

Compact Proofs of Retrievability

Hovav Shacham*
hovav@cs.ucsd.edu

Brent Waters†
bwaters@cs.utexas.edu

Abstract

In a proof-of-retrievability system, a data storage center must prove to a verifier that he is actually storing all of a client’s data. The central challenge is to build systems that are both efficient and *provably* secure—that is, it should be possible to extract the client’s data from any prover that passes a verification check. In this paper, we give the first proof-of-retrievability schemes with full proofs of security against arbitrary adversaries in the strongest model, that of Juels and Kaliski.

Our first scheme, built from BLS signatures and secure in the random oracle model, features a proof-of-retrievability protocol in which the client’s query and server’s response are both extremely short. This scheme allows public verifiability: anyone can act as a verifier, not just the file owner. Our second scheme, which builds on pseudorandom functions (PRFs) and is secure in the standard model, allows only private verification. It features a proof-of-retrievability protocol with an even shorter server’s response than our first scheme, but the client’s query is long. Both schemes rely on homomorphic properties to aggregate a proof into one small authenticator value.

1 Introduction

Proofs of storage. Recent visions of “cloud computing” and “software as a service” call for data, both personal and commercial, to be stored by third parties, but deployment has lagged. Users of outsourced storage are at the mercy of their storage providers for the continued availability of their data. Even Amazon’s S3, the best-known storage service, has experienced significant downtime.

The solution, as Shah et al. argue [30], is storage auditing: cryptographic systems that would allow users of outsourced storage services (or their agents) to verify that their data is still available and ready for retrieval if needed. Such a capability can be important to storage providers as well. Users may be reluctant to entrust their data to an unknown startup; an auditing mechanism can reassure them that their data is indeed still available.

Early proof-of-storage systems were proposed by Deswarte, Quisquater, and Saïdane [14], Gazzoni Filho and Barreto [17], and Schwarz and Miller [29].

Evaluation: formal security models. Such proof-of-storage systems should be evaluated by both “systems” and “crypto” criteria. Systems criteria include: (1) the system should be as *efficient*

*Supported by the MURI program under AFOSR Grant No. FA9550-08-1-0352 and by the UCSD Center for Networked Systems

†Supported by NSF CNS-0749931, CNS-0524252, CNS-0716199; the US Army Research Office under the CyberTA Grant No. W911NF-06-1-0316; and the U.S. Department of Homeland Security under Grant Award Number 2006-CS-001-000001.

as possible in terms of both computational complexity and communication complexity of the proof-of-storage protocol, and the storage overhead on the server should be as small as possible; (2) the system should allow *unbounded use* rather than imposing a priori bound on the number of audit-protocol interactions; (3) verifiers should be *stateless*, and not need to maintain and update state between audits, since such state is difficult to maintain if the verifier’s machine crashes or if the verifier’s role is delegated to third parties or distributed among multiple machines.¹ Statelessness and unbounded use are required for proof-of-storage systems with *public verifiability*, in which anyone can undertake the role of verifier in the proof-of-storage protocol, not just the user who originally stored the file. Public verifiability for proof-of-storage schemes was first proposed by Ateniese et al. [3].

The most important crypto criterion is this: Whether the protocol actually establishes that any server that passes a verification check for a file—even a malicious server that exhibits arbitrary, Byzantine behavior—is *actually storing the file*. The early cryptographic papers lacked a formal security model, let alone proofs. But provable security matters. Even reasonable-looking protocols could in fact be insecure; see Appendix B for an example.

The first papers to consider formal models for proofs of storage were by Naor and Rothblum, for “authenticators” [26], and by Juels and Kaliski, for “proofs of retrievability” [22]. Though the details of the two models are different, the insight behind both is the same: in a secure system if a server can pass an audit then a special extractor algorithm, interacting with the server, must be able (w.h.p.) to extract the file. This is, of course, similar to the intuition behind proofs of knowledge.

A simple MAC-based construction. In addition, the Naor-Rothblum and Juels-Kaliski papers describe similar proof-of-retrievability protocols. The insight behind both is that checking that *most* of a file is stored is easier than checking that *all* is. If the file to be stored is first encoded redundantly, and each block of the encoded file is authenticated using a MAC, then it is sufficient for the client to retrieve a few blocks together with their MACs and check, using his secret key, that these blocks are correct. Naor and Rothblum prove their scheme secure in their model. Juels and Kaliski do not give a proof of security against arbitrary adversaries, but this proof can be done straightforwardly using the techniques we develop in this paper; for completeness, we give the proof in Section 5. The simple protocol obtained here uses techniques similar to those proposed by Lillibridge et al. [23]. Signatures can be used instead of MACs to obtain public verifiability.

The downside to this simple solution is that the server’s response consists of λ block-authenticator pairs, where λ is the security parameter. If each authenticator is λ bits long, as required in the Juels-Kaliski model, then the response is $\lambda^2 \cdot (s+1)$ bits, where the ratio of file block to authenticator length is $s : 1$.²

Homomorphic authenticators. Ateniese et al. [3] describe a proof-of-storage scheme that improves on the response length of the simple MAC-based scheme using *homomorphic authenticators*. In their scheme, the authenticators σ_i on each file block m_i are constructed in such a way that

¹We note that the sentinel-based scheme of Juels and Kaliski [22], the scheme of Ateniese, Di Pietro, Mancini, and Tsudik [4], and the scheme of Shah, Swaminathan and Baker [31] lack both unbounded use and statelessness. We do not consider these schemes further in this paper.

²Naor and Rothblum show that one-bit MACs suffice for proving security in their less stringent model, for an overall response length of $\lambda \cdot (s + 1)$ bits. The Naor-Rothblum scheme is not secure in the Juels-Kaliski model.

a verifier can be convinced that a linear combination of blocks $\sum_i \nu_i m_i$ (with arbitrary weights $\{\nu_i\}$) was correctly generated using an authenticator computed from $\{\sigma_i\}$. In the Ateniese et al. construction, for example, the aggregate authenticator is $\prod_i \sigma_i^{\nu_i} \bmod N$.

When using homomorphic authenticators, the server can combine the blocks and λ authenticators in its response into a single aggregate block and authenticator, reducing the response length by a factor of λ . As an additional benefit, the Ateniese et al. scheme is the first with public verifiability. The homomorphic authenticators of Ateniese et al. are based on RSA and are thus relatively long.

Unfortunately, Ateniese et al. do not give a rigorous proof of security for their scheme. In particular, they do not show that one can extract a file (or even a significant fraction of one) from a prover that is able to answer auditing queries convincingly. The need for rigor in extraction arguments applies equally to both the proof-of-retrievability model we consider and the weaker proof of data possession model considered by Ateniese et al. For completeness, we give a correct and fully proven Ateniese-et-al.-inspired, RSA-based scheme, together with a full proof of security, in Section 6.

Our contributions. In this paper, we make two contributions.

1. We describe two new short, efficient homomorphic authenticators. The first, based on PRFs, gives a proof-of-retrievability scheme secure in the standard model. The second, based on BLS signatures [9], gives a proof-of-retrievability scheme with public verifiability secure in the random oracle model.
2. We prove both of the resulting schemes secure in a variant of the Juels-Kaliski model. Our schemes are the first with a security proof against arbitrary adversaries in this model.

The scheme with public retrievability features a proof-of-retrievability protocol in which the client’s query and server’s response are both extremely short: 20 bytes and 40 bytes, respectively, at the 80-bit security level. The scheme with private retrievability features a proof-of-retrievability protocol with an even shorter server’s response than our first scheme: 20 bytes at the 80-bit security level, matching the response length of the Naor-Rothblum scheme in a more stringent security model, albeit at the cost of a longer query.

1.1 Our Schemes

In our schemes, the user breaks an erasure encoded file into n blocks $m_1, \dots, m_n \in \mathbb{Z}_p$ for some large prime p . The erasure code should allow decoding in the presence of *adversarial* erasure. Erasure codes derived from Reed-Solomon codes have this property, but decoding and encoding are slow for large files. In Appendix A we discuss how to make use of more efficient codes secure only against random erasures.

The user authenticates each block as follows. She chooses a random $\alpha \in \mathbb{Z}_p$ and PRF key k for function f . These values serve as her secret key. She calculates an authentication value for each block i as

$$\sigma_i = f_k(i) + \alpha m_i \in \mathbb{Z}_p .$$

The blocks $\{m_i\}$ and authenticators $\{\sigma_i\}$ are stored on the server. The proof of retrievability protocol is as follows. The verifier chooses a random challenge set I of l indices along with l random coefficients in \mathbb{Z}_p .³ Let Q be the set $\{(i, \nu_i)\}$ of challenge index-coefficient pairs. The verifier sends

³Or, more generally, from a subset B of \mathbb{Z}_p of appropriate size; see Section 1.1.

Q to the prover. The prover then calculates the response, a pair (σ, μ) , as

$$\sigma \leftarrow \sum_{(i, \nu_i) \in Q} \nu_i \cdot \sigma_i \quad \text{and} \quad \mu \leftarrow \sum_{(i, \nu_i) \in Q} \nu_i \cdot m_i .$$

Now verifier can check that the response was correctly formed by checking that

$$\sigma \stackrel{?}{=} \alpha \cdot \mu + \sum_{(i, \nu_i) \in Q} \nu_i \cdot f_k(i) .$$

It is clear that our techniques admit short responses. But it is not clear that our new system admits a simulator that can extract files. Proving that it does takes some work, as we discuss below. In fact, unlike similar, seemingly correct schemes (see Appendix B), our scheme is provably secure. Moreover, our proofs are in the standard model.

A scheme with public verifiability. Our second scheme is publicly verifiable. It follows the same framework as the first, but instead uses BLS signatures [9] for authentication values that can be publicly verified. The structure of these signatures allows for them to be aggregated into linear combinations as above. We prove the security of this scheme under the Computational Diffie-Hellman assumption over bilinear groups in the random oracle model.

Let $e: G \times G \rightarrow G_T$ be a computable bilinear map with group G 's support being \mathbb{Z}_p . A user's private key is $x \in \mathbb{Z}_p$, and her public key is $v = g^x \in G$ along with another generator $u \in G$. The signature on block i is $\sigma_i = [H(i)u^{m_i}]^x$. On receiving query $Q = \{(i, \nu_i)\}$, the prover computes and sends back $\sigma \leftarrow \prod_{(i, \nu_i) \in Q} \sigma_i^{\nu_i}$ and $\mu \leftarrow \sum_{(i, \nu_i) \in Q} \nu_i \cdot m_i$. The verification equation is:

$$e(\sigma, g) \stackrel{?}{=} e\left(\prod_{(i, \nu_i) \in Q} H(i)^{\nu_i} \cdot u^\mu, v\right) .$$

This scheme has public verifiability: the private key x is required for generating the authenticators $\{\sigma_i\}$ but the *public* key v is sufficient for the verifier in the proof-of-retrievability protocol. As we note below, the query can be generated from a short seed using a random oracle, and this short seed can be transmitted instead of the longer query.

Parameter selection. Let λ be the security parameter; typically, $\lambda = 80$. For the scheme with private verification, p should be a λ -bit prime. For the scheme with public verification, p should be a 2λ -bit prime, and the curve should be chosen so that discrete logarithm is 2^λ -secure. For values of λ up to 128, Barreto-Naehrig curves [6] are the right choice; see the survey by Freeman, Scott, and Teske [16].

Let n be the number of blocks in the file. We assume that $n \gg \lambda$. Suppose we use a rate- ρ erasure code, i.e., one in which any ρ -fraction of the blocks suffices for decoding. (Encoding will cause the file length to grow approximately $(1/\rho) \times$.) Let l be the number of indices in the query Q , and $B \subseteq \mathbb{Z}_p$ be the set from which the challenge weights ν_i are drawn.

Our proofs—see Section 4.2 for the details—guarantee that extraction will succeed from any adversary that convincingly answers an ϵ -fraction of queries, provided that $\epsilon - \rho^l - 1/\#B$ is non-negligible in λ . It is this requirement that guides the choice of parameters.

A conservative choice is $\rho = 1/2$, $l = \lambda$, and $B = \{0, 1\}^\lambda$; this guarantees extraction against any adversary. For applications that can tolerate a larger error rate these parameters can be reduced. For example, if a 1-in-1,000,000 error is acceptable, we can take B to be the set of 22-bit strings and l to be 22; alternatively, the coding expansion $1/\rho$ can be reduced.

A tradeoff between storage and communication. As we have described our schemes above, each file block is accompanied by an authenticator of equal length. This gives a $2\times$ overhead beyond that imposed by the erasure code, and the server’s response in the proof-of-retrievability protocol is $2\times$ the length of an authenticator. In the full schemes of Section 3, we introduce a parameter s that gives a tradeoff between storage overhead and response length. Each block consists of s elements of \mathbb{Z}_p that we call *sectors*. There is one authenticator per block, reducing the overhead to $(1 + 1/s)\times$. The server’s response is one aggregated block and authenticator, and is $(1 + s)\times$ as long as an authenticator. Thus, a larger value of s gives less storage overhead at the cost of higher communication. The choice $s = 1$ corresponds to our schemes as we described them above and to the scheme given by Ateniese et al. [3].⁴

Compressing the request. A request, as we have seen, consists of an l element subset of $[1, n]$ together with l elements of the coefficient set B , chosen uniformly and independently at random. In the conservative parametrization above, a request is thus $\lambda \cdot (\lceil \lg n \rceil + \lambda)$ bits long. One can reduce the randomness required to *generate* the request using standard techniques⁵. This would reduce the size of the client’s request only if the PRF keys were sent in place of the computed PRF output, but we do not know how to prove this method secure in the standard model. By contrast, in the random oracle model, the verifier *can* send a short (2λ bit) seed for the random oracle from which the prover will generate the full query. Using this technique we can make the queries compact as well as responses in our publicly verifiable scheme, which already relies on random oracles. Apply the same trick to our PRF-based scheme would introduce a reliance on the random oracle heuristic.

We note that, by techniques similar to those discussed above, a PRF can be used to generate the per-file secret values $\{\alpha_j\}$ for our privately verifiable scheme and a random oracle seed can be used to generate the per-file public generators $\{u_j\}$ in our publicly verifiable scheme. This allows file tags for both schemes to be short: $O(\lambda)$, asymptotically.

Followup work. The major problem left open by our work was to obtain short queries for schemes whose analysis does not rely on the random oracle heuristic. Dodis, Vadhan, and Wichs [15] and, independently, Bowers, Juels, and Oprea [10] both observed that the “ B coefficients” in our queries are a Hadamard code in disguise, and that more efficient error-correcting codes can substantially reduce the query size. In addition, Dodis, Vadhan, and Wichs showed that it is possible to reduce the query size further using hitting samplers [18]. The result is a proof-of-retrievability protocol that is essentially optimal, with query and response size both linear in the security parameter, with a formal proof in the standard model.

Ateniese, Kamara, and Katz [5] gave a framework for constructing a proof-of-retrievability scheme with public verifiability (in the random oracle model) from any homomorphic identification protocol. They showed how to fit the scheme of Ateniese et al. [3] (based on the RSA problem; see Section 6) and our publicly verifiable scheme (based on the Diffie-Hellman problem in bilinear groups; see Section 3.3) into their framework, and gave a new instantiation based on the factoring problem.

⁴It would be possible to shorten the response further using knowledge-of-exponent assumptions, as Ateniese et al. do, but such assumptions are strong and nonstandard; more importantly, their use means that the extractor can never be implemented in the real world.

⁵For example, choose keys k' and k'' for PRFs with respective ranges $[1, n]$ and B . The query indices are the first l *distinct* values amongst $f_{k'}(1), f_{k'}(2), \dots$; the query coefficients are the l values $f_{k''}(1), \dots, f_{k''}(l)$, not necessarily all distinct.

1.2 Our Proofs

We provide a modular proof framework for the security of our schemes. Our framework allows us to argue about the systems’ unforgeability, extractability, and retrievability with these three parts based respectively on cryptographic, combinatorial, and coding-theoretical techniques. Only the first part differs between the four schemes we propose. The combinatorial techniques we develop are nontrivial and we believe they will be of independent interest.

It is interesting to compare both our security model and our proof methodology to those in related work.

The proof of retrievability model has two major distinctions from that used by Naor and Rothblum [26] (in addition to the public-key setting). First, the NR model assumes a checker can request and receive specific memory locations from the prover. In the proof of retrievability model, the prover consists of an arbitrary program as opposed to a simple memory layout and this program may answer these questions in an arbitrary manner. We believe that this realistically represents an adversary in the type of setting we are considering. In the NR setting the extractor needs to retrieve the file given the server’s memory; in the POR setting the analogy is that the extractor receives the adversary’s program.

