

DSKE: Digital Signatures with Key Extraction

Zhipeng Wang¹, Orestis Alpos², Alireza Kavousi³, Sze Yiu Chau⁴,
Duc V. Le⁵, and Christian Cachin²

¹*Imperial College London*

²*University of Bern*

³*University College London*

⁴*The Chinese University of Hong Kong*

⁵*VISA Research*

Abstract

This work introduces DSKE, digital signatures with key extraction. In a DSKE scheme, the private key can be extracted if more than a threshold of signatures on different messages are ever created while, within the threshold, each signature continues to authenticate the signed message. We give a formal definition of DSKE, as well as two provably secure constructions, one from hash-based digital signatures and one from polynomial commitments.

We demonstrate that DSKE is useful for various applications, such as spam prevention and deniability. First, we introduce the GroupForge signature scheme, leveraging DSKE constructions to achieve deniability in digital communication. GroupForge integrates DSKE with a Merkle tree and timestamps to produce a “short-lived” signature equipped with extractable sets, ensuring deniability under a fixed public key. We illustrate that GroupForge can serve as a viable alternative to Keyforge in the non-attributable email protocol of Specter, Park, and Green (USENIX Sec ’21), thereby eliminating the need for continuous disclosure of outdated private keys. Second, we leverage the inherent extraction property of DSKE to develop a Rate-Limiting Nullifier (RLN) scheme. RLN efficiently identifies and expels spammers once they exceed a predetermined action threshold, thereby jeopardizing their private keys.

Moreover, we implement both variants of the DSKE scheme to demonstrate their performance and show it is comparable to existing signature schemes. We also implement GroupForge from the polynomial commitment-based DSKE and illustrate the practicality of our proposed method.

1 Introduction

Digital signature schemes [GMR88] play an important role in protecting the integrity of data transmitted over the Internet. In some jurisdictions [Bly05, Mas16, Kar19], a digital signature applied to data can serve as evidence of the sender’s authorship of the data. Moreover, the signature of a message remains valid until either the underlying signature scheme is broken or the private key is compromised. However, as pointed out by Borisov, Goldberg, and Brewer in their work on off-the-record (OTR) communication work [BGB04], this “long-lived” property is unsuitable for certain types of messages. For instance, if Alice wishes to communicate privately with Bob, she can encrypt her messages using Bob’s public key (i.e., for confidentiality) and sign them with her private key (i.e., for authenticity). However, if Eve compromises Bob’s computer at some point in the future, Eve will be able to read all of Bob’s previous messages from Alice, and use the signatures to prove to Judy that the messages indeed originated from Alice.

To address this problem, OTR messaging requires an interactive key agreement protocol between the sender and the recipient to agree on session keys before exchanging messages. However, this pair-wise key agreement required in OTR is not scalable for applications such as email protocols, in which there is often no prior end-to-end interaction among the participating parties. Another way to achieve deniability is to simply require the sender to periodically rotate keys and publish their old private keys [Gre]. This method enables anyone to forge signatures using the published private keys and thus offers deniability to old transcripts. In fact, this method is being suggested to offer deniability in domain keys identified mail (DKIM) [ACD⁺07], where SMTP servers sign outgoing emails on behalf of the whole domain using a single key, as a way to safeguard against email spoofing.

A server sending an email cryptographically signs it, so that the recipient can verify that it has originated from the reported server. A side effect of this action is *email attributability* which stems from the fact that the digital signature remains valid for a long time, potentially forever [Gre]. As a result, a malicious actor, who at any time gains access to these emails, can provably link them to their sender, which in turn incentivizes extortion and retaliation, among others. In this paper, we focus on answering the following question that arises naturally from the limitations of existing attempts:

Is it possible to design a signature scheme that allows the recipients to verify the validity of the signature, while enabling the sender to gain plausible deniability, without requiring the constant publication of old key materials or any additional components?

It is worth noting that the question itself is, seemingly, a contradiction due to the non-repudiation property of digital signature schemes. In this

work, we propose digital signatures with key extraction (DSKE) scheme to circumvent the aforementioned question. We first propose a general framework, showing that any one-time hash-based signature scheme, such as Lamport [Lam79] and Winternitz OTS [Mer89, BDE⁺11], can be turned into a DSKE scheme, and present two concrete constructions. We then propose another DSKE construction based on polynomial commitments [KZG10] that offers properties like a flexible threshold and uniqueness. This type of signature scheme comes with an *extractable set*, a set of signatures from which the private key can be extracted *asynchronously*.

Independently, the key extraction property of DKSE is particularly useful when there is a need to disincentivize or penalize the creation of more than a certain number of signatures. As concrete examples, one may consider the issue of double-spending [RKS15] or double-signing on executable code [DRS18]. Both could be mitigated by putting the user under the threat of key leakage.¹

Contributions. Our contributions can be summarized as follows:

- We formally define the notion of DSKE. DSKE comes with an *extractable set*, a set (with a predefined set size) of signatures that can be used to extract the private key.
- We introduce a one-time hash-based DSKE, denoted as $\text{DSKE}_{\text{hash}}$. As demonstrated, generating multiple signatures while reusing a private key in one-time hash-based signatures results in a *probabilistic* extractable set with an overwhelming probability. We provide two concrete $\text{DSKE}_{\text{hash}}$ constructions based on Lamport [Lam79] and Winternitz signatures [Mer89, BDE⁺11].
- We propose a polynomial commitment-based DSKE, $\text{DSKE}_{\text{poly}}$. Compared to $\text{DSKE}_{\text{hash}}$, $\text{DSKE}_{\text{poly}}$ offers shorter signature size with *deterministic* extractable set. Additionally, it empowers signers to customize the size of the extractable set without compromising security. We further provide a concrete $\text{DSKE}_{\text{poly}}$ construction based on the KZG commitment scheme [KZG10].
- We use the $\text{DSKE}_{\text{poly}}$ construction to propose two concrete applications. First, we present GroupForge, a construction combining DSKE with a Merkle tree and timestamps. It yields a *short-lived* signature with extractable sets for deniability under a fixed public key. GroupForge can replace Keyforge in the non-attributable email protocol by Specter, Park, and Green [SPG21], eliminating the need for continual disclosure of outdated private keys. Second, we harness the inherent extraction property

¹We remark that DSKE in this sense resembles an existing primitive in the literature, named double-authentication-preventing signatures (DAPS). However, there are crucial differences between the two families of signatures, as elaborated in Section 8.

of DSKE to build a robust Rate-Limiting Nullifier (RLN) protocol, which can identify and expel spammers once they exceed a certain threshold in a certain action, placing their private keys in jeopardy.

- We implement and evaluate both DSKE constructions and the Group-Forge construction, thereby highlighting their practicality.

Paper Organization The remainder of the paper is organized as follows. Section 2 outlines the necessary background and building blocks for DSKE. In Section 3, we present the formal definition of DSKE. Sections 4 and 5 provide concrete constructions of DSKE, namely $\text{DSKE}_{\text{hash}}$, based on hash-based signature schemes, and $\text{DSKE}_{\text{poly}}$, based on polynomial commitment schemes, respectively. We show DSKE’s applications to RLN and Group-Forge in Section 6. We evaluate the performance of our constructions and discuss DSKE’s limitations and potential improvements in Section 7. Section 8 discusses the related work and Section 9 concludes our work.

2 Preliminaries

2.1 Notation

We denote by 1^λ the security parameter and by $\text{negl}(\lambda)$ a negligible function of λ . We express by (pk, sk) a pair of public and private keys. We let $[n]$ denote the set $\{1, \dots, n\}$. Moreover, we require that pk can always be efficiently derived from sk , and we denote $\text{extractPK}(sk) = pk$ to be the deterministic function for doing so. We denote as $\mathbb{Z}_{\geq a}$ the set of integers that are greater than or equal to a . For a field \mathbb{F} , we denote $\mathbb{F}^{\leq d}(X)$ the set of polynomials in $\mathbb{F}[X]$ with degree at most d . We denote by \mathcal{M} the message space and \mathcal{S} the signature space. For a non-negative integer j , we let $f^{(j)}$ be the j -th iterate of the function f , i.e., $f^{(j)}(x) = f(f(\dots f(x)\dots))$ where f is repetitively calculated j times.

2.2 Hash Functions

Our constructions employ the following standard properties of cryptographic hash functions. We use $H : \mathcal{K} \times \mathcal{M} \rightarrow \{0, 1\}^\lambda$ to denote a family of hash functions that is parameterized by a key $k \in \mathcal{K}$ and message $m \in \mathcal{M}$ and outputs a binary string of length λ . For this work, we consider cryptographic hash functions [RS04], satisfying preimage resistance and collision resistance properties.

Definition 1 (Preimage Resistance). *A family H of hash functions is preimage-resistant, if for any PPT adversary \mathcal{A} , the adversary’s advantage in finding*

the preimage of a given hash value is:

$$\Pr \left[\begin{array}{l} k \xleftarrow{\$} \mathcal{K}, x \xleftarrow{\$} \mathcal{M}, y \leftarrow H(k, x) : H(k, x') = y \\ x' \leftarrow \mathcal{A}(k, y) \end{array} \right] \leq \text{negl}(\lambda)$$

Definition 2 (Collision Resistance). *A family H of hash functions is collision-resistant, if for any PPT adversary \mathcal{A} , the adversary's advantage in finding collisions is:*

$$\Pr \left[\begin{array}{l} k \xleftarrow{\$} \mathcal{K} \\ (x, x') \leftarrow \mathcal{A}(k) \end{array} : (x \neq x') \wedge (H(k, x) = H(k, x')) \right] \leq \text{negl}(\lambda)$$

In practice, the key for standard hash functions is public; therefore, from this point, we refer to the cryptographic hash function h sampled from a family of hash functions as a fixed function $h : \mathcal{M} \rightarrow \{0, 1\}^\lambda$.

2.3 Polynomial Commitment Schemes

A polynomial commitment scheme (PCS) allows a prover to commit to a polynomial $f(X) \in \mathbb{F}^{\leq d}(X)$ and later open $f(X)$ at arbitrary points x , revealing only the value $f(x)$.

Definition 3 (Polynomial Commitment). *A PCS consists of the following algorithms.*

- $(ck, vk) \leftarrow \text{Setup}(1^\lambda, d)$: *The setup algorithm takes as input a security parameter λ , a maximum degree $d \in \mathbb{N}$, and outputs the public commitment key ck , which allows committing to polynomials in $\mathbb{F}^{\leq d}(X)$, and the public verification key vk .*
- $C_f \leftarrow \text{Com}(ck, f(X))$: *The commitment algorithm takes as input the commitment key ck , a polynomial $f(X) \in \mathbb{F}^{\leq d}(X)$, and outputs a commitment $C_f \in G$ to the polynomial $f(X)$.*
- $(\pi, y) \leftarrow \text{Open}(ck, C_f, x, f(X))$: *The algorithm takes as input a commitment key ck , a commitment C_f , an evaluation point x , the polynomial $f(X)$, and outputs $y = f(x) \in \mathbb{F}$ and a proof $\pi \in G$.*
- $0/1 \leftarrow \text{Check}(vk, C_f, x, y, \pi)$: *The algorithm takes as input the verification key vk , the commitment C_f , a point x , the claimed evaluation y , the opening proof π , and outputs 1 iff $y = f(x)$.*

PCS can guarantee *correctness*, *computational hiding*, *evaluation binding*, and *polynomial binding* properties from the polynomial commitment scheme.

Definition 4 (Correctness [KZG10]). Let $(ck, vk) \leftarrow \text{Setup}(1^\lambda, k)$, $f(X) \in \mathbb{F}^{\leq d}(X)$, and $C_f \leftarrow \text{Com}(ck, f(X))$. Then for any (π, y) output by $\text{Open}(ck, C_f, x, f(X))$, we have $\text{Check}(vk, C_f, x, y, \pi) = 1$.

Definition 5 (Computational Hiding [KZG10]). Given (ck, vk) , the commitment C_f , and \hat{d} valid openings (y_i, π_i) for points x_i , where $i \in \{1, \dots, \hat{d}\}$ and $\hat{d} \leq d$, no PPT adversary can determine the value $f(x')$, for $x' \notin \{x_1, \dots, x_{\hat{d}}\}$, except with a negligible probability.

Definition 6 (Evaluation Binding [KZG10]). Given (ck, vk) , no PPT adversary can compute commitment C_f , point x , and two openings (π_1, y_1) , (π_2, y_2) for x , such that $\text{Check}(vk, C, x, y_1, \pi_1) = 1$, $\text{Check}(vk, C, x, y_2, \pi_2) = 1$, and $y_1 \neq y_2$.

Definition 7 (Polynomial Binding [KZG10]). Given (ck, vk) , no PPT adversary can compute polynomials $f(X)$ and $f'(X)$, such that $f(X) \neq f'(X)$ and $\text{Com}(ck, f(X)) = \text{Com}(ck, f'(X))$.

KZG Polynomial Commitment Scheme. We now revisit KZG scheme [KZG10] as a concrete polynomial commitment construction. It works over a bilinear pairing group $\mathbb{G} = \langle e, G, G_t \rangle$, where G is a group of prime order p , e is a symmetric pairing $e : G \times G \rightarrow G_t$, and g and \hat{h} are generators of G .

