Signature Schemes with Bounded Leakage Resilience

Jonathan Katz[*]

Abstract

A *leakage-resilient* cryptosystem remains secure even if arbitrary, but bounded, information about the secret key (or possibly other internal state information) is leaked to an adversary. Denote the length of the secret key by n. We show a signature scheme tolerating (optimal) leakage of up to $n - n^{\epsilon}$ bits of information about the secret key, and a more efficient one-time signature scheme that tolerates leakage of $(\frac{1}{4} - \epsilon) \cdot n$ bits of information about the signer's entire state. The latter construction extends to give a leakage-resilient t-time signature scheme. All these constructions are in the standard model under general assumptions.

1 Introduction

Proofs of security for cryptographic primitives traditionally treat the primitive as a "black box" which an adversary is able to access in a relatively limited fashion. For example, in the usual model for proving security of signature schemes an adversary is given the public key and allowed to request signatures on any messages of its choice, but is unable to get any *other* information about the secret key or any internal randomness or state information used during signature generation.

In real-world implementations of cryptographic primitives, on the other hand, an adversary may be able to recover a significant amount of additional information not captured by standard security models. Examples include information leaked by side-channel cryptanalysis [18, 19], fault attacks [6, 4], or timing attacks [5], or even bits of the secret key itself in case this key is improperly stored or erased [16]. Potentially, schemes can also be attacked when they are implemented using poor random number generation [25] (which can be viewed as giving the adversary additional information on the internal state, beyond that which would be available were the output truly random), or when the same key is used in multiple contexts (e.g., for decryption and signing).

In the past few years cryptographers have made tremendous progress toward modeling security in the face of such information leakage [22, 31], and in constructing *leakage-resilient* cryptosystems secure even in case such leakage occurs. (There has also been corresponding work on reducing unwanted leakage by, e.g., building tamper-proof hardware; this is not the focus of our work.) Most relevant to the current work is a recent series of results [11, 1, 28, 10, 32, 23, 2] showing cryptosystems that guarantee security even when *arbitrary* information about the secret key is leaked (under suitable restrictions); we discuss this work, along with other related results, in further detail below. This prior work gives constructions of stream ciphers [11, 28] (and hence stateful symmetric-key encryption and MACs), symmetric-key encryption schemes [10], public-key encryption schemes [1, 32, 23], and signature schemes [2] achieving various notions of leakage resilience.

[*]Dept. of Computer Science, University of Maryland. Work done while visiting IBM. Email: `jkatz@cs.umd.edu`.

Most prior work has focused on primitives for ensuring *secrecy.* The only work of which we are aware that deals with *authenticity* is that of Alwen et al. [2] which shows, among other results, leakage-resilient signature schemes based on number-theoretic assumptions in the random oracle model that tolerate leakage of up to half the secret key.[1] In this work, we give constructions of leakage-resilient signature schemes based on general assumptions in the standard model; our first construction tolerates (optimal) leakage of all but n^ϵ bits of information. We postpone a more thorough discussion of our results until after we define leakage resilience in more detail.

1.1 Modeling Leakage Resilience

We provide a brief overview of the framework in which leakage resilience is defined (specialized to signatures), providing a discussion of previous related work along the way.

At a high level, definitions of leakage resilience take the following form: Begin with a "standard" security notion (e.g., existential unforgeability under adaptive chosen message attacks [14]) and modify this definition by allowing the adversary to (adaptively) specify a series of *leakage functions* $f_1, \ldots$. The adversary, in addition to getting whatever other information is specified by the original security definition, is given the result of applying f_i to the secret key and possibly other internal state of the honest party (e.g., the signer). We then require that the adversary's success probability — in the case of signature schemes, the probability with which it can output a forged signature on a previously unsigned message — remain negligible. It should be clear that this is a general methodology that can be applied to many different primitives.