Second, in the proof of retrievability model we allow the attacker to execute a polynomial number of proof attempts before committing to how it will store memory. In the NR model the adversary does not get to execute the protocol before committing its memory. This weaker model is precisely what allows for the use of 1-bit MACs with error correcting codes in one NR variant. One might argue that in many situations this is sufficient. If a storage server responds incorrectly to an audit request we might assume that it is declared to be cheating and there is no need to go further. However, this limited view overlooks several scenarios. In particular, we want to be able to handle setups where there are several verifiers that do not communicate or several storage servers handling the same encoded file that are audited independently. Only our stronger model can correctly reflect these situations. In general, we believe that the strongest security model allows for a system to be secure in the most contexts including those not previously considered.

One of the distinctive and challenging parts of our work is to argue extraction from homomorphically accumulated blocks. Extractability issues arise in several natural constructions. Proving extraction from aggregated authenticator values can be challenging. In Appendix B we show an attack on a natural but incorrect system that is very similar to the “E-PDP” efficient alternative scheme given by Ateniese et al. For their E-PDP scheme, Ateniese et al. claim only that the protocol establishes that a cheating prover has the sum $\sum_{i \in I} m_i$ of the blocks. Our attack suggests that this guarantee is insufficient for recovering file contents.

Finally, we argue that the POR is the “right” model for considering practical data storage problems, since it provides a successful audit guarantees that *all* the data can be extracted. Other work has advocated for a weaker model, Proof of Data Possession [3]. In this model, one only wants to guarantee that a certain percentage (e.g., 90%) of data blocks are available. By offering this weaker guarantee one might hope to avoid the overhead of applying erasure codes. However, this weaker condition is unsatisfactory for most practical application demands. One might consider how happy a user would be were 10% of an accounting data file lost. Or if, for a compressed file, the compression tables were lost — and with them all useful data. Instead of hoping that there is enough redundancy left to reconstruct important data in an ad-hoc way, it is much more desirable to have a model that inherently provides this.

In another difference from previous work, we insist that files be recoverable from an adversary

that correctly answers any small (but nonnegligible) fraction of queries. We believe that this frees systems implementers from having to worry about whether a substantial error rate (for example, due to an intermittent connection between auditor and server) invalidates the assumptions of the underlying cryptographic protocol.

Our proofs show that a system that provably allows recovery of a constant fraction of file blocks gives a secure POR scheme when combined with a suitable erasure code; the question is whether the erasure coding can be omitted. We believe that provable full retrievability is crucial, especially when when cryptographic storage is one building block in a larger system.

2 Security Model

We recall the security definition of Juels and Kaliski [22]. Our version differs from the original definition in several details:

- we rule out any state (“ α ”) in key generation and in verification, because (as explained in Section 1) we believe that verifiers in proof-of-retrievability schemes should be stateless;
- we allow the proof protocol to be arbitrary, rather than two-move, challenge-response; and
- our key generation emits a public key as well as a private key, to allow us to capture the notion of public verifiability.

Note that any stateless scheme secure in the original Juels-Kaliski model will be secure in our variant, and any scheme secure in our variant whose proof protocol can be cast as two-move, challenge-response protocol will be secure in the Juels-Kaliski definition. In particular, our scheme with private verifiability is secure in the original Juels-Kaliski model.⁶

A proof of retrievability scheme defines four algorithms, Kg , St , \mathcal{V} , and \mathcal{P} , which behave thus:

Kg(\cdot). This randomized algorithm generates a public-private keypair (pk, sk) .

St(sk, M). This randomized file-storing algorithm takes a secret key sk and a file $M \in \{0, 1\}^*$ to store. It processes M to produce and output M^* , which will be stored on the server, and a tag τ . The tag contains information that names the file being stored; it could also contain additional secret information encrypted under the secret key sk .

\mathcal{P} , \mathcal{V} . The randomized proving and verifying algorithms define a protocol for proving file retrievability. During protocol execution, both algorithms take as input the public key pk and the file tag τ output by St . The prover algorithm also takes as input the processed file description M^* that is output by St , and the verifier algorithm takes as input the secret key. At the end of the protocol run, \mathcal{V} outputs 0 or 1, where 1 means that the file is being stored on the server. We can denote a run of two machines executing the algorithms as: $\{0, 1\} \stackrel{\text{R}}{\leftarrow} (\mathcal{V}(pk, sk, \tau) \Rightarrow \mathcal{P}(pk, \tau, M^*))$.

⁶In an additional minor difference, we do not specify the extraction algorithm as part of a scheme, because we do not expect that the extract algorithm will be deployed in outsourced storage applications. Nevertheless, the extract algorithm used in our proofs (cf. Section 4.2) is quite simple: undertake many random \mathcal{V} interactions with the cheating prover; keep track of those queries for which \mathcal{V} accepts the cheating prover’s reply as valid; and continue until enough information has been gathered to recover file blocks by means of linear algebra. The adversary \mathcal{A} could implement this algorithm by means of its proof-of-retrievability protocol access.

We would like a proof-of-retrievability protocol to be correct and sound. Correctness requires that, for all keypairs (pk, sk) output by Kg , for all files $M \in \{0, 1\}^*$, and for all (M^*, τ) output by $\text{St}(sk, M)$, the verification algorithm accepts when interacting with the valid prover:

$$(\mathcal{V}(pk, sk, \tau) \Rightarrow \mathcal{P}(pk, \tau, M^*)) = 1 .$$

A proof-of-retrievability protocol is sound if any cheating prover that convinces the verification algorithm that it is storing a file M is actually storing that file, which we define in saying that it yields up the file M to an extractor algorithm that interacts with it using the proof-of-retrievability protocol. We formalize the notion of an extractor and then give a precise definition for soundness.

An extractor algorithm $\text{Extr}(pk, sk, \tau, \mathcal{P}')$ takes the public and private keys, the file tag τ , and the description of a machine implementing the prover's role in the proof-of-retrievability protocol: for example, the description of an interactive Turing machine, or of a circuit in an appropriately augmented model. The algorithm's output is the file $M \in \{0, 1\}^*$. Note that Extr is given non-black-box access to \mathcal{P}' and can, in particular, rewind it. The extraction algorithm must be efficient: It must run in time polynomial in n and $(1/\epsilon)$. In an asymptotic formalization, Extr 's running time must also be polynomial in the security parameter λ .

Consider the following setup game between an adversary \mathcal{A} and an environment:

1. The environment generates a keypair (pk, sk) by running Kg , and provides pk to \mathcal{A} .
2. The adversary can now interact with the environment. It can make queries to a store oracle, providing, for each query, some file M . The environment computes $(M^*, \tau) \stackrel{\text{R}}{\leftarrow} \text{St}(sk, M)$ and returns both M^* and τ to the adversary.
3. For any M on which it previously made a store query, the adversary can undertake executions of the proof-of-retrievability protocol, by specifying the corresponding tag τ . In these protocol executions, the environment plays the part of the verifier and the adversary plays the part of the prover: $\mathcal{V}(pk, sk, \tau) \Rightarrow \mathcal{A}$. When a protocol execution completes, the adversary is provided with the output of \mathcal{V} . These protocol executions can be arbitrarily interleaved with each other and with the store queries described above.
4. Finally, the adversary outputs a challenge tag τ returned from some store query, and the description of a prover \mathcal{P}' .

The cheating prover \mathcal{P}' is ϵ -admissible if it convincingly answers an ϵ fraction of verification challenges, i.e., if $\Pr[(\mathcal{V}(pk, sk, \tau) \Rightarrow \mathcal{P}') = 1] \geq \epsilon$. Here the probability is over the coins of the verifier and the prover. Let M be the message input to the store query that returned the challenge tag τ (along with a processed version M^* of M).

Definition 2.1. We say a proof-of-retrievability scheme is ϵ -sound if there exists an efficient extraction algorithm Extr such that, for every adversary \mathcal{A} , whenever \mathcal{A} , playing the setup game, outputs an ϵ -admissible cheating prover \mathcal{P}' for a file M , the extraction algorithm recovers M from \mathcal{P}' —i.e., $\text{Extr}(pk, sk, \tau, \mathcal{P}') = M$ —except possibly with negligible probability.

Note that it is okay for \mathcal{A} to have engaged in the proof-of-retrievability protocol for M in its interaction with the environment. Note also that each run of the proof-of-retrievability protocol is independent: the verifier implemented by the environment is stateless.

Finally, note that we require that extraction succeed (with all but negligible probability) from an adversary that causes \mathcal{V} to accept with any nonnegligible probability ϵ . An adversary that passes the verification even a very small but nonnegligible fraction of the time—say, once in a million interactions—is fair game. Intuitively, recovering enough blocks to reconstruct the original file from such an adversary should take $O(n/\epsilon)$ interactions; our proofs achieve essentially this bound.

Concrete or asymptotic formalization. A proof-of-retrievability scheme is secure if no efficient algorithm wins the game above except rarely, where the precise meaning of “efficient” and “rarely” depends on whether we employ a concrete or asymptotic formalization.

In a concrete formalization, we require that each algorithm defining the proof-of-retrievability scheme run in at most some number of steps, and that for any algorithm \mathcal{A} that runs in time t steps, that makes at most q_s store queries, and that undertakes at most q_p proof-of-retrievability protocol executions, extraction from an ϵ -admissible prover succeeds except with some small probability δ . In an asymptotic formalization, every algorithm is provided with an additional parameter 1^λ for security parameter λ , we require each algorithm to run in time polynomial in λ , and we require that extraction fail from an ϵ -admissible prover with only negligible probability in λ , provided ϵ is nonnegligible.

Public or private verification, public or private extraction. In the model above, the verifier and extractor are provided with a secret that is not known to the prover or other parties. This is a secret-verification, secret-extraction model. If the verification algorithm does not use the secret key, any third party can check that a file is being stored, giving public verification. Similarly, if the extraction algorithm does not use the secret key, any third party can extract the file from a server, giving public extraction.

3 Constructions

In this section we give formal descriptions for both our private and public verification systems. The systems here follow the constructions outlined in the introduction with a few added generalizations. First, we allow blocks to contain $s \geq 1$ elements of \mathbb{Z}_p . This allows for a tradeoff between storage overhead and communication overhead. Roughly the communication complexity grows as $s + 1$ elements of \mathbb{Z}_p and the ratio of authentication overhead to data stored (post encoding) is $1 : s$. Second, we describe our systems where the set of coefficients sampled from B can be smaller than all of \mathbb{Z}_p . This enables us to obtain more efficient systems in certain situations.

3.1 Common Notation

We will work in the group \mathbb{Z}_p . When we work in the bilinear setting, the group \mathbb{Z}_p is the support of the bilinear group G , i.e., $\#G = p$. In queries, coefficients will come from a set $B \subseteq \mathbb{Z}_p$. For example, B could equal \mathbb{Z}_p , in which case query coefficients will be randomly chosen out of all of \mathbb{Z}_p .

After a file undergoes preliminary processing, the processed file is split into *blocks*, and each block is split into *sectors*. Each sector is one element of \mathbb{Z}_p , and there are s sectors per block. If the processed file is b bits long, then there are $n = \lceil b/s \lg p \rceil$ blocks. We will refer to individual file sectors as $\{m_{ij}\}$, with $1 \leq i \leq n$ and $1 \leq j \leq s$.

Queries. A query is an l -element set $Q = \{(i, \nu_i)\}$. Each entry $(i, \nu_i) \in Q$ is such that i is a block index in the range $[1, n]$, and ν_i is a multiplier in B . The size l of Q is a system parameter, as is the choice of the set B .

The verifier chooses a random query as follows. First, she chooses, uniformly at random, an l -element subset I of $[1, n]$. Then, for each element $i \in I$ she chooses, uniformly at random, an element $\nu_i \xleftarrow{R} B$. We observe that this procedure implies selection of l elements from $[1, n]$ *without* replacement but a selection of l elements from B *with* replacement.

Although the set notation $Q = \{(i, \nu_i)\}$ is space-efficient and convenient for implementation, we will also make use of a vector notation in the analysis. A query Q over indices $I \subset [1, n]$ is represented by a vector $\mathbf{q} \in (\mathbb{Z}_p)^n$ where $\mathbf{q}_i = \nu_i$ for $i \in I$ and $\mathbf{q}_i = 0$ for all $i \notin I$. Equivalently, letting $\mathbf{u}_1, \dots, \mathbf{u}_n$ be the usual basis for $(\mathbb{Z}_p)^n$, we have $\mathbf{q} = \sum_{(i, \nu_i) \in Q} \nu_i \mathbf{u}_i$.⁷

If the set B does not contain 0 then a random query (according to the selection procedure defined above) is a random weight- l vector in $(\mathbb{Z}_p)^n$ with coefficients in B . If B does contain 0, then a similar argument can be made, but care must be taken to distinguish the case “ $i \in I$ and $\nu_i = 0$ ” from the case “ $i \notin I$.”

Aggregation. For its response, the server responds to a query Q by computing, for each j , $1 \leq j \leq s$, the value

$$\mu_j \leftarrow \sum_{(i, \nu_i) \in Q} \nu_i m_{ij} .$$

That is, by combining sectorwise the blocks named in Q , each with its multiplier ν_i . Addition, of course, is modulo p . The response is $(\mu_1, \dots, \mu_s) \in (\mathbb{Z}_p)^s$.

Suppose we view the message blocks on the server as an $n \times s$ element matrix $M = (m_{ij})$, then, using the vector notation for queries given above, the server’s response is given by $\mathbf{q}M$.

3.2 Construction for Private Verification

Let $f: \{0, 1\}^* \times \mathcal{K}_{\text{prf}} \rightarrow \mathbb{Z}_p$ be a PRF.⁸ The construction of the private verification scheme Priv is:

Priv.Kg(). Choose a random symmetric encryption key $k_{\text{enc}} \xleftarrow{R} \mathcal{K}_{\text{enc}}$ and a random MAC key $k_{\text{mac}} \xleftarrow{R} \mathcal{K}_{\text{mac}}$. The secret key is $sk = (k_{\text{enc}}, k_{\text{mac}})$; there is no public key.

Priv.St(sk, M). Given the file M , first apply the erasure code to obtain M' ; then split M' into n blocks (for some n), each s sectors long: $\{m_{ij}\}_{\substack{1 \leq i \leq n \\ 1 \leq j \leq s}}$. Now choose a PRF key $k_{\text{prf}} \xleftarrow{R} \mathcal{K}_{\text{prf}}$ and s random numbers $\alpha_1, \dots, \alpha_s \xleftarrow{R} \mathbb{Z}_p$. Let τ_0 be $n \|\text{Enc}_{k_{\text{enc}}}(k_{\text{prf}} \|\alpha_1\| \dots \|\alpha_s)\|$; the file tag is $\tau = \tau_0 \|\text{MAC}_{k_{\text{mac}}}(\tau_0)\|$. Now, for each i , $1 \leq i \leq n$, compute

$$\sigma_i \leftarrow f_{k_{\text{prf}}}(i) + \sum_{j=1}^s \alpha_j m_{ij} .$$

The processed file M^* is $\{m_{ij}\}$, $1 \leq i \leq n$, $1 \leq j \leq s$ together with $\{\sigma_i\}$, $1 \leq i \leq n$.

⁷We are using subscripts to denote vector elements (for \mathbf{q}) and to choose a particular vector from a set (for \mathbf{u}); but no confusion should arise.

⁸In fact, the domain need only be $\lceil \lg N \rceil$ -bit strings, where N is a bound on the number of blocks in a file.

Priv. $\mathcal{V}(pk, sk, \tau)$. Parse sk as $(k_{\text{enc}}, k_{\text{mac}})$. Use k_{mac} to verify the MAC on τ ; if the MAC is invalid, reject by emitting 0 and halting. Otherwise, parse τ and use k_{enc} to decrypt the encrypted portions, recovering n , k_{prf} , and $\alpha_1, \dots, \alpha_s$. Now pick a random l -element subset I of the set $[1, n]$, and, for each $i \in I$, a random element $\nu_i \xleftarrow{R} B$. Let Q be the set $\{(i, \nu_i)\}$. Send Q to the prover.

Parse the prover's response to obtain μ_1, \dots, μ_s and σ , all in \mathbb{Z}_p . If parsing fails, fail by emitting 0 and halting. Otherwise, check whether

$$\sigma \stackrel{?}{=} \sum_{(i, \nu_i) \in Q} \nu_i f_{k_{\text{prf}}}(i) + \sum_{j=1}^s \alpha_j \mu_j ;$$

if so, output 1; otherwise, output 0.