- $(\mathbb{G}, ck, vk) \leftarrow \text{Setup}(1^\lambda, d)$: The algorithm outputs a representation of the bilinear group \mathbb{G} , commitment key $ck = \{g, g^\alpha, \dots, g^{\alpha^d}\}$, and verification key $vk = \hat{h}^\alpha$, for an $\alpha \in \mathbb{Z}_p$.
- $C_f \leftarrow \text{Com}(ck, f(X))$: The algorithm computes $C_f = g^{f(\alpha)}$ using ck and outputs C_f .
- $(\pi, y) \leftarrow \text{Open}(ck, C_f, x, f(X))$: The algorithm computes $y = f(x)$ and the quotient polynomial $q(X) = \frac{f(X) - y}{X - x}$, and outputs y and $\pi = C_q = \text{Com}(ck, q(X))$.
- $0/1 \leftarrow \text{Check}(vk, C_f, x, y, \pi)$: The algorithm outputs 1 if $e(C_f \cdot g^{-y}, \hat{h}) = e(C_q, \hat{h}^\alpha \cdot \hat{h}^{-x})$, and 0 otherwise.

KZG scheme satisfies *correctness*, *computational hiding*, *evaluation binding*, and *polynomial binding* properties, provided the DL and d -SDH assumptions hold in \mathbb{G} [KZG10].

3 Digital Signatures with Key Extraction (DSKE)

In this section, we formally define the notion of digital signatures with key extraction (DSKE). We adopt the standard digital signature definition and introduce a new algorithm to capture the capability of extracting the private key from a set of signatures.

Definition 8 ((k, δ) -Digital Signature with Key Extraction). A signature scheme Σ , with key extraction consists of five algorithms:

- $par \leftarrow \text{Setup}(1^\lambda)$: The setup algorithm takes a security parameter 1^λ and outputs a set of public parameters par . This algorithm runs once, and the public parameters are implicitly input to all subsequent algorithms.
- $(pk, sk) \leftarrow \text{KeyGen}(par)$: The key generation is a probabilistic algorithm that outputs a pair (pk, sk) of public and private keys.
- $\sigma \leftarrow \text{Sign}(sk, m)$: The signing algorithm is a probabilistic algorithm that takes a private key sk and a message $m \in \mathcal{M}$ as input and outputs a signature σ in the signature space \mathcal{S} .
- $b \leftarrow \text{Verify}(pk, m, \sigma)$: The verification algorithm is a deterministic algorithm that takes a public key pk , a message m , a signature σ as input, and outputs the validity of the signature, $b \in \{0, 1\}$.
- $sk \leftarrow \text{Extract}(\{(m_i, \sigma_i)\}_{i \in [k]}, pk)$: The extraction algorithm is a probabilistic algorithm that takes as input a set of distinct message-signature pairs $(m_i, \sigma_i)_{i \in [k]}$, such that $\sigma_i \leftarrow \text{Sign}(sk, m_i)$, the public key pk , and outputs the underlying private key sk with probability δ and \perp with probability $1 - \delta$.

Apart from the straightforward correctness definition, we consider two other properties of DSKE: existential unforgeability, and the existence of an extractable set. The security of digital signatures is defined through the following experiment.

d -times signature experiment $\text{SignExp}_{\mathcal{A}, \Sigma}^d(\lambda)$.

1. $\text{Setup}(1^\lambda)$ and $\text{KeyGen}()$ are run to obtain keys (pk, sk) .
2. \mathcal{A} is given pk and can ask up to d queries to the signing oracle $\text{Sign}(sk, \cdot)$. Let $Q_{\mathcal{A}}^{\text{Sign}(sk, \cdot)} = \{m_i\}_{i \in [d]}$ be the set of all messages for which \mathcal{A} queries $\text{Sign}(sk, \cdot)$, where the i^{th} query is a message $m_i \in \mathcal{M}$. Eventually, \mathcal{A} outputs a pair $(m^*, \sigma^*) \in \mathcal{M} \times \mathcal{S}$.
3. The output of the experiment is defined to be 1 if and only if $m^* \notin Q_{\mathcal{A}}^{\text{Sign}(sk, \cdot)}$ and $\text{Verify}(pk, m^*, \sigma^*) = 1$.

Definition 9 (Existential Unforgeability). A digital signature scheme Σ is existentially unforgeable under a d -times adaptive chosen-message attack, or d -times-secure, if for all PPT adversaries \mathcal{A} the success probability in the previous experiment is negligible:

$$\Pr[\text{SignExp}_{\mathcal{A}, \Sigma}^d(\lambda) = 1] \leq \text{negl}(\lambda).$$

Definition 10 (Extractable Set). *A digital signature scheme has a (k, δ) -extractable set when the extraction algorithm $\text{Extract}(\cdot)$ on input k distinct message-signature pairs $\{(m_i, \sigma_i)\}_{i \in [k]}$ and the public key pk , such that each σ_i is a valid signature on m_i under pk , outputs the private key sk with probability δ . That is:*

$$\Pr \left[\begin{array}{l} m_i \leftarrow \mathcal{M}, \text{ s.t. } m_i \neq m_j, \text{ for } i, j \in [k], i \neq j \\ par \leftarrow \text{Setup}(), (pk, sk) \leftarrow \text{KeyGen}() \\ \sigma_i \leftarrow \text{Sign}(sk, m_i) \\ sk' \leftarrow \text{Extract}(\{(m_i, \sigma_i)\}_{i \in [k]}, pk) \end{array} : pk = \text{extractPK}(sk') \right] = \delta$$

4 DSKE from Hash-Based Signature Schemes

Hash-based signature schemes (such as Lamport [Lam79] and Winternitz OTS [Mer89, BDE⁺11]) leverage the security of one-way functions to construct digital signatures. In hash-based signature schemes, the private key is often a list that is derived from a succinct seed, and based on the message, the signature usually reveals “partial” information about the private key. In essence, having a *sufficiently* large number of signatures allows us to obtain enough information for the reconstruction of the private key.

In this section, we provide DSKE constructions based on the hash-based signature schemes, denoted as $\text{DSKE}_{\text{hash}}$. We first give a generic definition to capture the private key leakage of hash-based signature schemes and then provide two DSKE constructions that are based on the Lamport signature scheme and Winternitz OTS.

Definition 11 (Leakage of Hash-based Signature). *A hash-based signature, Σ_{hash} , consists of the four algorithms Setup , KeyGen , Sign , and Verify as defined in Definition 8, as well as a leakage algorithm defined as follows:*

- $S \leftarrow \text{Leak}((m, \sigma), pk)$: *This leakage algorithm is a deterministic algorithm that takes as input a message-signature pairs (m, σ) , such that $\sigma \leftarrow \text{Sign}(sk, m)$, the public key pk , such that $pk = \text{extractPK}(sk)$, and outputs a list S , containing a fraction of the private key.*

We can then propose a hash signature-based DSKE:

Definition 12 (Hash-based DSKE). *A hash-based DSKE scheme, $\text{DSKE}_{\text{hash}}$, consists of the four algorithms, Setup , KeyGen , Sign , and Verify which are same as a hash-based signature, as well as an Extract algorithm defined as follows:*

- $sk/\perp \leftarrow \text{Extract}(\{(m_i, \sigma_i)\}_{i \in [k]}, pk)$: *The extraction algorithm is a probabilistic algorithm that receives a set of distinct message-signature pairs $\{m_i, \sigma_i\}_{i \in [k]}$, such that $\sigma_i = \text{Sign}(sk, m_i)$, and the public key pk as inputs. For each (m_i, σ_i) , this algorithm runs $S_i \leftarrow \text{Leak}((m_i, \sigma_i), pk)$, and*

computes $sk' = \bigcup_{i=1}^k (S_i)$. The algorithm outputs the private key sk' if $pk = \text{extractPK}(sk')$; otherwise, it outputs \perp .

We subsequently introduce a definition aimed at quantifying the probability that an element of the private key is not revealed even after receiving k hash-based signatures. This definition serves as a tool to compute the probability of a successful output of the private key by running the $\text{Extract}()$ function.

Definition 13. Given k distinct messages $\{m_i\}_{i \in [k]}$ which are randomly sampled from \mathcal{M} , a key pair (sk, pk) generated by $\text{KeyGen}()$, k signatures $\{\sigma_i\}_{i \in [k]}$ such that each σ_i is a valid signature on m_i under pk , and an element sk_j which are sampled from the secret key sk , we define p_{leak} as the probability that sk_j is not leaked after running $\text{Leak}()$ on all message-signature pairs $\{(m_i, \sigma_i)\}_{i \in [k]}$.

$$p_{\text{leak}} = \Pr \left[\begin{array}{l} m_i \xleftarrow{\$} \mathcal{M}, \text{ for } i \in [k] \\ par \leftarrow \text{Setup}(), (pk, sk) \leftarrow \text{KeyGen}() \\ \sigma_i \leftarrow \text{Sign}(sk, m_i) \\ S_i \leftarrow \text{Leak}((m_i, \sigma_i), pk) \\ sk_j \xleftarrow{\$} sk \end{array} : sk_j \notin \bigcup_{i=1}^k S_i \right]$$

The following theorem shows that the lower bound of δ in any hash-based DSKE scheme can be determined by p_{leak} .

Theorem 1 (Extractable Set). Given a hash-based DSKE scheme with a private key of size λ , if the underlying hash function is modeled as a random oracle, then $\text{DSKE}_{\text{hash}}$ has a (k, δ) -extractable set, where:

$$\delta \geq 1 - \lambda \cdot p_{\text{leak}} - \text{negl}(\lambda)$$

Proof. Given a private key of a hash function, which is a list of λ elements, denoted as $sk = (sk_j)_{j \in [\lambda]}$. We denote Bad_0 to be the event that at least one element of sk is not included in $\bigcup_{i=1}^k \text{Leak}((m_i, \sigma_i), pk)$. We have:

$$\begin{aligned} \Pr[\text{Bad}_0] &= \Pr \left[\bigvee_{j=1}^{\lambda} \left(sk_j \notin \bigcup_{i=1}^k \text{Leak}((m_i, \sigma_i), pk) \right) \right] \\ &\leq \sum_{j=1}^{\lambda} \Pr \left[sk_j \notin \bigcup_{i=1}^k \text{Leak}((m_i, \sigma_i), pk) \right] = \lambda \cdot p_{\text{leak}} \end{aligned}$$

Moreover, we denote Bad_1 to be the event that there are two private key elements that map to the same public key element. However, since the underlying hash is modeled as a random oracle, this probability is negligible; hence, $\Pr[\text{Bad}_1] = \text{negl}(\lambda)$.

Therefore, the probability that the extraction algorithm $\text{Extract}(\cdot)$ outputs a private key sk' that $pk = \text{extractPK}(sk')$ is $\delta \geq 1 - \Pr[\text{Bad}_0 \vee \text{Bad}_1] \geq 1 - \lambda \cdot p_{\text{leak}} - \text{negl}(\lambda)$. \square

In the following, we provide two DSKE constructions that are based on the Lamport signature scheme and Winternitz OTS, respectively. We will show that, for a Lamport signature scheme, \mathfrak{p}_{leak} is determined by the number of given signatures (i.e., k), $\mathfrak{p}_{leak} = \frac{1}{2^{k-1}}$ (see Section 4.1); and for a Winternitz OTS, $\mathfrak{p}_{leak} = \frac{w^k}{(w+1)^k}$, where w are a parameter for the Winternitz OTS (see Section 4.2).

4.1 Lamport Signature-Based Construction

The Lamport signature scheme is parameterized by a preimage-resistant hash function f and a collision-resistant hash function H . The private key SK contains 2λ binary strings uniformly sampled from $\{0, 1\}^\lambda$, where λ is the security parameter. The public key PK consists of 2λ binary strings that are evaluations of f on each element in the private key. To form a signature σ on a message m , the signer reveals components of the private key SK , according to the binary representation of $H(m)$ as the signature. Verification is carried out naturally; the verifier uses the binary representation of the digest, $H(m)$, to validate if each element contained in the signature is the actual preimage of the public key elements.

Extracting Private Key from Message-Signature Pairs. As shown in Figure 1, the intuition behind the extracting function is that each signature σ_i is a subset of the private key SK . To extract the complete private key, one needs to compute the union of k different signatures $\{\sigma_i\}_{i \in [k]}$. Since the underlying hash function is unpredictable, the probability of successfully extracting the private key depends on the number of distinct message-signature pairs (i.e., k). Based on the above intuitions, we propose a hash-based construction, $\text{DSKE}_{\text{Lamp}}$, as follows.

Lamport-Based DSKE. A Lamport-based DSKE scheme, $\text{DSKE}_{\text{Lamp}}$, consists of:

- **Setup**(1^λ): On input the security parameter λ , the algorithm outputs a parameter, par , containing a hash function, $f : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$, chosen from a family of preimage-resistant hash functions, and a hash function, $H : \mathcal{M} \rightarrow \{0, 1\}^\lambda$, chosen from a family of collision-resistant hash function. The public parameters are implicitly input to all subsequent algorithms.
- **KeyGen**(\cdot): For each $i \in [\lambda]$, $b \in \{0, 1\}$, sample $sk_i[b] \xleftarrow{\$} \{0, 1\}^\lambda$; output the private key, $SK = (sk_i[b])_{i \in [\lambda], b \in \{0, 1\}}$, and the public key, $PK = (pk_i[b])_{i \in [\lambda], b \in \{0, 1\}}$, where $pk_i[b] = f(sk_i[b])$.
- **Sign**(SK, m): Parse $SK = (sk_i[b])_{i \in [\lambda], b \in \{0, 1\}}$, compute $d = H(m) = (d_i)_{i \in [\lambda]}$, and output $\sigma = (sk_i[d_i])_{i \in [\lambda]}$.
- **Verify**(PK, m, σ): Parse PK , σ , and compute $d = H(m) = (d_i)_{i \in [\lambda]}$. For all $i \in [\lambda]$, if $f(\sigma_i) \neq pk_i[d_i]$, return 0. Otherwise, return 1.