The exact model is then determined by the restrictions placed on the leakage function(s) f_i. We may distinguish the following possibilities:

Limited or arbitrary information A first issue to be resolved is whether the $\{f_i\}$ are allowed to be arbitrary (polynomial-time computable) functions, or whether they are restricted to be in some more limited class. Early work considered the latter case, for example where the adversary is restricted to learning *specific bits* of the secret key [8], or the values on *specific wires* of the circuit implementing the primitive [17]. More recent work [11, 1, 28, 10, 32, 23, 2] allows arbitrary $\{f_i\}$.

Bounded or unbounded information leakage. Let n denote the length of the secret key. If the secret key does not change over time, and the $\{f_i\}$ are allowed to be arbitrary, then security in the traditional sense cannot be achieved once the total length of the *leakage* — that is, the outputs of all the $\{f_i\}$ — is n bits or more. For the case of signatures, the length of the leakage must also be less than the signature length. This inherent restriction is used in [1, 32, 23]. (Alwen et al. [2] do not impose this restriction, but as a consequence can only achieve a weaker notion of security.)

One can avoid this restriction, and potentially tolerate an unbounded amount of leakage overall, if the secret key is updated over time; even in this case, one must somehow limit the amount of leakage between successive key updates. This approach to leakage resilience was considered in [11, 28] in the context of stateful symmetric-key primitives.

One can also avoid imposing a bound on the leakage by restricting the $\{f_i\}$, as discussed next.

Computational min-entropy of the secret key. If the leakage is shorter than the secret key (as discussed above), then the secret key will have high min-entropy conditioned on the leakage. This setting is considered in [1, 23, 32, 2], and is also enforced on a per-period basis in the work

[1] The results of [2] were obtained independently of our own work. Also, the primary focus of [2] was on identification schemes in the bounded retrieval model.

of [11, 28] (i.e., the leakage per time period is required to be shorter than the secret key). More recent work [10, 32] shows schemes that remain secure for leakage of arbitrary length, as long as the secret key remains exponentially hard to compute given the leakage (but even if the secret key is fully determined by the leakage in an information-theoretic sense). A drawback of this guarantee is that given some collection of functions $\{f_i\}$ (say, as determined experimentally for some particular set of side-channel attacks) there is no way to tell, in general, whether they satisfy the stated requirement or not. Furthermore, existing results in this direction currently require super-polynomial hardness assumptions.

Inputs to the leakage functions. A final consideration is the allowed inputs of the leakage functions. Work of [11, 28] assumes, following [22], that *only computation leaks information*; this is modeled by letting each f_i take as input only those portions of the secret key that are accessed during the ith phase of the scheme. Halderman et al. [16], however, show that memory contents can be leaked even when they are not being accessed. Motivated (in part) by this result, the schemes of [1, 10, 32, 23, 2] allow the $\{f_i\}$ to take the entire secret key as input at all times.

For the specific primitives considered in [11, 1, 28, 10, 32, 23], the secret key sk is the only internal state maintained by the party holding the secret key, and so allowing the $\{f_i\}$ to depend on sk is (almost) the most general choice.[2] For signature schemes, however, any randomness used during signing might also be leaked to an adversary. The strongest definition of leakage resilience is thus obtained by allowing the $\{f_i\}$ to depend on *all* the state information used by the honest signer during the course of the experiment.

All these variants may be meaningful depending on the particular attacks one is trying to model. Memory attacks [16, 1], which probe long-term secret information during a time when computation is *not* taking place, can be faithfully modeled by allowing the leakage functions to take only sk as input. On the other hand, side-channel attacks that collect information while computation is occurring might be more accurately captured by allowing the leakage functions to take as input only those portions of the internal state that are actively be accessed.

1.2 Our Results

With the preceding discussion in mind, we can now describe our results in further detail. In all cases, we allow the leakage function(s) to be *arbitrary* as long as the total leakage is *bounded* as some function of the secret key length n; recall that such a restriction on the leakage is essential if the secret key is unchanging, as it is in all our schemes. Our results can be summarized as follows:

1. We show a construction of a leakage-resilient signature scheme in the standard model, based on general (as opposed to number-theoretic) assumptions. This scheme tolerates leakage of $n - n^\epsilon$ bits of information about the secret key for any $\epsilon > 0$, which is optimal (unless one is willing to make super-polynomial hardness assumptions).