Priv. $\mathcal{P}(pk, \tau, M^*)$. Parse the processed file M^* as $\{m_{ij}\}$, $1 \leq i \leq n$, $1 \leq j \leq s$, along with $\{\sigma_i\}$, $1 \leq i \leq n$. Parse the message sent by the verifier as Q , an l -element set $\{(i, \nu_i)\}$, with the i 's distinct, each $i \in [1, n]$, and each $\nu_i \in B$. Compute

$$\mu_j \leftarrow \sum_{(i, \nu_i) \in Q} \nu_i m_{ij} \quad \text{for } 1 \leq j \leq s, \quad \text{and} \quad \sigma \leftarrow \sum_{(i, \nu_i) \in Q} \nu_i \sigma_i .$$

Send to the prover in response the values μ_1, \dots, μ_s and σ .

A note on the field \mathbb{Z}_p . In our description, we specified that the output of the PRF and the size of the file sectors $\{m_{ij}\}$ be \mathbb{Z}_p for a prime p . In fact, any finite field will do, and \mathbb{F}_{2^k} may be a more convenient choice for some implementations.

Correctness. It is easy to see that the scheme is correct. Let the PRF key be k_{prf} and the secret coefficients be $\alpha_1, \dots, \alpha_s \xleftarrow{R} \mathbb{Z}_p$. Let the file sectors be $\{m_{ij}\}$, so that the block authenticators are $\sigma_i = f_{k_{\text{prf}}}(i) + \sum_{j=1}^s \alpha_j m_{ij}$. For a prover who responds honestly to a query $\{(i, \nu_i)\}$, so that each $\mu_j = \sum_{(i, \nu_i) \in Q} \nu_i m_{ij}$ and $\sigma = \sum_{(i, \nu_i) \in Q} \nu_i \sigma_i$, we have

$$\begin{aligned} \sigma &= \sum_{(i, \nu_i) \in Q} \nu_i \sigma_i = \sum_{(i, \nu_i) \in Q} (\nu_i) \left(f_{k_{\text{prf}}}(i) + \sum_{j=1}^s \alpha_j m_{ij} \right) = \sum_{(i, \nu_i) \in Q} \nu_i f_{k_{\text{prf}}}(i) + \sum_{(i, \nu_i) \in Q} \nu_i \sum_{j=1}^s \alpha_j m_{ij} \\ &= \sum_{(i, \nu_i) \in Q} \nu_i f_{k_{\text{prf}}}(i) + \sum_{j=1}^s \alpha_j \sum_{(i, \nu_i) \in Q} \nu_i m_{ij} = \sum_{(i, \nu_i) \in Q} \nu_i f_{k_{\text{prf}}}(i) + \sum_{j=1}^s \alpha_j \mu_j , \end{aligned}$$

So the verification equation is satisfied.

3.3 Construction for Public Verification

Let $e: G \times G \rightarrow G_T$ be a bilinear map, let g be a generator of G , and let $H: \{0, 1\}^* \rightarrow G$ be the BLS hash, treated as a random oracle.⁹ The construction of the public verification scheme Pub is:

⁹For notational simplicity, we present our scheme using a symmetric bilinear map, but efficient implementations will use an asymmetric map $e: G_1 \times G_2 \rightarrow G_T$. Translating our scheme to this setting is simple. User public keys v will live in G_2 ; file generators u_j will live in G_1 , as will the output of H ; and security will be reduced to co-CDH [9].

Pub.Kg(). Generate a random signing keypair $(spk, ssk) \stackrel{R}{\leftarrow} \text{SKg}$. Choose a random $\alpha \stackrel{R}{\leftarrow} \mathbb{Z}_p$ and compute $v \leftarrow g^\alpha$. The secret key is $sk = (\alpha, ssk)$; the public key is $pk = (v, spk)$.

Pub.St (sk, M) . Given the file M , first apply the erasure code to obtain M' ; then split M' into n blocks (for some n), each s sectors long: $\{m_{ij}\}_{\substack{1 \leq i \leq n \\ 1 \leq j \leq s}}$. Now parse sk as (α, ssk) . Choose a random file name $name$ from some sufficiently large domain (e.g., \mathbb{Z}_p). Choose s random elements $u_1, \dots, u_s \stackrel{R}{\leftarrow} G$. Let τ_0 be “ $name||n||u_1||\dots||u_s$ ”; the file tag τ is τ_0 together with a signature on τ_0 under private key ssk : $\tau \leftarrow \tau_0 || \text{SSig}_{ssk}(\tau_0)$. For each i , $1 \leq i \leq n$, compute

$$\sigma_i \leftarrow \left(H(name||i) \cdot \prod_{j=1}^s u_j^{m_{ij}} \right)^\alpha .$$

The processed file M^* is $\{m_{ij}\}$, $1 \leq i \leq n$, $1 \leq j \leq s$ together with $\{\sigma_i\}$, $1 \leq i \leq n$.

Pub.V (pk, sk, τ) . Parse pk as (v, spk) . Use spk to verify the signature on τ ; if the signature is invalid, reject by emitting 0 and halting. Otherwise, parse τ , recovering $name$, n , and u_1, \dots, u_s . Now pick a random l -element subset I of the set $[1, n]$, and, for each $i \in I$, a random element $\nu_i \stackrel{R}{\leftarrow} B$. Let Q be the set $\{(i, \nu_i)\}$. Send Q to the prover.

Parse the prover’s response to obtain $(\mu_1, \dots, \mu_s) \in (\mathbb{Z}_p)^s$ and $\sigma \in G$. If parsing fails, fail by emitting 0 and halting. Otherwise, check whether

$$e(\sigma, g) \stackrel{?}{=} e\left(\prod_{(i, \nu_i) \in Q} H(name||i)^{\nu_i} \cdot \prod_{j=1}^s u_j^{\mu_j}, v \right) ;$$

if so, output 1; otherwise, output 0.

Pub.P (pk, τ, M^*) . Parse the processed file M^* as $\{m_{ij}\}$, $1 \leq i \leq n$, $1 \leq j \leq s$, along with $\{\sigma_i\}$, $1 \leq i \leq n$. Parse the message sent by the verifier as Q , an l -element set $\{(i, \nu_i)\}$, with the i ’s distinct, each $i \in [1, n]$, and each $\nu_i \in B$. Compute

$$\mu_j \leftarrow \sum_{(i, \nu_i) \in Q} \nu_i m_{ij} \in \mathbb{Z}_p \quad \text{for } 1 \leq j \leq s, \quad \text{and} \quad \sigma \leftarrow \prod_{(i, \nu_i) \in Q} \sigma_i^{\nu_i} \in G .$$

Send to the prover in response the values μ_1, \dots, μ_s and σ .

Correctness. Again, it is easy to see that the scheme is correct. Let the secret key be α and the corresponding public key be $v = g^\alpha$. Let the public generators be u_1, \dots, u_s . Let the file sectors be $\{m_{ij}\}$, so that the block authenticators are $\sigma_i = (H(name||i) \cdot \prod_{j=1}^s u_j^{m_{ij}})^\alpha$. For a prover who responds honestly to a query $\{(i, \nu_i)\}$, so that each $\mu_j = \sum_{(i, \nu_i) \in Q} \nu_i m_{ij}$ and $\sigma = \prod_{(i, \nu_i) \in Q} \sigma_i^{\nu_i}$, we have

$$\begin{aligned} \sigma &= \prod_{(i, \nu_i) \in Q} \sigma_i^{\nu_i} = \prod_{(i, \nu_i) \in Q} \left(H(name||i) \cdot \prod_{j=1}^s u_j^{m_{ij}} \right)^{\nu_i \alpha} = \left(\prod_{(i, \nu_i) \in Q} H(name||i)^{\nu_i} \times \prod_{(i, \nu_i) \in Q} \prod_{j=1}^s u_j^{\nu_i m_{ij}} \right)^\alpha \\ &= \left(\prod_{(i, \nu_i) \in Q} H(name||i)^{\nu_i} \times \prod_{j=1}^s u_j^{\left(\sum_{(i, \nu_i) \in Q} \nu_i m_{ij} \right)} \right)^\alpha = \left(\prod_{(i, \nu_i) \in Q} H(name||i)^{\nu_i} \times \prod_{j=1}^s u_j^{\mu_j} \right)^\alpha , \end{aligned}$$

which means that

$$e(\sigma, g) = e\left(\left(\prod_{(i, \nu_i) \in Q} H(\text{name} \| i)^{\nu_i} \times \prod_{j=1}^s u_j^{\mu_j}\right)^\alpha, g\right) = e\left(\prod_{(i, \nu_i) \in Q} H(\text{name} \| i)^{\nu_i} \times \prod_{j=1}^s u_j^{\mu_j}, v\right),$$

so the verification equation is satisfied.

4 Security Proofs

In this section we prove that both of our systems are secure under the model we provided. We break our proof into three parts. Intuitively, the first part shows that the attacker can never give a forged response back to the a verifier. The second part of the proof shows that from any adversary that passes the check a non-negligible amount of the time we will be able to extract a constant fraction of the encoded blocks. The second step uses the fact that (w.h.p.) all verified responses must be legitimate. Finally, we show that if this constant fraction of blocks is recovered we can use the erasure code to reconstruct the original file.

The proof, for both schemes, is in three parts:

1. Prove that the verification algorithm will reject except when the prover's $\{\mu_j\}$ are correctly computed, i.e., are such that $\mu_j = \sum_{(i, \nu_i) \in Q} \nu_i m_{ij}$. This part of the proof uses cryptographic techniques.
2. Prove that the extraction procedure can efficiently reconstruct a ρ fraction of the file blocks when interacting with a prover that provides correctly-computed $\{\mu_j\}$ responses for a non-negligible fraction of the query space. This part of the proof uses combinatorial techniques.
3. Prove that a ρ fraction of the blocks of the erasure-coded file suffice for reconstructing the original file. This part of the proof uses coding theoretic techniques.

The crucial point is the second and third parts of the proof are *identical* for our two schemes; only the first part is different.

4.1 Part-One Proofs

4.1.1 Scheme with Private Verifiability

Theorem 4.1. *If the MAC scheme is unforgeable, the symmetric encryption scheme is semantically secure, and the PRF is secure, then (except with negligible probability) no adversary against the soundness of our private-verification scheme ever causes \mathcal{V} to accept in a proof-of-retrievability protocol instance, except by responding with values $\{\mu_j\}$ and σ that are computed correctly, i.e., as they would be by Priv.P .*

We prove the theorem in a series of games. Note that the reductions are not tight. The reduction to PRF security, for example, loses a factor of $1/(Nq_S)$, where N is a bound on the number of blocks in the encoding of any file the adversary requests to have stored. In the proof below, we interleave the game descriptions and the analysis limiting the difference in adversary behavior between successive games.

Game 0. The first game, Game 0, is simply the challenge game defined in Section 2.

Game 1. Game 1 is the same as Game 0, with one difference. The challenger keeps a list of all MAC-authenticated tags ever issued as part of a store query. If the adversary ever submits a tag τ either in initiating a proof-of-retrievability protocol or as the challenge tag, that (1) verifies as valid under k_{mac} but (2) is not on the list of tags authenticated by the challenger, the challenger declares failure and aborts.

Analysis. Clearly, if the adversary causes the challenger in Game 1 to abort with nonnegligible probability, we can use the adversary to construct a forger against the MAC scheme.

If the adversary does not cause the challenger to abort, his view is identical in Game 0 and in Game 1. With the modification made in Game 1, the verification and extraction algorithms will never attempt to decrypt a tag except those generated by the challenger. To see why this is so, observe that the first thing algorithm \mathcal{V} does, given a tag τ , is to check that the MAC on the tag is valid. If the MAC isn't valid, \mathcal{V} rejects immediately, without attempting to decrypt. Tags with a valid MAC will be decrypted, and these could either (a) have been produced by the challenger or (b) somehow mauled by the adversary; but, in Game 1, the challenger will abort if the adversary ever produces a tag with a valid MAC but different from all tags generated by the challenger itself, meaning that the verification and extraction algorithms will never deal with case (b). From now on, we need not worry about decrypting adversarially generated tags.

Game 2. In Game 2, the challenger includes in the tags not the encryption of $k_{\text{prf}}\|\alpha_1\|\cdots\|\alpha_s$ but a random bit-string of the same length. When given a tag by the adversary whose MAC verifies as correct, the challenger uses the values that would (in previous games) have been encrypted in the tag, rather than attempting to decrypt the ciphertext.

Analysis. The changes made in Game 1 guarantee that the challenger will never attempt to decrypt any ciphertext it did not generate, because the only tags with valid MACs the challenger will see are those it itself generated. The challenger can thus keep a table of plaintext values $k_{\text{prf}}\|\alpha_1\|\cdots\|\alpha_s$ values and the corresponding bit string it emitted as their tags. Decryption is replaced with table lookup.

If there is a difference in the adversary's success probability between Games 1 and 2, we can use the adversary to break the semantic security of the symmetric encryption scheme. Note that the reduction so obtained will suffer a $1/q_s$ security loss, where q_s is the number of **St** queries made by the adversary, because we must use a hybrid argument between "all valid encryptions" and "no valid encryptions."

Specifically, consider a challenger interacting with the adversary according to the game in Definition 2.1. The challenger keeps track of the files stored by the adversary. If the adversary succeeds in any proof-of-retrievability protocol interaction but sends values $\{\mu_j\}$ and σ that are different from those values that would be by the (deterministic) Priv.P algorithm, the challenger halts and outputs 1. Otherwise, the challenger outputs 0.

If this challenger's behavior in interacting with adversary \mathcal{A} is as specified in Game 0, then by assumption it will output 1 with some nonnegligible probability ϵ_0 . By the analysis of Game 1, if the challenger's behavior is as specified by Game 1, then it will output 1 with some nonnegligible probability ϵ_1 , because the difference between ϵ_0 and ϵ_1 is negligible assuming the MAC is secure.

If the challenger’s behavior is as specified in Game 2, then it will output 1 with some probability ϵ_3 . We will show that the difference between ϵ_2 and ϵ_3 is negligible assuming the symmetric encryption scheme is secure.

In Game 2, the challenger includes the encryption of the values $k_{\text{prf}}\|\alpha_1\|\cdots\|\alpha_s$ in each tag it generates in response to a store query by \mathcal{A} . In Game 3, the challenger encrypts a random string of the same length instead in each tag it generates. Suppose that $|\epsilon_3 - \epsilon_2|$ is nonnegligible. Then consider the hybrids in which the challenger encrypts a random string in the first i tags, and encrypts a random value in the remaining $q_S - i$ tags. Then there must be a value of i such that the difference between the challenger’s output in hybrid i and hybrid $i + 1$ is at least $|\epsilon_3 - \epsilon_2|/q_S$, which is nonnegligible. We will use this to construct an algorithm \mathcal{B} that breaks the security of the symmetric encryption scheme.

Algorithm \mathcal{B} is given access to an encryption oracle for a key k_{enc} , as well as a left-or-right oracle that, given strings m_0 and m_1 of the same length, outputs the encryption of m_b , where b is a randomly chosen bit [7]. Algorithm \mathcal{B} plays the part of the challenger, interacting with adversary \mathcal{A} . In answering \mathcal{A} ’s first i store queries, \mathcal{B} uses its encryption oracle to obtain the encryption of $k_{\text{prf}}\|\alpha_1\|\cdots\|\alpha_s$, which it includes in the tag. In answering \mathcal{A} ’s $(i + 1)$ st query, \mathcal{B} computes the correct plaintext $m_0 = k_{\text{prf}}\|\alpha_1\|\cdots\|\alpha_s$ and a random plaintext m_1 of the same length and submits both to its left-or-right oracle, including the oracle’s response in the tag. In answering \mathcal{A} ’s remaining store queries, \mathcal{B} computes the correct plaintext, generates a random plaintext of the same length, encrypts this random plaintext using its encryption oracle, and includes the result in the tag. Algorithm \mathcal{B} keeps track of the files stored by the adversary. If the adversary succeeds in any proof-of-retrievability protocol interaction but sends values $\{\mu_j\}$ and σ that are different from those values that would be by the (deterministic) Priv.P algorithm, \mathcal{B} outputs 1, otherwise 0.

If the left-or-right oracle encrypts its left input, \mathcal{B} is interacting with \mathcal{A} according to hybrid i . If the left-or-right oracle encrypts its right input, \mathcal{B} is interacting with \mathcal{A} according to hybrid $i + 1$. There is a nonnegligible difference in \mathcal{A} ’s behavior and therefore in \mathcal{B} ’s, which breaks the security of the symmetric encryption scheme. Note that, because the values $k_{\text{prf}}\|\alpha_1\|\cdots\|\alpha_s$ are chosen independently at random for each file, the values given by algorithm \mathcal{B} to its left-or-right oracle coincide with a query it makes to its encryption oracle only with negligible probability.