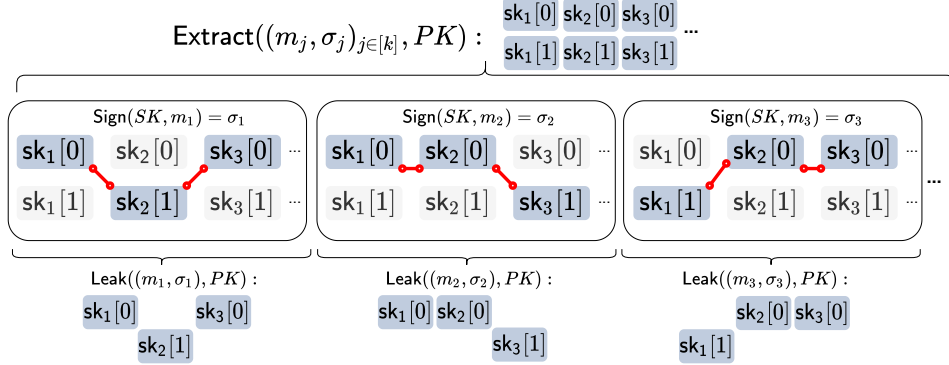


Figure 1: Example of the algorithm `Extract` for $\text{DSKE}_{\text{lamp}}$. The elements of a signature σ are linked via red lines. In this example, the original private key can be collectively reconstructed from $\{(m_i, \sigma_i)\}_{i \in [2]}$.

- $\text{Extract}(\{m_j, \sigma_j\}_{j \in [k]}, PK)$: For each message-signature pair (m_j, σ_j) , compute $d_j = H(m_j) = (d_{ji})_{i \in [\lambda]}$, and parse $\sigma_j = (\sigma_{ji})_{i \in [\lambda]}$. For each $j \in [k], i \in [\lambda], b \in \{0, 1\}$, if $\exists d_{ji} = b$, then let $sk_i[b] = \sigma_{ji}$. Set $SK = (sk_i[b])_{i \in [\lambda], b \in \{0, 1\}}$ and parse $PK = (pk_i[b])_{i \in [\lambda], b \in \{0, 1\}}$. If for all $i \in [\lambda]$ and $b \in \{0, 1\}$, $pk_i[b] = f(sk_i[b])$, then output SK . Otherwise, the algorithm outputs \perp .

Leakage Function for Lamport Signatures. As illustrated in Figure 1, within a Lamport signature scheme, each signature encapsulates half of the information about the privacy key SK . The leakage function is employed to extract the divulged information of the privacy key from the signature.

- $\text{Leak}((m, \sigma), PK)$: Check if $\text{Verify}(PK, m, \sigma) = 1$; otherwise, output \perp . Compute $d = H(m) = (d_i)_{i \in [\lambda]}$, and parse $\sigma = (\sigma_i)_{i \in [\lambda]}$. For each $i \in [\lambda], b \in \{0, 1\}$, if $d_i = b$, then let $S_i[b] = \sigma_i$ and $S_i[1 - b] = \perp$. Output $S = (S_i[b])_{i \in [\lambda], b \in \{0, 1\}}$.

In the following, we prove that $\text{DSKE}_{\text{lamp}}$ satisfies unforgeability and has an extractable set depending on the number of distinct message-signature pairs.

Theorem 2 (Existential Unforgeability). *$\text{DSKE}_{\text{lamp}}$ is existentially unforgeable under a 1-time adaptive chosen-message attack, that is,*

$$\Pr[\text{SignExp}_{\mathcal{A}, \text{DSKE}_{\text{lamp}}}^1(\lambda) = 1] \leq \text{negl}(\lambda)$$

Proof. **The t -Repeated One-Way Problem.** We first define the attack game of t -repeated one-way problem. Let f be a one-way function over $(\mathcal{X}, \mathcal{Y})$. For a given positive integer t and a given adversary \mathcal{A}_1 , the game runs as follows:

- The challenger samples $x_1, \dots, x_t \xleftarrow{\$} \mathcal{X}$, computes $y_1 \leftarrow f(x_1), \dots, y_t \leftarrow f(x_t)$, and sends (y_1, \dots, y_t) to the adversary \mathcal{A}_1 .
- The adversary \mathcal{A}_1 makes a sequence of reveal queries by sending indexes from $(1, \dots, t)$ to the challenger. Upon receiving an index j , the challenger returns x_j to \mathcal{A}_1 .
- \mathcal{A}_1 outputs (j^*, x) , where $j^* \in (1, \dots, t)$ and $x \in \mathcal{X}$.

We say that the adversary \mathcal{A}_1 wins the game if index j^* is not among \mathcal{A}_1 's reveal queries, and $f(x) = y_{j^*}$. We denote $\text{rOWadv}[\mathcal{A}_1, f, t]$ as the probability that \mathcal{A}_1 wins the game.

According to Lemma 13.5 in [BS20], for every t -repeated one-way problem adversary \mathcal{A}_1 , there exists an adversary \mathcal{A}_0 that can leverage \mathcal{A}_1 as a subroutine to break the preimage resistance of the function f . Moreover, let $\text{OWadv}[\mathcal{A}_0, f]$ be the probability that \mathcal{A}_0 breaks the preimage resistance, we have $\text{rOWadv}[\mathcal{A}, f, t] \leq t \cdot \text{OWadv}[\mathcal{A}_0, f]$.

We now consider the adversary \mathcal{A} that wins the **1-time signature experiment** $\text{SignExp}_{\mathcal{A}, \Sigma}^1(\lambda)$, and show that \mathcal{A} can be used to solve the repeated one-way problem for f .

We construct an adversary \mathcal{B} that uses \mathcal{A} to win the repeated one-way game as follows:

- The repeated one-way challenger \mathcal{C} gives \mathcal{B} a list of λ elements $y_1, \dots, y_\lambda \in \{0, 1\}^\lambda$. \mathcal{B} aims to invert one of the elements.
- \mathcal{B} sends (y_1, \dots, y_λ) as the public key pk to \mathcal{A} . Note that pk is indistinguishable from a public key generated by $\text{KeyGen}()$.
- Upon receiving the signing request of a message m from \mathcal{A} , \mathcal{B} requests from \mathcal{C} the preimages of all i that $i \in H(m)$, and sends these preimages as the signature to \mathcal{A} .
- Eventually, \mathcal{A} outputs a forgery σ^* for a message m^* , which is not already requested by \mathcal{A} , i.e., $m^* \neq m$. Let bad_1 be the event that $H(m^*) \neq H(m)$. In this event, there exists some $j \in (1, \dots, \lambda)$ which is not requested by \mathcal{B} , and if σ^* is a valid signature on m^* then σ^* contains a preimage x_j of y_j . \mathcal{B} then outputs (j, x_j) as its solution to the repeated one-way problem.

Therefore, we have

$$\begin{aligned}
& \Pr[\text{SignExp}_{\mathcal{A}, \text{DSKE}_{\text{Iamp}}}^1(\lambda) = 1] \\
& \leq (1 - \Pr[\text{bad}_1]) \cdot \text{rOWadv}[\mathcal{A}, f, t] + \Pr[\text{bad}_1] \\
& \leq \lambda \cdot \text{OWadv}[\mathcal{A}_0, f] + \text{negl}(\lambda) \\
& \leq \text{negl}(\lambda)
\end{aligned}$$

□

Lemma 1. *In a Lamport-based DSKE scheme, if $H(\cdot)$ is modeled as a random oracle, the probability that an element of the private key will not be leaked after receiving k signatures is $\mathbf{p}_{leak} = \frac{1}{2^{k-1}}$.*

Proof. We define $d_i = H(m_i)$. All d_i for $i \in \{1, \dots, k\}$ forms a $k \times \lambda$ matrix:

$$\begin{pmatrix} H(m_1) \\ \dots \\ H(m_k) \end{pmatrix} = \begin{pmatrix} d_1 \\ \dots \\ d_k \end{pmatrix} = \begin{pmatrix} d_{1,1} & \dots & d_{1,\lambda} \\ \dots & \ddots & \dots \\ d_{k,1} & \dots & d_{k,\lambda} \end{pmatrix}$$

where $d_{i,j} \in \{0, 1\}$.

For $j \in [\lambda], b \in \{0, 1\}$, to extract the value of $sk_j[b]$, there should be at least one d_{ij} that satisfies $d_{ij} = b$, i.e., for $j \in [\lambda]$, d_{ij} should not all be $1 - b$. Therefore, to be able to extract an element sk_j , we do not want any columns in the matrix that contain all 0s or all 1s.

Since $H(\cdot)$ is modeled as a random oracle, \mathbf{p}_{leak} is the probability that one single column consists of all 0s or all 1s:

$$\mathbf{p}_{leak} = \Pr \left[x \stackrel{\$}{\leftarrow} \{0, 1\}^k : x = 0^k \vee x = 1^k \right] = \frac{1}{2^{k-1}}$$

□

By combining Theorem 1 and Lemma 1, we can deduce the implications outlined in Theorem 3.

Theorem 3 (DSKE_{lamp} Extractable Set). *If $H(\cdot)$ is modeled as a random oracle, then DSKE_{lamp} has a (k, δ) -extractable set, where:*

$$\delta \geq 1 - \frac{\lambda}{2^{k-1}} - \text{negl}(\lambda)$$

4.2 Winternitz OTS-Based Construction

In the following, we demonstrate how hash-based Winternitz can also be a good candidate for DSKE construction.

Domination Free Function. A domination free function $P : \mathcal{M} \rightarrow I_w^\lambda$ determines the existential unforgeability and extractable set of our Winternitz Hash-based DSKE_{wots}. We use the definition of the domination free function construction from [BS20].

Given a message $m \in \mathcal{M}$ and a hash function $H : \mathcal{M} \rightarrow \{0, 1\}^\nu$, we compute $H(m)$ and convert it to an integer in $[0, 2^\nu)$. Given the public parameter w , let λ_0 be the smallest number satisfying $2^\nu \leq (w + 1)^{\lambda_0}$, set $\lambda_1 = \log_{w+1}(w \cdot \lambda_0) + 1$ and $\lambda = \lambda_0 + \lambda_1$. Given an input message m , the function $P(m)$ works as follows:

- Compute $H(m)$, convert $H(m)$ to a λ_0 -digit number in base $(w + 1)$: $(s_1, \dots, s_{\lambda_0})$.

- Compute the checksum $c = w \cdot \lambda_0 - (s_1 + \dots + s_{\lambda_0})$.
- Convert c to a λ_1 -digit number in base $(w + 1)$: $(c_1, \dots, c_{\lambda_1})$.
- Output $(s_1, \dots, s_{\lambda_0}, c_1, \dots, c_{\lambda_1})$.

In this case, the function P is domination free [BS20] and can map the message m to a vector $P(m) = (s_1, \dots, s_{\lambda_0}, c_1, \dots, c_{\lambda_1}) \in I_w^\lambda$.

Winternitz-Based DSKE. The Winternitz signature scheme is another hash-based signature scheme, that allows a trade-off between computation and the size of the signature. Winternitz-based DSKE, $\text{DSKE}_{\text{wots}}$, consists of the following algorithms:

- **Setup**(1^λ): On input the security parameter λ , the algorithm outputs the public parameter, par , which contains: (1) a hash function, $f : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$, chosen from a family of preimage-resistant hash functions, (2) a hash function, $H(\mathcal{M} \rightarrow \{0, 1\}^\nu)$, chosen from a family of collision-resistant hash function (3) a parameter integer w , and (4) a domination free function P parameterized by g and w , which maps a message m to a vector \vec{s} of length λ , and each component of s is a number in $\{0, \dots, w\}$, namely $P : \mathcal{M} \rightarrow I_w^\lambda$, where $I_w^\lambda = (\{0, \dots, w\})^\lambda$. The public parameters are implicitly input to all subsequent algorithms.
- **KeyGen**(\cdot): For each $i \in [\lambda]$, sample $sk_i \xleftarrow{\$} \{0, 1\}^\lambda$; compute $pk_i = f^{(w)}(sk_i)$; and output the private key, $SK = (sk_i)_{i \in [\lambda]}$, and the public key, $PK = (pk_i)_{i \in [\lambda]}$.
- **Sign**(SK, m): Parse $SK = (sk_i)_{i \in [\lambda]}$; compute $\vec{s} = P(m) = (s_1, \dots, s_\lambda) \in I_w^\lambda$; and output $\sigma = (f^{(s_1)}(sk_1), \dots, f^{(s_\lambda)}(sk_\lambda))$.
- **Verify**(PK, m, σ): Parse $PK = (pk_i)_{i \in [\lambda]}$, $\sigma = (\sigma_i)_{i \in [\lambda]}$, and compute $\vec{s} = P(m) = (s_1, \dots, s_\lambda) \in I_w^\lambda$. For all $i \in [\lambda]$, if $f^{(w-s_i)}(\sigma_i) \neq pk_i$, return 0. Otherwise, return 1.
- **Extract**($\{m_j, \sigma_j\}_{j \in [k]}, PK$): For each message-signature pair (m_j, σ_j) , compute $\vec{s}_j = P(m_j) = (s_{ji})_{i \in [\lambda]}$, and parse $\sigma_j = (\sigma_{ji})_{i \in [\lambda]}$. For each $j \in [k], i \in [\lambda]$, if $\exists s_{ji} = 0$, then let $sk_i = \sigma_{ji}$. Set $SK = (sk_i)_{i \in [\lambda]}$ and parse $PK = (pk_i)_{i \in [\lambda]}$. If for all $i \in [\lambda]$, $pk_i = f^{(w)}(sk_i)$, then output SK . Otherwise, output \perp .

