

Provable Data Possession at Untrusted Stores

GIUSEPPE ATENIESE* RANDAL BURNS* REZA CURTMOLA*
JOSEPH HERRING* LEA KISSNER † ZACHARY PETERSON* DAWN SONG ‡

Abstract

We introduce a model for *provable data possession* (PDP) that allows a client that has stored data at an untrusted server to verify that the server possesses the original data without retrieving it. The model generates probabilistic proofs of possession by sampling random sets of blocks from the server, which drastically reduces I/O costs. The client maintains a constant amount of metadata to verify the proof. The challenge/response protocol transmits a small, constant amount of data, which minimizes network communication. Thus, the PDP model for remote data checking supports large data sets in widely-distributed storage systems. Previous work offers guarantees weaker than data possession, or requires prohibitive overhead at the server.

We present two provably-secure PDP schemes that are more efficient than previous solutions, even when compared with schemes that achieve weaker guarantees. In particular, the overhead at the server is low (or even constant), as opposed to linear in the size of the data. Experiments using our implementation verify the practicality of PDP and reveal that the performance of PDP is bounded by disk I/O and not by cryptographic computation.

1 Introduction

Verifying the authenticity of data has emerged as a critical issue in storing data on untrusted servers. It arises in peer-to-peer storage systems [26, 32], network file systems [27, 23], long-term archives [29], web-service object stores [42], and database systems [28]. Such systems prevent storage servers from misrepresenting or modifying data by providing authenticity checks when accessing data.

However, archival storage requires guarantees about the authenticity of data on storage, *i.e.*, that storage servers possess data. It is insufficient to detect that data have been modified or deleted when accessing the data, because it may be too late to recover lost or damaged data. Archival storage servers retain tremendous amounts of data, little of which are accessed. They also hold data for long periods of time during which there may be exposure to data loss from administration errors as the physical implementation of storage evolves (*e.g.*, backup and restore, data migration to new systems, and changing memberships in peer-to-peer systems).

Archival network storage presents unique performance demands. Given that file data are large and are stored at remote sites, accessing an entire file is expensive in I/O costs to the storage server and in transmitting the file across a network. Reading an entire archive, even periodically, greatly limits the scalability of network stores. (The growth in storage capacity has far outstripped the growth in storage access times and bandwidth [40]). Furthermore, I/O incurred to establish

*Department of Computer Science, Johns Hopkins University

†Google, Inc.

‡Carnegie Mellon University

data possession interferes with on-demand bandwidth to store and retrieve data. We conclude that clients need to be able to verify that a server has retained file data *without retrieving the data from the server* and *without having the server access the entire file*.

Previous solutions do not meet these requirements for proving data possession. Some schemes [18] provide a weaker guarantee by enforcing *storage complexity*: The server has to store an amount of data at least as large as the client’s data, but not necessarily the same exact data. Moreover, all previous techniques require the server to access the entire data file, which is not feasible when dealing with large amounts of data.

We define a model for provable data possession (PDP) that provides probabilistic proof that a third party stores a file. The model is unique in that it allows the server to access small portions of the file in generating the proof; all other techniques must access the entire file. Within this model, we give the first provably-secure scheme for remote data checking. The scheme conforms to our notion of an outsourced storage relationship. The client stores a small $O(1)$ amount of metadata to verify the server’s proof. Also, the scheme uses $O(1)$ bandwidth¹. The challenge and the response are each slightly more than 1 Kilobit. We also present a more efficient version of this scheme that proves data possession using a single modular exponentiation at the server, even though it provides a weaker guarantee.

Both schemes use *homomorphic verifiable tags*. The homomorphic properties of these tags allow one to combine tags computed from multiple file blocks into a single value and to verify these combined tags. To outsource the storage of a file, the client pre-computes tags for each block of the file and stores the file and its tags with a server. To verify that the server possesses that file at a later time, the client generates a random challenge against a randomly selected set of file blocks. Using the queried blocks and their corresponding tags, the server generates a proof of possession which ensures the client of data possession without retrieving the actual file blocks. This construct enables the client to *sample* the server’s storage.

The efficient PDP scheme is the fundamental construct underlying an archival introspection system that we are developing for the long-term preservation of Astronomy data. We are taking possession of multi-terabyte Astronomy databases at a University library in order to preserve the information long after the research projects and instruments used to collect the data are gone. The database will be replicated at multiple sites. Sites include resource-sharing partners that exchange storage capacity to achieve reliability and scale. As such, the system is subject to freeloading in which partners attempt to use storage resources and contribute none of their own [18]. The location and physical implementation of these replicas are managed independently by each partner and will evolve over time. Partners may even outsource storage to third-party storage server providers [21]. Efficient PDP schemes will ensure that the computational requirements of remote data checking do not unduly burden the remote storage sites.

We implemented our more efficient scheme (E-PDP) and two other remote data checking protocols (B-PDP and MHT-SC) and evaluated their performance. Experiments show that probabilistic possession guarantees make it practical to verify possession of large data sets. With sampling, E-PDP verifies a 64MB file in about 0.4 seconds as compared to 1.8 seconds without sampling. Further, I/O bounds the performance of E-PDP; it generates proofs as quickly as the disk produces data. Finally, E-PDP is 185 times faster than the previous secure protocol on 768 KB files.

¹Storage overhead and network overhead are constant in the size of the file, but depend on the chosen security parameter.

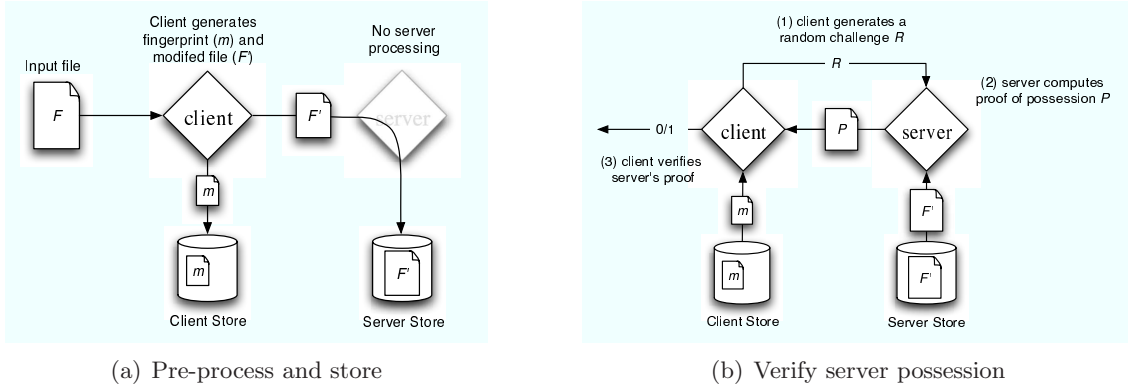


Figure 1: Protocol for provable data possession.

Contributions. In this paper we:

- formally define protocols for provable data possession (PDP) that provide probabilistic proof that a third party stores a file.
- introduce the first provably-secure and practical PDP schemes that guarantee data possession.
- implement one of our PDP schemes and show experimentally that probabilistic possession guarantees make it practical to verify possession of large data sets.

Paper Organization. The rest of the paper is organized as follows. In Section 2 we describe a framework for provable data possession, emphasizing the features and parameters that are relevant for PDP. Section 3 overviews related work. In Section 4 we introduce homomorphic verifiable tags, followed by definitions for PDP schemes; We then give our constructions (S-PDP and E-PDP). We support our theoretical claims with experiments that show the practicality of our schemes in Section 5 and conclude in Section 6.

2 Provable Data Possession (PDP)

We describe a framework for provable data possession. This provides background for related work and for the specific description of our schemes. A PDP protocol (Fig. 1) checks that an outsourced storage site retains a file, seen as a collection of n blocks. The client C (data owner) pre-processes the file, generating a piece of metadata (e.g., a fingerprint of the file) that is stored locally, transmits the file to the server S , and deletes its local copy. The server stores the file and responds to challenges issued by the client. Storage at the server is in $\Omega(n)$ and storage at the client is in $O(1)$, conforming to our notion of an outsourced storage relationship.

As part of pre-processing, the client may alter the file to be stored at the server. The client may expand the file or include additional metadata to be stored at the server. Before deleting its local copy of the file, the client may execute a data possession challenge to make sure the server has successfully stored the file. Clients often encrypt a file prior to out-sourcing the storage. For our purposes, encryption is an orthogonal issue; the “file” may consist of encrypted data and our metadata does not include encryption keys.

At a later time, the client issues a challenge to the server to establish that the server has retained the file. The client requests that the server compute a function of the stored file, which it sends back to the client. Using its local metadata, the client verifies the response.

