the oracle of same-input-PAKE functionality, which outputs k_i to each

C_i . Then the simulator invokes \mathcal{F}_{dedup} as \mathcal{S} , and then run same procedure as it is in the case of a corrupt \mathcal{S} . \square

Based on Theorem 1, we conclude that compromised parties cannot run the computation in a way that is different than is defined by (relaxed) \mathcal{F}_{dedup} . This satisfies requirement **S2**. We now discuss several extensions to the basic protocol to account for the types of issues we alluded to in Section 4.

5.1 Randomized Threshold

The protocol in Figure 6 is for server-side deduplication. To save bandwidth, we transform it to support client-side deduplication. In order to satisfy requirement **S1** and protect against a corrupt uploader, we use the randomized threshold approach from Harnik et al. [18]: for each file F , \mathcal{S} maintains a random threshold t_F ($t_F \geq 2$), and a counter c_F that indicates the number of \mathcal{C} s that have previously uploaded F .

In step 6 of the deduplication protocol,

- In the case of a match (6a), if $c_{F_i} < t_{F_i}$, \mathcal{S} tells \mathcal{C} to upload $E(k_F, F)$ as if no match occurred (but \mathcal{S} does not store this copy). Otherwise, \mathcal{S} informs \mathcal{C} that the file is duplicated and there is no need to upload it;
- In the case of a no match (6b), \mathcal{S} asks \mathcal{C} to upload $E(k_F, F)$.

5.2 Rate Limiting

A compromised active \mathcal{S} can apply online brute-force attacks against \mathcal{C} or C_i . Specifically, if F_i is predictable, \mathcal{S} can pretend to be an uploader and send PAKE requests to C_i attempting to guess F_i . \mathcal{S} can also pretend to be a checker and send PAKE responses to \mathcal{C} to guess F . Therefore both uploaders and checkers should limit the number of PAKE runs for each file in their respective roles. This per-file rate limiting strategy can both improve security (see below) and reduce overhead (namely the number of PAKE runs) without damaging the deduplication effectiveness (as shown in Section 6).

We use RL_c to denote the rate limit for checkers, i.e., each C_i can process at most RL_c PAKE requests for F_i and will ignore further requests. Similarly, RL_u is the rate limit for uploaders, i.e., a \mathcal{C} will send at most RL_u PAKE requests to upload F . Suppose that n is the length of the short hash and m is the min-entropy of a predictable file F , then \mathcal{A} can uniquely identify F only if

$$2^m \leq 2^n \cdot x \cdot (RL_u + RL_c) \quad (1)$$

where x is the number of clients who potentially own F .

As a comparison, DupLESS [4] uses a per-client rate limiting strategy to prevent such online brute-force attacks from any single \mathcal{C} . The rate limit must not degrade usability for a \mathcal{C} that needs to legitimately upload a large number of files within a brief time interval (such as backing up a local file system). Therefore the authors of DupLESS chose a large bound (825 000) for the total number of requests a single \mathcal{C} can make during one week. That is to say, for a predictable file with entropy $2^m > 825\,000$, an attack by a single compromised \mathcal{C} is expected to take more than one week. But if \mathcal{A} has compromised multiple \mathcal{C} s, the attack time in DupLESS decreases linearly. Recall also, that a compromised active \mathcal{S} can masquerade as any number of \mathcal{C} s that is needed.

To prevent a compromised active \mathcal{S} from masquerading multiple C s, the authors of [26] introduce another independent server called *identity server*. When C s first join the system, the identity server is responsible for verifying their identity and issuing credentials to them. However, it is hard to deploy an *independent* identity server in a real world setting. As far as we know, our scheme is the first deduplication protocol that can prevent online brute-force attacks (i.e., satisfying requirement **S3**) without the aid of an identity server.

5.3 Checker Selection

If an uploader is required to run only a few (or none) PAKE instances, due to no short hash matches, it will learn that it is less likely that its file is already in \mathcal{S} 's storage. To avoid this leakage of information, \mathcal{S} fixes the number of PAKE runs (i.e., RL_u) for an upload request to be constant. For a requested short hash sh , checkers are selected according to the following procedure:

1. \mathcal{S} selects the most popular file among the files whose short hash is sh and which were not already selected with respect for the current upload (popularity is measured in terms of the number of C s who own the file).
2. \mathcal{S} selects a checker for that file in ascending order of engagement among the C s that are currently online (in terms of the number of PAKE requests they have serviced so far for that specific file).
3. If the number of selected files is less than RL_u , repeat Step 1-3.
4. If the total number of selected files for which there are online clients is smaller than RL_u , \mathcal{S} uses additional dummy files and clients, until reaching RL_u files.

Then, \mathcal{S} lets the uploader run PAKE instances with the selected RL_u clients (\mathcal{S} itself runs as dummy clients).

6. SIMULATION

Our use of rate limiting can impact deduplication effectiveness. In this section, we use realistic simulations to study the effect of various parameter choices in our protocol on deduplication effectiveness.

Datasets. We want to consider two types of storage environments. The first consists predominantly of media files, such as audio and video files from many users. We did not have access to such a dataset. Instead, we use a dataset comprising of Android application prevalence data to represent an environment with media files. This is based on the assumption that the popularity of Android applications is likely to be similar to that of media files: both are created and published by a small number of authors (artists/developers), made available on online stores or other distribution sites, and are acquired by consumers either for free or for a fee. We call this the *media dataset*. We use a publicly available dataset⁴. It consists of data collected from 77 782 Android devices. For each device, the dataset identifies the set of (anonymized) application identifiers found on that device. We treat each application identifier as a “file” and consider the presence of an app on a device as an “upload request” to add the corresponding “file” to the storage. This dataset

⁴<https://se-sy.org/projects/malware/>

has 7 396 235 “upload requests” in total, of which 178 396 are for distinct files.