Game 3. In Game 3, the challenger uses truly random values in \mathbb{Z}_p instead of PRF outputs, remembering these values to use when verifying the adversary’s responses in proof-of-retrievability protocol instances. More specifically, the challenger evaluates $f_{k_{\text{prf}}}(i)$ not by applying the PRF algorithm but by generating a random value $r \xleftarrow{R} \mathbb{Z}_p$ and inserting an entry (k_{prf}, i, r) in a table; it consults this table when evaluating the PRF to ensure consistency.

Analysis. If there is a difference in the adversary’s success probability between Games 2 and 3, we can use the adversary to break the security of the PRF. It is important to note that, because of the change made in Game 2, the tags given to the adversary no longer contain k_{prf} , so the simulator does not need to know this value. The adversary will therefore see only PRF outputs; if it can distinguish these from random values it can be used to break the security of the PRF.

As in the analysis of Game 2, the difference in behavior we use to break the PRF security is the event that the adversary succeeds in a proof-of-retrievability protocol interaction but sends values $\{\mu_j\}$ and σ that are different from those values that would be by the (deterministic) Priv.P algorithm.

As before, a hybrid argument necessitates a security loss in the reduction; this time, the loss is $1/(Nq_s)$, where N is a bound on the number of blocks in the encoding of any file the adversary requests to have stored.

Game 4. In Game 4, the challenger handles proof-of-retrievability protocol executions initiated by the adversary differently than in Game 3.

In each such proof-of-retrievability protocol execution, the challenger issues a challenge as before. However, the challenger verifies the adversary’s response differently than is specified in algorithm \mathcal{V} .

The challenger keeps a table of the St queries made by the adversary, and of its responses to those queries; based on that table, the challenger knows the values $\{\mu_j\}$ and σ that the honest prover \mathcal{P} would have produced in response to the query it issued. (The honest prover is deterministic, so there is no ambiguity about the response it would have generated.) If the values the adversary sent were exactly these values, the challenger accepts the adversary’s response, returning a 1. If the values the adversary sent were different from these honest values, the challenger rejects the adversary’s response, returning a 0.

Analysis. The adversary’s view is different in Game 3 and Game 4 only when, in one of the proof-of-retrievability protocol interactions, the adversary responds in a way that (1) passes the verification algorithm but (2) is not what would have been computed by an honest prover, the challenger. We will now show that the probability that this happens is negligible.

We first establish some notation. Suppose a protocol instance involves an n -block file with secret values $\alpha_1, \dots, \alpha_s$ and content sectors $\{m_{ij}\}$, and that the block signatures issued by St are $\{\sigma_i\}$. Suppose $Q = \{(i, \nu_i)\}$ is the query issued by the challenger, and that the adversary’s response to that query was μ'_1, \dots, μ'_s together with σ' . Let the expected response—i.e., the one that would have been obtained from an honest prover—be μ_1, \dots, μ_s and σ , where $\sigma = \sum_{(i, \nu_i) \in Q} \nu_i \sigma_i$ and $\mu_j = \sum_{(i, \nu_i) \in Q} \nu_i m_{ij}$ for $1 \leq j \leq s$. If the adversary’s response satisfies the verifier—i.e., if $\sigma' = \sum_{(i, \nu_i) \in Q} \nu_i r_{k_{\text{prf}}, i} + \sum_{j=1}^s \alpha_j \mu'_j$, where $r_{k_{\text{prf}}, i}$ is the random value substituted by Game 2 for $f_{k_{\text{prf}}}(i)$, but $\mu'_j \neq \mu_j$ for at least one j , the challenger aborts. (If $\mu'_j = \mu_j$ for all j but $\sigma' \neq \sigma$, it is impossible that the verification equation holds, so we need not worry about this case.)

By the correctness of the scheme the expected values σ along with $\{\mu_j\}$ also satisfy the verification equation, so we have $\sigma = \sum_{i \in I} r_{k_{\text{prf}}, i} + \sum_{j=1}^s \alpha_j \mu_j$. Letting $\Delta\sigma \stackrel{\text{def}}{=} \sigma' - \sigma$ and $\Delta\mu_j \stackrel{\text{def}}{=} \mu'_j - \mu_j$ for $1 \leq j \leq s$ and subtracting the verification equation for σ from that for σ' , we have

$$\Delta\sigma = \sum_{j=1}^s \alpha_j \Delta\mu_j . \tag{1}$$

The bad event we are trying to rule out—the adversary’s submitting a convincing response different from an honest prover’s response—occurs exactly when some $\Delta\mu_j$ isn’t zero yet (1) holds.

However, with the Game 4 challenger, the values $\alpha_1, \dots, \alpha_s$ for every file are independent of the adversary’s view. They are no longer encrypted in the tag, and their only other appearance is in computing $\sigma_i = r_{k_{\text{prf}}, i} + \sum_{j=1}^s \alpha_j m_{ij}$ for $1 \leq i \leq n$; but the random value $r_{k_{\text{prf}}, i}$ replacing $f_{k_{\text{prf}}}(i)$ (and used only there) means that σ_i is independent of $\alpha_1, \dots, \alpha_s$.

Accordingly, the probability that the bad event happens if the simulator first picks the values $\{\alpha_j\}$ for each stored file and then undertakes the proof-of-retrievability interactions is the

same as the probability that the bad event happens if the simulator first undertakes the proof-of-retrievability interactions and only then chooses the values $\{\alpha_j\}$ for each file.

Fix the sequence of values $\Delta\mu_j$ and $\Delta\sigma$ in proof-of-retrievability responses by the adversary. The probability (over the choice of $\{\alpha_j\}$) that (1) holds for a specific entry in this sequence is $1/p$. The probability that (1) holds for a nonzero number of entries is at most q_P/p , where q_P is the number of proof-of-retrievability protocol interactions initiated by the adversary. (This upper bound is achieved only if all these interactions are for the same file.)

If the bound of q_P/p holds for any fixed sequence of values $\Delta\mu_j$ and $\Delta\sigma$, it holds also over a random choice of these values by the adversary. Except with negligible probability q_P/p , then, the adversary never generates a convincing response different from an honest prover’s response, so the adversary’s view in Game 4 is identical to its view in Game 3 except with negligible probability.

Game 5. In Game 5, the challenger observes each instance of the proof-of-retrievability protocol with the adversary — whether because of a proof-of-retrievability query made by the adversary, or in the test made of \mathcal{P}' , or as part of the extraction attempt by Extr . It compares the response made by the adversary to the response that would have been made by an honest prover. If in any of these interactions the adversary responds in a way that (1) passes the verification algorithm but (2) is not what would have been computed by an honest prover, the challenger sets a flag. At the end of the game, if the flag is set, the challenger declares failure and aborts.

Analysis. In the analysis of Game 4, we argued that the secret values $\{\alpha_j\}$ for each file are independent of the adversary’s view until the adversary outputs the cheating prover \mathcal{P}' for the challenge file. Although the values $\{\alpha_j\}$ for the challenge file *are* used by the extractor (in particular, to make the adversary “polite,” as defined below), \mathcal{P}' is rewound after each protocol interaction, meaning that it cannot learn information about the values $\{\alpha_j\}$, which thus remain independent of the adversary’s view for the entire game.

By the analysis of Game 4, the probability that any proof-of-retrievability interaction initiated by the adversary causes an abort is at most q_P/p , which is negligible. If there are k subsequent proof-of-retrievability interactions initiated by the extraction algorithm, the probability that any of these causes the challenger to abort is at most k/p . This probability is also negligible, since the extractor may make only polynomially many queries. The Game 5 challenger will thus abort only with negligible probability.

(This argument is inspired by Cramer and Shoup’s analysis of their encryption scheme [13]. The present version is simpler than the one we originally supplied, and was proposed by an anonymous *Journal of Cryptology* reviewer.)

Wrapping up. In Game 5, the adversary is constrained from answering any verification query with values other than those that would have been computed by $\text{Priv}\mathcal{P}$. Yet we have argued that, assuming the MAC, encryption scheme, and PRF are secure, there is only a negligible difference in the success probability of the adversary in this game compared to Game 0, where the adversary is not constrained in this manner. This completes the proof of Theorem 4.1. \square

4.1.2 Scheme with Public Verifiability

Theorem 4.2. *If the signature scheme used for file tags is existentially unforgeable and the computational Diffie-Hellman problem is hard in bilinear groups, then, in the random oracle model, except*

with negligible probability no adversary against the soundness of our public-verification scheme ever causes \mathcal{V} to accept in a proof-of-retrievability protocol instance, except by responding with values $\{\mu_j\}$ and σ that are computed correctly, i.e., as they would be by Pub.P .

Once again, we prove the theorem as a series of games with interleaved analysis. In this case, the reductions *are* tight.

Game 0. The first game, Game 0, is simply the challenge game defined in Section 2, with the changes for public verifiability sketched at the end of that section.

Game 1. Game 1 is the same as Game 0, with one difference. The challenger keeps a list of all signed tags ever issued as part of a store-protocol query. If the adversary ever submits a tag τ either in initiating a proof-of-retrievability protocol or as the challenge tag, that (1) has a valid signature under ssk but (2) is not a tag signed by the challenger, the challenger declares failure and aborts.

Analysis. Clearly, if the adversary causes the challenger in Game 1 to abort with nonnegligible probability, we can use the adversary to construct a forger against the signature scheme.

If the adversary does not cause the challenger to abort, his view is identical in Game 0 and in Game 1. With the modification made in Game 1, the verification and extraction algorithms will never attempt to make use of values u_1, \dots, u_s from a tag, except those generated by the challenger. To see why this is so, observe that the first thing algorithm \mathcal{V} does, given a tag τ , is to check that the signature on the tag is valid. If the signature isn't valid, \mathcal{V} rejects immediately. Those tags with a valid signature could either (a) have been produced by the challenger or (b) somehow mauled by the adversary; but, in Game 1, the challenger will abort if the adversary ever produces a tag with a valid signature but different from all tags generated by the challenger itself, meaning that the verification and extraction algorithms will never deal with case (b). From now on, we can be sure that any values u_1, \dots, u_s used in proof-of-retrievability interactions with the adversary will have been generated by the challenger.

Game 2. Game 2 is the same as Game 1, with one difference. The challenger keeps a list of its responses to St queries made by the adversary. Now the challenger observes each instance of the proof-of-retrievability protocol with the adversary — whether because of a proof-of-retrievability query made by the adversary, or in the test made of \mathcal{P}' , or as part of the extraction attempt by Extr . If in any of these instances the adversary is successful (i.e., \mathcal{V} outputs 1) but the adversary's aggregate signature σ is not equal to $\prod_{(i, \nu_i) \in Q} \sigma_i^{\nu_i}$ (where Q is the challenge issued by the verifier and σ_i are the signatures on the blocks of the file considered in the protocol instance) the challenger declares failure and aborts.

Analysis. Before analyzing the difference in success probabilities between Games 1 and 2, we will establish some notation and draw a few conclusions. Suppose the file that causes the abort is n blocks long, has name name , has generating exponents $\{u_j\}$, and contains sectors $\{m_{ij}\}$, and that the block signatures issued by St are $\{\sigma_i\}$. Suppose $Q = \{(i, \nu_i)\}$ is the query that causes the challenger to abort, and that the adversary's response to that query was μ'_1, \dots, μ'_s together with σ' . Let the expected response — i.e., the one that would have been obtained from an honest

prover—be μ_1, \dots, μ_s and σ , where $\sigma = \prod_{(i, \nu_i) \in Q} \sigma_i^{\nu_i}$ and $\mu_j = \sum_{(i, \nu_i) \in Q} \nu_i m_{ij}$ for $1 \leq j \leq s$. By the correctness of the scheme, we know that the expected response satisfies the verification equation, i.e., that

$$e(\sigma, g) = e\left(\prod_{(i, \nu_i) \in Q} H(\text{name}||i)^{\nu_i} \cdot \prod_{j=1}^s u_j^{\mu_j}, v\right) ;$$

Because the challenger aborted, we know that $\sigma \neq \sigma'$ and that σ' passes the verification equation, i.e., that

$$e(\sigma', g) = e\left(\prod_{(i, \nu_i) \in Q} H(\text{name}||i)^{\nu_i} \cdot \prod_{j=1}^s u_j^{\mu'_j}, v\right) ,$$

where $v = g^\alpha$ is part of the challenger's public key. Observe that if $\mu'_j = \mu_j$ for each j , it follows from the verification equation that $\sigma' = \sigma$, which contradicts our assumption above. Therefore, if we define $\Delta\mu_j \stackrel{\text{def}}{=} \mu'_j - \mu_j$ for $1 \leq j \leq s$, it must be the case that at least one of $\{\Delta\mu_j\}$ is nonzero.

With this in mind, we now show that if the adversary causes the challenger in Game 2 to abort with nonnegligible probability we can construct a simulator that solves the computational Diffie-Hellman problem.

The simulator is given as inputs values $g, g^\alpha, h \in G$; its goal is to output h^α . The simulator behaves like the Game 1 challenger, with the following differences:

- In generating a key, it sets the public key v to g^α received in the challenge. This means that it does not know the corresponding secret key α .
- The simulator programs the random oracle H . It keeps a list of queries and responses to answers consistently. In answering the adversary's queries it chooses a random $r \xleftarrow{R} \mathbb{Z}_p$ and responds with $g^r \in G$. It also answers queries of the form $H(\text{name}||i)$ in a special way, as we will see below.
- When asked to store some file whose coded representation comprises the n blocks $\{m_{ij}\}$, $1 \leq i \leq n$, $1 \leq j \leq s$, the simulator behaves as follows. It chooses a name name at random. Because the space from which names are drawn is large, it follows that, except with negligible probability, the simulator has not chosen this name before for some other file and a query has not been made to the random oracle at $\text{name}||i$ for any i .

For each j , $1 \leq j \leq s$, the simulator chooses random values $\beta_j, \gamma_j \xleftarrow{R} \mathbb{Z}_p$ and sets $u_j \leftarrow g^{\beta_j} \cdot h^{\gamma_j}$. For each i , $1 \leq i \leq n$, the simulator chooses a random value $r_i \xleftarrow{R} \mathbb{Z}_p$, and programs the random oracle at i as

$$H(\text{name}||i) = g^{r_i} / \left(g^{\sum_{j=1}^s \beta_j m_{ij}} \cdot h^{\sum_{j=1}^s \gamma_j m_{ij}}\right) .$$

Now the simulator can compute σ_i , since we have

$$\begin{aligned} H(\text{name}||i) \cdot \prod_{j=1}^s u_j^{m_{ij}} &= \left(\prod_{j=1}^s u_j^{m_{ij}}\right) \cdot \left(g^{r_i} / g^{\sum_{j=1}^s \beta_j m_{ij}} \cdot h^{\sum_{j=1}^s \gamma_j m_{ij}}\right) \\ &= g^{\sum_{j=1}^s \beta_j m_{ij}} \cdot h^{\sum_{j=1}^s \gamma_j m_{ij}} \cdot g^{r_i} / g^{\sum_{j=1}^s \beta_j m_{ij}} \cdot h^{\sum_{j=1}^s \gamma_j m_{ij}} = g^{r_i} , \end{aligned}$$

so the simulator computes $\sigma_i = (H(\text{name}||i) \cdot \prod_{j=1}^s u_j^{m_{ij}})^\alpha = (g^\alpha)^{r_i}$.

- The simulator continues interacting with the adversary until the condition specified in the definition of Game 2 occurs: the adversary, as part of a proof-of-retrievability protocol, succeeds in responding with a signature σ' that is different from the expected signature σ .

The change made from Game 0 to Game 1 establishes that the parameters associated with this protocol instance — $name$, n , $\{u_j\}$, $\{m_{ij}\}$, and $\{\sigma_i\}$ — were generated by the simulator as part of a St query; otherwise, execution would have already aborted. This means that these parameters were generated according to the simulator's procedure described above. Now, dividing the verification equation for the forged signature σ' by the verification equation for the expected signature σ , we obtain

$$e(\sigma'/\sigma, g) = e\left(\prod_{j=1}^s u_j^{\Delta\mu_j}, v\right) = e\left(\prod_{j=1}^s (g^{\beta_j} \cdot h^{\gamma_j})^{\Delta\mu_j}, v\right) .$$

Rearranging terms yields

$$e(\sigma' \cdot \sigma^{-1} \cdot v^{-\sum_{j=1}^s \beta_j \Delta\mu_j}, g) = e(h, v)^{\sum_{j=1}^s \gamma_j \Delta\mu_j} ,$$

Noting that v equals g^α , we see that we have found the solution to the computational Diffie-Hellman problem,

$$h^\alpha = \left(\sigma' \cdot \sigma^{-1} \cdot v^{-\sum_{j=1}^s \beta_j \Delta\mu_j}\right)^{\frac{1}{\sum_{j=1}^s \gamma_j \Delta\mu_j}} ,$$

unless evaluating the exponent causes a divide-by-zero. However, we noted already that not all of $\{\Delta\mu_j\}$ can be zero, and the values of $\{\gamma_j\}$ are information theoretically hidden from the adversary,¹⁰ so the denominator is zero only with probability $1/p$, which is negligible.