Threat model. The server S must answer challenges from the client C ; failure to do so represents

	[18]	[18]-Wagner (MHT-SC)	[15, 17] (B-PDP)	[38]	[37]	S-PDP	E-PDP
Data possession	No	No	Yes	Yes	Yes*	Yes	Yes
Supports sampling	No	No	No	No	No [†]	Yes	Yes
Type of guarantee	deterministic				probabilistic / deterministic		probabilistic
Server block access	$O(n)$	$O(\log n)$	$O(n)$	$O(n)$	$O(n)$	$O(1)$	$O(1)$
Server computation	$O(n)$	$O(1)$	$O(n)$	$O(1)$	$O(n)$	$O(1)$	$O(1)$
Client computation	$O(1)$	$O(1)$	$O(1)$	$O(1)$	$O(1)$	$O(1)$	$O(1)$
Communication	$O(1)$	$O(\log n)$	$O(1)$	$O(1)$	$O(n)$	$O(1)$	$O(1)$
Client storage	$O(1)$	$O(1)$	$O(1)$	$O(n)$	$O(1)$	$O(1)$	$O(1)$

Table 1: Features and parameters (per challenge) of various PDP schemes when the server misbehaves by deleting a fraction of n (e.g., 1% of n), where n is the number of file blocks. The server and client computation is expressed as the total cost of performing modular exponentiation operations. For simplicity, the security parameter is not included as a factor for the relevant costs.

* No security proof is given for this scheme, so assurance of data possession is not confirmed.

[†] The client can ask proof for select symbols inside a block, but cannot sample across blocks.

a data loss. However, the server is not trusted to store the client’s file and may try to convince the client that it possesses the file even though it is totally or partially missing the file. The server’s motivation for misbehavior can be diverse and includes: To reclaim storage by discarding data that has not been or is rarely accessed (for monetary reasons); Or to hide a data loss incident (due to management errors, hardware failure, compromise by outside or inside attacks etc). The goal of a PDP scheme that achieves probabilistic proof of data possession is to detect server misbehavior when the server has deleted a percentage of the file.

Requirements and Parameters. The important performance parameters of a PDP scheme include:

- *Computation complexity:* The computational cost to pre-process a file (at C), to generate a proof of possession (at S) and to verify such a proof (at C);
- *Block access complexity:* The number of file blocks accessed to generate a proof of possession (at S);
- *Communication complexity:* The amount of data transferred (between C and S).

For a scalable solution, the amount of computation and block accesses at the server should be minimized, because the server may be involved in concurrent interactions with many clients. *We stress that in order to minimize bandwidth, an efficient PDP scheme cannot consist of retrieving entire file blocks.* While relevant, the computation complexity at the client is of less importance, even though our schemes minimize that as well. Also, a PDP scheme requires a small, constant amount of storage at the client, conforming to our notion of an outsourced storage relationship.

To meet these performance goals, our PDP schemes sample the server’s storage, accessing a random subset of blocks. In doing so, the PDP schemes provide a probabilistic guarantee of possession; a deterministic guarantee cannot be provided without accessing all blocks. In fact, as a special case of our PDP scheme, the client may ask proof for all the file blocks, making the data possession guarantee deterministic. Sampling proves data possession with high probability based on

accessing few blocks in the file, which radically alters the performance of proving data possession. Interestingly, when the server deletes a fraction of the file, the client can detect server misbehavior with high probability by asking proof for a constant amount of blocks, independently of the total number of file blocks. As an example, for a file with $n = 10,000$ blocks, if S has deleted 1% of the blocks, then C can detect server misbehavior with probability greater than 99% by asking proof of possession for only 460 randomly selected blocks (representing 4.6% of n). For more details see Sec. 5.1.

We list the features of our PDP schemes (S-PDP and E-PDP) in Table 1. We also include a comparison of related techniques [15, 17](B-PDP), [18], [38] and [37]. The scheme [18]-Wagner (MHT-SC) refers to the variant suggested by David Wagner [18], based on Merkle hash trees. Both schemes in [18] do not provide a data possession guarantee, but only enforce storage complexity. Indeed, as noted by Golle *et al.* [18], the server could pass the verification procedure by using (and storing) a different file, which is at least as big as the original one. We emphasize that PDP schemes that offer an inherently deterministic guarantee by accessing all the blocks of the file ([18, 15, 17]) cannot offer *both* sampling across blocks and constant storage on the client; fundamental changes would be required in these schemes in order to avoid storing $O(n)$ metadata on the client. The scheme in [37] allows the client to request proof of possession of segments from each file block. However, the server’s response includes one “signature” per each block, which makes its size linear with respect to the number of blocks.

A “strawman” solution. We notice that if one is willing to allow the server to send entire file blocks back to the client as proof of possession, then a simple solution is possible. Initially, the client computes a message authentication code (MAC) for each file block and stores the MACs at the server. Later, to obtain a probabilistic proof of data possession, the client retrieves a number of randomly selected file blocks and their corresponding MACs. However, this solution requires an arbitrarily large amount of bandwidth, which is linear with respect to the number of queried blocks. In contrast, our solutions have constant network communication.

3 Related Work

Deswarte *et al.* [15] and Filho *et al.* [17] provide similar techniques to verify that a remote server stores a file using RSA-based hash functions. Unlike other hash-based approaches, it allows a client to perform multiple challenges using the same metadata. In this protocol, communication and client storage complexity are both $O(1)$. The limitation of the algorithm lies in the computational complexity at the server, which must exponentiate the entire file, accessing all of the file’s blocks. Further, RSA over the entire file is extremely slow — 20 seconds per Megabyte for 1024-bit keys on a 3.0 GHz processor [17]. In fact, these limitations led us to study algorithms that allowed for sub-file access (sampling). We implement this protocol for comparison with our PDP scheme and refer to it as B-PDP (basic PDP). A description of B-PDP is provided in Appendix B.

Schwarz and Miller [37] propose a scheme that allows a client to verify the storage of m/n erasure-coded data across multiple sites even if sites collude. The data possession guarantee is achieved using a special construct, called “algebraic signature”: a function that fingerprints a block and has the property that the signature of the parity block equals the parity of the signatures of the data blocks. The parameters of the scheme limit its applicability: the file access and computation complexity at the server and the communication complexity are all linear in the number of file blocks (n) per challenge. Additionally, the security of the scheme is not proven and remains in question.

Sebe *et al.* [38] give a protocol for remote file integrity checking, based on the Diffie-Hellman problem in \mathbb{Z}_N . The client has to store N bits per block, where N is the size of an RSA modulus, so the total storage on the client is $O(n)$ (which does not conform to our notion of an outsourced storage relationship). Indeed, the authors state that this solution only makes sense if the size of a block is much larger than N . Moreover, the protocol requires the server to access the entire file.

Homomorphic hash functions have been used in *source authentication*, *e.g.* to verify data transfer in a pay-for-bandwidth content distribution network. Krohn *et al.* [25] provide the best example and give a good review of other applications of homomorphic hashing. They use homomorphism to compose multiple blocks inputs into a single value. However, this protocol can only compose specific subsets of blocks, based on erasure coding. Source authentication techniques do not apply to provable data possession.

Related to provable data possession is the enforcement of storage complexity, which shows that a server retains *an amount of information at least as large as the file* received from the client; The server does not necessarily retain the original file. To the best of our knowledge, Golle *et al.* [18] were the first to propose a scheme that enforces storage complexity, based on a new assumption (n -Power Decisional Diffie-Hellman). Their scheme considers a different setting that involves three parties and provides two additional features (binding and concealing) that resemble commitment schemes. Golle *et al.* also briefly mention a scheme suggested by David Wagner, based on Merkle hash trees, which lowers the computational requirements for the server at the expense of increased communication. We implement Wagner’s suggestion for comparison with our PDP scheme and refer to it as MHT-SC. A description of MHT-SC is provided in in Appendix B. Note that MHT-SC does not guarantee data possession (it only enforces storage complexity).

Oprea *et al.* [36] propose a scheme based on tweakable block ciphers that allows a client to detect the modification of data blocks by an untrusted server. The scheme does not require additional storage at the server and if the client’s data has low entropy then the client only needs to keep a relatively low amount of state. However, verification requires the entire file to be retrieved, which means that the server file access and communication complexity are both linear with the file size per challenge. The scheme is targeted for data retrieval, but is impractical for verifying data possession.

Memory checking protocols [10, 34] verify that all reads and writes to a remote memory behave identically to reads and writes to a local memory. PDP is a restricted form of memory checking in that memory checking verifies every read and write generated from a program at any granularity. Due to this restriction, memory checking is much more difficult and expensive than proving data possession.

An alternative to checking remote storage is to make data resistant to undetectable deletion through entanglement [41, 2], which encodes data to create dependencies among unrelated data throughout the storage system. Thus, deleting any data reveals itself as it deletes other unrelated data throughout the system.

Our homomorphic verifiable tags are related with the concept of homomorphic signatures introduced by Johnson *et al.* [22]. As pointed out by the same authors, fully-additive homomorphic signature schemes cannot be secure; Our homomorphic verifiable tags are able to achieve a form of message homomorphism under addition because of their special structure that uses one-time indices.