2. We also construct a leakage-resilient *one-time* signature scheme in the standard model. In contrast to the previous result, this scheme is more efficient and can be based on the minimal assumption that one-way functions exist; it also tolerates leakage that may depend on the entire state of the signer throughout the experiment. (This follows from the fact that signing

[2] More generally, one could also allow the $\{f_i\}$ to depend on the *randomness* used to generate the (public and) secret key(s); this possibility is mentioned in [23, Section 8.2]. (For the specific schemes considered in [11, 1, 28, 10, 32, 23], however, this makes no substantive difference.)

is deterministic.) Here security holds only as long as the leakage is bounded by $(\frac{1}{4} - \epsilon) \cdot n$ bits, for any $\epsilon > 0$. This construction extends to give an t-time signature scheme tolerating leakage of $\Theta(n/t^2)$ bits.

In the appendix, we also discuss an efficient scheme in the random oracle model that is secure as long as the leakage is bounded by $(\frac{1}{2} - \epsilon) \cdot n$ bits for any $\epsilon > 0$. This scheme was discovered independently by Alwen et al. [2]. Our analysis offers some advantages as compared to theirs; see the appendix for further discussion. (We include the result for completeness, but do not claim any significant novelty with regard to the existing work of [2].)

1.3 Overview of Our Techniques

Our constructions all rely on the same basic idea. Roughly, we consider signature schemes with the following properties:

- A given public key pk corresponds to a set S_{pk} of *exponentially many* secret keys. Furthermore, given (sk, pk) with $sk \in S_{pk}$ it remains hard to compute any other $sk' \in S_{pk}$.
- The secret key sk used by the signer has high min-entropy (at least in a computational sense) even for an adversary who observes signatures on messages of its choice. (For our one-time scheme, this is only required to hold for an adversary who observes a single signature.)
- A signature forgery can be used to compute a secret key in S_{pk}.

To prove that any such signature scheme is leakage resilient, we show how to use an adversary $\mathcal{A}$ attacking the scheme to find distinct $sk, sk' \in S_{pk}$ given (sk, pk) (in violation of the assumed hardness of doing so). Given (sk, pk), we simply run $\mathcal{A}$ on input pk and respond to its signing queries using the given key sk. Leakage queries can also be answered using sk. If the adversary forges a signature, we extract some $sk' \in S_{pk}$; it remains only to show that $sk' \neq sk$ with high probability. Let $n = \log |S_{pk}|$ be the (computational) min-entropy of sk conditioned on pk and the signatures seen by the adversary. (We assume that all secret keys in S_{pk} are equally likely, which will be the case in our constructions.) A standard argument shows that if the leakage is bounded by ℓ bits, then the conditional min-entropy of the secret key is still at least $n - \ell - t$ bits except with probability 2^{-t}. So as long as the leakage is bounded away from n, with high probability the min-entropy of sk conditioned on $\mathcal{A}$'s entire view is still at least 2. But then $sk' \neq sk$ with probability at least $1/2$. This concludes the outline of the proof. We remark, however, that various subtleties arise in the formal proofs of security.

Some existing signature schemes in the random oracle model already satisfy the requirements stated above. In particular, these include schemes constructed using the Fiat-Shamir transform [12] applied to a witness-indistinguishable Σ-protocol where there are an *exponential* number of witnesses corresponding to a given statement. Concrete examples include the signature schemes of Okamoto [26] (extending the Schnorr [30] and Guillou-Quisquater [15] schemes) based on the discrete logarithm or RSA assumptions, as well as the signature scheme of Fischlin and Fischlin [13] (extending the Ong-Schnorr [27] scheme) based on the hardness of factoring. This class of schemes was also considered by Alwen et al. [2]. See Appendix A for further discussion.