Thus if there is a nonnegligible difference between the adversary's probabilities of success in Games 1 and 2, we can construct a simulator that uses the adversary to solve computational Diffie-Hellman, as required.

Game 3. Game 3 is the same as Game 2, with one difference. As before, the challenger tracks St queries and observes proof-of-retrievability protocol instances. This time, if in any of these instances the adversary is successful (i.e., \mathcal{V} outputs 1) but at least one of the aggregate messages m_j is not equal to the expected $\sum_{(i, \nu_i) \in Q} \nu_i m_{ij}$ (where, again, Q is the challenge issued by the verifier) the challenger declares failure and aborts.

Analysis. Again, let us establish some notation. Suppose the file that causes the abort is n blocks long, has name $name$, has generating exponents $\{u_j\}$, and contains sectors $\{m_{ij}\}$, and that the block signatures issued by St are $\{\sigma_i\}$. Suppose $Q = \{(i, \nu_i)\}$ is the query that causes the challenger to abort, and that the adversary's response to that query was μ'_1, \dots, μ'_s together with σ' . Let the expected response — i.e., the one that would have been obtained from an honest prover — be μ_1, \dots, μ_s and σ , where $\sigma = \prod_{(i, \nu_i) \in Q} \sigma_i^{\nu_i}$ and $\mu_j = \sum_{(i, \nu_i) \in Q} \nu_i m_{ij}$ for $1 \leq j \leq s$. Game 2 already guarantees that we have $\sigma' = \sigma$; it is only the values $\{\mu'_j\}$ and $\{\mu_j\}$ that can differ. Define $\Delta\mu_j \stackrel{\text{def}}{=} \mu'_j - \mu_j$ for $1 \leq j \leq s$; again, at least one of $\{\Delta\mu_j\}$ is nonzero.

¹⁰Hidden because they are used to compute only the values $\{u_j\}$ in the adversary's view, and these are Pedersen commitments and so information-theoretically hiding.

We now show that if the adversary causes the challenger in Game 3 to abort with nonnegligible probability we can construct a simulator that solves the discrete logarithm problem.

The simulator is given as inputs values $g, h \in G$; its goal is to output x such that $h = g^x$. The simulator behaves like the Game 2 challenger, with the following differences:

- When asked to store some file whose coded representation comprises the n blocks $\{m_{ij}\}$, $1 \leq i \leq n$, $1 \leq j \leq s$, the simulator behaves according to St , except that For each j , $1 \leq j \leq s$, the simulator chooses random values $\beta_j, \gamma_j \xleftarrow{R} \mathbb{Z}_p$ and sets $u_j \leftarrow g^{\beta_j} \cdot h^{\gamma_j}$.
- The simulator continues interacting with the adversary until the condition specified in the definition of Game 3 occurs: the adversary, as part of a proof-of-retrievability protocol, succeeds in responding with aggregate messages $\{\mu'_j\}$ that are different from the expected aggregate messages $\{\mu_j\}$.

As before, we know because of the change made in Game 1 that the parameters associated with this protocol instance were generated by the simulator as part of a St query. Because of the change made in Game 2 we know that $\sigma' = \sigma$. Equating the verification equations using $\{\mu'_j\}$ and $\{\mu_j\}$ gives us

$$e\left(\prod_{(i,\nu_i) \in Q} H(\text{name}||i)^{\nu_i} \cdot \prod_{j=1}^s u_j^{\mu_j}, v\right) = e(\sigma, g) = e(\sigma', g) = e\left(\prod_{(i,\nu_i) \in Q} H(\text{name}||i)^{\nu_i} \cdot \prod_{j=1}^s u_j^{\mu'_j}, v\right),$$

from which we conclude that

$$\prod_{j=1}^s u_j^{\mu_j} = \prod_{j=1}^s u_j^{\mu'_j}$$

and therefore that

$$1 = \prod_{j=1}^s u_j^{\Delta\mu_j} = \prod_{j=1}^s (g^{\beta_j} \cdot h^{\gamma_j})^{\Delta\mu_j} = g^{\sum_{j=1}^s \beta_j \Delta\mu_j} \cdot h^{\sum_{j=1}^s \gamma_j \Delta\mu_j}$$

We see that we have found the solution to the discrete logarithm problem,

$$h = g^{-\frac{\sum_{j=1}^s \beta_j \Delta\mu_j}{\sum_{j=1}^s \gamma_j \Delta\mu_j}},$$

unless the denominator is zero. However, not all of $\{\Delta\mu_j\}$ can be zero, and the values of $\{\gamma_j\}$ are information theoretically hidden from the adversary, so the denominator is zero only with probability $1/p$, which is negligible.

Thus if there is a nonnegligible difference between the adversary's probabilities of success in Games 2 and 3, we can construct a simulator that uses the adversary to compute discrete logarithms, as required.

Wrapping up. In Game 3, the adversary is constrained from answering any verification query with values other than those that would have been computed by Pub.P . Yet we have argued that, assuming the signature scheme is secure and computational Diffie-Hellman and discrete logarithm are hard in bilinear groups, there is only a negligible difference in the success probability of the adversary in this game compared to Game 0, where the adversary is not constrained in this manner. Moreover, the hardness of the CDH problem implies the hardness of the discrete logarithm problem. This completes the proof of Theorem 4.2. \square

4.2 Part-Two Proof

We say that a cheating prover \mathcal{P}' is *well-behaved* if it never causes \mathcal{V} to accept in a proof-of-retrievability protocol instance except by responding with values $\{\mu_j\}$ and σ that are computed correctly, i.e., as they would be by Pub.P . The part-one proofs above guarantee that all adversaries that win the soundness game with nonnegligible probability output cheating provers that are well-behaved, provided that the cryptographic primitives we employ are secure. The part-two theorem shows that extraction always succeeds against a well-behaved cheating prover:

Theorem 4.3. *Suppose a cheating prover \mathcal{P}' on an n -block file M is well-behaved in the sense above, and that it is ϵ -admissible: i.e., convincingly answers an ϵ fraction of verification queries. Let $\omega = 1/\#B + (\rho n)^l/(n-l+1)^l$. Then, provided that $\epsilon - \omega$ is positive and nonnegligible, it is possible to recover a ρ fraction of the encoded file blocks in $O(n / (\epsilon - \omega))$ interactions with \mathcal{P}' and in $O(n^2s + (1 + \epsilon n^2)(n) / (\epsilon - \omega))$ time overall.*

We first make the following definition.

Definition 4.4. Consider an adversary \mathcal{B} , implemented as a probabilistic polynomial-time Turing machine, that, given a query Q on its input tape, outputs either the correct response ($\mathbf{q}M$ in vector notation) or a special symbol \perp to its output tape. Suppose \mathcal{B} responds with probability ϵ , i.e., on an ϵ fraction of the query-and-randomness-tape space. We say that such an adversary is ϵ -polite.

The proof of our theorem depends upon the following lemma that is proved below.

Lemma 4.5. *Suppose that \mathcal{B} is an ϵ -polite adversary as defined above. Let ω equal $1/\#B + (\rho n)^l/(n-l+1)^l$. If $\epsilon > \omega$ then it is possible to recover a ρ fraction of the encoded file blocks in $O(n / (\epsilon - \omega))$ interactions with \mathcal{B} and in $O(n^2s + (1 + \epsilon n^2)(n) / (\epsilon - \omega))$ time overall.*

To apply Lemma 4.5, we need only show that a well-behaved ϵ -admissible cheating prover \mathcal{P}' , as output by a setup-game adversary \mathcal{A} , can be turned into an ϵ -polite adversary \mathcal{B} . But this is quite simple. Here is how \mathcal{B} is implemented. We will use the \mathcal{P}' to construct the ϵ -adversary \mathcal{B} . Given a query Q , interact with \mathcal{P}' according to $(\mathcal{V}(pk, sk, \tau) \rightleftharpoons \mathcal{P}')$, playing the part of the verifier. If the output of the interaction is 1, write (μ_1, \dots, μ_s) to the output tape; otherwise, write \perp . Each time \mathcal{B} runs \mathcal{P}' , it provides it with a clean scratch tape and a new randomness tape, effectively rewinding it. Since \mathcal{P}' is well-behaved, a successful response will compute (μ_1, \dots, μ_s) as prescribed for an honest prover. Since \mathcal{P}' is ϵ -admissible, on an ϵ fraction of interactions it answers correctly. Thus algorithm \mathcal{B} that we have constructed is an ϵ -polite adversary.

The only use for \mathcal{V} above is to check that \mathcal{P}' 's responses are convincing. For schemes with private verification, this requires the secret key sk . For schemes with public verification, however, the secret key is not needed.

All that remains to guarantee that $\omega = 1/\#B + (\rho n)^l/(n-l+1)^l$ is such that $\epsilon - \omega$ is positive—indeed, nonnegligible. But this simply requires that each of $1/\#B$ and $(\rho n)^l/(n-l+1)^l$ be negligible in the security parameter; see Section 1.1. \square

To prove Lemma 4.5, we must first introduce some arguments in linear algebra.

For a subspace \mathbb{D} of $(\mathbb{Z}_p)^n$, denote the dimension of \mathbb{D} by $\dim \mathbb{D}$. Furthermore, let the free variables of a space, $\text{free } \mathbb{D}$, be the indices of the basis vectors $\{\mathbf{u}_i\}$ included in \mathbb{D} , i.e.,

$$\text{free } \mathbb{D} \stackrel{\text{def}}{=} \{i \in [1, n] : \mathbf{u}_i \in \mathbb{D}\} .$$

Observe that if we represent \mathbb{D} by means of a basis matrix in row-reduced echelon form, then we can efficiently compute $\dim \mathbb{D}$ and $\text{free } \mathbb{D}$.

Next, we give two claims.

Claim 4.6. *Let \mathbb{D} be a subspace of $(\mathbb{Z}_p)^n$, and let I be an l -element subset of $[1, n]$. If $I \not\subseteq \text{free } \mathbb{D}$, then a random query over indices I with coefficients in B is in \mathbb{D} with probability at most $1/\#B$.*

Proof. Let \mathbb{I} be the subspace spanned by the unit vectors in I , i.e., by $\{\mathbf{u}_i\}_{i \in I}$. Clearly, $\dim \mathbb{D} \cap \mathbb{I}$ is at most $l - 1$; if it equalled l , then we would have $\mathbb{D} \cap \mathbb{I} = \mathbb{I}$ and each of the vectors $\{\mathbf{u}_i\}_{i \in I}$ would be in \mathbb{D} , contradicting the lemma statement. Suppose $\dim \mathbb{D} \cap \mathbb{I}$ equals r . Then there exist r indices in I such that a choice of values for the coordinates at these indices determines the values of the remaining $l - r$ coordinates. This means that there are at most $(\#B)^r$ vectors in $\mathbb{D} \cap \mathbb{I}$ with coordinate values in B : a choice of one of $\#B$ values for each of the r coordinates above determines the value to each of the other $l - r$ coordinates; if the values of these coordinates are all in B , then this vector contributes 1 to the count; otherwise it contributes 0. The maximum possible count is thus $(\#B)^r$. By contrast, there are $(\#B)^l$ vectors in \mathbb{I} with coordinates in B , and these are exactly the vectors corresponding to each random query with indices I . Thus the probability that a random query is in \mathbb{D} is at most $1 / (\#B)^{l-r} \leq 1 / (\#B)$, which proves the lemma. \square

Claim 4.7. *Let \mathbb{D} be a subspace of $(\mathbb{Z}_p)^n$, and suppose that $\#(\text{free } \mathbb{D}) = m$. Then for a random l -element subset I of $[1, n]$ the probability that $I \subseteq \text{free } \mathbb{D}$ is at most $m^l / (n - l + 1)^l$.*

Proof. Color the m indices included in $\text{free } \mathbb{D}$ black; color the remaining $n - m$ indices white. A query I corresponds to a choice of l indices out of all these, without replacement. A query satisfies the condition that $I \subseteq \text{free } \mathbb{D}$ exactly if every element of I is in $\text{free } \mathbb{D}$, i.e., is colored black. Thus the probability that a random query satisfies the condition is just the probability of drawing l black balls, without replacement, from a jar containing m black balls and $n - m$ white balls; and this probability is

$$\binom{m}{l} / \binom{n}{l} = \frac{(m! / (m - l)!)}{(n! / (n - l)!)} < \frac{m^l}{(n - l + 1)^l} ,$$

as required. \square

Note that the bound established in 4.7 is not particularly tight. For example, if $m < l$ then it is impossible that $I \subseteq \text{free } \mathbb{D}$, but the probability bound is still positive; and if $m > n - l$ the probability bound is larger than 1 and therefore vacuous.

Lemma 4.5. *Suppose that \mathcal{B} is an ϵ -polite adversary as defined above. Let ω equal $1/\#B + (\rho n)^l / (n - l + 1)^l$. If $\epsilon > \omega$ then it is possible to recover a ρ fraction of the encoded file blocks in $O(n / (\epsilon - \omega))$ interactions with \mathcal{B} and in $O(n^2 s + (1 + \epsilon n^2)(n) / (\epsilon - \omega))$ time overall.*

Proof. We say the extractor's knowledge at each point is a subspace \mathbb{D} , represented by a $t \times n$ matrix A in row-reduced echelon form. Suppose that the query-response pairs contributing to the extractor's knowledge are

$$\mathbf{q}^{(1)} M = (\mu_1^{(1)}, \dots, \mu_s^{(1)}) \quad \dots \quad \mathbf{q}^{(t)} M = (\mu_1^{(t)}, \dots, \mu_s^{(t)}) ,$$

or $VM = W$, where V is the $t \times n$ matrix whose rows are $\{\mathbf{q}^{(i)}\}$ and W is the $t \times s$ matrix whose rows are $(\mu_1^{(i)}, \dots, \mu_s^{(i)})$. The row-reduced echelon matrix A is related to V by $A = UV$, where U is a $t \times t$ matrix with nonzero determinant computed in applying Gaussian elimination to V .

The extractor's knowledge is initially empty, i.e., $\mathbb{D} = \emptyset$.

The extractor repeats the following behavior until $\#(\text{free } \mathbb{D}) \geq \rho n$:

The extractor chooses a random query Q . It runs \mathcal{B} on Q . Suppose \mathcal{B} chooses to respond, giving answer (μ_1, \dots, μ_s) ; clearly this happens with probability ϵ . Let Q be over indices $I \in [1, n]$, and denote it in vector notation as \mathbf{q} . Now we classify Q into three types:

1. $\mathbf{q} \notin \mathbb{D}$;
2. $\mathbf{q} \in \mathbb{D}$ but $I \not\subseteq \text{free } \mathbb{D}$; or
3. $I \subseteq \text{free } \mathbb{D}$.

For queries of the first type, the extractor adds Q to its knowledge \mathbb{D} , obtaining new knowledge \mathbb{D}' , as follows. It adds a row corresponding to the query to V , obtaining V' , and a row corresponding to the response to W , obtaining W' ; it modifies the transform matrix U , obtaining U' , so that $A' = U'V'$ is again in row-reduced echelon form and spans \mathbf{q} . The primed versions \mathbb{D}' , A' , U' , V' , and W' replace the unprimed versions in the extractor's state. For queries of type 2 or 3, the extractor does not add to its knowledge. Regardless, the extractor continues with another query.

Clearly, a type-1 query increases $\dim \mathbb{D}$ by 1. If $\dim \mathbb{D}$ equals n then $\text{free } \mathbb{D} = [1, n]$ and $\#(\text{free } \mathbb{D}) = n \geq \rho n$, so the extractor's query phase is guaranteed to terminate by the time it has encountered n type-1 queries.

We now observe that any time the simulator is in its query phase, type-1 queries make up at least a $1 - \omega$ fraction of the query space. By Claim 4.6, type-2 queries make up at most a $(1/\#B)$ fraction of the query space, since

$$\begin{aligned} \Pr_Q[Q \text{ is type-2}] &= \Pr_Q[\mathbf{q} \in \mathbb{D} \wedge I \not\subseteq \text{free } \mathbb{D}] \\ &= \Pr_Q[\mathbf{q} \in \mathbb{D} \mid I \not\subseteq \text{free } \mathbb{D}] \cdot \Pr_Q[I \not\subseteq \text{free } \mathbb{D}] \\ &\leq \Pr_Q[\mathbf{q} \in \mathbb{D} \mid I \not\subseteq \text{free } \mathbb{D}] \\ &\leq 1/\#B, \end{aligned}$$

where it is the last inequality that follows from Claim 4.6.¹¹ Here the probability expressions are all over a random choice of query Q , and I and \mathbf{q} are the index set and vector form corresponding to the chosen query.