The PDP solutions we propose are the first schemes that securely prove the possession of data on an untrusted server and are computationally efficient, *i.e.*, require a constant number of modular exponentiations and have constant I/O complexity. This makes our PDP solutions the first schemes

suitable for implementation in large data systems.

4 Provable Data Possession Schemes

4.1 Preliminaries

The client C wants to store on the server S a file F which is a finite ordered collection of n blocks: $F = (m_1, \dots, m_n)$.

We denote the output x of an algorithm \mathcal{A} by $x \leftarrow \mathcal{A}$.

Homomorphic Verifiable Tags (HVTs). We introduce the concept of a homomorphic verifiable tag that will be used as a building block for our PDP schemes.

Given a message m (corresponding to a file block), an HVT is a pair $(T_{W,m}, W)$, where W is a random value. The client must ensure that the values W are not reused, *i.e.*, that no two tags use the same random value W . For example, these different random values can be generated by applying a pseudo-random function on a set of *one-time indices*, *i.e.*, a different index is used each time. A simple way to ensure that the values W are different and unpredictable is to use a counter for the index. The tags will be stored on the server together with the file F . Homomorphic verifiable tags act as verification metadata for the file blocks and have the following properties:

- *Blockless verification:* Using HVTs the server can construct a proof that allows the client to verify if the server possesses certain file blocks, even when the client does not have access to the actual file blocks. HVTs and their corresponding proofs have a fixed constant size and are (much) smaller than the actual file blocks.
- *Homomorphic tags:* Given two tags T_{W_i, m_i} and T_{W_j, m_j} , anyone can combine them into a tag $T_{\{W_i, W_j\}, m_i + m_j}$ corresponding to the set of random values $\{W_i, W_j\}$ and the sum of the messages $m_i + m_j$.

We stress that techniques based on *aggregate signatures* [11], *multi-signatures* [30, 35], batch RSA [16], batch verification of RSA [20, 3], condensed RSA [33], etc. would all fail to provide *blockless verification*, which is needed by our PDP scheme. Indeed, the client has to have the ability to verify the tags on *specific* file blocks even though he *does not* possess any of those blocks.

4.2 Definitions

We start with the precise definition of a provable data possession scheme, followed by a security definition that captures the data possession property.

Definition 4.1 (PROVABLE DATA POSSESSION SCHEME (PDP)) *A PDP scheme is a collection of four polynomial-time algorithms (KeyGen, TagBlock, GenProof, CheckProof) such that:*

$\text{KeyGen}(1^k) \rightarrow (\text{pk}, \text{sk})$ *is a probabilistic key generation algorithm that is run by the client to setup the scheme. It takes a security parameter k as input, and returns a pair of matching public and secret keys (pk, sk) .*

$\text{TagBlock}(\text{sk}, m) \rightarrow (T_{W,m}, W)$ *is run by the client to generate the verification metadata. It takes as inputs a secret key sk and a file block m , and returns the verification metadata $(T_{W,m}, W)$ corresponding to the tuple (W, m) .*

$\text{GenProof}(\text{pk}, F, \text{chal}, \Sigma) \rightarrow \mathcal{V}$ *is run by the server in order to generate a proof of possession. It takes as inputs a public key pk , an ordered collection F of blocks, a challenge chal and the verification metadata set Σ corresponding to the blocks in F . It returns a proof of possession \mathcal{V} for the blocks in F that are determined by the challenge chal .*

$\text{CheckProof}(\text{pk}, \text{sk}, \text{chal}, \mathcal{V}) \rightarrow \{\text{“success”}, \text{“failure”}\}$ is run by the client in order to validate a proof of possession. It takes as inputs a public key pk , a secret key sk , a challenge chal and a proof of possession \mathcal{V} . It returns whether \mathcal{V} is a correct proof of possession for the blocks determined by chal .

A PDP system can be constructed from a PDP scheme in two phases, **Setup** and **Challenge**:

Setup: The client C is in possession of the file F and runs $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^k)$, followed by $(\text{T}_{W_i, m_i}, W_i) \leftarrow \text{TagBlock}(\text{sk}, m_i)$ for all $1 \leq i \leq n$. C stores the pair (sk, pk) . C then sends pk, F and $\Sigma = \{\text{T}_{W_1, m_1}, \dots, \text{T}_{W_n, m_n}\}$ to S for storage and deletes F and Σ from its local storage.

Challenge: C generates a challenge chal that, among other things, indicates the specific blocks for which C wants a proof of possession. C then sends chal to S . S runs $\mathcal{V} \leftarrow \text{GenProof}(\text{pk}, F, \text{chal}, \Sigma)$ and sends to C the proof of possession \mathcal{V} . Finally, C can check the validity of the proof \mathcal{V} by running $\text{CheckProof}(\text{pk}, \text{sk}, \text{chal}, \mathcal{V})$.

In the **Setup** phase, C computes tags for each file block and stores them together with the file at S . In the **Challenge** phase, C requests proof of possession for a subset of the blocks in F . This phase can be executed an unlimited number of times in order to ascertain whether S still possesses the selected blocks.

We state the security for a PDP system using a game that captures the data possession property. Intuitively, the Data Possession Game captures that an adversary cannot successfully construct a valid proof without possessing all the blocks corresponding to a given challenge, unless it guesses all the missing blocks. We assume w.l.o.g. that F contains random data (see Remark 3 in Sec. 4.3).

Data Possession Game:

- **Setup:** The challenger runs $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^k)$, sends pk to the adversary and keeps sk secret.
- **Query:** The adversary makes tagging queries adaptively: It selects a block m_1 and sends it to the challenger. The challenger picks a random value W_1 and computes the verification metadata $(\text{T}_{W_1, m_1}, W_1) \leftarrow \text{TagBlock}(\text{sk}, m_1)$. The challenger then sends back to the adversary the values T_{W_1, m_1} and W_1 . The adversary continues to query the challenger for the verification metadata $(\text{T}_{W_2, m_2}, W_2), \dots, (\text{T}_{W_n, m_n}, W_n)$ on the blocks of its choice m_2, \dots, m_n . As a general rule, the challenger generates T_{W_j, m_j} for some $1 \leq j \leq n$, by picking a random value W_j and computing $(\text{T}_{W_j, m_j}, W_j) \leftarrow \text{TagBlock}(\text{sk}, m_j)$. The adversary then stores all the blocks as an ordered collection $F = (m_1, \dots, m_n)$, together with the corresponding verification metadata $(\text{T}_{W_1, m_1}, W_1), \dots, (\text{T}_{W_n, m_n}, W_n)$.
- **Challenge:** The challenger selects a challenge chal and requests the adversary to provide a proof of possession of any blocks of the adversary’s choice.
- **Forge:** The adversary computes \mathcal{V} without knowledge of at least α bits of information about some blocks of its choice m_{i_1}, \dots, m_{i_c} , $1 \leq i_j \leq n, 1 \leq c \leq n$, and returns \mathcal{V} as a proof of possession for the blocks m_{i_1}, \dots, m_{i_c} .

If $\text{CheckProof}(\text{pk}, \text{sk}, \text{chal}, \mathcal{V}) = \text{“success”}$, then the adversary has won the Data Possession Game.

Definition 4.2 A PDP system (**Setup, Challenge**) built on a PDP scheme ($\text{KeyGen}, \text{TagBlock}, \text{GenProof}, \text{CheckProof}$) guarantees data possession if no (probabilistic polynomial-time) adversary can win the Data Possession Game with probability more than $2^{-\alpha} + \epsilon(k)$, where α is the number of bits that the adversary is missing and ϵ is negligible in the security parameter k .

4.3 Efficient and Secure PDP Schemes

In this section we present our PDP constructions: the first (S-PDP) provides a strong data possession guarantee, while the second (E-PDP) achieves better efficiency at the cost of weakening the data possession guarantee.

We start by introducing some additional notation used by the constructions. Let $p = 2p' + 1$ and $q = 2q' + 1$ be safe primes and let $N = pq$ be an RSA modulus. Let g be a generator of QR_N , the unique cyclic subgroup of \mathbb{Z}_N^* of order $p'q'$ (i.e., QR_N is the set of quadratic residues modulo N). We can obtain g as $g = a^2$, where $a \stackrel{R}{\leftarrow} \mathbb{Z}_N^*$ such that $\gcd(a \pm 1, N) = 1$. All exponentiations are performed modulo N , and for simplicity we sometimes omit writing it explicitly. Let $h : \{0, 1\}^* \rightarrow QR_N$ be a secure deterministic hash-and-encode function² that maps strings uniformly to QR_N .

The schemes are based on the KEA1 assumption which was introduced by Damgard in 1991 [12] and subsequently used by several others, most notably in [19, 4, 5, 24, 13]. In particular, Bellare and Palacio [4] provided a formulation of KEA1, that we follow and adapt to work in the RSA ring:

KEA1-r (Knowledge of Exponent Assumption): For any adversary \mathbf{A} that takes input (N, g, g^s) and returns (C, Y) such that $Y = C^s$, there exists an “extractor” $\bar{\mathbf{A}}$, which given the same inputs as \mathbf{A} , returns x such that $g^x = C$.