The second is the type of storage environments that are found in enterprise backup systems. We use data gathered by the Debian Popularity Contest⁵ to approximate such an environment. The *popularity-contest* package on a Debian device regularly reports the list of packages installed on that device. The resulting data consists of a list of debian packages along with the number of devices which reported that package. We took a snapshot of this data on Nov 27, 2014. It consists of data collected from 175 903 Debian users. From this data we generated our *enterprise dataset*: it has 217 927 332 “upload requests” (debian package installations) of which 143 949 are for distinct files (unique packages).

Figure 7 shows the file popularity distribution (i.e., the number of upload requests for each file) in logarithmic scale for both datasets. We map each dataset to a stream of upload requests by generating the requests in random order, where a file that has x copies generates x upload requests at random time intervals.

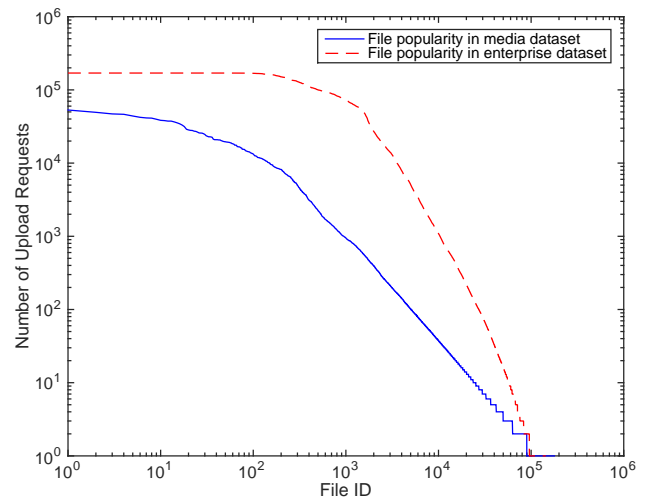


Figure 7: File popularity in both datasets.

Parameters. To facilitate comparison with DupLESS [4], we set the min-entropy to $\log(825000)$. We then set the length of the short hash $n = 13$, and $(RL_u + RL_c) = 100$ (i.e., a C will run PAKE at most 100 times for a certain file as both uploader and checker), so that we achieve the bound in inequality 1 in Section 5.2: \mathcal{A} cannot uniquely identify a file within the rate limit. We use these parameters in our simulations.

We measure overhead as the average number of PAKE runs⁶, which can be calculated as:

$$\mu = \frac{\text{Total number of PAKE runs}}{\text{Total number of upload requests}} \quad (2)$$

We measure deduplication effectiveness using the deduplication percentage (Section 2.1). We assume that all files are

⁵<http://popcon.debian.org>

⁶We do not include fake PAKE runs by \mathcal{S} (Section 5.3) since we are interested in estimating the average number of real PAKE runs.

of equal size so that the deduplication percentage ρ is:

$$\rho = \left(1 - \frac{\text{Number of all files in storage}}{\text{Total number of upload requests}}\right) \cdot 100\% \quad (3)$$

Rate limiting. We first assume that all C s are online during the simulation, and study the impact of rate limits. Having selected $RL_u + RL_c$ to be 100, we now see how selecting specific values for RL_u and RL_c affects the average number of PAKE runs and the deduplication effectiveness. Figure 8 shows the average number of PAKE runs resulting from different values of RL_u (and hence RL_c) in both datasets. Both values are very low, in the range 1.3-1.75. We also ran the simulation without any rate limits, which led to an average of 26.88 PAKE runs in the media dataset and 13.19 PAKE runs in the enterprise dataset. These numbers are significantly larger than the results with rate limiting. Figure 9 shows ρ resulting from different rate limit choices in both datasets. We see that setting $RL_u = 30$ (and hence setting $RL_c = 70$), maximizes ρ to be (97.58% and 99.9332%, respectively). These values are extremely close to the perfect deduplication percentages in both datasets (97.59% and 99.9339% respectively). A major conclusion is that rate lim-

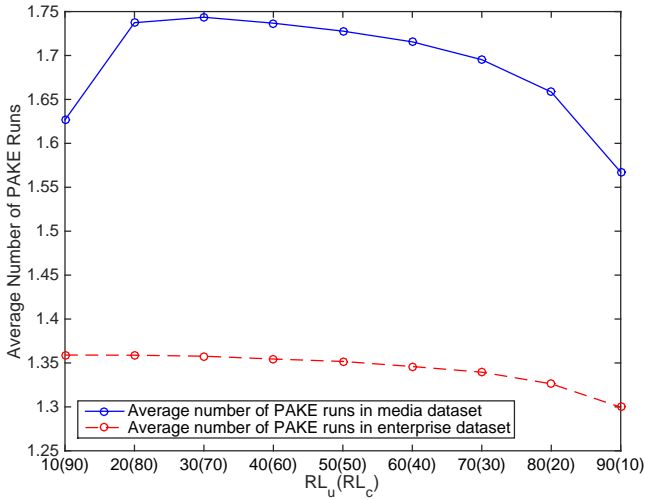


Figure 8: Average number of PAKE runs VS. rate limits.

iting can improve security and reduce overhead without negatively impacting deduplication effectiveness.