Similarly, suppose that $\#(\text{free } \mathbb{D}) = m$. Then by Claim 4.7, type-3 queries make up at most an $m^l/(n-l+1)^l$ fraction of the query space, and since $m < \rho n$ (otherwise the extractor would have ended the query phase) this fraction is at most $(\rho n)^l/(n-l+1)^l$.

Therefore the fraction of the query space consisting of type-2 and type-3 queries is at most $1/\#B + (\rho n)^l/(n-l+1)^l = \omega$. Since query type depends on the query and not on the randomness supplied to \mathcal{B} , it follows that the fraction of query-and-randomness-tape space consisting of type-2 and type-3 queries is also at most ω . Now, \mathcal{B} must respond correctly on an ϵ fraction of the query-and-randomness-tape space. Even if the adversary is as unhelpful as it can be and this ϵ fraction includes the entire ω fraction of type-2 and type-3 queries, there remains at least an $(\epsilon - \omega)$ fraction

¹¹The claim gives a condition for a single I satisfying the condition $I \not\subseteq \text{free } \mathbb{D}$; the inequality here is over *all* such I ; but if the probability never exceeds $1/\#B$ for any specific I then it doesn't exceed $1/\#B$ over a random choice of I , either.

of the query-and-randomness-tape space to which the adversary will respond correctly and in which the query is of type 1 and therefore helpful to the extractor. (By assumption $\epsilon > \omega$, so this fraction is nonempty.)

Since the extractor needs at most n successful type-1 queries to complete the query phase and it obtains a successful type-1 query from an interaction with \mathcal{B} with probability $O(\epsilon - \omega)$, it follows that the extractor will require at most $O(n/(\epsilon - \omega))$ interactions in expectation.

With \mathbb{D} represented by a basis matrix A in row-reduced echelon form, it is possible, given a query \mathbf{q} to which the adversary has responded, to determine efficiently which type it is. The extractor appends \mathbf{q} to A and runs the Gaussian elimination algorithm on the new row, a process that takes $O(n^2)$ time [11, Section 2.3].¹² If the reduced row is not all zeros then the query is type 1; the reduction also means that the augmented matrix A' is again in row-reduced echelon form, and the steps of the reduction also give the appropriate updates to the transform matrix U' . Since the reduction need only be performed for the ϵ fraction of queries to which \mathcal{B} correctly responds, the overall running time of the query phase is $O((1 + \epsilon n^2)(n) / (\epsilon - \omega))$.

Once the query phase is complete, the extractor has matrices A , U , V , and W such that $VM = W$ (where $M = (m_{ij})$ is the matrix consisting of encoded file blocks), $A = UV$, and A is in row-reduced echelon form. Moreover, there are at least ρn free dimensions in the subspace \mathbb{D} spanned by A and by V . Suppose i is in $\text{free}\mathbb{D}$. Since A is in row-reduced echelon form, there must be a row in A , say row t , that equals the i th basis vector \mathbf{u}_i . Multiplying both sides of $VM = W$ by U on the left gives the equation $AM = UW$. For any j , $1 \leq j \leq s$, consider the entry at row t and column j in the matrix AM . It is equal to $\mathbf{u}_i \cdot (m_{1,j}, m_{2,j}, \dots, m_{n,j}) = m_{i,j}$. If we compute the matrix product UW , we can thus read off from it every block of every sector for $i \in \text{free}\mathbb{D}$. Computing the matrix product takes $O(n^2 s)$ time. The extractor computes the relevant rows, outputs them, and halts. \square

Note that while we have described the extraction algorithm as performing row reduction operations, it could instead collect n successful interactions with the cheating prover and then perform a single Gaussian elimination using an algorithm specialized for sparse matrices, reducing the asymptotic runtime substantially. We do not expect that the extraction algorithm will be used in actual outsourced storage deployments, so this improvement is not important in practice. This completes the proof of Lemma 4.5. \square

4.3 Part-Three Proof

Theorem 4.8. *Given a ρ fraction of the n blocks of an encoded file M^* , it is possible to recover the entire original file M with all but negligible probability.*

Proof. For rate- ρ Reed-Solomon codes this is trivially true, since any ρ fraction of encoded file blocks suffices for decoding; see Appendix A. For rate- ρ linear-time codes the additional measures described in Appendix A guarantee that the ρ fraction of blocks retrieved will allow decoding with overwhelming probability. Note, however, that these measures do not protect the user if the pattern of block accesses she makes, in reading or reconstructing her file, reveals correlations between the plaintext blocks. If proofs of retrievability are used as part of a larger system where individual file blocks will be accessed, then Reed-Solomon codes should be used instead. \square

¹²More specifically, $O(tn)$ time if A is a $t \times n$ matrix; but of course $t \leq n$.

5 Proof for the Simple MAC Scheme

In this section we recall the simple MAC scheme described by Naor and Rothblum [26] and Juels and Kaliski [22] and give a formal proof for its security in the proof-of-retrievability model. We use the same common notation as in Section 3.1.

5.1 The Construction

Let $f: \{0, 1\}^* \times \mathcal{K}_{\text{prf}} \rightarrow \mathbb{Z}_p$ be a PRF. The construction of the simple scheme Simple is:

Simple.Kg(). Choose a random MAC key $k_{\text{mac}} \xleftarrow{\text{R}} \mathcal{K}_{\text{mac}}$. The secret key is $sk = (k_{\text{mac}})$; there is no public key.

Simple.St(sk, M). Given the file M , first apply the erasure code to obtain M' ; then split M' into n blocks (for some n), each s sectors long: $\{m_{ij}\}_{\substack{1 \leq i \leq n \\ 1 \leq j \leq s}}$. Choose a random file name $name$ from some sufficiently large domain (e.g., \mathbb{Z}_p). The file tag is $\tau = (name)$. Now, for each i , $1 \leq i \leq n$, compute

$$\sigma_i \leftarrow \text{MAC}_{k_{\text{mac}}}(name \| i \| m_{i1} \| \cdots \| m_{is}) .$$

The processed file M^* is $\{m_{ij}\}$, $1 \leq i \leq n$, $1 \leq j \leq s$ together with $\{\sigma_i\}$, $1 \leq i \leq n$.

Simple.V(pk, sk, τ). Parse sk as (k_{mac}) . Parse τ as $(name)$. Pick a random l -element subset I of the set $[1, n]$. Send I to the prover.

Parse the prover's response to obtain m_{i1}, \dots, m_{is} and σ_i , all in \mathbb{Z}_p , for each $i \in I$. If parsing fails, fail by emitting 0 and halting. Otherwise, check for each $i \in I$ whether

$$\sigma_i \stackrel{?}{=} \text{MAC}_{k_{\text{mac}}}(name \| i \| m_{i1} \| \cdots \| m_{is}) ;$$

if all l equations hold, output 1; otherwise, output 0.

Simple.P(pk, τ, M^*). Parse the processed file M^* as $\{m_{ij}\}$, $1 \leq i \leq n$, $1 \leq j \leq s$, along with $\{\sigma_i\}$, $1 \leq i \leq n$. Parse the message sent by the verifier as I , an l -element subset of $[1, n]$. Send to the prover, for each $i \in I$, the values m_{i1}, \dots, m_{is} and σ_i .

The correctness of the scheme is trivial to establish. Note that it is easy to modify the scheme and the proof to use a signature scheme instead of a MAC to obtain public verifiability.

5.2 The Proof

Theorem 5.1. *If the MAC scheme is unforgeable then (except with negligible probability) no adversary against the soundness of the simple scheme ever causes \mathcal{V} to accept in a proof-of-retrievability protocol instance, except by responding with values $\{m_{ij}\}$ and $\{\sigma_i\}$ that are computed correctly, i.e., as they would be by **Simple.P**.*

Proof. The simulator is given oracle access to the MAC; its goal is to create a forgery. The simulator plays the part of the environment in interacting with the attacker, using its MAC-generation oracle to create the $\{\sigma_i\}$ MACs. Whenever the adversary responds in a proof-of-storage protocol instance where $name$ is not one of the names issued by the simulator in a store query,

the simulator uses its MAC verification oracle to check whether any σ_i sent by the adversary, for $i \in I$, is a valid MAC.¹³ Such a valid MAC would be a forgery, since the simulator never requests a MAC on a *name* not chosen in a store query. Whenever the adversary responds in a proof-of-storage protocol instance on a file with tag *name* whose blocks are $\{m_{ij}\}$ and where, for some $i \in I$, the values $\{m'_{ij}\}_j$ sent by the adversary are different from the values $\{m_{ij}\}_j$ in the file, the simulator uses its MAC verification oracle to check whether the corresponding authenticator is valid. Such a valid MAC would be a forgery, since the simulator never requested a MAC on any string beginning “*name*||*i*|| \dots ” except for “*name*||*i*|| m_{i1} || \dots || m_{is} ”, and $(m_{i1}, \dots, m_{is}) \neq (m'_{i1}, \dots, m'_{is})$ by assumption. (Because *name* is drawn from a large space, each file storage query will use a different value for *name*, except with negligible probability.) We see that if the adversary ever causes \mathcal{V} to accept in a proof-of-retrievability protocol instance without responding with values $\{m_{ij}\}$ and $\{\sigma_i\}$ computed as they would be by Simple.P , the simulator finds a MAC forgery. \square

As before, we say that a cheating prover \mathcal{P}' is *well-behaved* if it never causes \mathcal{V} to accept in a proof-of-retrievability protocol instance except by responding with values $\{m_{ij}\}$ and $\{\sigma_i\}$ that are computed correctly, i.e., as they would be by Simple.P . The theorem above guarantee that all adversaries that win the soundness game with nonnegligible probability output cheating provers that are well-behaved, provided that the MAC we employ is secure. The next theorem shows that extraction always succeeds against a well-behaved cheating prover:

Theorem 5.2. *Suppose a cheating prover \mathcal{P}' on an n -block file M is well-behaved in the sense above, and that it is ϵ -admissible: i.e., convincingly answers an ϵ fraction of verification queries. Then, provided that $\epsilon - (\rho n)^l / (n - l + 1)^l$ is positive and nonnegligible, it is possible to recover a ρ fraction of the encoded file blocks in $O(\rho n / (\epsilon - (\rho n)^l / (n - l + 1)^l))$ interactions with \mathcal{P}' .*

Proof. We turn the ϵ -admissible, well-behaved cheating prover \mathcal{P}' into an ϵ -polite adversary \mathcal{B} as in the proof of Theorem 4.3, by interacting with \mathcal{P}' , checking the MACs $\{\sigma_i\}$ on each block $i \in I$, and emitting $\{m_{ij}\}$ if all l MACs are valid, \perp otherwise.

Against an ϵ -polite adversary the extractor works as follows. Its knowledge is a subset $S \subseteq [1, n]$, initially empty. If $\#S$ ever reaches ρn , the extractor halts. The extractor repeatedly chooses a random l -element query $I \subset [1, n]$ and sends I to the polite adversary \mathcal{B} . If the adversary does not output \perp , the extractor updates its knowledge as $S' \leftarrow S \cup I$. Regardless, the extractor continues with another query.

A query is answered as not \perp with probability ϵ . For such a query, $\#S$ increases by at least 1 provided that $I \not\subseteq S$. But if the extractor has not yet halted we have $\#S < \rho n$, and the probability that a random l -element subset I of $[1, n]$ is such that $I \subseteq S$ is at most $(\rho n)^l / (n - l + 1)^l$, by reasoning identical to that used in the proof of Lemma 4.7. This means that on an $\epsilon - (\rho n)^l / (n - l + 1)^l$ fraction of the query–randomness space the adversary \mathcal{B} will give a response that increases $\#S$. Thus $O(\rho n / (\epsilon - (\rho n)^l / (n - l + 1)^l))$ interactions with the adversary suffice to grow $\#S$ to ρn elements, at which point the extractor halts.

But each element $i \in S$ corresponds to a block i for which the extractor has learned the values m_{i1}, \dots, m_{is} (since the adversary is polite), so the extractor will have recovered ρn blocks of the file, as required. \square

Theorem 5.3. *Given a ρ fraction of the n blocks of an encoded file M^* , it is possible to recover the entire original file M with all but negligible probability.*

¹³See Bellare, Goldreich, and Mityagin [8] for why the MAC security definition incorporates verification queries.

The proof is identical to the proof of Theorem 4.8 in Section 4.3.

6 Construction with RSA Signatures

In this section, we show how the RSA construction of Ateniese et al. [3] can be considered an instantiation of our framework for proofs of retrievability. The construction closely follows that of Ateniese et al., and the part-one proof also uses RSA techniques similar to those used in their Theorem 3.3. The benefit of this section is to show that an RSA-based construction very similar to that of Ateniese et al. admits a full and rigorous proof of security.

6.1 Construction

Let λ be the security parameter, and let λ_1 be a bitlength such that the difficulty of factoring a $(2\lambda_1 - 1)$ -bit modulus is appropriate to the security parameter λ . Let $\max B$ be the largest element in B , and let λ_2 be a bitlength equal to $\lceil \lg(l \cdot \max B) \rceil + 1$.

The construction of the public verification scheme PubRSA is:

PubRSA.Kg(). Generate a random signing keypair $(spk, ssk) \xleftarrow{R} \text{SKg}$. Choose two random primes p and q in the range $[2^{\lambda_1-1}, 2^{\lambda_1} - 1]$. Let $N = pq$ be the RSA modulus; we have $2^{2\lambda_1-2} < N < 2^{2\lambda_1}$. Let $H: \{0, 1\}^* \rightarrow \mathbb{Z}_N^*$ be a full-domain hash, which we treat as a random oracle.¹⁴ Choose a random $2\lambda_1 + \lambda_2$ -bit prime e , and set $d = e^{-1} \bmod \phi(N)$. The secret key is $sk = (N, d, H, ssk)$; the public key is $pk = (N, e, H, spk)$.

PubRSA.St (sk, M) . Given the file M , first apply the erasure code to obtain M' ; then split M' into n blocks (for some n), each s sectors: $\{m_{ij}\}_{\substack{1 \leq i \leq n \\ 1 \leq j \leq s}}$. Each sector m_{ij} is an element of \mathbb{Z}_N . Now parse sk as (N, d, H, ssk) . Choose a random file name $name$ from some sufficiently large domain (e.g., \mathbb{Z}_N). Choose s random elements $u_1, \dots, u_s \xleftarrow{R} \mathbb{Z}_N^*$. Let τ_0 be “ $name||n||u_1||\dots||u_s$ ”; the file tag τ is τ_0 together with a signature, on τ_0 under private key ssk : $\tau \leftarrow \tau_0 || \text{SSig}_{ssk}(\tau_0)$.

Now, for each i , $1 \leq i \leq n$, compute

$$\sigma_i \leftarrow \left(H(name||i) \cdot \prod_{j=1}^s u_j^{m_{ij}} \right)^d \bmod N .$$

The processed file M^* is $\{m_{ij}\}$, $1 \leq i \leq n$, $1 \leq j \leq s$ together with $\{\sigma_i\}$, $1 \leq i \leq n$.

PubRSA.V (pk, sk, τ) . Parse pk as (N, e, H, spk) . Use spk to verify the signature on τ ; if the signature is invalid, reject by emitting 0 and halting. Otherwise, parse τ , recovering $name$, n , and u_1, \dots, u_s . Now pick a random l -element subset I of the set $[1, n]$, and, for each $i \in I$, a random element $\nu_i \xleftarrow{R} B$. Let Q be the set $\{(i, \nu_i)\}$. Send Q to the prover.

Parse the prover’s response to obtain μ_1, \dots, μ_s and $\sigma \in \mathbb{Z}_N$. Check that each μ_j is in the range $[0, l \cdot N \cdot \max B]$. If parsing fails or the $\{\mu_j\}$ values are not in range, fail by emitting 0

¹⁴The hash H can be instantiated using a hash onto all of \mathbb{Z}_N ; a value not in \mathbb{Z}_N^* would disclose the factorization of N and thus will never appear in practice.

and halting. Otherwise, check whether

$$\sigma^e \stackrel{?}{=} \prod_{(i,\nu_i) \in Q} H(\text{name}||i)^{\nu_i} \cdot \prod_{j=1}^s u_j^{\mu_j} \pmod{N} ;$$

if so, output 1; otherwise, output 0.