Recently, KEA1 has been shown to hold in generic groups (i.e., it is secure in the generic group model) by A. Dent [14] and independently by Abe and Fehr [1].

S-PDP overview. We first give an overview of our provable data possession scheme that supports sampling. In the **Setup** phase, the client computes a homomorphic verifiable tag (T_{W_i, m_i}, W_i) for each block m_i of the file. In order to maintain constant storage, the client generates the random values W_i using a pseudo-random function; thus, **TagBlock** has an extra parameter, i . Each value T_{W_i, m_i} includes information about the index i of the block m_i , in the form of a value $h(W_i)$, as W_i is obtained from i by using a pseudo-random function. This binds the tag on a block to that specific block, and prevents using the tag to obtain a proof for a different block. These tags are stored on the server together with the file F . The extra storage at the server is the price to pay for allowing thin clients that only store a small, constant amount of data, regardless of the file size.

In the **Challenge** phase, the client asks the server for proof of possession of c file blocks whose indices are randomly chosen using a pseudo-random permutation keyed with a fresh randomly-chosen key for each challenge. This prevents the server from anticipating which blocks will be queried in each challenge. C also generates a fresh (random) challenge $g_s = g^s$ to ensure that S does not reuse any values from a previous **Challenge** phase. The server returns a proof of possession that consists of two values: T and ρ . T is obtained by combining into a single value the individual tags T_{W_i, m_i} corresponding to each of the requested blocks. ρ is obtained by raising the challenge g_s to a function of the requested blocks. The value T contains information about the indices of the blocks requested by the client (in the form of the $h(W_i)$ values). C can remove all the $h(W_i)$ values from T because it has both the key for the pseudo-random permutation (used to determine the indices of the requested blocks) and the key for the pseudo-random function (used to generate the values W_i). C can then verify the validity of the server’s proof by checking if a certain relation

² h is computed by squaring the output of the full-domain hash function for the provably secure FDH signature scheme [6, 7] based on RSA. We refer the reader to [6] for ways to construct an FDH function out of regular hash functions, such as SHA-1. Alternatively, h can be the deterministic encoding function used in RSA-PSS[8]

As previously defined, let f and w be pseudo-random functions, let π be a pseudo-random permutation and let H be a cryptographic hash function.

KeyGen(1^k): Generate $\text{pk} = (N, e, g)$ and $\text{sk} = (N, d, v)$, such that (N, e) and (N, d) are matching RSA public and secret keys, and $v \xleftarrow{R} \{0, 1\}^\kappa$. Output (pk, sk) .

TagBlock(sk, m, i):

1. Let $(N, d, v) = \text{sk}$. Generate $W_i = w_v(i)$. Compute a tag $(T_{W_i, m}, W_i)$ on the tuple (W_i, m) :
 $T_{W_i, m} = (h(W_i) \cdot g^m)^d \bmod N$.
2. Output $(T_{W_i, m}, W_i)$.

GenProof($\text{pk}, F = (m_1, \dots, m_n), \text{chal}, \Sigma = \{T_{W_1, m_1}, \dots, T_{W_n, m_n}\}$):

1. Let $(N, e, g) = \text{pk}$ and $(c, k_1, k_2, g_s) = \text{chal}$.
 For $1 \leq j \leq c$:
 - compute the indices of the blocks for which the proof is generated: $i_j = \pi_{k_1}(j)$
 - compute coefficients: $a_j = f_{k_2}(j)$.
2. Compute $T = T_{W_{i_1}, m_{i_1}}^{a_1} \cdot \dots \cdot T_{W_{i_c}, m_{i_c}}^{a_c} = (h(W_{i_1})^{a_1} \cdot \dots \cdot h(W_{i_c})^{a_c} \cdot g^{a_1 m_{i_1} + \dots + a_c m_{i_c}})^d \bmod N$.
3. Compute $\rho = H(g_s^{a_1 m_{i_1} + \dots + a_c m_{i_c}} \bmod N)$.
4. Output $\mathcal{V} = (T, \rho)$.

CheckProof($\text{pk}, \text{sk}, \text{chal}, \mathcal{V}$):

1. Let $(N, e, g) = \text{pk}$, $(N, d, v) = \text{sk}$, $(c, k_1, k_2, s) = \text{chal}$ and $(T, \rho) = \mathcal{V}$.
2. Let $\tau = T^e$. For $1 \leq j \leq c$:
 - compute $i_j = \pi_{k_1}(j)$, $W_{i_j} = w_v(i_j)$, $a_j = f_{k_2}(j)$, and $\tau = \frac{\tau}{h(W_{i_j})^{a_j}} \bmod N$
 (As a result, one should get $\tau = g^{a_1 m_{i_1} + \dots + a_c m_{i_c}} \bmod N$)
3. If $H(\tau^s) = \rho \bmod N$, then output “success”. Otherwise output “failure”.

We construct a PDP system from a PDP scheme in two phases, **Setup** and **Challenge**:

Setup: The client C runs $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(1^k)$, followed by $(T_{W_i, m_i}, W_i) \leftarrow \text{TagBlock}(\text{sk}, m_i, i)$ for all $1 \leq i \leq n$ (note that **TagBlock** has an extra parameter, i , since the random values W_i are generated using a PRF). C stores the pair (sk, pk) . C then sends pk , F and $\Sigma = \{T_{W_1, m_1}, \dots, T_{W_n, m_n}\}$ to S for storage and deletes F and Σ from its local storage.

Challenge: C requests proof of possession for c distinct blocks of the file F (with $1 \leq c \leq n$):

1. C generates the challenge $\text{chal} = (c, k_1, k_2, g_s)$, where $k_1 \xleftarrow{R} \{0, 1\}^\kappa$, $k_2 \xleftarrow{R} \{0, 1\}^\kappa$, $g_s = g^s \bmod N$ and $s \xleftarrow{R} \mathbb{Z}_N^*$. C sends chal to S .
2. S runs $\mathcal{V} \leftarrow \text{GenProof}(\text{pk}, F, \text{chal}, \Sigma = \{T_{W_1, m_1}, \dots, T_{W_n, m_n}\})$ and sends to C the proof of possession \mathcal{V} .
3. C sets $\text{chal} = (c, k_1, k_2, s)$ and checks the validity of the proof \mathcal{V} by running **CheckProof**($\text{pk}, \text{sk}, \text{chal}, \mathcal{V}$).

Figure 2: S-PDP: a PDP scheme with strong data possession guarantee

holds between T and ρ .

S-PDP in detail. We are now ready to proceed to the details of the construction. Let κ, ℓ be

security parameters and let H be a cryptographic hash function. In addition, we make use of two pseudo-random functions (PRF) f, w , and a pseudo-random permutation (PRP) π with the following parameters:

- $f : \{0, 1\}^\kappa \times \{0, 1\}^{\log_2(n)} \rightarrow \{0, 1\}^\ell$;
- $w : \{0, 1\}^\kappa \times \{0, 1\}^{\log_2(n)} \rightarrow \{0, 1\}^\ell$;
- $\pi : \{0, 1\}^\kappa \times \{0, 1\}^{\log_2(n)} \rightarrow \{0, 1\}^{\log_2(n)}$

We write $f_k(x)$ to denote f keyed with key k applied on input x . Our S-PDP scheme is described in Fig. 2.

In the B-PDP scheme, C is storing a single piece of metadata for the entire file. A direct extension to a scheme capable of block sampling would require C to store a piece of metadata for each block of the file. We are able to reduce C 's storage requirement by having the client tag each piece of metadata and store these tags on the server. We include $h(W_i)$ in the tag for a block m_i in order to ensure that the values returned by the server are not computed on a set of blocks different than the set of blocks requested by the client. The purpose of including the a_j coefficients in the values for ρ and T computed by S is to ensure that S possesses each one of the requested blocks. These coefficients are determined by a PRF keyed with a fresh randomly-chosen key for each challenge, which prevents S from storing combinations (*e.g.*, sums) of the original blocks instead of the original file blocks themselves. Also, we are able to maintain constant communication cost because tags on blocks can be combined into a single value.

In our analysis we assume w.l.o.g. that the indices for the blocks picked by the client in a challenge are different. One way to achieve this is to implement π using the techniques proposed by Black and Rogaway [9]. In a practical deployment, our protocol can tolerate collisions of these indices.