We are not aware of any existing signature scheme in the standard model that meets our requirements. We construct one as follows. Let H be a universal one-way hash function (UOWHF) [24] mapping n-bit inputs to n^ϵ-bit outputs. The secret key of the signature scheme is $x \in \{0,1\}^n$, and the public key is $(y = H(x), pk, r)$ where pk is a public key for a CPA-secure public-key encryption scheme, and r is a common reference string for an unbounded simulation-sound NIZK proof system [29, 9]. A signature on a message m consists of an encryption $C \leftarrow \mathsf{Enc}_{pk}(m\|x)$ of both m and x, along with a proof π that C is an encryption of $m\|x'$ with $H(x') = y$. Observe that, with high probability over choice of x, there are exponentially many pre-images of $y = H(x)$ and hence exponentially many valid secret keys; furthermore, finding another such secret key $sk' \neq sk$ requires finding a collision in H. Details are given in Section 3.

Our leakage-resilient one-time signature scheme is constructed using a similar idea, applied to the Lamport signature scheme [21]. That is, the secret key is $\{(x_{i,0}, x_{i,1})\}_{i=1}^{k}$ and the public key is $\{(y_{i,0}, y_{i,1})\}_{i=1}^{k}$ where k is the message length and $y_{i,b} = H(x_{i,b})$ for H a UOWHF as above. Once again, there are exponentially many secret keys associated with any public key and finding any two such keys yields a collision in H. Adapting the standard Lamport scheme in this way yields a signature scheme secure against leakage if $n^{1-\epsilon}$ bits. By encoding the message using an error-correcting code with high minimum distance, it is possible to "boost" the leakage resilience to $(\frac{1}{4} - \epsilon) \cdot n$ bits. Finally, using cover-free families this approach extends also to give a leakage-resilient t-time signature scheme. These constructions are all described in Section 4.

2 Definitions and Preliminaries

We provide a formal definition of leakage resilience for signature schemes, and state a technical lemma that will be used repeatedly in our ananlsis. We denote the security parameter by k, and let PPT stand for "probabilistic polynomial time".

2.1 Definitions

Definition 1 A *signature scheme* is a tuple of PPT algorithms (Gen, Sign, Vrfy) such that:

- Gen is a randomized algorithm that takes as input 1^k and outputs (pk, sk), where pk is the public key and sk is the secret key.
- Sign is a (possibly) randomized algorithm that takes as input the secret key sk, the public key pk, and a message m, and outputs a signature σ. We denote this by $\sigma \leftarrow \mathsf{Sign}_{sk}(m)$, leaving the public key implicit.[3]
- Vrfy is a deterministic algorithm that takes as input a public key pk, a message m, and a purported signature σ. It outputs a bit b indicating acceptance or rejection, and we write this as $b := \mathsf{Vrfy}_{pk}(m, \sigma)$.

It is required that for all k, all (pk, sk) output by $\mathsf{Gen}(1^k)$, and all messages m in the message space, we have $\mathsf{Vrfy}_{pk}(m, \mathsf{Sign}_{sk}(m)) = 1$. $\diamondsuit$

Our definition of leakage resilience is the standard notion of existential unforgeability under adaptive chosen-message attacks [14], except that we additionally allow the adversary to specify arbitrary leakage functions $\{f_i\}$ and obtain the value of these functions applied to the secret key (and possibly other state information).

[3] Usually, one assumes without loss of generality that the public key is included as part of the secret key. Since we measure leakage as a function of the secret-key length, however, we seek to minimize the size of the secret key.