Offline rate. The possibility of some C s being offline may adversely impact deduplication effectiveness. To estimate this impact, we assign an *offline rate* to each C as its probability to be offline during one upload request. Using the chosen rate limits ($RL_u = 30$ and $RL_c = 70$), we measured ρ by varying the offline rate. The results for both datasets are shown in Figure 10. It shows that ρ is still reasonably high if the offline rate is lower than 70%. But drops quickly beyond that. We can solve this by introducing *deferred check*. Specifically, we split RL_u to $RL_{u1} + RL_{u2}$. \mathcal{S} will let the uploader run RL_{u1} times PAKE before uploading, and later ask it to run further RL_{u2} PAKE instances when some C s who are previously offline, come online. If \mathcal{S} finds a match after uploading, it checks the counter and random threshold for the matched file. If the counter has exceeded the threshold, \mathcal{S} deletes the previously uploaded file and asks the uploader to change the encryption key to

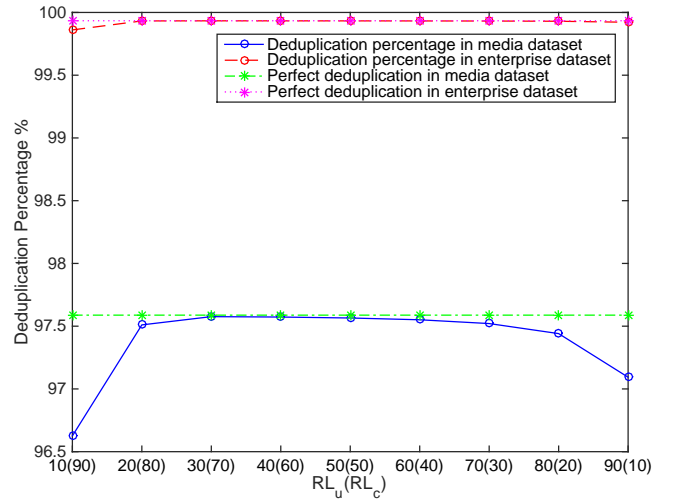


Figure 9: Dedup. percentage VS. rate limits.

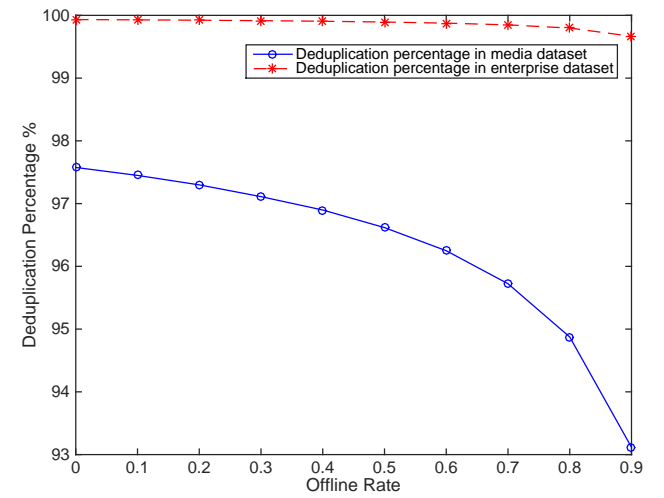


Figure 10: Dedup. percentage VS. offline rates.

match the detected duplicate. The only issue for this solution is that the uploader needs to keep the randomness of all PAKE runs of offline check. Otherwise, \mathcal{S} keeps the messages for that PAKE instance until the threshold being crossed. Figure 11 shows that this method can significantly improve the deduplication effectiveness when offline rate is high.

Evolution of deduplication effectiveness. Figure 12 shows that the ρ achieved by our scheme increases as more files are added to the storage, and it meets the realistic expectation (95%) quickly: after receiving 304 160 (4%) upload requests in the media dataset, and 121 110 (0.05%) upload requests in the enterprise dataset. Given that the deduplication effectiveness of our scheme is close to that of perfect deduplication and exceeds typical expected values, we can conclude that it satisfies functionality goal **F1**. Using rate limits implies that ρ increases more slowly in our scheme than in perfect deduplication. Figure 13 shows that this difference stabilizes as the number of upload requests increases. Figure 16 (in Appendix B) shows the second order difference in deduplication effectiveness compared to perfect dedupli-

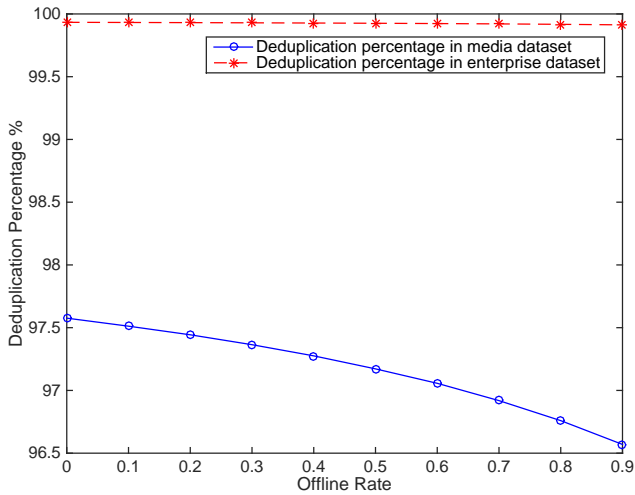


Figure 11: Dedup. percentage VS. offline rates.

cation. The second order difference vanishes as more files are uploaded to the storage.

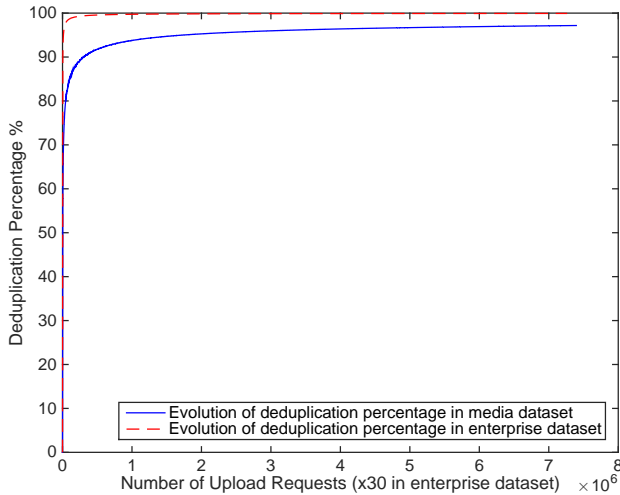


Figure 12: Dedup. percentage VS. number of upload requests.