In Appendix A, we prove:

Theorem 4.3 *Under the RSA and KEA1-r assumptions, S-PDP guarantees data possession in the random oracle model.*

Regarding efficiency, we remark that each challenge requires a small, constant amount of communication between C and S (the challenge and the response are each slightly more than 1 Kilobit). In terms of server block access, the demands are c accesses for S , while in terms of computation we have c exponentiations for both C and S . When S deletes a fraction of the file blocks, c is a relatively small, constant value, which gives the $O(1)$ parameters in Table 1 (for more details, see Sec. 5.1). Since the size of the file is $O(n)$, accommodating the additional tags does not change (asymptotically) the storage requirements for the server.

A concrete example of using S-PDP. For a concrete example of using S-PDP, we consider a 1024-bit modulus N and a 4 GB file F which has $n = 1,000,000$ 4KB blocks. During **Setup**, C stores the file and the tags at S . The tags require additional storage of 128 MB. The client stores less than 2048 bits (N has 1024 bits; e, d, v have less than 1024 bits). During the **Challenge** phase, C and S use AES for π (used to select the random block indices i), HMAC for w (used to compute the random values W), HMAC for f (used to determine the random coefficients a) and SHA1 for H . In a challenge, C sends to S four values which total 168 bytes (c has 4 bytes, k_1 has 16 bytes, k_2 has 20 bytes, g_s has 1024 bits). Assuming that S deletes at least 1% of F , then C can detect server misbehavior with probability over 99% by asking proof for $c = 460$ randomly selected blocks. The server's response contains two values which total 148 bytes (T has 1024 bits, ρ has 20 bytes). We

emphasize that the server’s response to a challenge consists of a small, constant value; in particular, the server does not send back to the client any of the file blocks and not even their sum.

Remark 1. We have described a PDP scheme for the case when a client stores a single file F on the server. In the **TagBlock** algorithm, for each block m_i the client computes a tag over the tuple (W_i, m_i) . We emphasize that the values W_i cannot be reused. Since W_i is obtained from i using a PRF, this implies that the indices i must be different across all tags. In other words, the client should not use the same index twice for computing tags. This condition holds because in our scheme an index i is simply the position of the block m_i in the file.

In order to store multiple files on the server, the client must ensure that indices used to compute tags are distinct not only across the tags corresponding to the blocks of each file, but also across the tags corresponding to the blocks of all files. One simple method to achieve this is to pre-pend the file’s identifier to the index. For example, if the identifier of a file $F = (m_1, \dots, m_n)$ is given by $\text{id}(F)$, then for each block m_i , $1 \leq i \leq n$, C computes the tag $(T_{W_{\text{id}(F)}||i}, m_i)$ on input $(W_{\text{id}(F)}||i, m_i)$, where $W_{\text{id}(F)}||i = w_v(\text{id}(F)||i)$. The uniqueness of indices is ensured under the assumption that each file has a unique identifier. Another simple way to ensure that indices are only used once is to use a global counter for the index, which is incremented by the client each time after a tag is computed.

Remark 2. In step 3 of the **GenProof** algorithm, S has two ways of computing ρ : either sum the values $a_j m_{i_j}$ (as integers) and then exponentiate g_s to this sum, or exponentiate g_s to each value $a_j m_{i_j}$ and then multiply all values. We observed that the former choice takes considerable less time, as it only involves one exponentiation to a $(|m_i| + \ell + \log_2(c))$ -bit number, as opposed to c exponentiations to a $(|m_i| + \ell)$ -bit number (typically, $\ell = 160$).

Remark 3. Note that the server is free to store the client’s file F however it sees fit, as long as it is able to recover the file when answering a challenge. For example, it is allowed to compress F (e.g., if all the blocks of F are identical, then the server may only store one full block and the information that all the blocks are equal). Alternatively, w.l.o.g., one could assume that F has been optimally compressed by the client and the size of the F is equal to F ’s information entropy function.

A more efficient scheme, with weaker guarantees (E-PDP). Our S-PDP scheme provides the guarantee that S possesses each one of the c blocks for which C requested proof of possession in a challenge. We now describe a more efficient variant of S-PDP, which we call E-PDP, that achieves better performance at the cost of offering weaker guarantees. E-PDP differs from S-PDP only in that all the coefficients a_j are equal to 1:

- In **GenProof** (steps 2 and 3) the server computes $T = T_{W_{i_1}, m_{i_1}} \cdot \dots \cdot T_{W_{i_c}, m_{i_c}}$ and $\rho = H(g_s^{m_{i_1} + \dots + m_{i_c}} \text{ mod } N)$.
- In **CheckProof** (step 2) the client computes $\tau = \frac{T^e}{h(W_{i_1}) \cdot \dots \cdot h(W_{i_c})} \text{ mod } N$.

The E-PDP scheme reduces the computation on both the server and the the client to one exponentiation, as described in Table 1 (see Remark 3 above, the server computes ρ as one exponentiation to a value whose size in bits is slightly larger than $|m_i|$). E-PDP only guarantees possession of the sum of the blocks $m_{i_1} + \dots + m_{i_c}$ and not necessarily possession of each one of the blocks for which the client requests proof of possession. However, we argue next that this is not a practical concern.

If the server pre-computes and stores the sums of all possible combinations of c blocks out of the n blocks ($\binom{n}{c}$ values), then the server could successfully pass any challenge. When using E-PDP,

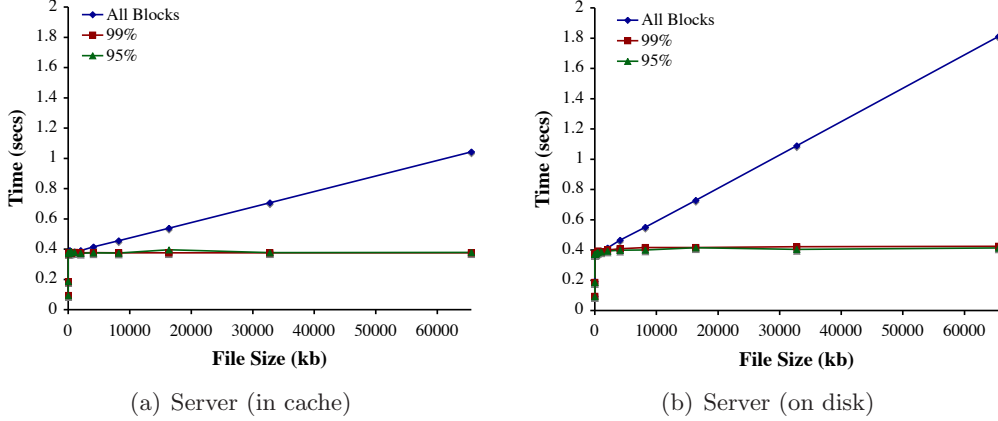


Figure 4: Performance of sampling at multiple confidence levels.

Since $\frac{n-i-t}{n-i} \geq \frac{n-i-1-t}{n-i-1}$, it follows that:

$$1 - \left(\frac{n-t}{n}\right)^c \leq P_X \leq 1 - \left(\frac{n-c+1-t}{n-c+1}\right)^c.$$

P_X indicates the probability that, if S deleted a fraction t of the file, then C will detect server misbehavior after a challenge in which it asks proof for c blocks. Fig. 3 plots P_X for different values of n , t , c . Interestingly, when S deletes a fraction t of the file, C can detect server misbehavior with a certain probability by asking proof for a constant amount of blocks, independently of the total number of file blocks: If $t = 1\%$, then C asks for 460 blocks and 300 blocks in order to achieve P_X of at least 99% and 95%, respectively.

5.2 Implementation and Experimental Results

We measure the performance of E-PDP and the benefits of sampling based on our implementation of E-PDP in Linux. As a basis for comparison, we have also implemented the scheme of Deswarte *et al.* [15] and Filho *et al.* [17] (B-PDP), and the more efficient scheme in [18] (MHT-SC) suggested by David Wagner (these schemes are described in Appendix B).

All experiments were conducted on an Intel 2.8 GHz Pentium IV system with a 512 KB cache, an 800 MHz EPCI bus, and 1024 MB of RAM. The system runs Red Hat Linux 9, kernel version 2.4.22. Algorithms use the crypto library of OpenSSL version 0.9.8b with a modulus N of size 1024 bits and files have 4KB blocks. Experiments that measure disk I/O performance do so by storing files on an ext3 file system on a Seagate Barracuda 7200.7 (ST380011A) 80GB Ultra ATA/100 drive. All experimental results represent the mean of 20 trials. Because results varied little across trials, we do not present confidence intervals.

Sampling. To quantify the performance benefits of sampling for E-PDP, we compare the client and server performance for detecting 1% missing or faulty data at 95% and 99% confidence (Fig. 4). These results are compared with using E-PDP over all blocks of the file at large file sizes, up to 64 MB. We measure both the computation time only (in memory) as well as the overall time (on disk), which includes I/O costs.