Definition 2 Let $\Pi = (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Vrfy})$ be a signature scheme, and let λ be a function. Given an adversary $\mathcal{A}$, define the following experiment parameterized by k:

1. Choose $r \leftarrow \{0,1\}^*$ and compute $(pk, sk) := \mathsf{Gen}(1^k; r)$. Set $\mathsf{state} := \{r\}$.
2. Run $\mathcal{A}(1^k, pk)$. The adversary may then adaptively access a *signing oracle* $\mathsf{Sign}_{sk}(\cdot)$ and a *leakage oracle* $\mathsf{Leak}(\cdot)$ that have the following functionality:
    - In response to the ith query $\mathsf{Sign}_{sk}(m_i)$, this oracle chooses random $r_i \leftarrow \{0,1\}^*$, computes $\sigma_i := \mathsf{Sign}_{sk}(m_i; r_i)$, and returns σ_i to $\mathcal{A}$. It also sets $\mathsf{state} := \mathsf{state} \cup \{r_i\}$.
    - In response to the ith query $\mathsf{Leak}(f_i)$ (where f_i is specified as a circuit), this oracle gives $f_i(\mathsf{state})$ to $\mathcal{A}$. (To make the definition meaningful in the random oracle model, the $\{f_i\}$ are allowed to be oracle circuits that depend on the random oracle H.)

      The $\{f_i\}$ can be arbitrary, subject to the restriction that the total output length of all the f_i is at most $\lambda(|sk|)$.
3. At some point, $\mathcal{A}$ outputs (m, σ).

We say $\mathcal{A}$ *succeeds* if (1) $\mathsf{Vrfy}_{pk}(m, \sigma) = 1$ and (2) m was not previously queried to the $\mathsf{Sign}_{sk}(\cdot)$ oracle. We denote the probability of this event by $\Pr[\mathsf{Succ}^{\lambda\text{-leakage}^*}_{\mathcal{A},\Pi}(k)]$. We say Π is *fully λ-leakage resilient* if $\Pr[\mathsf{Succ}^{\lambda\text{-leakage}^*}_{\mathcal{A},\Pi}(k)]$ is negligible for every PPT adversary $\mathcal{A}$.

If state is not updated after each signing query (and so always contains only the randomness r used to generate the secret key), we denote the probability of success by $\Pr[\mathsf{Succ}^{\lambda\text{-leakage}}_{\mathcal{A},\Pi}(k)]$ and say Π is *λ-leakage resilient* if $\Pr[\mathsf{Succ}^{\lambda\text{-leakage}}_{\mathcal{A},\Pi}(k)]$ is negligible for every PPT adversary $\mathcal{A}$. ◊

Leakage resilience in the definition above corresponds to the *memory attacks* of [1] (except that we allow the leakage to depend also on the random coins used to generate the secret key). Other variations of the definition are, of course, also possible: state could include only sk (and not the random coins r used to generate it), or could include only the most recently used random coins r_i.

2.2 A Technical Lemma

Let X be a random variable taking values in $\{0,1\}^n$. The *min-entropy* of X is given by

$$H_\infty(X) \stackrel{\text{def}}{=} \min_{x \in \{0,1\}^n} \{-\log_2 \Pr[X = x]\}.$$

The conditional min-entropy of X given an event E is defined as:

$$H_\infty(X \mid E) \stackrel{\text{def}}{=} \min_{x \in \{0,1\}^n} \{-\log_2 \Pr[X = x \mid E]\}.$$

Lemma 1 *Let X be a random variable with $H \stackrel{\text{def}}{=} H_\infty(X)$, and fix $\delta \in [0, H]$. Let f be an arbitrary function with range $\{0,1\}^\lambda$, and set $Y \stackrel{\text{def}}{=} \left\{y \in \{0,1\}^\lambda \mid H_\infty(X \mid y = f(X)) \leq H - \Delta\right\}$. Then*

$$\Pr[f(X) \in Y] \leq 2^{\lambda - \Delta}.$$

In words: the probability that knowledge of $f(X)$ decreases the min-entropy of X by Δ or more is at most $2^{\lambda-\Delta}$. Put differently, the min-entropy of X after observing the value of $f(X)$ is greater than H' except with probability at most $2^{\lambda-H+H'}$.