Explanation. The fact that our scheme achieves close to perfect deduplication even in the presence of rate limits may appear counter-intuitive at first glance. But this phenomenon can be explained by *Zipf's law* [30]. As seen from Figure 7, beyond the initial plateau, the file popularity distribution is a straight line and thus follows a *power law distribution* (also known as *Zipf distribution*). The initial plateau does not impact our system. This is evident when we account for the use of short hash function. Figure 14 shows the file popularity in both datasets for some specific, but randomly selected, short hash values (of length 10).

Even though we use rate limits, \mathcal{S} always selects files based on descending order of popularity (step 1 in Section 5.3). Since file popularity follows the Zipf distribution, selecting files based on popularity ensures that popular uploaded files have a much higher likelihood of being selected and thus deduplicated. There are other examples of using the Zipf

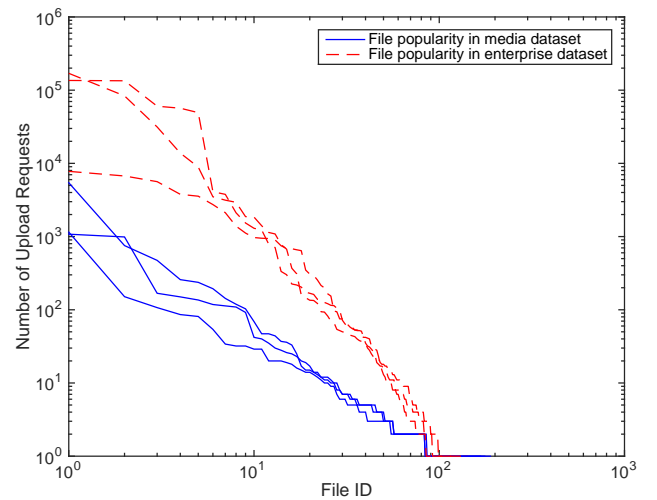


Figure 14: File popularity for six short hashes.

distribution to design surprisingly efficient systems. *Web proxy caching proxies* are such an example [9]. Breslau et al. observe that the distribution of page requests follows Zipf's law. Consequently, proxies use their limited storage to only cache popular files but still achieve significant bandwidth savings. The frequency of a request for the m^{th} most popular page can be calculated as $\frac{1/m^\alpha}{\sum_{i=1}^N (1/i^\alpha)}$, where N is the size of the cache, and α is the value of the exponent characterising the distribution [9]. As a result, most of the requested pages can be found in the cache. Similarly, in our case, most of the upload requests for files that have already been uploaded can find a matched file within the rate limit.

7. PERFORMANCE EVALUATION

Our deduplication scheme incurs some extra computation and communication due to the number of PAKE runs. In this section, we demonstrate that the overhead is negligible for large files by implementing a proof-of-concept prototype. **Prototype.** Our prototype consists of two parts: (1) a server program which simulates \mathcal{S} and (2) a client program which simulates \mathcal{C} (performing file uploading/downloading, encryption/decryption, and assisting \mathcal{S} in deduplication). We used *Node.js*⁷ for the implementation of both parties, and *Redis*⁸ for the implementation of \mathcal{S} 's data structure. We used SHA-256 as the cryptographic hash function and AES with 256-bit keys as the symmetric encryption scheme, both of which are provided by the *Crypto* module in *Node.js*. We truncated the cryptographic hash to get the short hash. We used the *GNU multiple precision arithmetic library*⁹ to implement the public key operations.

Test setting and methodology. We ran the server-side program on a remote server (Intel Xeon with 4 2.66 GHz cores) and the client-side program on an Intel Core i7 machine with 4 2.2 GHz cores. We measured the running time using the *Date* module in Javascript and measured the bandwidth usage using *TCPdump*.

As the downloading phase in our protocol is simply down-

⁷<http://nodejs.org>

⁸<http://redis.io>

⁹<https://gmplib.org>

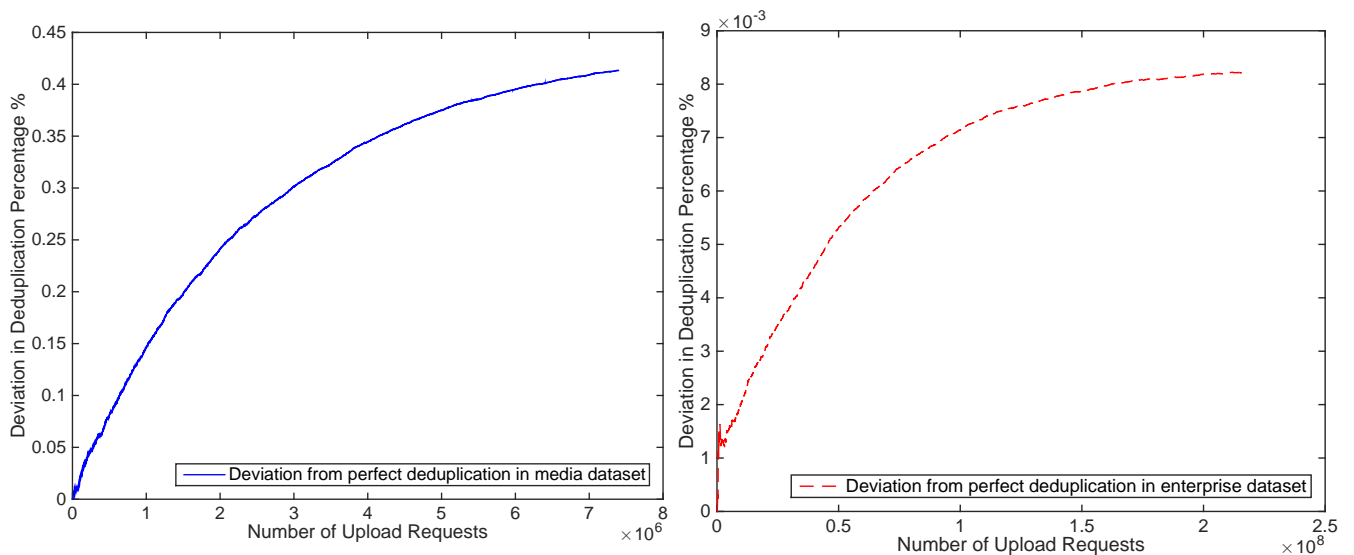


Figure 13: Deviation from perfect deduplication VS. Number of upload requests.