Examining all blocks uses time linear in the file size for files larger than 4MB. This is the point at which the computation becomes bound from either memory or disk throughput. Larger

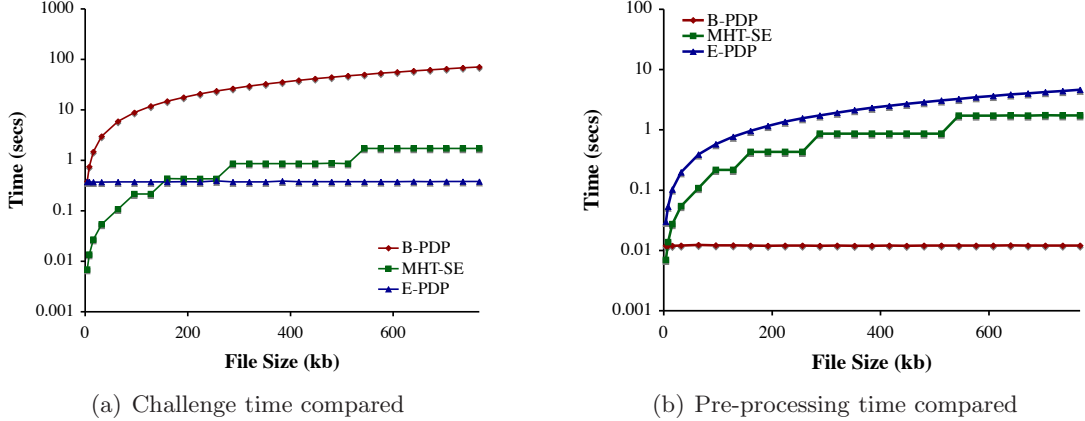


Figure 5: Computation performance.

inputs amortize the cost of the single exponentiation required by E-PDP. This is also the point at which the performance of sampling diverges. The number of blocks needed to achieve the target confidence level governs performance.

For larger files, E-PDP generates data as fast as it can be accessed from memory and summed, because it only computes a single exponentiation. In E-PDP, the server generates $\sum_{i=1}^c m_i$, which it exponentiates. The maximum size of this quantity in bits is $|m_i| + \log_2(c)$; its maximum value is $c \cdot 2^{|m_i|}$. Thus, the cryptographic costs grows logarithmically in the file size. The linear cost of accessing all data blocks and computing the sum dominate this logarithmic growth.

Comparing results when data are on disk versus in cache shows that disk throughput bounds E-PDP’s performance when accessing all blocks. With the exception of the first blocks of a file, I/O and the challenge computation occur in parallel. Thus, E-PDP generates proofs faster than the disk can deliver data: 1.0 second versus 1.8 seconds for a 64 MB file. Because I/O bounds performance, no protocol can outperform E-PDP by more than the startup costs. While faster, multiple-disk storage may remove the I/O bound today. Over time increases in processor speeds will exceed those of disk bandwidth and the I/O bound will hold.

Sampling breaks the linear scaling relationship between time to generate a proof of data possession and the size of the file. At 99% confidence, E-PDP can build a proof of possession for any file, up to 64 MB in size in about 0.4 seconds. Disk I/O incurs about 0.04 seconds of additional runtime for larger file sizes over the in-memory results. Sampling performance characterizes the benefits of E-PDP. Probabilistic guarantees make it practical to use public-key cryptography constructs to verify possession of very large data sets.

Server Computation. The next experiments look at the worst-case performance of generating a proof of possession, which is useful for planning purposes to allow the server to allocate enough resources. For E-PDP, this means sampling every block in the file, while for MHT-SC this means computing the entire hash tree. We compare the computation complexity of E-PDP with other algorithms, which do not support sampling. All schemes perform an equivalent number of disk and memory accesses.

Fig. 5(a) shows the computation time as a function of file size used at the server when computing a proof for B-PDP, MHT-SC and E-PDP. Note the logarithmic scale. Computation time includes the time to access the memory blocks that contain file data in cache. We restrict this experiment

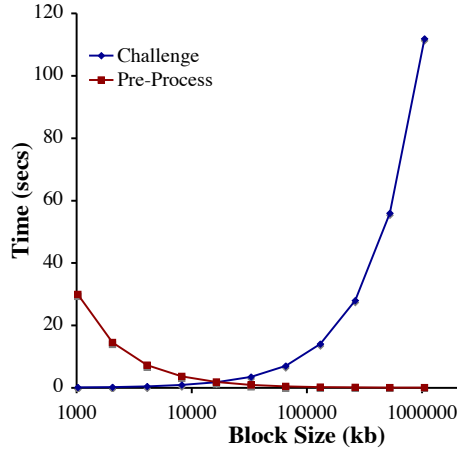


Figure 6: E-PDP pre-processing vs. challenge trade-offs with block size for a 1 GB file.

to files of 768 KB or less, because of the amount of time consumed by B-PDP.

E-PDP radically alters the complexity of data possession protocols and even outperforms protocols that provide weaker guarantees, specifically MHT-SC. For files of 768 KB, E-PDP is more than 185 times faster than B-PDP and more than 4.5 times as fast as MHT-SC. These performance ratios become arbitrarily large for larger file sizes. For B-PDP performance grows linearly with the file size, as it exponentiates the entire file. For MHT-SC, performance also grows linearly, but in disjoint clusters or equivalence classes, which represent the height of the Merkle-tree needed to represent a file of that size.

Pre-Processing: In preparing a file for outsourced storage, the client generates its local metadata. In this experiment, we measure the processor time for metadata generation only. This does not include the I/O time to load data to the client or store metadata to disk, nor does it include the time to transfer the file to the server. Fig. 5(b) shows the pre-processing time as a function of file size for B-PDP, MHT-SC and E-PDP.

E-PDP exhibits slower pre-processing performance. The costs grow linearly with the file size at 162 KB/s. E-PDP performs an exponentiation on every block of the file in order to create the per-block tags. For MHT-SC, preprocessing performance mirrors challenge performance, because both protocol steps perform the same computation. It generates data at about 433 KB/s on average.

The preprocessing performance of B-PDP differs from the challenge phase even though both steps compute the exact same signature. This is because the client has access to $\phi(N)$ and can reduce the file modulo $\phi(N)$ before exponentiating. In contrast, the security of the protocol depends on $\phi(N)$ being a secret, unavailable to the server. The preprocessing costs comprise a single exponentiation and computing a modulus against the entire file.

E-PDP also exponentiates data that was reduced modulo $\phi(N)$ but does not reap the same speed up, because it must do so for every block. This creates a natural trade-off between preprocessing time and challenge time by varying the block size; *e.g.*, the protocol devolves to B-PDP for files of a single block. Fig. 6 shows this trade-off and indicates that the best balance occurs at natural file system and memory blocks sizes of 4-64 KB. This trade-off may be used to distribute computational across the client and server, depending upon application requirements. We choose a block size of 4K in order to minimize the server's effort.

Given the efficiency of computing challenges, pre-processing represents the limiting performance

factor for E-PDP. The rate at which clients can generate data to outsource bounds the overall system performance perceived by the client. However, there are several mitigating factors. (1) Outsourcing data is a one time task, as compared to challenging outsourced data, which will be done repeatedly. (2) The process is completely parallelizable. Each file can be processed independently at a different processor. A single file can be parallelized trivially if processors share key material.

6 Conclusion

We focused on the problem of verifying if an untrusted server stores a client's data. We introduced a model for provable data possession, in which it is desirable to minimize the file block accesses and the computation on the server, and the client-server communication. Our solutions for PDP fit this model: They incur a low (or even constant) overhead at the server, and require a small, constant amount of communication per challenge. Key components of our schemes are the homomorphic verifiable tags, which allow to verify data possession without having access to the actual data file.

Experiments show that our schemes, which offer a probabilistic possession guarantee by sampling the server's storage, make it practical to verify possession of large data sets. Previous schemes that do not allow sampling are not practical when PDP is used to prove possession of large amounts of data. Our experiments show that such schemes also impose a large computational burden on the server.

7 Acknowledgments

This work was supported in part by NSF awards CCF-0238305 and IIS-0456027 and by the IBM Corporation. The third author would like to thank Răzvan Musăloiu-E. for fruitful discussions in the early stages of this work.

References

- [1] M. Abe and S. Fehr. Perfect NIZK with adaptive soundness. In *Proc. of Theory of Cryptography Conference (TCC '07)*, 2007. Full version available on Cryptology ePrint Archive, Report 2006/423.
- [2] J. Aspnes, J. Feigenbaum, A. Yampolskiy, and S. Zhong. Towards a theory of data entanglement. In *Proc. of Euro. Symp. on Research in Computer Security*, 2004.
- [3] M. Bellare, J. Garay, and T. Rabin. Fast batch verification for modular exponentiation and digital signatures. In *Proc. of EUROCRYPT '98*, LNCS, pages 236–250, 1998.
- [4] M. Bellare and A. Palacio. The knowledge-of-exponent assumptions and 3-round zero-knowledge protocols. In *Proc. of CRYPTO '04*, Lecture Notes in Computer Science, pages 273–289. Springer, 2004.
- [5] M. Bellare and A. Palacio. Towards plaintext-aware public-key encryption without random oracles. In *Proc. of ASIACRYPT '04*, volume 3329 of LNCS, pages 48–62. Springer, 2004.
- [6] M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *First Conference on Computer and Communications Security*, pages 62–73. ACM, 1993.