Proof Fix $y \in \{0,1\}^\lambda$ and $x \in \{0,1\}^n$ with $f(x) = y$. Since

$$\Pr[X = x \mid y = f(X)] = \frac{\Pr[X = x]}{\Pr[y = f(X)]}\,,$$

we have that $y \in Y$ only if $\Pr[y = f(X)] \leq 2^{-\Delta}$. The fact that the range of f is $\{0,1\}^\lambda$ means that $|Y| \leq 2^\lambda$, and it follows that $\Pr[f(X) \in Y] \leq 2^{\lambda-\Delta}$ as claimed. ■

3 A Leakage-Resilient Signature Scheme

We construct a leakage-resilient signature scheme in the standard model, following the intuition described in Section 1.2. Let (Gen_H, H) be a UOWHF [24] mapping n-bit inputs to $\frac{1}{2} \cdot n^\epsilon$-bit outputs for $n = \mathsf{poly}(k)$ and $\epsilon \in (0,1)$; we assume that finding second pre-images is hard even given the randomness used to generate the hash key. (Standard constructions of UOWHFs have this property.) Let $(\mathsf{Gen}_E, \mathsf{Enc}, \mathsf{Dec})$ be a CPA-secure, dense[4] public-key encryption scheme, and let $(\ell, \mathcal{P}, \mathcal{V}, \mathcal{S}_1, \mathcal{S}_2)$ be an unbounded simulation-sound NIZK proof system [9] for the following language L:

$$L = \{(s, y, pk, m, C) : \exists x, \omega \text{ s.t. } C = \mathsf{Enc}_{pk}(x;\omega) \text{ and } H_s(x) = y\}\,.$$

The signature scheme is defined as follows:

Key generation: Choose random $x \leftarrow \{0,1\}^n$ and compute $s \leftarrow \mathsf{Gen}_H(1^k)$. Obliviously sample a public key pk for the encryption scheme, and choose a random string $r \leftarrow \{0,1\}^{\ell(k)}$. The public key is $(s, y := H_s(x), pk, r)$ and the secret key is x.

Signing: To sign message m using secret key x and public key (s, y, pk, r), first choose random ω and compute $C := \mathsf{Enc}_{pk}(x;\omega)$. Then compute $\pi \leftarrow \mathcal{P}_r((s, y, pk, m, C), (x, \omega))$; i.e., π is a proof that $(s, y, pk, m, C) \in L$ using witness (x, ω). The signature is (C, π).

Verification: Given a signature (C, π) on the message m with respect to the public key (s, y, pk, r), output 1 iff $\mathcal{V}_r((s, y, pk, m, C), \pi) = 1$.

Theorem 1 *Under the stated assumptions, the signature scheme above is $(n - n^\epsilon)$-leakage resilient.*

Proof (Sketch) Let Π denote the scheme given above, and let $\mathcal{A}$ be a PPT adversary with $\delta = \delta(k) \stackrel{\text{def}}{=} \Pr[\mathsf{Succ}^{\lambda\text{-leakage}}_{\mathcal{A},\Pi}(k)]$. We consider a sequence of experiments, and let $\Pr_i[\cdot]$ denote the probability of an event in experiment i. We abbreviate $\mathsf{Succ}^{\lambda\text{-leakage}}_{\mathcal{A},\Pi}(k)$ by Succ.

Experiment 0: This is the experiment of Definition 2. Given the public key (s, y, pk, r) defined by the experiment, Succ denotes the event that $\mathcal{A}$ outputs $(m, (C, \pi))$ where $\mathcal{V}_r((s, y, pk, m, C), \pi) = 1$ and m was never queried to the signing oracle. By assumption, we have $\Pr_0[\mathsf{Succ}] = \delta$.

[4] This means it is possible to sample a public key "obliviously," without knowing the corresponding secret key.

Experiment 1: We introduce the following differences with respect to the preceding experiment: when setting up the public key, we now generate the common random string r of the simulation-sound NIZK by computing $(r, \tau) \leftarrow \mathcal{S}_1(1^k)$. Furthermore, signing queries are now answered as follows: to sign m, generate $C \leftarrow \mathsf{Enc}_{pk}(x)$ as before but compute π as $\pi \leftarrow \mathcal{S}_2((s, y, pk, m, C), \tau)$.