PubRSA. $\mathcal{P}(pk, \tau, M^*)$. Parse the processed file M^* as $\{m_{ij}\}$, $1 \leq i \leq n$, $1 \leq j \leq s$, along with $\{\sigma_i\}$, $1 \leq i \leq n$. Parse the message sent by the verifier as Q , an l -element set $\{(i, \nu_i)\}$, with the i 's distinct, each $i \in [1, n]$ and each $\nu_i \in B$.

For each j , $1 \leq j \leq s$, compute

$$\mu_j \leftarrow \sum_{(i,\nu_i) \in Q} \nu_i m_{ij} \in \mathbb{Z} ,$$

where the sum is computed in \mathbb{Z} , without modular reduction. In addition, compute

$$\sigma \leftarrow \prod_{(i,\nu_i) \in Q} \sigma_i^{\nu_i} \pmod{N} .$$

Send to the prover in response the values μ_1, \dots, μ_s and σ .

Correctness. It is easy to see that the scheme is correct. Let the modulus be N and the public and private exponents be e and d . Let the public generators be u_1, \dots, u_s . Let the file sectors be $\{m_{ij}\}$, so that the block authenticators are $\sigma_i = (H(\text{name}||i) \cdot \prod_{j=1}^s u_j^{m_{ij}})^d \pmod{N}$. For a prover who responds honestly to a query $\{(i, \nu_i)\}$, so that each $\mu_j = \sum_{(i,\nu_i) \in Q} \nu_i m_{ij}$ and $\sigma = \prod_{(i,\nu_i) \in Q} \sigma_i^{\nu_i} \pmod{N}$, we have, modulo N , that

$$\begin{aligned} \sigma &= \prod_{(i,\nu_i) \in Q} \sigma_i^{\nu_i} = \prod_{(i,\nu_i) \in Q} \left(H(\text{name}||i) \cdot \prod_{j=1}^s u_j^{m_{ij}} \right)^{\nu_i d} = \left(\prod_{(i,\nu_i) \in Q} H(\text{name}||i)^{\nu_i} \times \prod_{(i,\nu_i) \in Q} \prod_{j=1}^s u_j^{\nu_i m_{ij}} \right)^d \\ &= \left(\prod_{(i,\nu_i) \in Q} H(\text{name}||i)^{\nu_i} \times \prod_{j=1}^s u_j^{\left(\sum_{(i,\nu_i) \in Q} \nu_i m_{ij} \right)} \right)^d = \left(\prod_{(i,\nu_i) \in Q} H(\text{name}||i)^{\nu_i} \times \prod_{j=1}^s u_j^{\mu_j} \right)^d , \end{aligned}$$

which means that

$$\sigma^e = \prod_{(i,\nu_i) \in Q} H(\text{name}||i)^{\nu_i} \times \prod_{j=1}^s u_j^{\mu_j} \pmod{N} ,$$

so the verification equation is satisfied.

6.2 Part-One Proof

We now give the part one proof of our scheme.

We begin with technical observations about \mathbb{Z}_N^* that will be of use below. For e relatively prime to $\phi(N)$, the map $x \mapsto x^e \pmod{N}$ is an isomorphism of \mathbb{Z}_N^* ; since e as chosen above is prime and larger than N , it must be relatively prime to $\phi(N)$, as required. For $c \in \mathbb{Z}_N^*$, the map $x \mapsto cx \pmod{N}$ is also an isomorphism of \mathbb{Z}_N^* . Thus for $x \in \mathbb{Z}_N^*$, the value cx for a random $c \in \mathbb{Z}_N^*$ is information-theoretically independent of x . In addition, we will use the following lemma (see [19] and Lemma 1 of [12]):

Lemma 6.1. *Given $x, y \in \mathbb{Z}_N$, along with $a, b \in \mathbb{Z}$ such that $x^a = y^b$ and $\gcd(a, b) = 1$, one can efficiently compute $\bar{x} \in \mathbb{Z}_N$ such that $\bar{x}^a = y$.*

Theorem 6.2. *If the signature scheme used for file tags is existentially unforgeable and the RSA problem with large public exponents is hard, then, in the random oracle model, except with negligible probability no adversary against the soundness of our public-verification scheme ever causes \mathcal{V} to accept in a proof-of-retrievability protocol instance, except by responding with values $\{\mu_j\}$ and σ that are computed correctly, i.e., as they would be by PubRSA.P .*

Once more, we prove the theorem as a series of games with interleaved analysis.

Game 0. The first game, Game 0, is simply the challenge game defined in Section 2. By assumption, the adversary \mathcal{A} wins with nonnegligible probability.

Game 1. Game 1 is the same as Game 0, with one difference. The challenger keeps a list of all signed tags ever issued as part of a store-protocol query. If the adversary ever submits a tag τ either in initiating a proof-of-storage protocol or as the challenge tag, that (1) has a valid signature under ssk but (2) is not a tag signed by the challenger, the challenger declares failure and aborts.

Analysis. Clearly, if there is a difference in the adversary's success probability between Games 0 and 1, we can use the adversary to construct a forger against the signature scheme.

Game 2. Game 2 is the same as Game 1, with one difference. The challenger keeps a list of its responses to St queries made by the adversary. Now the challenger observes each instance of the proof-of-storage protocol with the adversary — whether because of a proof-of-storage query made by the adversary, or in the test made of \mathcal{P}' , or as part of the extraction attempt by Extr . If in any of these instances the adversary is successful (i.e., \mathcal{V} outputs 1) but either

1. the adversary's aggregate signature σ is not equal to $\prod_{(i, \nu_i) \in Q} \sigma_i^{\nu_i} \bmod N$ (where Q is the challenge issued by the verifier and σ_i are the signatures on the blocks of the file considered in the protocol instance) or
2. at least one the adversary's aggregate block values μ'_1, \dots, μ'_s is not equal to the expected block value $\mu_j = \sum_{(i, \nu_i) \in Q} \nu_i m_{ij}$,

the challenger declares failure and aborts.

Analysis. Before analyzing the difference in success probabilities between Games 1 and 2, we will establish some notation and draw a few conclusions. Suppose the file that causes the abort is n blocks long, has name name , has generating exponents $\{u_j\}$, and contains sectors $\{m_{ij}\}$, and that the block signatures issued by St are $\{\sigma_i\}$. Suppose $Q = \{(i, \nu_i)\}$ is the query that causes the challenger to abort, and that the adversary's response to that query was μ'_1, \dots, μ'_s together with σ' . Let the expected response — i.e., the one that would have been obtained from an honest prover — be μ_1, \dots, μ_s and σ , where $\sigma = \prod_{(i, \nu_i) \in Q} \sigma_i^{\nu_i} \bmod N$ and $\mu_j = \sum_{(i, \nu_i) \in Q} \nu_i m_{ij}$ for $1 \leq j \leq s$. By the correctness of the scheme, we know that the expected response satisfies the verification equation, i.e., that

$$\sigma^e = \prod_{(i, \nu_i) \in Q} H(\text{name} || i)^{\nu_i} \cdot \prod_{j=1}^s u_j^{\mu_j} ;$$

Because the challenger aborted, we know that σ' and μ'_1, \dots, μ'_s passed the verification equation, i.e., that

$$(\sigma')^e = \prod_{(i, \nu_i) \in Q} H(\text{name}||i)^{\nu_i} \cdot \prod_{j=1}^s u_j^{\mu'_j} .$$

(Note that if σ' is in $\mathbb{Z}_N \setminus \mathbb{Z}_N^*$ then so is $(\sigma')^e$, whereas the right-hand side of the verification equation is in \mathbb{Z}_N^* . Thus the verification equation will not hold unless σ' is in \mathbb{Z}_N^* , which is why no separate check in \mathcal{V} that is required that σ' is relatively prime to N .)

Now observe that condition 1, above, implies condition 2, which means that having the simulator abort on either condition 1 or 2 is the same as having it abort on just condition 2: if condition 2 doesn't hold then $\mu'_j = \mu_j$ for each j , and it follows from the verification equation that $(\sigma')^e = \sigma^e$. Because \mathcal{V} checked that σ' is in \mathbb{Z}_N and because, as noted, the verification equation requires that σ' , like σ , is in \mathbb{Z}_N^* , the fact that exponentiation by e is an isomorphism of \mathbb{Z}_N^* means that $(\sigma')^e = \sigma^e$ implies $\sigma' = \sigma$, so condition 1 doesn't hold, either.

Therefore, if we define $\Delta\mu_j \stackrel{\text{def}}{=} \mu'_j - \mu_j$ for $1 \leq j \leq s$, it must be the case that if the simulator aborts at least one of $\{\Delta\mu_j\}$ is nonzero.

With this in mind, we now show that if there is a nonnegligible difference in the adversary's success probabilities between Games 1 and 2 we can construct a simulator that solves the RSA problem when the public exponent e is large.

The simulator is given as inputs a $2\lambda_1$ -bit modulus N and a $(2\lambda_1 + \lambda_2)$ -bit public exponent e , along with a value $y \in \mathbb{Z}_N^*$; its goal is to output $x \in \mathbb{Z}_N^*$ such that $x^e = y$. The simulator behaves like the Game 1 challenger, with the following differences:

- In generating a public key, it sets the modulus and public exponent to N and e ; it does not know the corresponding secret modulus d .
- The simulator programs the random oracle H . It keeps a list of queries and responses to answers consistently. In answering the adversary's queries it responds with a random $g \stackrel{\text{R}}{\leftarrow} \mathbb{Z}_N^*$. The simulator also answers queries of the form $H(\text{name}||i)$ in a special way, as we will see below.
- When asked to store some file whose coded representation comprises the n blocks $\{m_{ij}\}$, $1 \leq i \leq n$, $1 \leq j \leq s$, the simulator behaves as follows. It chooses a name name at random. Because the space from which names are drawn is large, it follows that, except with negligible probability, the simulator has not chosen this name before for some other file and a query has not been made to the random oracle at $\text{name}||i$ for any i .

For each j , $1 \leq j \leq s$, the simulator chooses a random $g_j \stackrel{\text{R}}{\leftarrow} \mathbb{Z}_N^*$ and $\beta_j \stackrel{\text{R}}{\leftarrow} [1, 2^\lambda]$ and sets $u_j \leftarrow g_j^e y^{\beta_j}$. For each i , $1 \leq i \leq n$, the simulator chooses a random value $h_i \stackrel{\text{R}}{\leftarrow} \mathbb{Z}_N^*$, and programs the random oracle at i as

$$H(\text{name}||i) = h_i^e / \prod_{j=1}^s u_j^{m_{ij}} .$$

Now the simulator can compute σ_i , since we have

$$H(\text{name}||i) \cdot \prod_{j=1}^s u_j^{m_{ij}} = h_i^e ;$$

if the simulator sets $\sigma_i = h_i$, we will have $\sigma_i^e = h_i^{e \cdot r_i} = H(\text{name}||i) \cdot \prod_{j=1}^s u_j^{m_{ij}}$, as required.

- The simulator continues interacting with the adversary until the condition specified in the definition of Game 2 occurs: the adversary, as part of a proof-of-storage protocol, succeeds in responding with a signature σ' that is different from the expected signature σ .

The change made from Game 0 to Game 1 establishes that the parameters associated with this protocol instance — name , n , $\{u_j\}$, $\{m_{ij}\}$, and $\{\sigma_i\}$ — were generated by the simulator as part of a St query; otherwise, execution would have already aborted. This means that these parameters were generated according to the simulator's procedure described above. Now, dividing the verification equation for the forged signature σ' by the verification equation for the expected signature σ , we obtain

$$(\sigma'/\sigma)^e = \prod_{j=1}^s u_j^{\Delta\mu_j} = \left[\prod_{j=1}^s (g_j^e)^{\beta_j \Delta\mu_j} \right] \cdot y^{\sum_{j=1}^s \beta_j \Delta\mu_j} ;$$

rearranging terms yields

$$\left[(\sigma'/\sigma) \cdot \prod_{j=1}^s g_j^{\beta_j \Delta\mu_j} \right]^e = y^{\sum_{j=1}^s \beta_j \Delta\mu_j} ; \quad (2)$$

Now, provided that $\gcd(e, \sum_{j=1}^s \beta_j \Delta\mu_j) = 1$, we can compute, using Lemma 6.1, a value x from (2) such that $x^e = y$.

It remains only to argue that $\gcd(e, \sum_{j=1}^s \beta_j \Delta\mu_j) \neq 1$ occurs with negligible probability. First, we noted already that not all of $\{\Delta\mu_j\}$ can be zero. Second, the values of $\{\beta_j\}$ are statistically hidden from the adversary.¹⁵ Third, the verification equation checks that each μ'_j is in the range $[0, l \cdot N \cdot \max B]$, and each μ_j is also in the same range; thus for each j we have

$$|\Delta\mu_j| = |\mu'_j - \mu_j| \leq l \cdot N \cdot \max B < 2^{\lceil \lg N \rceil} \cdot 2^{\lceil \lg(l \cdot \max B) \rceil} < 2^{2\lambda_1} \cdot 2^{\lambda_2} < e ,$$

and since e is prime this means that $\gcd(\Delta\mu_j, e)$ must equal 1. Now, because e is prime, $\gcd(e, \sum_{j=1}^s \beta_j \Delta\mu_j) \neq 1$ means that e divides $\sum_{j=1}^s \beta_j \Delta\mu_j$, i.e., that $\sum_{j=1}^s \beta_j \Delta\mu_j \equiv 0 \pmod{e}$. For any particular fixed choice of $\{\Delta\mu_j\}$ values, the probability that this happens, over the independent random choices of each β_j from $[1, 2^\lambda]$, is at most $2^{-\lambda}$, which is negligible. (Let j^* be some index such that $\Delta\mu_{j^*} \neq 0$ and fix $\{\beta_j\}_{j \neq j^*}$. Then let $c \equiv \sum_{j \neq j^*} \beta_j \Delta\mu_j \pmod{e}$; then $\sum_{j=1}^s \beta_j \Delta\mu_j \equiv c + \beta_{j^*} \Delta\mu_{j^*}$; and this is congruent to 0 modulo e for exactly one value of β_{j^*} , namely $\beta_{j^*} = -(c)(\Delta\mu_{j^*}^{-1}) \pmod{e}$; since β_{j^*} is drawn from the range $[1, 2^\lambda]$, the probability that it takes on this value is at most $1/2^\lambda$.)

Thus if there is a nonnegligible difference between the adversary's probabilities of success in Games 1 and 2, we can construct a simulator that uses the adversary to solve the RSA problem, as required.

Wrapping up. Assuming the signature scheme used for file tags is secure, and that the RSA problem with large public exponent is hard, we see that any adversary that wins the soundness game against our public-verification scheme responds in proof-of-storage protocol instance with values $\{\mu_j\}$ and σ that are computed according to $\text{PubRSA}.\mathcal{P}$, which completes the proof of Theorem 6.2. \square

¹⁵Hidden because they are used to compute only the values $\{u_j\}$ in the adversary's view, and thus are hidden by the $\{g_j^e\}$ multipliers.

6.3 Part-Two and Part Three Proofs

It is easy to see that the Part-Two proof of Section 4.2 carries over unchanged to the case where blocks are drawn from \mathbb{Z}_N instead of \mathbb{Z}_p . The matrix operations used there require only that inversion be efficiently computable, and this is, of course, the case in \mathbb{Z}_N using Euclid’s algorithm, provided we never encounter values in $\mathbb{Z}_N \setminus \mathbb{Z}_N^*$; but such a value would allow us to factor N , so they occur with negligible probability provided the RSA problem—and therefore factoring—is hard.

Similarly, erasure decoding works just as well when blocks are drawn from \mathbb{Z}_N ; and because nothing in the proof requires that blocks be distributed uniformly in all of \mathbb{Z}_N , we could treat each m_{ij} as an element of $\mathbb{Z}_{p_0}^k$ where p_0 is some prime convenient for whatever erasure code we employ and k is the largest integer such that $p_0^k < N$.

Acknowledgements

We thank Dan Boneh, Guy Rothblum, and Moni Naor for helpful discussions about this work; Eric Rescorla for detailed comments on the manuscript; Giuseppe Ateniese for his helpful comments, and in particular for his suggested improvements to our RSA construction; attendees of the MIT Cryptography and Information Security Seminar and the UC Irvine Crypto Seminar for their questions and comments; and the anonymous conference reviewers and Journal of Cryptology referees.