loading an encrypted file, we only consider the uploading phase. We set the length of short hash to be 13, and set $RL_u = 30$. We considered the case where the uploader C runs PAKE with 30 checkers C_i . So we simulate the uploading phase in our protocol as:

1. C sends the short hash of the file it wants to upload to the server S ;
2. S forwards requests to 30 checkers C_i and lets them run PAKE with C ;
3. S waits for responses in all instances back from C and $\{C_i\}$;
4. S chooses one instance and sends the result to C ;
5. C uses the resulting key to encrypt its file with AES and uploads it to S .

We measured both running time and bandwidth usage during the whole procedure above. We compare the results to two baselines: (1) uploading without encryption and (2) uploading with AES encryption. As in [4], we repeat our experiment using files of size 2^{2i} KB for $i \in \{0, 1, \dots, 8\}$, which provides a file size range of 1KB to 64 MB. For each file, we upload it 100 times and calculate the mean. For files that are larger than the computer buffer, we do loading, encryption and uploading at the same time by pipelining the data stream. As a result, uploading encrypted files uses almost the same amount of time as uploading plain files.

Results. Figure 15 reports the uploading time and bandwidth usage in our protocol compared to the two baselines. For files that are smaller than 1 MB, the overhead introduced by our deduplication protocol is relatively high. For example, it takes 15 ms (2 508 bytes) to upload a 1 KB encrypted file, while it takes 319 ms (145 359 bytes) to upload the same file in our protocol. However, the overhead introduced by our protocol is independent of the file size (about 10 ms for each PAKE run), and becomes negligible when the file is large enough. For files that are larger than 64 MB file, the time overhead is below 2%, and the bandwidth overhead is below 0.16%. So our scheme meets **F2**.

8. RELATED WORK

There are several types of schemes that enable deduplication with client-side encrypted data. The simplest approach (which is used by most commercial products) is to encrypt C 's files using a global key which is encoded in the client-side software. As a result, different copies of F result in the same ciphertext and can therefore be deduplicated. This approach is, of course, insecure if S is untrusted.

Another approach is *convergent encryption* [13], which uses $H(F)$ as a key to encrypt F , where $H()$ is a publicly known cryptographic hash function. This approach ensures that different copies of F result in the same ciphertext. However, a compromised passive S can perform an offline brute-force attack if F has a small (or medium) entropy. Bellare et al. proposed *message-locked encryption* (MLE), which uses a semantically secure encryption scheme but produces a deterministic tag [5]. So it still suffers from the same attack.

Other solutions are based on the aid of additional independent servers (\mathcal{IS} s). For example, *Cloudedup* is a deduplication system that introduces an \mathcal{IS} that is responsible for encryption and decryption [24]. Specifically, C first encrypts each block with convergent encryption and sends the ciphertexts to \mathcal{IS} , who then encrypts them again with a key only known by itself. During file retrieval, blocks are first decrypted by \mathcal{IS} and sent back to C . In this scheme, a compromised active S can easily perform an online brute-force attack by uploading guessing files and see if deduplication happens.

Stanek et al. propose a scheme that only deduplicates popular files [26]. C s encrypt their files with two layers of encryption: the inner layer is obtained through convergent encryption, and the outer layer is obtained through a semantically secure threshold encryption scheme with the aid of an \mathcal{IS} . S can decrypt the outer layer of F iff the number of C s who have uploaded F reaches the threshold, and thus perform a deduplication. In addition, they introduce another \mathcal{IS} as an identity server to prevent online brute-force attacks by multiple compromised C s.

Both [24] and [26] are vulnerable to offline brute-force attacks by compromised \mathcal{IS} s. To prevent this, Bellare et al.