Leakage Function for Winternitz Signatures. As illustrated in Figure 2, given a Winternitz signature $\sigma = (\sigma_1, \dots, \sigma_\lambda)$, the **Leak**(\cdot) function of Winternitz signature will output all the signature elements $(\sigma_i)_i$ that are equal to the corresponding elements in the private key list, which satisfy $f^{(w)}(\sigma_i) = pk_i$.

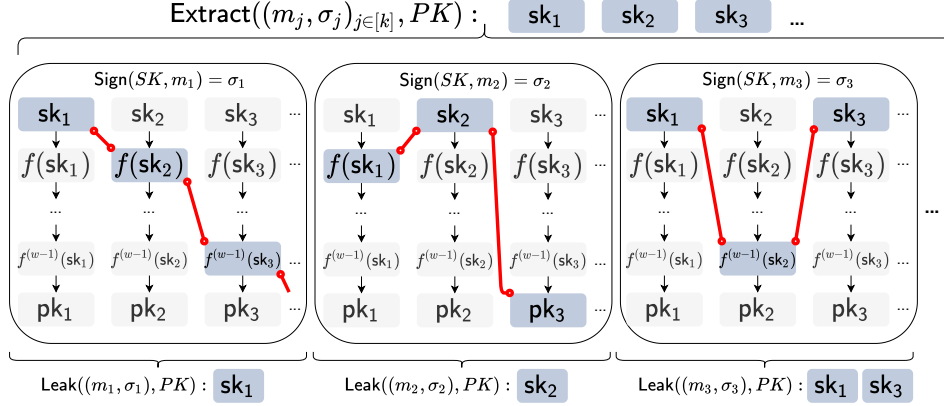


Figure 2: Example of the algorithm **Extract** for $\text{DSKE}_{\text{wotS}}$. The elements of a signature σ are linked via red lines. In this example, the original private key can be collectively reconstructed from $\{(m_i, \sigma_i)\}_{i \in [3]}$.

- $\text{Leak}((m, \sigma), PK)$: Check if $\text{Verify}(PK, m, \sigma) = 1$; otherwise, output \perp . Compute $\vec{s} = P(m) = (s_i)_{i \in [\lambda]}$, and parse $\sigma = (\sigma_i)_{i \in [\lambda]}$. For each $i \in [\lambda]$, if $s_i = 0$, then let $S_i = \sigma_i$; otherwise, let $S_i = \perp$. Output $S = (S_i)_{i \in [\lambda]}$.

The intuition of the domination free function is to ensure that if a forger can obtain a valid signature for m' from one signature for m , domination property of P ensures that there exists one element s'_i in $P(m')$ that is smaller than $s_i \in P(m)$. Therefore, it implies the fact that such a forger can be used to invert the preimage-resistant hash function f .

Security Analysis. In this part, we prove that $\text{DSKE}_{\text{wotS}}$ satisfies unforgeability and has an extractable set depending on the number of distinct message-signature pairs.

Theorem 4 (Existential Unforgeability). *$\text{DSKE}_{\text{wotS}}$ is existentially unforgeable under a 1-time adaptive chosen-message attack, that is,*

$$\Pr[\text{SignExp}_{\mathcal{A}, \text{DSKE}_{\text{wotS}}}^1(\lambda) = 1] \leq \text{negl}(\lambda)$$

Proof. We defer the proof of Theorem 4 to Theorem 14.4 in [BS20]. \square

In the following analysis for the extractable set, we assume that the checksum computed by P is independent and uniformly random. This is not really the case, because there is a dependency between the checksum elements and other elements in the vector s computed by P . However, as mentioned in [GBH17], without such an assumption, the analysis will become significantly more complex.

Lemma 2. *In a Winternitz-based DSKE scheme, if $H(\cdot)$ is modeled as a random oracle and the domination-free function $P(\cdot)$ outputs a uniformly*

random vector in I_w^λ , then the probability that an element of the private key will not be leaked after receiving k signatures is $\mathbf{p}_{leak} = \frac{w^k}{(w+1)^k}$.

Proof. Following the existing work [GBH17], we assume that in a domination free function P , the checksum c is independent of the digests $(s_1, \dots, s_{\lambda_0})$. Therefore, given a message m to the function P , we can denote its output as $(s_1, \dots, s_{\lambda_0}, s_{\lambda_0+1}, \dots, s_\lambda) \in I_w^\lambda$. We define $s_i = P(m_i)$. All s_i for $i \in \{1, \dots, k\}$ forms a $k \times \lambda$ matrix:

$$\begin{pmatrix} P(m_1) \\ \dots \\ P(m_k) \end{pmatrix} = \begin{pmatrix} s_1 \\ \dots \\ s_k \end{pmatrix} = \begin{pmatrix} s_{1,1} & \dots & s_{1,\lambda} \\ \dots & \ddots & \dots \\ s_{k,1} & \dots & s_{k,\lambda} \end{pmatrix}$$

where $s_{i,j}$ is a number in $\{0, \dots, w\}$.

For $j \in [\lambda]$, to extract the value of sk_j , there should be at least one s_{ij} that satisfies $s_{ij} = 0$, i.e., for $j \in [\lambda]$, s_{ij} should not all be non-zero. Therefore, to be able to extract SK , we do not want any columns in the matrix that contain all non-zero values.

We further assume that the output of the function P is uniformly random, then \mathbf{p}_{leak} is the probability that one single column consists of all non-zero values:

$$\mathbf{p}_{leak} = \Pr \left[x \stackrel{\$}{\leftarrow} \{0, \dots, w\}^k : x \in \{1, \dots, w\}^k \right] = \frac{w^k}{(w+1)^k}$$

□

By combining Theorem 1 and Lemma 2, we can deduce the implications outlined in Theorem 5.

Theorem 5 (DSKE_{wots} Extractable Set). *If $H(\cdot)$ is modeled as a random oracle and the domination free function $P(\cdot)$ outputs a uniformly random vector in I_w^λ , then DSKE_{wots} has an (k, δ) -extractable set, where:*

$$\delta \geq 1 - \frac{\lambda \cdot w^k}{(w+1)^k} - \text{negl}(\lambda)$$

Generalization to Other Many-Time Signature Schemes. Our constructions of DSKE can be generalized to other many-time signature schemes. The basic intuition behind DSKE_{hash} is that each signature σ_j discloses partial information of the private key SK . Therefore, besides Lamport signature and Winternitz OTS, other signature schemes such as optimized Lamport's scheme and BiBa [Per01], which directly derive signatures from the private key, can also be used to construct DSKE.

5 DSKE from Polynomial Commitment Schemes

Although the scheme $\text{DSKE}_{\text{hash}}$ proposed in Section 4 satisfies our definition of a signature scheme with key extraction, it has two inherent weaknesses: (1) The key and signature sizes are linear with the security parameter λ , making it inefficient in practice, and (2) **Extract** is not deterministic, hence the probability δ of extracting the private key depends on the size of the extractable set k ². In this section, we construct a DSKE scheme $\text{DSKE}_{\text{poly}}$ from a polynomial commitment scheme Π without aforementioned drawbacks.

In the following, we assume a degree bound $d \in \mathbb{Z}$ and a polynomial degree $\ell \in \mathbb{Z}$ that satisfy $1 \leq \ell \leq d$. The idea is to use the polynomial $f(X)$ of degree ℓ as the private key. Then, the signature on a message m is the evaluation of $f(X)$ at point $x = h(m)$, where h is a collision-resistant hash function. For key extraction we employ polynomial interpolation: any set of $\ell + 1$ valid message-signature pairs (m_i, σ_i) can reconstruct $f(X)$.

Definition 14 (Polynomial Commitment-based DSKE). *A polynomial commitment-based DSKE scheme, $\text{DSKE}_{\text{poly}}$, consists of:*

- **Setup** $(1^\lambda, d)$: *On input the security parameter λ and degree $d \in \mathbb{N}$, it runs $\Pi.\text{Setup}(1^\lambda, d)$ to obtain (ck, vk) , which allows a signer to commit to polynomials in $\mathbb{F}^{\leq d}(X)$. It samples a collision-resistant hash function $h : \mathcal{M} \rightarrow \mathbb{F}$. The public parameters, par , contain ck, vk, d , and the specification of h .*
- **KeyGen** (ℓ) : *On input $1 \leq \ell \leq d$, it samples $f(X) \xleftarrow{\$} \mathbb{F}^\ell(X)$ as an ℓ -degree polynomial and computes $\Pi.\text{Com}(ck, f(X)) \rightarrow C_f$. Let $sk = f(X)$, $pk = C_f$.*
- **Sign** (sk, m) : *It parses $sk = f(X)$, computes $x = h(m)$, and runs $\Pi.\text{Open}(ck, C_f, x, f(X)) \rightarrow (\pi, y)$. It outputs the signature $\sigma = (\pi, y)$.*
- **Verify** (pk, m, σ) : *It parses $pk = C_f$ and $\sigma = (\pi, y)$, computes $x = h(m)$, and outputs $\Pi.\text{Check}(vk, C_f, x, y, \pi) \in \{0, 1\}$.*
- **Extract** $(\{(m_i, \sigma_i)\}_{i \in [k]}, pk)$: *If $k \leq \ell$, or if m_i are not all distinct, then return \perp . If $\text{Verify}(pk, m_i, \sigma_i) = 0$ for some $i \in [k]$, then return \perp . Otherwise, compute $x_i = h(m_i)$ and parse $\sigma_i = (\pi_i, y_i)$, for $i \in [k]$. At least $\ell + 1$ pairs (x_i, y_i) interpolate the unique polynomial $\phi(X) \in \mathbb{F}^\ell(X)$, where $\lambda_i(X)$ are the Lagrange coefficients. That is, $\phi(X) = \sum_{i \in [k]} y_i \lambda_i(X)$*

$$\text{and } \lambda_i(X) = \prod_{\substack{m \in [k] \\ m \neq i}} \frac{X - x_m}{x_i - x_m}.$$

²This drawback can be circumvented by requiring signers to add an extra nonce to k different messages so that the set of these signatures guarantees the existence of the extractable set.

Remark on the degree of $f(X)$. The extraction of the private key from $k \geq \ell + 1$ points requires the signer to commit to a polynomial of degree at most ℓ . As the publicly available information in ck allows the signer to commit to any polynomial in $\mathbb{F}^{\leq d}(X)$, stronger properties, such as *strong correctness* [KZG10], *bounded-polynomial extractability* [MBKM19], and *knowledge soundness* [BDFG21], have been formulated in the literature to enforce the claimed degree on $f(X)$. However, the signer in our scheme is allowed to choose any $\ell \in [1, d]$ and has no incentive to commit to a polynomial of degree larger than ℓ , as that would cost them the deniability, as we discuss in the following sections. Hence, we can assume that $f(X)$ is indeed of degree ℓ and does not require Π to satisfy any stronger property.

Theorem 6 (Existential Unforgeability). *Assuming the underlying polynomial commitment scheme Π satisfies the computational hiding, evaluation binding, and polynomial binding properties, and that h is a collision-resistant hash function, the DSKE scheme $\text{DSKE}_{\text{poly}}$ is existentially unforgeable under an ℓ -times adaptive chosen-message attack. That is,*

$$\Pr[\text{SignExp}_{\mathcal{A}, \text{DSKE}_{\text{poly}}}^{\ell}(\lambda) = 1] \leq \text{negl}(\lambda)$$

Proof. Let \mathcal{A} be a PPT adversary breaking the signature scheme $\text{DSKE}_{\text{poly}}$. We construct a PPT algorithm \mathcal{B} that runs \mathcal{A} as a subroutine and attacks the hiding property of the polynomial commitment scheme Π . Specifically, \mathcal{B} receives from its challenger up to ℓ openings $(i_j, f(i_j), w_{i_j})$, for $i_j \in \mathbb{F}$ and $j \in [\ell]$, and for $f(X)$ not known to \mathcal{B} . It outputs (i^*, y^*) for unqueried i^* and wins if $y^* = f(i^*)$. Algorithm \mathcal{B} also receives C_f, d , and ck, vk . \mathcal{B} works as follows:

- Initiate \mathcal{A} with input $pk = C_f$, and create an empty set S_{quer} .
- Whenever \mathcal{A} requests a signature on message m , compute $x = h(m)$ and check whether $x \in S_{\text{quer}}$. If this is the case, then \mathcal{B} has already asked his challenger for the opening of point x . Otherwise, add x to S_{quer} and obtain the opening (π, y) of point x from the challenger of \mathcal{B} . Return $\sigma = (\pi, y)$ to \mathcal{A} .
- If \mathcal{A} fails to output a valid forgery on an unqueried message, then abort. Otherwise \mathcal{A} has output a message m^* and a forgery $\sigma^* = (\pi^*, y^*)$ on m^* . We assume w.l.o.g. \mathcal{A} has made ℓ signature queries (if not, \mathcal{B} queries these values itself) and hence \mathcal{B} has openings (π_i, y_i) for points x_i , with $i \in [\ell]$. Calculate $x^* = h(m^*)$. If $x^* \in S_{\text{quer}}$, then set $\text{bad}_1 \leftarrow 1$ and abort. Otherwise interpolate $f'(X) \in \mathbb{F}^{\leq \ell}(X)$ from the $\ell + 1$ points $\{(x_1, y_1), \dots, (x_\ell, y_\ell), (x^*, y^*)\}$ and compute $C_{f'} = \Pi.\text{Com}(ck, f'(X))$. If $C_{f'} \neq C_f$, then set $\text{bad}_2 \leftarrow 1$ and abort. Otherwise, output (x^*, y^*) . Observe that, even though $C_{f'} = C_f$, it could still be the case that $f(X) \neq f'(X)$.