References

- [1] M. Aigner and G. Ziegler. *Proofs from The Book*. Springer-Verlag, 3rd edition, 2004.
- [2] N. Alon and M. Luby. A linear time erasure-resilient code with nearly optimal recovery. *IEEE Trans. Info. Theory*, 42(6):1732–6, Nov. 1996.
- [3] G. Ateniese, R. Burns, R. Curtmola, J. Herring, O. Khan, L. Kissner, Z. Peterson, and D. Song. Remote data checking using provable data possession. *ACM Trans. Info. & System Security*, 14(1), May 2011.
- [4] G. Ateniese, R. Di Pietro, L. Mancini, and G. Tsudik. Scalable and efficient provable data possession. In P. Liu and R. Molva, editors, *Proceedings of SecureComm 2008*, pages 1–10. ICST, Sept. 2008.
- [5] G. Ateniese, S. Kamara, and J. Katz. Proofs of storage from homomorphic identification protocols. In M. Matsui, editor, *Proceedings of Asiacrypt 2009*, volume 5912 of *LNCS*, pages 319–33. Springer-Verlag, Dec. 2009.
- [6] P. Barreto and M. Naehrig. Pairing-friendly elliptic curves of prime order. In B. Preneel and S. Tavares, editors, *Proceedings of SAC 2005*, volume 3897 of *LNCS*, pages 319–31. Springer-Verlag, 2006.
- [7] M. Bellare, A. Desai, E. Jorjipii, and P. Rogaway. A concrete security treatment of symmetric encryption. In A. R. Karlin, editor, *Proceedings of FOCS 2007*, pages 394–403. IEEE Computer Society, Oct. 1997.

- [8] M. Bellare, O. Goldreich, and A. Mityagin. The power of verification queries in message authentication and authenticated encryption. Cryptology ePrint Archive, Report 2004/309, 2004. <http://eprint.iacr.org/>.
- [9] D. Boneh, B. Lynn, and H. Shacham. Short signatures from the Weil pairing. *J. Cryptology*, 17(4):297–319, Sept. 2004.
- [10] K. D. Bowers, A. Juels, and A. Oprea. Proofs of retrievability: Theory and implementation. In R. Sion and D. Song, editors, *Proceedings of CCSW 2009*, pages 43–54. ACM Press, Nov. 2009.
- [11] H. Cohen. *A Course in Computational Algebraic Number Theory*, volume 138 of *Graduate Texts in Mathematics*. Springer-Verlag, 1993.
- [12] R. Cramer and V. Shoup. Signature schemes based on the strong RSA assumption. *ACM Trans. Info. & System Security*, 3(3):161–85, 2000.
- [13] R. Cramer and V. Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. *SIAM J. Computing*, 33(1):167–226, 2003.
- [14] Y. Deswarte, J.-J. Quisquater, and A. Saïdane. Remote integrity checking. In S. Jajodia and L. Strous, editors, *Proceedings of IICIS 2003*, volume 140 of *IFIP*, pages 1–11. Kluwer Academic, Jan. 2004.
- [15] Y. Dodis, S. Vadhan, and D. Wichs. Proofs of retrievability via hardness amplification. In O. Reingold, editor, *Proceedings of TCC 2009*, volume 5444 of *LNCS*, pages 109–27. Springer-Verlag, Mar. 2009.
- [16] D. Freeman, M. Scott, and E. Teske. A taxonomy of pairing-friendly elliptic curves. *J. Cryptology*, 23(2):224–80, Apr. 2010.
- [17] D. Gazzoni Filho and P. Barreto. Demonstrating data possession and uncheatable data transfer. Cryptology ePrint Archive, Report 2006/150, 2006. <http://eprint.iacr.org/>.
- [18] O. Goldreich. A sample of samplers: A computational perspective on sampling. In *Studies in Complexity and Cryptography: Miscellanea on the Interplay between Randomness and Computation*, volume 6650 of *LNCS*, pages 302–32. Springer-Verlag, 2011.
- [19] L. Guillou and J.-J. Quisquater. A practical zero-knowledge protocol fitted to security micro-processor minimizing both transmission and memory. In C. Günther, editor, *Proceedings of Eurocrypt 1988*, volume 330 of *LNCS*, pages 123–8. Springer-Verlag, May 1988.
- [20] V. T. Hoang, B. Morris, and P. Rogaway. An enciphering scheme based on a card shuffle. In R. Safavi-Naini, editor, *Proceedings of Crypto 2012*, LNCS. Springer-Verlag, Aug. 2012. To appear.
- [21] C. Huang, H. Simitci, Y. Xu, A. Ogus, B. Calder, P. Gopalan, J. Li, , and S. Yekhanin. Erasure coding in Windows Azure storage. In G. Heiser and W. Hsieh, editors, *Proceedings of USENIX ATC 2012*. USENIX, June 2012.

- [22] A. Juels and B. Kaliski. PORs: Proofs of retrievability for large files. In S. De Capitani di Vimercati and P. Syverson, editors, *Proceedings of CCS 2007*, pages 584–97. ACM Press, Oct. 2007. Full version: <http://www.rsa.com/rsalabs/node.asp?id=3357>.
- [23] M. Lillibridge, S. Elnikety, A. Birrell, M. Burrows, and M. Isard. A cooperative Internet backup scheme. In B. Noble, editor, *Proceedings of USENIX Technical 2003*, pages 29–41. USENIX, June 2003.
- [24] M. Liskov, R. Rivest, and D. Wagner. Tweakable block ciphers. *J. Cryptology*, 24(3):588–613, July 2011.
- [25] M. Mitzenmacher. Digital fountains: A survey and look forward. In R. Calderbank and A. Orlitsky, editors, *Proceedings of ITW 2004*, pages 271–76. IEEE Information Theory Society, Oct. 2004.
- [26] M. Naor and G. Rothblum. The complexity of online memory checking. *J. ACM*, 56(1), Jan. 2009.
- [27] M. Rabin. Efficient dispersal of information for security, load balancing, and fault tolerance. *J. ACM*, 36(2):335–48, Apr. 1989.
- [28] L. Rizzo. Effective erasure codes for reliable computer communication protocols. *ACM SIGCOMM Computer Communication Rev.*, 27(2):24–36, Apr. 1997.
- [29] T. Schwarz and E. Miller. Store, forget, and check: Using algebraic signatures to check remotely administered storage. In M. Ahamad and L. Rodrigues, editors, *Proceedings of ICDCS 2006*. IEEE Computer Society, July 2006.
- [30] M. Shah, M. Baker, J. Mogul, and R. Swaminathan. Auditing to keep online storage services honest. In G. Hunt, editor, *Proceedings of HotOS 2007*. USENIX, May 2007.
- [31] M. Shah, R. Swaminathan, and M. Baker. Privacy-preserving audit and extraction of digital contents. Cryptology ePrint Archive, Report 2008/186, 2008. <http://eprint.iacr.org/>.

A Erasure Codes

The proof-of-retrievability schemes we have presented allow a client of a storage server to be sure that he can recover an ϵ fraction of the stored blocks. Of course, clients would like to recover all their data, not a fraction of it, so stored files must be redundantly encoded such that any ϵ fraction of the redundant encoding allows reconstruction of the file’s contents. Erasure codes are the codes that provide this property [27, 2]. In this section, we briefly note the properties we require from erasure codes. For more on erasure codes, see the brief survey by Mitzenmacher [25]; for more on their use in storage systems, see the recent paper by Huang et al. [21].

Because an adversarial server can choose what blocks to “remember” and what blocks to “forget,” it is crucial that the erasure code be resilient against *adversarial erasure*. Reed-Solomon–style erasure codes can be constructed for arbitrary rates allowing recovery of the original file from any fraction of the encoded file blocks [28]. The code matrix used can be made public and any user can apply the decoding procedure. This provides public retrievability.

The downside of Reed-Solomon codes is the time required for encoding and decoding. For an n -block file, both of these procedure will take $O(n^2)$ time. For outsourced storage, n can be very large. For example, if a block is 1000 bytes and the file being stored is 1 GB, then we have $n \approx 2^{20}$.

Although one would like decoding to have performance linear in n , no codes are known that provide linear decoding time in the presence of adversarial erasure. Instead, current linear-time erasure codes are secure against *random erasure*: they allow reconstruction of the original file from an ϵ fraction of the encoded blocks with overwhelming probability [25]. To make use of these codes, we scramble the encoded file blocks so the server can do no better than randomly erasing blocks. It is crucial that the server not learn the secrets used for this scrambling step, which unfortunately makes public retrievability impossible.

Our proposed scrambling operation is essentially the same as that proposed by Ateniese et al. [3, Section 4.2.2], and we refer the reader to that paper for more details and a full analysis.

First, encode the file using the linear-time code. Second, permute the blocks of the file using a pseudorandom permutation over the domain $[1, n]$, where n is the number of blocks in the encoded file. (See Hoang, Morris, and Rogaway [20] for details on how to construct such a permutation.) Third, encrypt each block independently using a tweakable block cipher [24], with the the block index as tweak. Store the blocks output by this procedure on the server.

Now consider an adversary \mathcal{A} that is given a file scrambled according to this procedure, and can erase all but an ϵ fraction of the blocks. We argue that each block in the original encoded file is retained with probability ϵ , assuming the pseudorandom permutation and block cipher are secure. That is, \mathcal{A} can do no better than random erasure.

In Game 0, we play the erasure game with the adversary. In Game 1, we replace the pseudorandom permutation with a truly random permutation over $[1, n]$. If \mathcal{A} behaves differently in Game 1 and Game 2, we can construct an adversary that breaks the security of the pseudorandom permutation. In Game 3, we replace the encrypted block with truly random blocks. If \mathcal{A} behaves differently in Game 2 and Game 3, we can construct an adversary that breaks the security of the block cipher, by means of a hybrid argument. Note that, without the tweak, identical plaintext blocks would encrypt to identical ciphertext blocks, so this argument wouldn't apply. But now in Game 3 the permutation applied to the encoded blocks is independent of \mathcal{A} 's view. Thus no adversary that erases blocks can do better than random erasure, which is exactly the property we require for decoding to work with overwhelming probability.

It is important to note that our model, above, does not consider the access pattern for the file blocks. It is possible that the block accesses made by a user, in reading or in reconstructing her file, leak information about the correlation between plaintext blocks. In this case, the server might be able to do better than guessing in choosing which blocks to delete. Note that our proof-of-retrievability protocol queries blocks at random, as does the extraction algorithm used in our proofs, so neither leaks information to the adversary. If proofs of retrievability are used as part of a larger system where individual file blocks will be accessed, then codes secure against adversarial erasure should be used instead.

B Are Our Schemes Still Secure Without B Coefficients?

In both schemes proposed in Section 3, the verifier sends with each index i of a query a coefficient ν_i from a set B . If we could avoid sending these coefficients—equivalently, if we could set $B = \{1\}$ —then we would obtain a scheme that is more efficient in several respects:

- the verifier would need to flip fewer coins in generating a query;
- the query would be shorter by $l \cdot \lg \#B$ bits;
- the prover’s computation would be greatly reduced: essentially, one multiplication instead of $l + 1$ multiplications in the first scheme; one exponentiation instead of $l + 1$ exponentiations in the second scheme¹⁶; and
- the verifier’s computation would also be reduced, though not so dramatically.

Unfortunately, it is clear from Lemma 4.5 that the proof techniques of Section 4.2 cannot apply when $\#B = 1$, since we will not then have $\epsilon > 1/\#B$, however large the adversary’s success probability ϵ is.

This is not just a proof problem. Below, we present an attack on the schemes of Section 3 when $B = \{1\}$. In the attack, the server stores $n - 1$ blocks instead of n , can answer a nonnegligible fraction of all queries, yet no extraction technique can recover *any* of the original blocks.

Note that our argument is relevant only to those schemes, like those we presented in Section 3, in which the server’s response consists of a linear combination of file blocks. If individual blocks are returned, as in the simple scheme of Section 5, then no coefficients are necessary.

Note also that the scheme we attack is closely related to the “E-PDP” efficient alternative scheme given by Ateniese et al. [3]. For their E-PDP scheme, Ateniese et al. claim only that the protocol establishes that a cheating prover has the sum $\sum_{i \in I} m_i$ of the blocks. Our attack suggests that this guarantee is insufficient for recovering file contents.

A note on notation. In this section, we will make some simplifications to the notation for the sake of brevity and clarity. First, observe that the Part-1 proofs of Section 4.1 *do* apply in the case that $B = \{1\}$. We will thus elide the authenticators $\{\sigma_i\}$ in our attack; this allows us to address both the scheme with private verification and the scheme with public verification simultaneously. Second, we will set the number of sectors per block, s , to 1. Our attack easily generalizes to the case $s > 1$, but this simplification allows us to eliminate a subscript and simplify the presentation.

The attack. With the simplifications above, consider an n -block file with blocks (m_1, \dots, m_n) . A query will consist of l indices $I \subset [1, n]$; the response will be $\mu = \sum_{i \in I} m_i$. We assume that l is even.¹⁷

The adversary chooses an index i^* at random from $[1, n]$. For each $i \neq i^*$, he chooses $\tau_i \xleftarrow{R} \{-1, +1\}$ and sets

$$m'_i \leftarrow m_i + \zeta_i m_{i^*} .$$

Now the adversary remembers $(m'_1, \dots, m'_{i^*-1}, m'_{i^*+1}, m'_n)$. Clearly, the adversary needs to store one less block than an honest server would.

Now, consider a query I . If $i^* \notin I$, the adversary responds with $\mu' = \sum_{i \in I} m'_i$. Otherwise, $i^* \in I$, and the adversary responds with $\mu' = \sum_{i \in I \setminus \{i^*\}} m'_i$.

In our analysis, we will use the following simple lemma:

¹⁶The l element summation computed by the prover in the first scheme requires work comparable to a multiplication when $l \approx \lg p$.

¹⁷If l is odd, the adversary can set $m'_i \leftarrow m_i + (l^{-1} \bmod p)(m_{i^*})$, which allows him to respond to that l/n fraction of queries where $i^* \in I$.

Lemma B.1 ([1], p. 12). For $k \geq 2$ we have $\binom{k}{\lfloor k/2 \rfloor} \geq 2^k/k$.

Proof. $\binom{k}{\lfloor k/2 \rfloor}$ is the largest of the k values $\binom{k}{1}, \binom{k}{2}, \dots, \binom{k}{k-1}$, and $\binom{k}{0} + \binom{k}{l}$; and so it must be at least as large as their average, $2^k/k$. \square

In the case $i^* \notin I$, μ' will be correct provided that we have $\sum_{i \in I} \zeta_i = 0$. But this happens when the number of +1s and the number of -1s are equal in $\{\zeta_i\}_{i \in I}$, and this happens with probability

$$\binom{l}{l/2} \cdot \left(\frac{1}{2}\right)^{l/2} \cdot \left(\frac{1}{2}\right)^{l/2} \geq \left(\frac{2^l}{l}\right) \cdot \left(\frac{1}{2^l}\right) = \frac{1}{l}.$$

In the case $i^* \in I$, μ' will be correct provided we have $\sum_{i \in I \setminus \{i^*\}} \zeta_i = 1$. This happens when there are $(l/2 - 1)$ -1s and $(l/2)$ +1s in $\{\zeta_i\}_{i \in I \setminus \{i^*\}}$, and this happens with probability

$$\binom{l-1}{l/2-1} \cdot \left(\frac{1}{2}\right)^{l/2-1} \cdot \left(\frac{1}{2}\right)^{l/2} = \binom{l-1}{\lfloor (l-1)/2 \rfloor} \cdot \left(\frac{1}{2}\right)^{l-1} \geq \left(\frac{2^{l-1}}{l-1}\right) \cdot \left(\frac{1}{2^{l-1}}\right) = \frac{1}{l-1} > \frac{1}{l}.$$

Thus the adversary can respond to $1/l$ fraction of queries where $i^* \in I$ and to a $1/l$ fraction of queries where $i^* \notin I$; so he can respond to a $1/l$ fraction of *all* queries, which is clearly nonnegligible.

But now it is impossible for any extraction strategy to recover *any* block, let alone a ρ fraction of all blocks. This is because the subspace known to the adversary is insufficient to determine any block. Indeed, the adversary's knowledge is consistent with any value for any block. Fix $(m'_1, \dots, m'_{i^*-1}, m'_{i^*+1}, m'_n)$ where $m'_i = m_i + \zeta_i m_{i^*}$. Suppose we believe that $m_{i^*} = a$ for any value a . This fixes m_i for each $i \neq i^*$, as $m_i = m'_i - \zeta_i m_{i^*}$. If we believe, for some index $\lambda \neq i^*$, $m_\lambda = a$, then m_{i^*} is fixed because $m'_\lambda = m_\lambda + \zeta_\lambda m_{i^*}$ implies $m_{i^*} = (\zeta_\lambda)(m'_\lambda - a)$, and the argument proceeds as before. Since the adversary's knowledge is consistent with any choice of value for any (single) block, it cannot be the case that it allows recovery of the value of any block.