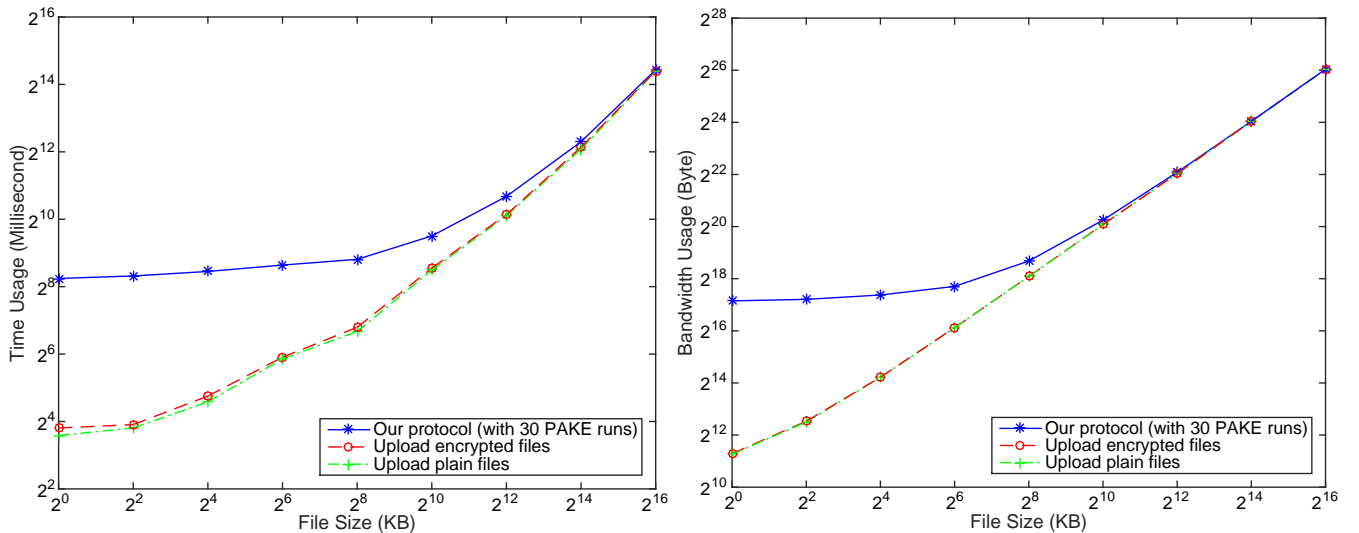


Figure 15: Time (left) and bandwidth usage (right) VS. file size.

propose DupLESS that enables C s to generate file keys by running an *oblivious pseudorandom function* (OPRF) with \mathcal{IS} . Specifically, in the key generation process of convergent encryption, they introduce another secret which is provided by \mathcal{IS} and identical for all C s. The OPRF enables C s to generate their keys without revealing their files to \mathcal{IS} , and without learning anything about \mathcal{IS} 's secret. To prevent the online brute-force attacks from compromised active \mathcal{S} . DupLESS uses a per-client rate limiting strategy to limit the number of requests that a C can send to \mathcal{IS} during each epoch. We have identified the limitations for this strategy in Section 5.2. In addition, if \mathcal{A} compromises both \mathcal{S} and \mathcal{IS} , it can get the secret from \mathcal{IS} , and the scheme is reduced to normal convergent encryption.

Duan proposes a scheme that uses the same idea as DupLESS, but distributes the task of \mathcal{IS} [14], where a C must interact with a threshold of other C s to generate the key. So this scheme is only suitable for peer-to-peer paradigm: a threshold number of C s must be online and interact with one another. While improving availability and security compared to DupLESS, this scheme is still susceptible to online brute-force attacks by compromised active \mathcal{S} , and it is unclear how to apply any rate-limiting strategy to it.

In Table 1, we summarize the resilience of these schemes with respect to the design goals from Section 3.3.

Schemes \ Threat	Compromised				
	C	\mathcal{S} (pas.)	\mathcal{S} (act.)	\mathcal{IS} s	$\mathcal{S}, \mathcal{IS}$ s
[13], [5]	✓	X	X	–	–
[24]	✓	✓	X	X	X
[26]	✓	✓	✓	X	X
[4]	✓	✓	X	✓	X
[14]	✓	✓	X	–	–
Our work	✓	✓	✓	–	–

Table 1: Resilience of deduplication schemes.

9. DISCUSSION

Incentives. In our scheme C s have to run several PAKE in-

stances as both uploaders and checkers. This imposes a cost on each C . \mathcal{S} is the direct beneficiary of deduplication. C s may indirectly benefit in that effective deduplication makes the storage system more efficient and can thus potentially lower the cost incurred by each C . Nevertheless, it is desirable to have more direct incentive mechanisms to encourage C s to do PAKE checks. For example, if a file uploaded by C is found to be shared by other C s, \mathcal{S} could reward the C s owning that file by giving them small increases in their respective storage quotas.

User involvement. A simple solution to the problem of deduplication vs. privacy is to have the user identify sensitive files. The storage client can then use convergent encryption for non-sensitive files and semantically secure encryption for sensitive files. This approach has three drawbacks: it is too burdensome for average users, it reveals which files are sensitive, and it foregoes deduplication of sensitive files altogether.

Deduplication effectiveness. We can improve deduplication effectiveness by introducing additional checks. For example, an uploader can indicate the (approximate) file size (which will be revealed anyway) so that \mathcal{S} can limit the selection of checkers to those whose files are of a similar size. Similarly, \mathcal{S} can keep track of similarities between C s based on the number of files they share and use this information while selecting checkers by prioritizing checkers who are similar to the uploader. Nevertheless, as discussed in Figure 12, our scheme exceeds what is considered as realistic levels of deduplication early in the life of the storage system. Whitehouse [28] reported that when selecting a deduplication scheme, enterprise administrators rated considerations such as ease of deployment and of use being more important than deduplication ratio. Therefore, we argue that the very small sacrifice in deduplication ratio is offset by the significant advantage of ensuring user privacy without having to use independent third party servers.

Block-level deduplication. Our scheme can be applied for both file-level and block-level deduplication. Applying it for block-level deduplication will incur more overhead.

Datasets. Deduplication effectiveness is highly dependent

on the dataset. Analysis using more realistic datasets can shed more light on the efficacy of our scheme.

Realistic modeling of offline status. In our analysis of how deduplication effectiveness is affected by the offline rate (Figure 10), we assumed a simple model where the offline status of clients is distributed uniformly at a specified rate. In practice the offline status is influenced by many factors like geography and time of day.

Acknowledgments

This work was supported in part by the “Cloud Security Services” project funded by the Academy of Finland (283135), the EU 7th Framework Program (FP7/2007-2013) under grant agreement n. 609611 (PRACTICE), and by a grant from the Israel Ministry of Science and Technology. We thank Ivan Martinovic for suggesting the analogy between our system and web-caching proxies. We thank Billy Brumley and Kaitai Liang for the kind review.