Denote $\Pr[\text{HidExp}_{\mathcal{B},\Pi}^\ell(\lambda) = 1]$ the probability that \mathcal{B} wins the game, **Output** the event that \mathcal{B} outputs some (x^*, y^*) , and **bad₃** the event that $f(X) \neq f'(X)$. \mathcal{B} wins if and only if **Output** happens and **bad₃** does not happen. Moreover, **Output** happens if \mathcal{A} forges a valid signature and **bad₁** and **bad₂** do not happen. Therefore,

$$\begin{aligned}
& \Pr[\text{HidExp}_{\mathcal{B},\Pi}^\ell(\lambda) = 1] \\
&= \Pr[\text{Output} \wedge \overline{\text{bad}_3}] \geq \Pr[\text{Output}] - \Pr[\text{bad}_3] \\
&= \Pr[\text{SignExp}_{\mathcal{A},\text{DSKE}_{\text{poly}}}^\ell(\lambda) = 1 \wedge \overline{\text{bad}_1} \wedge \overline{\text{bad}_2}] - \Pr[\text{bad}_3] \\
&\geq \Pr[\text{SignExp}_{\mathcal{A},\text{DSKE}_{\text{poly}}}^\ell(\lambda) = 1] - \Pr[\text{bad}_1] - \Pr[\text{bad}_2] - \Pr[\text{bad}_3]
\end{aligned} \tag{1}$$

For the bad events, we have the following: The event **bad₁** implies that \mathcal{A} breaks the collision resistance property of \mathcal{A} , which is assumed secure, hence $\Pr[\text{bad}_1] = \text{negl}(\lambda)$. The event **bad₂** implies $f'(X) \neq f(X)$, and hence it must be that (x^*, y^*) is not a point of $f(X)$, i.e., $f(x^*) \neq y^*$. Since \mathcal{A} succeeded, point (x^*, y^*) and proof π^* satisfy $\Pi.\text{Check}(vk, C_f, x^*, y^*, \pi^*) = 1$. But in this case, \mathcal{B} can break the *evaluation binding* of Π in the following way. It asks for the opening at x^* , hence obtaining (π, y) , where $y = f(x^*)$. This destroys any hope of \mathcal{B} to break the hiding property, but it can attack evaluation binding, using (x^*, y^*) and (x^*, y) , for which $y \neq y^*$, $\Pi.\text{Check}(vk, C_f, x^*, y^*, \pi^*) = 1$, and $\Pi.\text{Check}(vk, C_f, x^*, y, \pi) = 1$. Since by assumption Π satisfies the evaluation-binding property, we get $\Pr[\text{bad}_2] = \text{negl}(\lambda)$. Finally, the event **bad₃** would violate the polynomial-binding property of Π , since $f(X) \neq f'(X)$ and $\text{Com}(ck, f(X)) = \text{Com}(ck, f'(X))$, thus $\Pr[\text{bad}_3] = \text{negl}(\lambda)$. Moreover, from the hiding property of Π we get $\Pr[\text{HidExp}_{\mathcal{B},\Pi}^\ell(\lambda) = 1] \leq \text{negl}(\lambda)$.

Putting it all together, (1) gives that $\Pr[\text{SignExp}_{\mathcal{A},\text{DSKE}_{\text{poly}}}^\ell(\lambda) = 1] \leq \text{negl}(\lambda)$. \square

Theorem 7 (Extractable Set). *Assuming the underlying polynomial commitment scheme Π satisfies the evaluation binding property, the DSKE scheme $\text{DSKE}_{\text{poly}}$ has a $(k, 1 - \text{negl}(\lambda))$ -extractable set for any $k \geq \ell + 1$. That is,*

$$\begin{aligned}
\delta &= \Pr \left[\begin{array}{l} m_i \in \mathcal{M}, \text{ for } i \in [k] \\ m_i \neq m_j, \text{ for } i, j \in [k], i \neq j, \text{ and } k \geq \ell + 1 \\ \sigma_i \leftarrow \text{Sign}(sk, m_i), \text{ for } i \in [k] \\ \text{Extract}(\{(m_i, \sigma_i)\}_{i \in [k]}, pk) \rightarrow sk' \end{array} : sk = sk' \right] \\
&= 1 - \text{negl}(\lambda)
\end{aligned}$$

Proof. The proof follows from two facts. First, by assuming that the signer does not commit to polynomials of degree bigger than ℓ . Second, from the *evaluation binding property*, and since the points (x_i, y_i) correspond to valid signatures, we know that $y_i = f(x_i)$, for some polynomial $f(X) \in \mathbb{F}^{\leq \ell}(X)$

and for all $i \in [k]$, except with negligible probability. Due to the uniqueness of polynomial interpolation, we know that any $\ell + 1$ distinct points (x_i, y_i) define a unique polynomial $\phi(X)$ of degree at most ℓ , hence $\phi(X)$ must be the same as $f(X)$, hence $sk = sk'$ with probability $1 - \text{negl}(\lambda)$. \square

5.1 KZG-Based DSKE Scheme

We now show DSKE_{kzg} , a concrete construction of $\text{DSKE}_{\text{poly}}$ from the Π_{kzg} polynomial commitment scheme [KZG10] (see Section 2.3).

- **Setup** $(1^\lambda, d)$: Run $\Pi_{\text{kzg}}.\text{Setup}(1^\lambda, d)$ to obtain the commitment key, $ck = \{g^{\alpha^i} \in G\}_{i \in [d]}$, and, the verification key, $vk = \hat{h}^\alpha \in G$. It samples a collision-resistant hash function $h : \mathcal{M} \rightarrow \mathbb{Z}_p$. The algorithm returns the public parameters, par , containing ck, vk, d , and the specification of h .
- **KeyGen** $(1^\lambda, \ell)$: On input, $1 \leq \ell \leq d$, sample $f(X) \xleftarrow{\$} \mathbb{F}^\ell(X)$ as an ℓ -degree polynomial. The algorithm returns the public key, $pk = C_f = \Pi_{\text{kzg}}.\text{Com}(ck, f(X)) = g^{f(\alpha)} \in G$ and the secret key, $sk = f(X)$.
- **Sign** (sk, m) : Parse $sk = f(X)$, compute $x = h(m)$, and output the signature $\sigma = (\pi, y) = \Pi_{\text{kzg}}.\text{Open}(ck, C_f, x, f(X)) \in G \times \mathbb{Z}_p$.
- **Verify** (pk, m, σ) : Parse $pk = C_f$ and $\sigma = (\pi, y)$, compute $x = h(m)$, and output $\Pi_{\text{kzg}}.\text{Check}(vk, C_f, x, y, \pi) \in \{0, 1\}$.
- **Extract** $(\{(m_i, \sigma_i)\}_{i \in [k]}, pk)$: As in the general construction, check the validity of σ_i , for $i \in [k]$, interpolate $\phi(X)$ from points (x_i, y_i) , and output $sk' = \phi(X) \in \mathbb{F}^\ell(X)$.

Theorem 8 (Unforgeability). *The scheme DSKE_{kzg} is existentially unforgeable under an ℓ -times adaptive chosen-message attack, under the same cryptographic assumptions as the Π_{kzg} polynomial commitment scheme.*

Theorem 9 (Extractable set). *The scheme DSKE_{kzg} has an extractable set of size k for any $k \geq \ell + 1$ with probability $\delta = 1 - \text{negl}(\lambda)$, under the same cryptographic assumptions as the Π_{kzg} polynomial commitment scheme.*

The proofs for Theorems 8 and 9 follow directly from Theorems 6 and 7, because Π_{kzg} satisfies the properties required from the polynomial commitment scheme Π in the generic construction.

6 Applications

In this section, we present two concrete applications of DSKE.

6.1 Non-Attributable Email

Problem Statement: Non-Attributable Email. The notion of non-attributable emails [SPG21] addresses the inherent drawbacks of DKIM by deploying a forward-forgable signature (FFS) scheme. An FFS scheme consists of the standard $\text{KeyGen}(1^\lambda)$, $\text{Sign}(sk, m)$, and $\text{Verify}(pk, m, \sigma)$ algorithms, and additionally an $\text{Expire}(sk, \mathcal{T})$ algorithm, that generates expiry information η for a private key sk , possibly using additional information \mathcal{T} , and a $\text{Forge}(\eta, m)$ algorithm, that, given the expiry information of a key, outputs a signature on a message m . The scheme satisfies the standard *correctness* and *unforgeability* properties, and the *forgeability on expiry* property.

Definition 15 (Forgeability on Expiry [Gre]). *A digital signature scheme satisfies the forgeability on expiry property if no PPT adversary can distinguish a signature created with private key sk from a signature created with the expiry information η of sk . Formally, for any $m \in \mathcal{M}$ and any PPT distinguisher \mathcal{D} , there is a negligible function $\text{negl}(\cdot)$, such that for all λ ,*

$$\Pr \left[\begin{array}{l} (pk, sk) \leftarrow \text{KeyGen}(1^\lambda), \sigma_0 \leftarrow \text{Sign}(sk, m) \\ \eta \leftarrow \text{Expire}(sk, \mathcal{T}), \sigma_1 \leftarrow \text{Forge}(\eta, m) \\ b \stackrel{\$}{\leftarrow} \{0, 1\}, b' \leftarrow \mathcal{D}(\sigma_b, \eta) \end{array} : b = b' \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

The authors in [SPG21] provide two constructions, KeyForge and Timeforge. Both satisfy the FFS properties by requiring signers either to publish old private keys or to issue signed updates to a publicly verifiable timekeeping service, respectively.

KeyForge is an FFS scheme where signatures expire after a predefined delay Δ . The expiry information is the private key itself, which has to be published by the server every Δ time. It is based on a hierarchical identity-based signature scheme (HIBS). In a HIBS every signer has an identity, from which the private key can be derived. The identities are organized in a tree structure, with the property that any node (encoding an identity) can derive the private keys of its children (sub-identities). The master private key (the root of the tree) can generate all the other private keys. KeyForge uses a hierarchical tree structure with four levels, where each level corresponds to distinct timespans, namely, years, months, days, and minutes. Each tag (i.e., leaf in the tree) represents a unique 15-minute time chunk, and the signer of HIBS can derive a private key for this time chunk from the master private key. The underlying HIBS enables email servers to publish succinct expiration information; instead of publishing all private keys for, say, one day, it is enough to publish the private key that corresponds to the node that encodes that day on the third level.




Figure 3: An example of GroupForge for $h_T = 2$ and between the interval $[t, t + 2^{h_T} \cdot \Delta]$. In this example, $\{(m_i, \sigma_i)\}_{i \in [2]}$ can potentially act as the deniable group for the message-signature pair (m_3, σ_3) , and $\{(m_i, \sigma_i)\}_{i \in \{4,5\}}$ can act as the deniable group for the message-signature pair (m_6, σ_6) .

Solution: GroupForge from DSKE. In the following, We show how a DSKE scheme, such as $\text{DSKE}_{\text{poly}}$, can be used in a non-attributable email protocol [SPG21]. DSKE removes the requirement for email servers to periodically publish expiration information. Looking ahead, we introduce GroupForge, an FFS scheme that builds on top of a DSKE and achieves the same properties as KeyForge, while it does not require the email servers to publish *any* expiry information, i.e., algorithm $\text{Expire}()$ requires only information already published by the email servers.

GroupForge uses a DSKE scheme Σ_{DSKE} with extractable sets of size k , a collision-resistant hash function $H()$, and a *local* clock $\text{Time}()$ that returns the current time. GroupForge works as follows³:

- $\text{KeyGen}(1^\lambda, \Delta, h_T, t)$: On input the security parameter λ , the length Δ , the height of the Merkle tree $h_T \in \mathbb{N}$, and the starting time t , it invokes $\Sigma_{\text{DSKE}}.\text{Setup}()$, and generates 2^{h_T} key pairs $\{(pk_i, sk_i)\}_{i \in [2^{h_T}]}$ using the algorithm $\Sigma_{\text{DSKE}}.\text{KeyGen}()$. The algorithm then uses the hash function $H()$ to construct a Merkle tree from $\{pk_i\}_{i \in [2^{h_T}]}$. Let R be the root of this Merkle tree. Then it sets $PK = (R, \Delta, h_T, t)$ and $SK = (R, \Delta, h_T, t, \{sk_i\}_{i \in [2^{h_T}]})$, and outputs (SK, PK) .
- $\text{Sign}(SK, m)$: On input the key, $SK = (R, \Delta, h_T, t, \{sk_i\}_{i \in [2^{h_T}]})$, it obtains the current time $t' \leftarrow \text{Time}()$, and computes $\text{leaf}_i = \lfloor (t' - t) / \Delta \rfloor$ (the leaf that corresponds to the current time chunk i). It invokes $s = \Sigma_{\text{DSKE}}.\text{Sign}(sk_i, m)$, the authenticity proof proof_i for the path from leaf_i to the root R , and outputs the signature $\sigma = (s, pk_i, \text{leaf}_i, \text{proof}_i)$.