- [7] M. Bellare and P. Rogaway. The exact security of digital signatures - How to sign with RSA and Rabin. In *EUROCRYPT*, pages 399–416, 1996.
- [8] M. Bellare and P. Rogaway. PSS: Provably secure encoding method for digital signatures. IEEE P1363a: Provably secure signatures, 1998. <http://grouper.ieee.org/groups/1363/P1363a/PSSigs.html>.
- [9] J. Black and P. Rogaway. Ciphers with arbitrary finite domains. In *Proc. of CT-RSA*, volume 2271 of *LNCS*, pages 114–130. Springer-Verlag, 2002.
- [10] M. Blum, W. Evans, P. Gemmell, S. Kannan, and M. Naor. Checking the correctness of memories. In *Proc. of the FOCS '95*, 1995.
- [11] D. Boneh, C. Gentry, B. Lynn, and H. Shacham. Aggregate and verifiably encrypted signatures from bilinear maps. In *Proc. of EUROCRYPT '03*, volume 2656 of *LNCS*, pages 416–432. Springer, 2003.
- [12] I. Damgård. Towards practical public key systems secure against chosen ciphertext attacks. In J. Feigenbaum, editor, *CRYPTO91*, volume 576, pages 445–456. Springer, 1992.
- [13] A. W. Dent. The Cramer-Shoup encryption scheme is plaintext aware in the standard model. In *Proc. of EUROCRYPT '06*, volume 4004 of *LNCS*, pages 289–307. Springer, 2006.
- [14] A. W. Dent. The hardness of the DHK problem in the generic group model. Cryptology ePrint Archive, Report 2006/156, 2006. <http://eprint.iacr.org/2006/156>.
- [15] Y. Deswarte, J.-J. Quisquater, and A. Saidane. Remote integrity checking. In *Proc. of Conference on Integrity and Internal Control in Information Systems (IICIS'03)*, November 2003.
- [16] A. Fiat. Batch RSA. In G. Brassard, editor, *Proc. CRYPTO 89*, pages 175–185. Springer-Verlag, 1990. Lecture Notes in Computer Science No. 435.
- [17] D. L. G. Filho and P. S. L. M. Baretto. Demonstrating data possession and uncheatable data transfer. IACR ePrint archive, 2006. Report 2006/150, <http://eprint.iacr.org/2006/150>.
- [18] P. Golle, S. Jarecki, and I. Mironov. Cryptographic primitives enforcing communication and storage complexity. In *Financial Cryptography*, pages 120–135, 2002.
- [19] S. Hada and T. Tanaka. On the existence of 3-round zero-knowledge protocols. In *Proc. of CRYPTO '98*, volume 1462 of *Lecture Notes in Computer Science*, pages 408–423. Springer, 1998.
- [20] L. Harn. Batch verifying multiple RSA digital signatures. *Electronics Letters*, 34(12):1219–1220, 1998.
- [21] R. Hasan, W. Yurcik, and S. Myagmar. The evolution of storage server providers: Techniques and challenges to outsourcing storage. In *Proc. of the Workshop on Storage Security and Survivability*, 2005.
- [22] R. Johnson, D. Molnar, D. Song, and D. Wagner. Homomorphic signature schemes. In *Proc. of CT-RSA*, volume 2271 of *LNCS*, pages 244–262. Springer, 2002.

- [23] M. Kallahalla, E. Riedel, R. Swaminathan, Q. Wang, and K. Fu. Plutus: Scalable secure file sharing on untrusted storage. In *Proc. of FAST*, 2003.
- [24] H. Krawczyk. HMQV: A high-performance secure Diffie-Hellman protocol. In *Proc. of CRYPTO '05*, volume 3621 of *LNCS*, pages 546–566. Springer, 2005.
- [25] M. N. Krohn, M. J. Freedman, and D. Mazières. On-the-fly verification of rateless erasure codes for efficient content distribution. In *Proceedings of the IEEE Symposium on Security and Privacy*, 2004.
- [26] J. Kubiatowicz, D. Bindel, Y. Chen, P. Eaton, D. Geels, R. Gummadi, S. Rhea, H. Weatherspoon, W. Weimer, C. Wells, and B. Zhao. Oceanstore: An architecture for global-scale persistent storage. In *Proceedings of ACM ASPLOS '00*. ACM, November 2000.
- [27] J. Li, M. Krohn, D. Mazières, and D. Shasha. Secure untrusted data repository (SUNDR). In *Proceedings of the Symposium on Operating Systems Design and Implementation*, 2004.
- [28] U. Maheshwari, R. Vingralek, and W. Shapiro. How to build a trusted database system on untrusted storage. In *Proc. of OSDI*, 2000.
- [29] P. Maniatis, M. Roussopoulos, T. Giuli, D. Rosenthal, M. Baker, and Y. Muliadi. The LOCKSS peer-to-peer digital preservation system. *ACM Transactions on Computing Systems*, 23(1):2–50, 2005.
- [30] S. Micali, K. Ohta, and L. Reyzin. Accountable-subgroup multisignatures: extended abstract. In *Proc of ACM CCS '01*, pages 245–254, 2001.
- [31] G. L. Miller. Riemann’s hypothesis and tests for primality. *JCSS*, 13(3):300–317, 1976.
- [32] A. A. Muthitacharoen, R. Morris., T. M. Gil, and B. Chen. Ivy: A read/write peer-to-peer file system. In *Proceedings of 5th Symposium on Operating Systems Design and Implementation*, 2002.
- [33] E. Mykletun, M. Narasimha, and G. Tsudik. Authentication and integrity in outsourced databases. In *Proceedings of NDSS*. The Internet Society, 2004.
- [34] M. Naor and G. N. Rothblum. The complexity of online memory checking. *ePrint Archive Report*, (2006/091), 2006.
- [35] T. Okamoto. A digital multisignature schema using bijective public-key cryptosystems. *ACM Transactions on Computer Systems*, 6(4):432–441, 1988.
- [36] A. Oprea, M. K. Reiter, and K. Yang. Space-efficient block storage integrity. In *Proc. of NDSS '05*, 2005.
- [37] T. S. J. Schwarz and E. L. Miller. Store, forget, and check: Using algebraic signatures to check remotely administered storage. In *Proceedings of ICDCS '06*. IEEE Computer Society, 2006.
- [38] F. Sebe, A. Martinez-Balleste, Y. Deswarte, J. Domingo-Ferrer, and J.-J. Quisquater. Time-bounded remote file integrity checking. Technical Report 04429, LAAS, July 2004.

- [39] A. Shamir. On the generation of cryptographically strong pseudorandom sequences. *ACM Trans. Comput. Syst.*, 1(1):38–44, 1983.
- [40] D. Thompson and J. Best. The future of magnetic data storage technology. *IBM Journal Research and Development*, 44(3):311–319, May 2000.
- [41] M. Waldman and D. Mazières. Tangler: a censorship-resistant publishing system based on document entanglements. In *Proc. of CCS*, 2001.
- [42] A. Y. Yumerefendi and J. Chase. Strong accountability for network storage. In *Proc. of FAST*, 2007.

A Proof of Theorem 4.3

Under the KEA1-r assumption, we reduce the security of our S-PDP scheme to the security of the RSA problem and the security of integer factoring. We assume there exists an adversary \mathcal{B} that can break the data possession guarantee of the S-PDP scheme with probability greater than $2^{-\alpha} + \epsilon(k)$ (given that he has lost α bits of information about some blocks of the file) and construct an adversary \mathcal{A} that uses \mathcal{B} in order to either break RSA or factor the product of two large primes.

For the RSA problem, \mathcal{A} is given (N, e, y) , with $y \xleftarrow{R} \mathbb{Z}_N^*$, and needs to find an integer $b \equiv y^{1/e} \pmod{N}$. \mathcal{A} will play the role of the challenger in the Data Possession Game and will interact with \mathcal{B} .

We first look at the case when in `GenProof` and `CheckProof` all the coefficients a_1, \dots, a_c are equal to 1. This corresponds to the case where the server proves it possesses the sum of the requested blocks. We then generalize the proof to the case where the coefficients are random and pairwise distinct, which corresponds to the case where the server proves it possesses each individual block.

\mathcal{A} simulates a PDP environment for \mathcal{B} as follows:

Setup: \mathcal{A} computes $g = y^2 \pmod{N}$, sets the public key $\text{pk} = (N, e, g)$ and sends pk to \mathcal{B} . \mathcal{A} also generates $v \xleftarrow{R} \{0, 1\}^\kappa$, which is kept secret.