10. REFERENCES

- [1] M. Abdalla and D. Pointcheval. Simple password-based encrypted key exchange protocols. In A. Menezes, editor, *Topics in Cryptology - CT-RSA 2005, The Cryptographers' Track at the RSA Conference 2005, San Francisco, CA, USA, February 14-18, 2005, Proceedings*, volume 3376 of *Lecture Notes in Computer Science*, pages 191–208. Springer, 2005.
- [2] A. Afshar, P. Mohassel, B. Pinkas, and B. Riva. Non-interactive secure computation based on cut-and-choose. In *Advances in Cryptology - EUROCRYPT 2014*, pages 387–404. Springer Berlin Heidelberg, 2014.
- [3] B. Barak, R. Canetti, Y. Lindell, R. Pass, and T. Rabin. Secure computation without authentication. In V. Shoup, editor, *Advances in Cryptology - CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 361–377. Springer Berlin Heidelberg, 2005.
- [4] M. Bellare, S. Keelveedhi, and T. Ristenpart. DupLESS: Server-aided encryption for deduplicated storage. In *Proceedings of the 22Nd USENIX Conference on Security, SEC'13*, pages 179–194, Berkeley, CA, USA, 2013. USENIX Association.
- [5] M. Bellare, S. Keelveedhi, and T. Ristenpart. Message-locked encryption and secure deduplication. In T. Johansson and P. Nguyen, editors, *Advances in Cryptology - EUROCRYPT 2013*, volume 7881 of *Lecture Notes in Computer Science*, pages 296–312. Springer Berlin Heidelberg, 2013.
- [6] M. Bellare, D. Pointcheval, and P. Rogaway. Authenticated key exchange secure against dictionary attacks. In Preneel [23], pages 139–155.
- [7] S. Bellare and M. Merritt. Encrypted key exchange: password-based protocols secure against dictionary attacks. In *Research in Security and Privacy, 1992. Proceedings., 1992 IEEE Computer Society Symposium on*, pages 72–84, May 1992.
- [8] V. Boyko, P. D. MacKenzie, and S. Patel. Provably secure password-authenticated key exchange using diffie-hellman. In Preneel [23], pages 156–171.
- [9] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker. Web caching and zipf-like distributions: evidence and implications. In *INFOCOM '99. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 1, pages 126–134 vol.1, Mar 1999.
- [10] R. Canetti, D. Dachman-Soled, V. Vaikuntanathan, and H. Wee. Efficient password authenticated key exchange via oblivious transfer. In M. Fischlin, J. Buchmann, and M. Manulis, editors, *Public Key Cryptography - PKC 2012*, volume 7293 of *Lecture Notes in Computer Science*, pages 449–466. Springer Berlin Heidelberg, 2012.
- [11] R. Canetti, S. Halevi, J. Katz, Y. Lindell, and P. D. MacKenzie. Universally composable password-based key exchange. In R. Cramer, editor, *Advances in Cryptology - EUROCRYPT 2005, 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Aarhus, Denmark, May 22-26, 2005, Proceedings*, volume 3494 of *Lecture Notes in Computer Science*, pages 404–421. Springer, 2005.
- [12] W. Diffie and M. E. Hellman. New directions in cryptography. *Information Theory, IEEE Transactions on*, 22(6):644–654, 1976.
- [13] J. R. Douceur, A. Adya, W. J. Bolosky, P. Simon, and M. Theimer. Reclaiming space from duplicate files in a serverless distributed file system. In *Distributed Computing Systems, 2002. Proceedings. 22nd International Conference on*, pages 617–624. IEEE, 2002.
- [14] Y. Duan. Distributed key generation for encrypted deduplication: Achieving the strongest privacy. In *Proceedings of the 6th Edition of the ACM Workshop on Cloud Computing Security, CCSW '14*, pages 57–68, New York, NY, USA, 2014. ACM.
- [15] M. Dutch. Understanding data deduplication ratios. SNIA Data Management Forum, 2008. <http://storage.ctocio.com.cn/imagelist/2009/222/13pm284d8r1s.pdf>.
- [16] T. ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In G. Blakley and D. Chaum, editors, *Advances in Cryptology*, volume 196 of *Lecture Notes in Computer Science*, pages 10–18. Springer Berlin Heidelberg, 1985.
- [17] O. Goldreich and Y. Lindell. Session-key generation using human passwords only. In J. Kilian, editor, *Advances in Cryptology - CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 408–432. Springer Berlin Heidelberg, 2001.
- [18] D. Harnik, B. Pinkas, and A. Shulman-Peleg. Side channels in cloud services: Deduplication in cloud storage. *Security Privacy, IEEE*, 8(6):40–47, Nov 2010.
- [19] I. Jeong, J. Katz, and D. Lee. One-round protocols for two-party authenticated key exchange. In M. Jakobsson, M. Yung, and J. Zhou, editors, *Applied Cryptography and Network Security*, volume 3089 of *Lecture Notes in Computer Science*, pages 220–232. Springer Berlin Heidelberg, 2004.
- [20] J. Katz and V. Vaikuntanathan. Round-optimal password-based authenticated key exchange. *Journal of Cryptology*, 26(4):714–743, 2013.
- [21] D. T. Meyer and W. J. Bolosky. A study of practical deduplication. In *Proceedings of the 9th USENIX Conference on File and Storage Technologies, FAST'11*, pages 1–1, Berkeley, CA, USA, 2011. USENIX Association.
- [22] P. Paillier. Public-key cryptosystems based on composite degree residuosity classes. In J. Stern, editor, *Advances in Cryptology - EUROCRYPT '99*, volume 1592 of *Lecture Notes in Computer Science*, pages 223–238. Springer Berlin Heidelberg, 1999.
- [23] B. Preneel, editor. *Advances in Cryptology - EUROCRYPT 2000, International Conference on the Theory and Application of Cryptographic Techniques, Bruges, Belgium, May 14-18, 2000, Proceeding*, volume 1807 of *Lecture Notes in Computer Science*. Springer, 2000.
- [24] P. Puzio, R. Molva, M. Önen, and S. Loureiro.