³For simplicity we ignore the parameters specific to the $\text{DSKE}_{\text{poly}}$ scheme, e.g., d and ℓ .

- **Verify**(PK, m, σ): On input the public key PK , the message m , and the signature σ , it parses $PK = (R, \Delta, h_T, t)$, and $\sigma = (s, pk_i, \text{leaf}_i, \text{proof}_i)$, obtains the current time $t' \leftarrow \text{Time}()$, and computes $\text{leaf} \leftarrow \lfloor (t' - t) / \Delta \rfloor$. If $\text{leaf}_i \neq \text{leaf}$, then returns 0. It validates proof_i starting from leaf_i ; if invalid, it returns 0. Otherwise, returns $\Sigma_{\text{DSKE}}.\text{Verify}(pk_i, m, s)$.
- **Expire**($PK, \{(m_j, \sigma_j)\}_{j \in [k]}$): On input the public key PK and k message-signature pairs $\{(m_j, \sigma_j)\}_{j \in [k]}$, the algorithm checks that the messages are pairwise different, i.e., for all $j_1, j_2 \in [k]$, $m_{j_1} \neq m_{j_2}$, that all signatures are valid, i.e., for all $j \in [k]$, $\text{Verify}(PK, m_j, \sigma_j) = 1$, and that all signatures are created in the same time chunk, i.e., there exists $i \in [2^{h_T}]$ such that for all $j \in [k]$, $\sigma_j = (s_j, pk_i, \text{leaf}_i, \text{proof}_i)$. If the conditions do not hold, it returns \perp . Otherwise, it calls $\Sigma_{\text{DSKE}}.\text{Extract}(\{(m_j, \sigma_j)\}_{j \in [k]}, pk_i)$. If this call returns \perp , then **Expire**() also returns \perp . If it returns an extracted key sk_i , then **Expire**() returns the expiry information $\text{info} = (\text{leaf}_i, \text{proof}_i, sk_i)$.
- **Forge**(PK, m, info): On input the public key PK , a message m , and the expiry information $\text{info} = (\text{leaf}_i, \text{proof}_i, sk_i)$, for some $i \in [2^{h_T}]$, the algorithm computes $s = \Sigma_{\text{DSKE}}.\text{Sign}(sk_i, m)$ and returns the signature $\sigma = (s, pk_i, \text{leaf}_i, \text{proof}_i)$, where $pk_i = \text{extractPK}(sk_i)$ is the public key of leaf_i .

Figure 3 provides a pictorial example of how GroupForge works. GroupForge satisfies the *correctness*, *unforgeability*, and *forgeability on expiry* properties of an FFS scheme (see Appendix A).

Remark on Threat Models. In non-attributable email protocol, there are two types of attack scenarios: *real-time attacks* and *after-the-fact attacks*. In the first attack model, a real-time attacker can be the recipient of the message that immediately discloses received messages to third parties. In the after-the-fact attack model, the recipient of the message is assumed to be honest at the time of email receipt, but can be compromised afterward by an *after-the-fact attacker*; hence, the attacker can take a snapshot of all emails received in the past. More importantly, Specter *et al.* emphasize that it is impossible to defend against a real-time attacker without interaction. Therefore, our GroupForge protocol specifically targets the after-the-fact attack scenario. Due to the need for interaction, we believe interactive protocols such as OTR [BGB04] are more suitable for achieving deniability in the real-time attack scenario.

6.2 Rate-Limiting Nullifier

Problem Statement: Rate Limiting Nullifier. A Rate-Limiting Nullifier (RLN) [Pri24] is a scheme aimed at restricting the quantity of user actions, thereby facilitating a robust mechanism for spam prevention. This

approach has been extensively discussed in a range of real-world applications [Bla21b, Res21, Bla21a, BB21]. Generally, RLN achieves spam deterrence by either increasing the difficulty of duplicating identities or risking the disclosure of a user’s private key once a certain number of actions are exceeded. In the following, we propose a RLN construction that is directly derived from our $\text{DSKE}_{\text{poly}}$ scheme. The formal study and construction of secure RLN schemes has, to the best of our knowledge, not appeared in academic literature before.

Solution: RLN from $\text{DSKE}_{\text{poly}}$. A RLN scheme typically consists of four parts, *setup*, *registration*, *interaction*, and *slashing*, run between two types of participants, a *user* \mathcal{U} and a *server* \mathcal{T} . Our $\text{DSKE}_{\text{poly}}$ -based RLN works as follows:

- **Setup** $(1^\lambda, d, \mathcal{U}, \mathcal{T})$: On input the security parameter λ and degree $d \in \mathbb{N}$, it runs $\text{DSKE}_{\text{poly}}.\text{Setup}(1^\lambda, d)$ to obtain (ck, vk) , which allows a signer to commit to polynomials in $\mathbb{F}^{\leq d}(X)$. It samples a collision-resistant hash function $h : \mathcal{M} \rightarrow \mathbb{F}$. The public parameters *par* contain ck, vk, d , and h . The parameters are publicly known for the user \mathcal{U} and the server \mathcal{T} .
- **Registration** $(1^\lambda, \mathcal{U}, \mathcal{T})$: The user \mathcal{U} runs $\text{DSKE}_{\text{poly}}.\text{KeyGen}(1^\lambda, d)$ to obtain the private key $sk_{\mathcal{U}} = f(X)$ and the public key $pk_{\mathcal{U}} = C_f$. The public key $pk_{\mathcal{U}}$ is registered and stored on the server \mathcal{T} . \mathcal{T} initializes a counter $\text{MsgNum}_{\mathcal{U}} = 0$ to record the number of messages signed by \mathcal{U} .
- **Interaction** $(m, \mathcal{U}, \mathcal{T})$: The interaction algorithm is run between \mathcal{U} and \mathcal{T} in order to sign and verify a message m . User \mathcal{U} parses $sk = f(X)$, computes $x = h(m)$ and the signature $\sigma = \text{DSKE}_{\text{poly}}.\text{Open}(ck, C_f, x, f(X)) = (\pi, y)$, and sends σ to the server \mathcal{T} . Upon receiving σ , server \mathcal{T} parses the user’s $pk_{\mathcal{U}} = C_f$ and $\sigma = (\pi, y)$, computes $x = h(m)$, and outputs $\text{DSKE}_{\text{poly}}.\text{Check}(vk, C_f, x, y, \pi) \in \{0, 1\}$. If $\text{DSKE}_{\text{poly}}.\text{Check}(vk, C_f, x, y, \pi) = 1$, then \mathcal{T} increases the counter $\text{MsgNum}_{\mathcal{U}}$ and stores the message-signature pair (m, σ) .
- **Slashing** $(\{(m_i, \sigma_i)\}_{i \in [\text{MsgNum}]}, pk_{\mathcal{U}}, \mathcal{T})$: If server \mathcal{T} has received more than l message-signature pairs from the user \mathcal{U} , i.e., $\text{MsgNum} > d$, then \mathcal{T} runs the algorithm $\text{Extract}(\{(m_i, \sigma_i)\}_{i \in [\text{MsgNum}]}, pk_{\mathcal{U}})$ of DSKE_{kzg} to extract $sk_{\mathcal{U}}$. Subsequently, \mathcal{T} publishes $sk_{\mathcal{U}}$ and removes \mathcal{U} from the system.

Spam Prevention in $\text{DSKE}_{\text{poly}}$ -based RLN. The key extraction property of a $\text{DSKE}_{\text{poly}}$ (see Theorem 7) empowers the server to extract the private key of a user who issues more than d messages with their signatures. Here, d serves as a predefined parameter, influencing the polynomial’s degree that the user can select. Consequently, in a $\text{DSKE}_{\text{poly}}$ -based RLN, a user is prevented from generating excessive spam (exceeding d messages) because their private key would be publicly disclosed, leading to their expulsion from the system.

Table 1: Performance of $\text{DSKE}_{\text{lamp}}$ and DSKE_{kzg} . Note that the probability of extracting the key in $\text{DSKE}_{\text{lamp}}$ is overwhelming in k , so the actual size of the extractable group can be smaller than $k = 64$.

Scheme	Group Size	Key Generation	Signing	Verification	Extraction
$\text{DSKE}_{\text{lamp}}(\text{SHA256})$	$k = 64$	451.583 μs	17.625 μs	160.458 μs	45.083 μs
	$k = 16$	2.795 ms	2.425 ms	4.949 ms	1.785 ms
$\text{DSKE}_{\text{kzg}}(\text{BLS12-377})$	$k = 32$	2.825 ms	2.523 ms	6.123 ms	7.813 ms
	$k = 64$	2.852 ms	2.756 ms	7.202 ms	35.341 ms
	$k = 128$	3.624 ms	3.404 ms	8.045 ms	166.433 ms

Table 2: Performance of GroupForge Constructions with $\Delta = 0.5$ hour.

Scheme	Parameters:	Key Generation	Signing	Verification
	Tree height (Duration)			
GroupForge _{hash} (SHA256)	$h_T = 14$ (0.93 years)	7.407 s	21.791 μs	176.541 μs
	$h_T = 15$ (1.87 years)	14.815 s	34.708 μs	177.416 μs
	$h_T = 16$ (3.74 years)	29.631 s	41.291 μs	178.083 μs
	$h_T = 17$ (7.48 years)	59.262 s	45.916 μs	179.374 μs
GroupForge _{kzg} (SHA256, BLS12-377, $k = 32$)	$h_T = 14$ (0.93 years)	46.294 s	1.554 ms	6.139 ms
	$h_T = 15$ (1.87 years)	92.588 s	1.591 ms	6.140 ms
	$h_T = 16$ (3.74 years)	185.175 s	1.611 ms	6.141 ms
	$h_T = 17$ (7.48 years)	370.350 s	1.722 ms	6.143 ms

Using RLN in Anonymous Setting. As discussed in existing works of RLN [Bla21b, Res21, Bla21a, BB21], we can utilize non-interactive zero-knowledge proofs [Gro16] or mixing services [LG21, WCQ⁺23, WCL⁺23] to obscure the public key $pk_{\mathcal{U}}$ of a user \mathcal{U} , thereby achieving anonymity for an RLN scheme. Nevertheless, a significant challenge persists in linking the signed messages of \mathcal{U} in an anonymous context without revealing any details about their public key $pk_{\mathcal{U}}$. A feasible approach might be leveraging linkable ring signatures [LW05]: this method integrates a linkable tag into the signed message, allowing the server to track the number of messages originating from the same user. Incorporating these techniques for privacy protection into our $\text{DSKE}_{\text{poly}}$ -enhanced RLN represents a promising direction for further work.

7 Evaluation and Discussion

In this section, we evaluate the performance of both hash-based and polynomial commitment-based DSKE constructions, as well as the performance of GroupForge constructions. Finally, we discuss potential future directions for DSKE.

7.1 Evaluation

We evaluate the performance of Lamport signature-based DSKE ($\text{DSKE}_{\text{lamp}}$) and KZG-based DSKE scheme (DSKE_{kzg}) as well as the performance of

Groupforce constructions.

Testbed. We evaluate our schemes on a macOS Monterey machine with an Apple M1 chip (8-core, 3.2 GHz), 8 GB of RAM.

DSKE Constructions. We have implemented prototypes of our proposed constructions of $\text{DSKE}_{\text{lamp}}$ and DSKE_{kzg} in Rust. For our $\text{DSKE}_{\text{lamp}}$, we use the SHA-256 implementation. For DSKE_{kzg} , we adapt the KZG polynomial commitment implementation using the curve BLS12-377 from `arkworks` library ⁴.

Table 1 provides a comprehensive overview of the performance metrics for both $\text{DSKE}_{\text{lamp}}$ and DSKE_{kzg} . Remarkably, the execution time of functions within $\text{DSKE}_{\text{lamp}}$ consistently remains below $1ms$, indicating efficient processing across various $\text{DSKE}_{\text{lamp}}$ functions. The extraction time of DSKE_{kzg} appears as approximately a quadratic function of the group size k . This is because the complexity of $\phi(X)$ in DSKE_{kzg} increases quadratically as k grows.

GroupForge Constructions. We also implement two GroupForge constructions based on $\text{DSKE}_{\text{lamp}}$ and DSKE_{kzg} respectively. We leverage the Merkle tree library using SHA-256 from `arkworks` ⁵ to construct our GroupForge. We implement $\text{GroupForge}_{\text{hash}}$ with $\text{DSKE}_{\text{lamp}}$, and $\text{GroupForge}_{\text{kzg}}$ with DSKE_{kzg} using BLS12-377 curve and the group size $k = 32$. Table 2 summarizes the performance of GroupForge constructions. We evaluate our constructions with Merkle tree heights of 14, 15, 16, and 17, which correspond to a duration of 0.93, 1.87, 3.74 and 7.48 years respectively (given a time chunk of $\Delta = 0.5$ hour). For both $\text{GroupForge}_{\text{hash}}$ and $\text{GroupForge}_{\text{kzg}}$, the most expensive computation is the key generation. This is because we need to generate 2^{h_T} key pairs, where h_T is the Merkle tree height.