Query: \mathcal{B} makes tagging queries adaptively: \mathcal{B} selects a block m_1 and is also allowed to select an index i_1 . \mathcal{B} sends m_1 and i_1 to \mathcal{A} . \mathcal{A} runs $(\mathbb{T}_{\mathbb{W}_{i_1}, m_1}, \mathbb{W}_{i_1}) \leftarrow \text{TagBlock}(\text{sk}, m_1, i_1)$ and sends back to \mathcal{B} the tag $(\mathbb{T}_{\mathbb{W}_{i_1}, m_1}, \mathbb{W}_{i_1})$. \mathcal{B} continues to query \mathcal{A} for the tags $(\mathbb{T}_{\mathbb{W}_{i_2}, m_2}, \mathbb{W}_{i_2}), \dots, (\mathbb{T}_{\mathbb{W}_{i_n}, m_n}, \mathbb{W}_{i_n})$ on the blocks m_2, \dots, m_n and indices i_1, \dots, i_n of its choice. The only restriction is that \mathcal{B} cannot make tagging queries for two different blocks using the same index. Once \mathcal{B} gets the tag $(\mathbb{T}_{\mathbb{W}_i, m_i}, \mathbb{W}_i)$ on a block m_i , it may check the tag's validity by checking if the relation $(\mathbb{T}_{\mathbb{W}_i, m_i})^e = h(\mathbb{W}_i) \cdot g^{m_i}$ holds.

When \mathcal{B} makes tagging oracle queries and hash oracle queries, \mathcal{A} will answer those queries itself:

– when \mathcal{B} makes a tagging query for a block m and index i , with $1 \leq i \leq n$:

- if a previous tagging query has been made for the same m and i , then \mathcal{A} retrieves the recorded tuple $(m, i, r_i, \mathbb{W}_i)$ and returns $(\mathbb{T}_{\mathbb{W}_i, m} = r_i, \mathbb{W}_i)$.
- else, \mathcal{A} picks $r_i \xleftarrow{R} QR_N$, computes $\mathbb{W}_i = w_v(i)$, records the tuple $(m, i, r_i, \mathbb{W}_i)$ and returns $(\mathbb{T}_{\mathbb{W}_i, m} = r_i, \mathbb{W}_i)$.

– when \mathcal{B} makes a hash query for a value \mathbb{W}_i :

- if the value \mathbb{W}_i is recorded as having been previously used for a tagging query, then \mathcal{A} retrieves the corresponding tuple $(m, i, r_i, \mathbb{W}_i)$ and returns $h(\mathbb{W}_i) = r_i^e \cdot g^{-m_i} \pmod{N}$.
- else, \mathcal{A} picks at random a value from QR_N and returns it.

Challenge: \mathcal{A} generates the challenge $\text{chal} = (g_s)$, where $g_s = g^s \bmod N$ and $s \xleftarrow{R} \mathbb{Z}_N^*$. \mathcal{A} sends chal to \mathcal{B} .

Forge: \mathcal{B} picks an integer c , with $1 \leq c \leq n$, picks some indices i_1, \dots, i_c , with $1 \leq i_j \leq n$ and generates a proof $\mathcal{V} = (\mathbf{T}, \rho)$, where $\mathbf{T} = \mathbf{T}_{\{\mathbf{w}_{i_1}, \dots, \mathbf{w}_{i_c}\}, m_{i_1} + \dots + m_{i_c}}$, without $\alpha > 0$ bits of information about the blocks m_{i_1}, \dots, m_{i_c} determined by i_1, \dots, i_c . Note that \mathcal{V} is a valid proof that passes $\text{CheckProof}(\text{pk}, \text{sk}, \text{chal}, \mathcal{V})$. \mathcal{B} returns \mathcal{V} to \mathcal{A} and \mathcal{A} checks the validity of \mathcal{V} . Let $M = m_{i_1} + \dots + m_{i_c}$.

As H is a random oracle, with overwhelming probability we can extract the pre-image value ρ_p that \mathcal{B} utilized to calculate ρ . (By the definition of a random oracle, \mathcal{B} can guess a valid value of ρ with only negligible probability.)

\mathcal{A} has given \mathcal{B} both g, g^s and \mathcal{B} has implicitly returned $\tau = \frac{\mathbf{T}^e}{\prod_{j=1}^c h(\mathbf{w}_{i_j})}, \rho_p$ by returning \mathbf{T}, ρ .

Because $\tau^s = \rho_p$, by KEA-1r, \mathcal{A} can utilize the extractor $\bar{\mathbf{B}}$ to extract a value $M^* \neq M$ such that $g^{M^*} = \tau$. Note that M^* is the “full-domain” value utilized by this calculation. (If the extractor $\bar{\mathbf{B}}$ is able to extract a value $M' \neq M^*$ such that $g^{M'} = g^{M^*} \bmod N$, this will allow to compute a multiple of $\phi(N)$, from which the factorization of N can be efficiently computed [31].)

From $\tau = g^{M^*}$ we get $\mathbf{T}^e = \prod_{j=1}^c h(\mathbf{w}_{i_j}) \cdot g^{M^*}$, where clearly $g^{M^*} \neq g^M$, and thus:

$$\begin{aligned} \mathbf{T} &= \left(\prod_{j=1}^c h(\mathbf{w}_{i_j}) \cdot g^{M^*} \right)^d \\ &= \left(\prod_{j=1}^c (r_{i_j}^e \cdot g^{-m_{i_j}}) \cdot g^{M^*} \right)^d \\ &= \prod_{j=1}^c r_{i_j} \cdot (g^{M^* - M})^d \end{aligned}$$

\mathcal{A} computes:

$$z = \frac{\mathbf{T}}{\prod_{j=1}^c r_{i_j}} = (g^{M^* - M})^d$$

We have $z^e = g^{M^* - M} = y^{2(M^* - M)}$, and $\gcd(e, 2(M^* - M)) = 1$ with overwhelming probability (in fact this always holds if e is prime). Applying Shamir’s “trick” [39], \mathcal{A} uses the extended Euclidian algorithm to efficiently compute integers u and v such that $u \cdot e + v \cdot 2(M^* - M) = 1$ and outputs $y^{1/e} = y^u z^v$.

Note that the interactions of \mathcal{A} with \mathcal{B} are indistinguishable to \mathcal{B} from interactions with an honest challenger in the Data Possession Game, as \mathcal{A} chooses all parameters according to our protocol.

The proof generalizes to the case where the coefficients a_1, \dots, a_c are random and pairwise distinct. Indeed, in this case it is enough to apply the same simulation shown above and in addition to notice that at the end of the simulation \mathcal{A} will be able to extract $\bar{M} = a_1 m_{i_1} + \dots + a_c m_{i_c}$. We now have to show that our protocol constitutes a proof of knowledge of the blocks m_{i_1}, \dots, m_{i_c} when a_1, \dots, a_c are pairwise distinct. We show that \mathcal{A} may extract the file blocks m_{i_1}, \dots, m_{i_c} . Note that

each time \mathcal{A} runs the PDP protocol, he obtains a linear equation of the form $\bar{M} = a_1 m_{i_1} + \dots + a_c m_{i_c}$. By choosing independent coefficients a_1, \dots, a_c in c executions of the protocol on the same blocks m_{i_1}, \dots, m_{i_c} , \mathcal{A} obtains c independent linear equations in the variables m_{i_1}, \dots, m_{i_c} . \mathcal{A} may then solve these equations to obtain the file blocks m_{i_1}, \dots, m_{i_c} .

B Implemented PDP Schemes

As a basis of comparison, we have implemented the following two PDP schemes in addition to our E-PDP scheme:

Basic RSA-based PDP Scheme (B-PDP)[15, 17]. Let N be an RSA modulus and let $g \in \mathbb{Z}_N^*$.

Setup: C stores $a = g^F \bmod N$. C sends F to S .

Challenge:

1. C generates a random value $r \xleftarrow{R} \mathbb{Z}_n^*$ and sends g^r to S .
2. S computes $b = (g^r)^F \bmod N$ and sends b to C .
3. C computes a^r and checks if $b = a^r \bmod N$.

Merkle Hash Tree-based Storage Enforcing Scheme (MHT-SC)[18]. Let ψ be a one-way length-increasing transformation and let h be a cryptographic hash function. In a binary tree, we denote by $PATH(i)$ the set of nodes on the path between the root of the tree and the i -th leaf.

Setup: C applies the transformation ψ on the file F and obtains the expanded file $F' = \psi(F)$. C stores h_{root} as the root of the Merkle hash tree³ which is computed using h on the blocks of F' .

Challenge:

1. C randomly picks the index i of a block F'_i of the expanded file F' and sends i to S .
2. Let H_i be the set of hashes corresponding to the nodes that “hang” off $PATH(i)$ in the Merkle hash tree determined by F' . S sends H_i and F'_i back to C .
3. C uses H_i and F'_i to recompute the root of the tree and compares this value with the stored value for h_{root} .

³Merkle hash trees compute a parent node by applying a hash function on the concatenation of its children nodes.