Cloudedup: Secure deduplication with encrypted data for cloud storage. In *Proceedings of the 2013 IEEE International Conference on Cloud Computing Technology and Science - Volume 01*, CLOUDCOM '13, pages 363–370, Washington, DC, USA, 2013. IEEE Computer Society.

- [25] S. Quinlan and S. Dorward. Venti: A new approach to archival storage. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies*, FAST'02, pages 7–7, Berkeley, CA, USA, 2002. USENIX Association.
- [26] J. Stanek, A. Sorniotti, E. Androulaki, and L. Kencl. A secure data deduplication scheme for cloud storage. In N. Christin and R. Safavi-Naini, editors, *Financial Cryptography and Data Security*, Lecture Notes in Computer Science, pages 99–118. Springer Berlin Heidelberg, 2014.
- [27] J. M. Wendt. Getting Real About Deduplication Ratios. <http://www.d cig.com/2011/02/getting-real-about-deduplication.html>, 2011.
- [28] L. Whitehouse. Understanding data deduplication ratios in backup systems. TechTarget article, May 2009. <http://searchdatabackup.techtarget.com/tip/Understanding-data-deduplication-ratios-in-backup-systems>.
- [29] A. C.-C. Yao. How to generate and exchange secrets. In *Foundations of Computer Science, 1986., 27th Annual Symposium on*, pages 162–167, Oct 1986.
- [30] G. K. Zipf. Relative frequency as a determinant of phonetic change. *Harvard studies in classical philology*, pages 1–95, 1929.

APPENDIX

A. NOTATION TABLE

A table of notations is shown in Table 2.

B. SECOND ORDER DEVIATION FROM PERFECT DEDUPLICATION

Figure 16 how the second order difference in deduplication percentage (between perfect deduplication and our scheme) evolves as more files are uploaded to the storage server. This second order difference essentially vanishes as more files are added to the server.

Notation	Description
<i>Entities</i>	
\mathcal{C}	Client
\mathcal{S}	Server
\mathcal{A}	Adversary
\mathcal{IS}	Independent Server
<i>Cryptographic Notations</i>	
$E()$	Symmetric key encryption
$D()$	Symmetric key decryption
k	Symmetric encryption/decryption key
$Enc()$	Additively homomorphic encryption
$Dec()$	Additively homomorphic decryption
\oplus	Additively homomorphic addition
\ominus	Additively homomorphic subtraction
$H()$	Cryptographic hash function
h	Cryptographic hash
$SH()$	Short hash function
sh	Short hash
PAKE	Password Authenticated Key Exchange
\mathcal{F}_{pake}	Ideal functionality of PAKE
\mathcal{F}_{dedup}	Ideal functionality of deduplication protocol
<i>Parameters</i>	
F	File
m	Entropy of a predictable file
n	Length of the short hash
RL_u	Rate limit by uploaders
RL_c	Rate limit by checkers
t_F	Random threshold for a file
c_F	Counter for a file
ρ	Deduplication Percentage

Table 2: Summary of notations

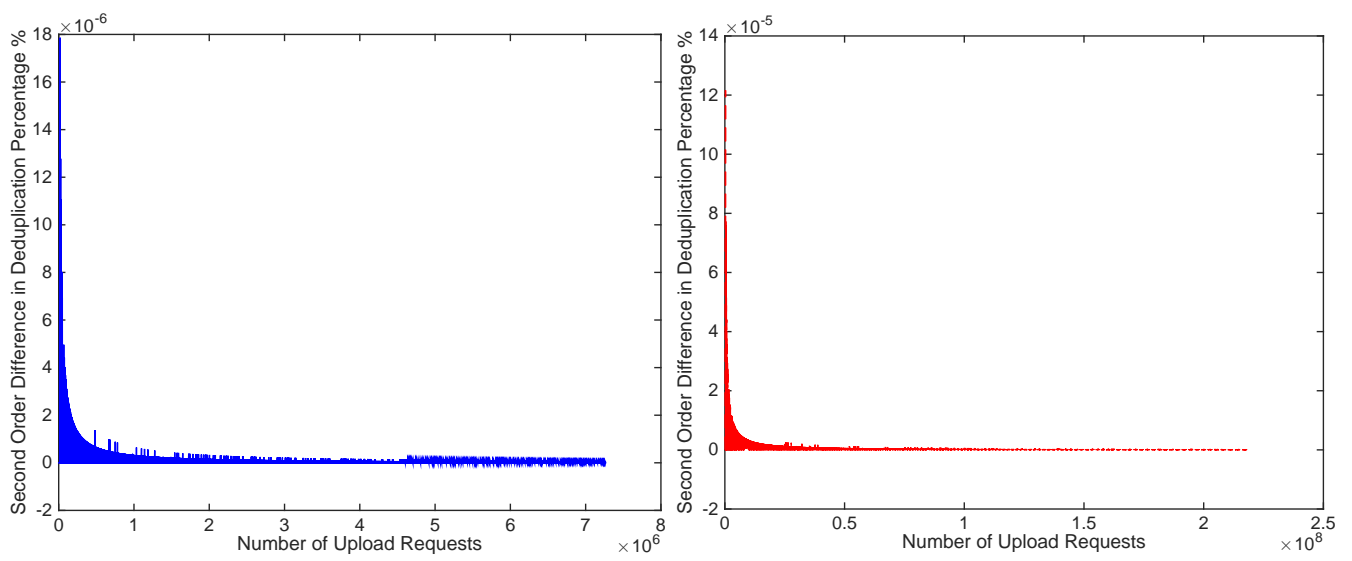


Figure 16: Second order difference in deduplication percentage: perfect deduplication VS. Number of upload requests. Left: Media dataset; Right: Enterprise dataset