7.2 Discussion

In this part, we discuss relevant properties and examine potential future directions and improvements for our constructions.

GroupForge with Forward-Secure Signature Schemes. Forward-secure signature schemes (FSS) ⁶ [BM99] allow signers to derive future keys from past keys while preventing users from deriving past keys from future keys. This property of FSS can be combined with GroupForge to demonstrate that if there is a deniable group at some point, the same group can be used as a deniable proof for all message-signature pairs before that point. In particular, to achieve this property, the signer can produce several FSS key pairs from a master private key, and use them in reverse order in the leaves of the Merkle hash tree. For instance, a private key sk_i used in the interval

⁴<https://github.com/arkworks-rs/poly-commit>

⁵<https://github.com/arkworks-rs/crypto-primitives>

⁶FSS is not to be confused with forward-forgeable signatures




Figure 4: An example of GroupForge using FSS. The group $\{(m_i, \sigma_i)\}_{i \in \{4,5\}}$ can act as a deniable group for all message-signature pairs in the past, $\{(m_i, \sigma_i)\}_{i \in [3]}$.

$[t_i, t_i + \Delta]$ can produce all previous private keys sk_j used before t_i . Hence, it is not difficult to see that, if a deniable group exists at some time t' , the same group can derive all private keys before t' . Figure 4 gives a high-level example of how FSS can help improve GroupForge.

Trusted Setup. The KZG scheme used in $\text{DSKE}_{\text{poly}}$ requires a trusted setup to generate the public parameters. Although the scheme can leverage a trusted third party to run the setup, such a reliance on a trusted party is typically not welcome in various settings. To remove this trusted assumption, practitioners and the academic community have developed practical solutions to securely generate these parameters, known as ceremony [NRBB22]. Also, one may deploy polynomial commitment schemes that do not require trusted setup, such as the one constructed from Bulletproofs [BBB⁺18].

Time-Agnostic Forgeability. Although not explicitly mentioned, making a synchronous assumption on the communication is required to ensure parties (i.e., group of size k) receive their signatures in the respective time chunk. Moreover, the proper operation of the system relies on parties having access to a local clock that advances at the same pace. The original work of [SPG21] requires the stronger assumption of a shared global clock. However, we observe that the forgeability aspect of the DSKE is generally independent of the notion of time as it comes from the availability of a sufficient number of signatures for key extraction. This essentially opens up the door to use DSKE for adding forgeability property in settings where reliance on time is not desirable.

DSKE and Unique Signatures. Apart from the aspect of key extractability, $\text{DSKE}_{\text{poly}}$ also satisfies the *uniqueness property*, i.e., there is only one valid signature for any combination of the message and public key. For reference, BLS [BLS01] is a unique-signature scheme while ECDSA [JMV01] is not. This property is of importance in applications that need to ensure the consistency of the produced signature, such as when generating shared randomness [HMW18].

Compacting the Keys and Signatures. We can also consider extending $\text{DSKE}_{\text{poly}}$ to *multi-signature schemes*, where a group of individuals, each holding their own key pairs, collectively sign a message and the resulting sig-

nature is verifiable under the respective set of public keys. In recent years, there has been a surge of interest in multi-signature schemes that support an aggregation property for the public key and the signature, aiming to reduce the size of blockchains [BDN18]. Leveraging the homomorphic property of the underlying KZG polynomial commitment scheme, these key and signature aggregation techniques are also applicable to DSKE_{kzg} , offering a high level of space efficiency.

Further Applications. DSKE can facilitate more applications where the need for authenticity is momentary and signers desire long-term deniability. Such examples are electronic voting and monetary donations. It has been argued that, even if future adversaries cannot go back to change the integrity or the result of a finished election, the privacy requirements of voters in the face of retaliation are everlasting [GCG⁺19]. Similarly, leaked donation records have reportedly caused individuals to be targeted and threatened [Ott22], and transaction deniability might be desirable in such scenarios as well.

8 Related Work

DSKE. Designing digital signature scheme with a key-extraction property has already been explored in the context of double authentication preventing signatures (DAPS). DAPS are genuinely designed with the purpose of double or multiple authentication prevention [PS14, BPS17, DRS18], and they particularly aim at messages of special form, namely $m = (a, p)$, where a is an address and p is a payload. In case a signer signs two (or more) messages with the same address but different payloads, its private key is leaked. However, DSKE essentially provides key extractability as an inherent feature without making assumptions about the type of message, increasing its applicability. In particular, the key extraction in our polynomial commitment-based DSKE (i.e., $\text{DSKE}_{\text{poly}}$) directly comes from the polynomial interpolation theorem. Note that one major downside of many DAPS schemes is their limited address space, i.e., exponentially large address space is not supported. Moreover, a notable difference between DSKE and DAPS is that the former has a sole design similar to the typical signature scheme in the literature like BLS [BLS01], while the latter has a hybrid design, containing different components. For instance, DAPS of [BPS17] builds on trapdoor identification schemes [MR02] and that of [DRS18] involves encryption scheme and secret sharing. This, in turn, results in DAPS having considerably larger key pairs and also signature sizes compared to the normal signatures [DRS18].

GroupForge. Specter, Park, and Green formally defined the notion of forward-forgeable signature (FFS) [SPG21] and showed how FFS can be used to achieve deniability in the email protocol. The main idea of their scheme is to make the signatures become forgeable after a fixed delay. They present two concrete constructions: KeyForge and TimeForge. In the KeyForge

construction, the email server needs to periodically publish expired keys to claim deniability over sent emails, and in the TimeForge construction, the signers need to rely on a trusted publicly verifiable time-keeping (PVTk) service to provide them with a verifiable clock time proof. The forgeability comes from the possibility of obtaining a valid proof by querying the PVTk service after a fixed delay.

Arun, Bonneau, and Clark proposed a similar notion to FFS called short-lived signatures [ABC23]. The main idea of their work is to leverage a disjunctive statement to achieve deniability, building up on the previous efforts in the literature, e.g., designated verifier proofs [JSI96, SWP04], proofs-of-work-or-knowledge (PoWorKs) [BKZZ16], ring signatures [RST01], and one-out-of-many proofs [GK15]. In particular, the construction in Arun *et al.*'s work [ABC23] deploys verifiable delay functions [BBBF18, Wes19] as its main building block together with a (trusted) randomness beacon [SJK⁺17]. They use a statement of the form: *I know the witness, x (e.g., private key), or someone solved a VDF on a beacon value derived from the trusted beacon before a time, t .* Hence, their construction offers deniability to the prover because anyone can produce a valid proof by evaluating the VDF through sequential computation. Moreover, a recent analysis report [LMPR] of the VDF implementation shows that VDFs in practice, such as Min-root [KMT22], the VDF for the Ethereum blockchain, might be prone to various potential attacks, which have yet not been explored thoroughly. This work, on the other hand, offers a simpler approach without requiring costly VDF evaluations and a trusted random beacon. We provide a comparison of our work and state-of-the-art signature schemes with deniability in Table 3.

Finally, there is a subtle difference between the aforementioned types of signatures. While forward-forgeable signatures provide non-attributability by selective release of some *expiry* information, short-lived signatures *automatically* become non-attributable after some time without further action. Interestingly, GroupForge, proposed in this work, can offer the best of both worlds. In other words, as the deniability property of GroupForge relies on the *number* of generated signatures, it essentially can be considered a short-lived signature in scenarios where the signer is supposed to generate a *certain* number of signatures, achieving deniability for free without requiring any trusted service or computing costly VDFs. Moreover, GroupForge can also be modeled as a forward-forgeable signature, achieving forgeability without constantly publishing key materials that are problematic in practice due to unreliable distribution [SPG21].

Rate-Limiting Nullifier. A Rate-Limiting Nullifier (RLN) [Pri24] is a mechanism designed to restrict users in the number of actions within a system, thereby facilitating an effective spam prevention mechanism. Typically, RLN accomplishes spam prevention by elevating the cost of identity replication or exposing a user's private key after a predefined number of actions.

Table 3: Comparison of signature schemes with deniability.

Scheme	Without Requiring Future Actions (e.g., Publishing Keys)	Without Requiring External Services (e.g., Random Beacon)	Without Requiring VDF
KeyForge [SPG21]	✗	✓	✓
TimeForge [SPG21]	✓	✗	✓
Short-lived signature [ABC23]	✓	✗	✗
GroupForge _{hash} (based on DSKE _{hash})	✓	✓	✓
GroupForge _{poly} (based on DSKE _{poly})	✓	✓	✓

RLN can be used in anonymous voting schemes [Bla21b], privacy-preserving P2P networks [Res21], and decentralized blockchain applications [Bla21a]. RLN has been extensively discussed by the blockchain community [Lim23] and implemented in practice ⁷. However, to our knowledge, no academic initiative has been undertaken to formally propose and establish secure RLN constructions as we do in this work with DSKE. In addition, it is feasible to extend DSKE-based RLN to an anonymous setting by employing non-interactive zero-knowledge proofs.

9 Conclusion

This paper introduces DSKE, a novel signature scheme featuring key extraction capabilities. We present concrete DSKE constructions based on the hash signatures and polynomial commitment schemes. We provide formal proof of security for our constructions, demonstrating that signers can consistently achieve deniability by presenting a set of signatures utilized to regenerate old private keys. Furthermore, we illustrate how DSKE can serve as the foundation for constructing both RLN schemes and GroupForge signatures. We posit that DSKE holds promise for application in other protocols that require short-term authenticity.

Acknowledgments

The authors would like to thank Harry W. H. Wong for helpful discussions and comments.

⁷<https://github.com/Rate-Limiting-Nullifier>,
<https://github.com/sec-bit/kzg-rln-go>,
<https://github.com/dynm/kzg-rln-contracts>

References

- [ABC23] Arasu Arun, Joseph Boneau, and Jeremy Clark. Short-lived zero-knowledge proofs and signatures. In *Advances in Cryptology–ASIACRYPT 2022: 28th International Conference on the Theory and Application of Cryptology and Information Security, Taipei, Taiwan, December 5–9, 2022, Proceedings, Part III*, pages 487–516. Springer, 2023.
- [ACD⁺07] E Allman, Jon Callas, M Delany, Miles Libbey, J Fenton, and M Thomas. Domainkeys identified mail (dkim) signatures. Technical report, RFC 4871, May, 2007.
- [BB21] Blagoj and WhiteHat Barry. Decentralised cloudflare - using rln and rich user identities, 2021. Available at: <https://ethresear.ch/t/decentralised-cloudflare-using-rln-and-rich-user-identities/10774>.
- [BBB⁺18] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE symposium on security and privacy (SP)*, pages 315–334. IEEE, 2018.
- [BBBF18] Dan Boneh, Joseph Boneau, Benedikt Bünz, and Ben Fisch. Verifiable delay functions. In *Annual international cryptology conference*, 2018.
- [BDE⁺11] Johannes Buchmann, Erik Dahmen, Sarah Ereth, Andreas Hülsing, and Markus Rückert. On the security of the winternitz one-time signature scheme. In *Progress in Cryptology–AFRICACRYPT 2011: 4th International Conference on Cryptology in Africa, Dakar, Senegal, July 5-7, 2011. Proceedings 4*, pages 363–378. Springer, 2011.
- [BDFG21] Dan Boneh, Justin Drake, Ben Fisch, and Ariel Gabizon. Halo infinite: Proof-carrying data from additive polynomial commitments. In *Annual International Cryptology Conference*, pages 649–680. Springer, 2021.
- [BDH11] Johannes Buchmann, Erik Dahmen, and Andreas Hülsing. XMSS - A Practical Forward Secure Signature Scheme Based on Minimal Security Assumptions. In Bo-Yin Yang, editor, *Post-Quantum Cryptography*, pages 117–129, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.
- [BDN18] Dan Boneh, Manu Drijvers, and Gregory Neven. Compact multi-signatures for smaller blockchains. In *International Con-*

- ference on the Theory and Application of Cryptology and Information Security*, pages 435–464. Springer, 2018.
- [BGB04] Nikita Borisov, Ian Goldberg, and Eric Brewer. Off-the-record communication, or, why not to use pgp. In *Proceedings of the 2004 ACM workshop on Privacy in the electronic society*, pages 77–84, 2004.
- [BKZZ16] Foteini Baldimtsi, Aggelos Kiayias, Thomas Zacharias, and Bingsheng Zhang. Indistinguishable proofs of work or knowledge. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 902–933. Springer, 2016.
- [Bla21a] Blagoj. Private message sharing for eth2 validators, 2021. Available at: <https://ethresear.ch/t/private-message-sharing-for-eth2-validators/10664>.
- [Bla21b] Blagoj. Rate limiting nullifier: A spam-protection mechanism for anonymous environments, 2021. Available at: <https://medium.com/privacy-scaling-explorations/rate-limiting-nullifier-a-spam-protection-mechanism-for-anonymous-environments-bbe4006a57d>.
- [BLS01] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the weil pairing. In *International conference on the theory and application of cryptology and information security*, pages 514–532. Springer, 2001.
- [Bly05] Stephen E Blythe. Digital signature law of the united nations, european union, united kingdom and united states: Promotion of growth in e-commerce with enhanced security. *Richmond Journal of Law & Technology*, 11(2):6, 2005.
- [BM99] Mihir Bellare and Sara K Miner. A forward-secure digital signature scheme. In *Annual international cryptology conference*, pages 431–448. Springer, 1999.
- [BPS17] Mihir Bellare, Bertram Poettering, and Douglas Stebila. Detering certificate subversion: efficient double-authentication-preventing signatures. In *Public-Key Cryptography–PKC 2017*, 2017.
- [BS20] Dan Boneh and Victor Shoup. A graduate course in applied cryptography. *Draft 0.5*, 2020.