It follows from the (adaptive) zero-knowledge property of $(\ell, \mathcal{P}, \mathcal{V}, \mathcal{S}_1, \mathcal{S}_2)$, that the difference $|\Pr_1[\mathsf{Succ}] - \Pr_0[\mathsf{Succ}]|$ must be negligible.

Experiment 2: We modify the preceding experiment in the following way: to answer a signing query for a message m, compute $C \leftarrow \mathsf{Enc}_{pk}(0^n)$ (and then compute π as in Experiment 1). CPA-security of the encryption scheme implies that $|\Pr_2[\mathsf{Succ}] - \Pr_1[\mathsf{Succ}]|$ is negligible.

Experiment 3: We now change the way the public key is generated. Namely, instead of obliviously sampling the encryption public key pk we compute it as $(pk, sk) \leftarrow \mathsf{Gen}_E(1^k)$. Note that this is only a syntactic change and so $\Pr_3[\mathsf{Succ}] = \Pr_2[\mathsf{Succ}]$. (This assumes perfect oblivious sampling; if an obliviously generated public key and a legitimately generated public key are only computationally indistinguishable, then the probability of Succ is affected by a negligible amount.)

Given the public key (s, y, pk, r) defined by the experiment, let Ext be the event that $\mathcal{A}$ outputs $(m, (C, \pi))$ such that Succ occurs and furthermore $H_s(\mathsf{Dec}_{sk}(C)) = y$. Unbounded simulation soundness of the NIZK proof system implies that $|\Pr_3[\mathsf{Ext}] - \Pr_3[\mathsf{Succ}]|$ is negligible. (Note that by definition of L the message m is included as part of the statement being proved, and so if $\mathcal{A}$ did not request a signature on m then it was never given a simulated proof of the statement (s, y, pk, m, C).)

To complete the proof, we show that $\Pr_3[\mathsf{Ext}]$ is negligible. Consider the following adversary $\mathcal{B}$ finding a second-preimage in the UOWHF: $\mathcal{B}$ chooses random $x \leftarrow \{0,1\}^n$ and is given key s (along with the randomness used to generate s). $\mathcal{B}$ then runs Experiment 3 with $\mathcal{A}$. In this experiment all signatures given to $\mathcal{A}$ are simulated (as described in Experiment 3 above); furthermore $\mathcal{B}$ can easily answer any leakage queries made by $\mathcal{A}$ since $\mathcal{B}$ knows a legitimate secret key. (Recall that here we allow the leakage functions to be applied only to [the randomness used to generate] the secret key, but not to any auxiliary state used during signing.) If event Ext occurs when $\mathcal{A}$ terminates, then $\mathcal{B}$ recovers a value $x' \stackrel{\text{def}}{=} \mathsf{Dec}_{sk}(C)$ for which $H_s(x') = y = H_s(x)$; i.e., $\mathcal{B}$ recovers such an x' with probability exactly $\Pr_3[\mathsf{Ext}]$. We now argue that $x' \neq x$ with high probability.

The only information about x revealed to $\mathcal{A}$ in Experiment 3 comes from the value y included in the public key and the leakage queries asked by $\mathcal{A}$; these total at most $\frac{1}{2} \cdot n^\epsilon + (n - n^\epsilon) = n - \frac{1}{2} \cdot n^\epsilon$ bits. Using Lemma 1 with $\Delta = H_\infty(x) = n$, the probability that $H_\infty(x \mid \mathcal{A}\text{'s view}) = 0$ (i.e., the probability that x is uniquely determined by the view of $\mathcal{A}$) is at most $2^{-n^\epsilon/2}$, which is negligible. When the conditional min-entropy of x is greater than 0 there are at least two (equally probable) possibilities for x and so $x' \neq x$ with probability at least $\frac{1}{2}$. Taken together, the probability that $\mathcal{B}$ recovers $x' \neq x$ with $H_s(x') = H_s(x)$ is at least

$$\frac{1}{2} \cdot \left(\Pr_3[\mathsf{Ext}] - 2^{-n^\epsilon/2}\right).$$

We thus see that if $\Pr_3[\mathsf{Ext}]$ is not negligible then $\mathcal{B}$ violates the security of the UOWHF with non-negligible probability, a contradiction. ■

4 A Fully Leakage-Resilient One-Time Signature Scheme

In this section we describe constructions of fully leakage-resilient one-time and t-time signature schemes. These results are incomparable to the result of the previous section: on the positive side, here we achieve *full* leakage resilience as well as better efficiency; on the downside, the schemes given here are only secure when the adversary obtains a bounded number of signatures, and the leakage that can be tolerated is lower.