- [DRS18] David Derler, Sebastian Ramacher, and Daniel Slamanig. Short double-and n-times-authentication-preventing signatures from ecdsa and more. In *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 273–287. IEEE, 2018.
- [GBH17] Leon Groot Bruinderink and Andreas Hülsing. “oops, i did it again”–security of one-time signatures under two-message attacks. In *International Conference on Selected Areas in Cryptography*, pages 299–322. Springer, 2017.
- [GCG⁺19] Huangyi Ge, Sze Yiu Chau, Victor E Gonsalves, Huian Li, Tianhao Wang, Xukai Zou, and Ninghui Li. Koinonia: verifiable e-voting with long-term privacy. In *Proceedings of the 35th Annual Computer Security Applications Conference*, pages 270–285, 2019.
- [GK15] Jens Groth and Markulf Kohlweiss. One-out-of-many proofs: Or how to leak a secret and spend a coin. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 253–280. Springer, 2015.
- [GMR88] Shafi Goldwasser, Silvio Micali, and Ronald L Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on computing*, 17(2):281–308, 1988.
- [Gre] Matthew Green. Ok Google: please publish your DKIM secret keys. <https://blog.cryptographyengineering.com/2020/11/16/ok-google-please-publish-your-dkim-secret-keys/>.
- [Gro16] Jens Groth. On the size of pairing-based non-interactive arguments. In *Advances in Cryptology–EUROCRYPT 2016: 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8–12, 2016, Proceedings, Part II 35*, pages 305–326. Springer, 2016.
- [HMW18] Timo Hanke, Mahnush Movahedi, and Dominic Williams. Dfinity technology overview series, consensus system. *arXiv preprint arXiv:1805.04548*, 2018.
- [JMV01] Don Johnson, Alfred Menezes, and Scott Vanstone. The elliptic curve digital signature algorithm (ecdsa). *International journal of information security*, 1(1):36–63, 2001.
- [JSI96] Markus Jakobsson, Kazue Sako, and Russell Impagliazzo. Designated verifier proofs and their applications. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 143–154. Springer, 1996.

- [Kar19] Nikolaos Karanikolas. Digital signature legality in different jurisdictions: Legally binding issues. 2019.
- [KMT22] Dmitry Khovratovich, Mary Maller, and Pratyush Ranjan Tiwari. Minroot: Candidate sequential function for ethereum vdf. *Cryptology ePrint Archive*, 2022.
- [KZG10] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In *Advances in Cryptology - ASIACRYPT 2010 - 16th International Conference on the Theory and Application of Cryptology and Information Security*, volume 6477, pages 177–194. Springer, 2010.
- [Lam79] Leslie Lamport. Constructing digital signatures from a one way function. 1979.
- [LG21] Duc Viet Le and Arthur Gervais. Amr: Autonomous coin mixer with privacy preserving reward distribution. In *AFT '21: 3rd ACM Conference on Advances in Financial Technologies*, pages 142–155, Arlington, Virginia, USA, 2021. ACM.
- [Lim23] Wanseob Lim. Rln on kzg polynomial commitment scheme, 2023. Available at: <https://zkresear.ch/t/rln-on-kzg-polynomial-commitment-scheme-cross-posted/114>.
- [LMPR] Gaetan Leurent, Bart Mennink, Krzysztof Pietrzak, and Vincent Rijmen. Analysis of minroot: Public report. Available at: <https://crypto.ethereum.org/events/minrootanalysis2023.pdf>.
- [LW05] Joseph K Liu and Duncan S Wong. Linkable ring signatures: Security models and new schemes. In *Computational Science and Its Applications—ICCSA 2005: International Conference, Singapore, May 9-12, 2005, Proceedings, Part II 5*, pages 614–623. Springer, 2005.
- [Mas16] Stephen Mason. *Electronic signatures in law*. University of London Press, 2016.
- [MBKM19] Mary Maller, Sean Bowe, Markulf Kohlweiss, and Sarah Meiklejohn. Sonic: Zero-knowledge snarks from linear-size universal and updatable structured reference strings. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 2111–2128, 2019.

- [Mer89] Ralph C Merkle. A certified digital signature. In *Conference on the Theory and Application of Cryptology*, pages 218–238. Springer, 1989.
- [MR02] Silvio Micali and Leonid Reyzin. Improving the exact security of digital signature schemes. *Journal of Cryptology*, 15(1):1–18, 2002.
- [NRBB22] Valeria Nikolaenko, Sam Ragsdale, Joseph Bonneau, and Dan Boneh. Powers-of-tau to the people: Decentralizing setup ceremonies. *Cryptology ePrint Archive*, 2022.
- [Ott22] Ottawa Citizen. Threats close stella luna gelato café after owner’s name appears in givesendgo data leak, 2022. <https://ottawacitizen.com/news/local-news/threats-close-stella-luna-gelato-cafe-after-owners-name-appears-in-givesendgo-data-leak>.
- [Per01] Adrian Perrig. The biba one-time signature and broadcast authentication protocol. In *Proceedings of the 8th ACM Conference on Computer and Communications Security*, pages 28–37, 2001.
- [Pri24] Privacy and Scaling Explorations team, Ethereum Foundation. Rate-limiting nullifier, 2024. Available at: <https://rate-limiting-nullifier.github.io/rln-docs/rln.html>.
- [PS14] Bertram Poettering and Douglas Stebila. Double-authentication-preventing signatures. In *Computer Security-ESORICS 2014: 19th European Symposium on Research in Computer Security, Wroclaw, Poland, September 7-11, 2014. Proceedings, Part I 19*, pages 436–453. Springer, 2014.
- [Res21] Vac Research. Privacy-preserving p2p economic spam protection in waku v2, 2021. Available at: <https://vac.dev/rlog/rln-relay/>.
- [RKS15] Tim Ruffing, Aniket Kate, and Dominique Schröder. Liar, liar, coins on fire! penalizing equivocation by loss of bitcoins. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 219–230, 2015.
- [RS04] Phillip Rogaway and Thomas Shrimpton. Cryptographic hash-function basics: Definitions, implications, and separations for preimage resistance, second-preimage resistance, and collision resistance. In *International workshop on fast software encryption*, pages 371–388. Springer, 2004.

- [RST01] Ronald L Rivest, Adi Shamir, and Yael Tauman. How to leak a secret. In *International conference on the theory and application of cryptology and information security*, pages 552–565. Springer, 2001.
- [SJK⁺17] Ewa Syta, Philipp Jovanovic, Eleftherios Kokoris Kogias, Nicolas Gailly, Linus Gasser, Ismail Khoffi, Michael J Fischer, and Bryan Ford. Scalable bias-resistant distributed randomness. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 444–460. Ieee, 2017.
- [SPG21] Michael A Specter, Sunoo Park, and Matthew Green. KeyForge:non-attributable email from Forward-Forgeable Signatures. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1755–1773, 2021.
- [SWP04] Ron Steinfeld, Huaxiong Wang, and Josef Pieprzyk. Efficient extension of standard schnorr/rsa signatures into universal designated-verifier signatures. In *International Workshop on Public Key Cryptography*, pages 86–100. Springer, 2004.
- [WCL⁺23] Zhipeng Wang, Marko Cirkovic, Duc V. Le, William Knottenbelt, and Christian Cachin. Pay Less for Your Privacy: Towards Cost-Effective On-Chain Mixers. In *5th Conference on Advances in Financial Technologies (AFT 2023)*, volume 282, pages 16:1–16:25, 2023.
- [WCQ⁺23] Zhipeng Wang, Stefanos Chaliasos, Kaihua Qin, Liyi Zhou, Lifeng Gao, Pascal Berrang, Benjamin Livshits, and Arthur Gervais. On how zero-knowledge proof blockchain mixers improve, and worsen user privacy. In *Proceedings of the ACM Web Conference 2023*, pages 2022–2032, 2023.
- [Wes19] Benjamin Wesolowski. Efficient verifiable delay functions. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2019.

A Proofs of GroupForge

GroupForge satisfies the *correctness*, *unforgeability*, and *forgeability on expiry* properties of an FFS scheme.

Correctness follows from the underlying Σ_{DSKE} scheme and from the correctness of the Merkle proofs: every signature $\sigma = (s, pk_i, \text{leaf}_i, \text{proof}_i)$ constructed with `Sign()` will be valid according to $\Sigma_{\text{DSKE}}.\text{Verify}()$, and proof_i will be a valid Merkle proof that pk_i is the public key that corresponds to leaf_i .

Unforgeability is also reduced to the unforgeability of Σ_{DSKE} through standard arguments on Merkle-based constructions [BDH11]. Assume an adversary \mathcal{A} that is given PK and attacks GroupForge. For each time chunk $i \in [2^h]$, \mathcal{A} is allowed to make up to k signing queries. The challenger of \mathcal{A} delegates these queries to Σ_{DSKE} . Assume a successful forgery $\sigma^* = (s^*, pk^*, \text{leaf}^*, \text{proof}^*)$, for some leaf leaf^* . Since Σ_{DSKE} is unforgeable, pk^* must be a public key under which s^* is a valid signature. In this case, however, the hash of pk^* and the Merkle proof proof^* can be used by the challenger to break the collision resistance of $H()$.

Forgeability on expiry is proven in the following theorem.

Theorem 10. *GroupForge achieves correctness, unforgeability, and forgeability on expiry properties.*

Proof. Let \mathcal{D} be an adversary that is given a signature σ_b , for $b \in \{0, 1\}$, and, as described in the definition of FFS, aims to distinguish the value of b . Signature σ_0 is created with the private key sk_i of some time chunk i , i.e., $\sigma_0 \leftarrow \text{Sign}(sk_i, m)$. Signature σ_1 is created using the expiry information η , which in turn is computed from the additional material \mathcal{T} , i.e., $\eta \leftarrow \text{Expire}(\mathcal{T})$ and $\sigma_1 \leftarrow \text{Forge}(\eta, m)$. Let us assume that \mathcal{T} contains k valid message-signature pairs, i.e., $\mathcal{T} = \{(m_j, \sigma_j)\}_{j \in [k]}$, that all m_j are pairwise different, and that all signatures are created in time chunk i , i.e., $\sigma_j = (s_j, pk_i, \text{leaf}_i, \text{proof}_i)$, for each $j \in [k]$. In this case, since the underlying DSKE scheme Σ_{DSKE} has (k, δ) -extractable sets, $\text{Expire}(\mathcal{T})$ returns $h \neq \perp$ with probability δ . For this $h = (\text{leaf}_i, \text{proof}_i, sk'_i)$ we know from the properties of Σ_{DSKE} that $sk_i = sk'_i$, and $\text{Forge}()$ will compute σ_1 by calling $s = \Sigma_{\text{DSKE}}.\text{Sign}(sk_i, m)$, which is identical to how σ_0 is created. Hence, \mathcal{D} can only guess the value of b with probability $1/2$.

Overall, assuming that the signer has created enough message-signature pairs \mathcal{T} with the properties described above, with probability δ , \mathcal{D} can only at random guess the value of b , while with probability $1 - \delta$ (when $\text{Expire}(\mathcal{T})$ returns \perp), the challenger of \mathcal{D} cannot produce a valid σ_1 and \mathcal{D} can correctly guess b . The overall success probability of \mathcal{D} is $P = 1 - \delta/2$. Observe that for the polynomial commitment-based DSKE scheme $\delta = 1 - \text{negl}(\lambda)$, hence $P = 1/2 + \text{negl}(\lambda)$, while for the hash-based DSKE scheme $\delta \geq 1 - \lambda/(2^{k-1}) + \text{negl}(\lambda)$, hence P asymptotically approaches $1/2$ as k grows. \square

Setting the Deniable Group Size k and Delay Δ . GroupForge requires that no recipients of k signatures collaborate within a time chunk Δ . This requirement can be made plausible by adjusting two variables. First, the size of the deniable group k can be set to a value large enough, for example, $\lceil n/2 \rceil$, where n is the number of recipient email servers. This would be equivalent to the assumption that no more than half of the nodes may be Byzantine, which is typical in distributed protocols. Second, time delay Δ can be set to a low value. Similar to that of KeyForge, we can assume an

upper bound $\hat{\Delta}$ in the time required for email delivery and then set $\Delta = \hat{\Delta}$. Hence, even if the signature recipients collaborate, the forged signature has a high probability of reaching the recipient email server in the next time chunk.

Moreover, the value of k and the choice of the underlying DSKE scheme affect the probability of the algorithm `Forge()` outputting a signature and not \perp . For $\text{DSKE}_{\text{poly}}$ we have $\delta = 1$ for all $k \geq \ell + 1$, and the signer can choose the size of the group by selecting the degree of the polynomial accordingly. For $\text{DSKE}_{\text{lamp}}$ the probability δ changes with k .