We describe a basic one-time signature scheme, and then present an extension that tolerates leakage of information up to a constant fraction of the secret key length. Let (Gen_H, H) be a UOWHF mapping k^c-bit inputs to k-bit outputs for some $c > 1$. (As before, we assume H is secure even given the randomness used to generate the hash key.) Our basic scheme is a variant on Lamport's signature scheme [21], using H as the one-way function:

Key generation: Choose random $x_{i,0}, x_{i,1} \leftarrow \{0,1\}^{k^c}$ for $i = 1, \ldots, k$, and generate $s \leftarrow \mathsf{Gen}_H(1^k)$. Compute $y_{i,b} := H_s(x_{i,b})$ for $i \in \{1, \ldots, k\}$ and $b \in \{0,1\}$. The public key is $(s, \{y_{i,b}\})$ and the secret key is $\{x_{i,b}\}$.

Signing: The signature on a k-bit message $m = m_1, \ldots, m_k$ consists of the k values $x_{1,m_1}, \ldots, x_{k,m_k}$.

Verification: Given a signature $x_1, \ldots, x_k$ on the k-bit message $m = m_1, \ldots, m_k$ with respect to the public key $(s, \{y_{i,b}\})$, output 1 iff $y_{i,m_i} \stackrel{?}{=} H_s(x_i)$ for all i.

It can be shown that the above scheme is fully $n^{(c-1)/(c+1)}$-leakage resilient (as a one-time signature scheme), where $n = 2k^{c+1}$ denotes the length of the secret key. (We omit the proof, since we will prove security for an improved scheme below.) Setting c appropriately, the above approach thus tolerates leakage $n^{1-\epsilon}$ for any desired $\epsilon > 0$. Note that the bound on the leakage is essentially tight, since an adversary who obtains the signature on the message 0^k and then leaks the value $x_{1,1}$ (which is only $k^c = (n/2)^{c/(c+1)}$ bits) can forge a signature on the message 10^{k-1}.

Tolerating leakage linear in the secret key length. An extension of the above scheme allows us to tolerate greater leakage: specifically, we apply Lamport's scheme to a high-distance *encoding* of the message. Details follow.

If A is a $k \times \ell$ matrix over $\{0,1\}$ (viewed as the field $\mathbb{F}_2$), then A defines a (linear) error-correcting code $\mathcal{C} \subset \{0,1\}^\ell$ where the message $m \in \{0,1\}^k$ (viewed as a row vector) is mapped to the codeword $m \cdot A$. It is well known that for every $\epsilon > 0$ there exists a constant R such that choosing $A \in \{0,1\}^{k \times Rk}$ uniformly at random defines a code with relative minimum distance $\frac{1}{2} - \epsilon$, except with probability negligible in k. (We will not need efficient decodability.)

Fix a constant $\epsilon \in (0,1)$ and let R be as above; set $\ell = Rk$. Let (Gen_H, H) be a UOWHF mapping ℓ_{in}-bit inputs to k-bit outputs where $\ell_{in} = 2k/\epsilon$. The signature scheme is defined as follows:

Key generation: Choose random $A \in \{0,1\}^{k \times \ell}$ and $x_{i,0}, x_{i,1} \leftarrow \{0,1\}^{\ell_{in}}$ for $i = 1, \ldots, \ell$. Generate $s \leftarrow \mathsf{Gen}_H(1^k)$. Compute $y_{i,b} := H_s(x_{i,b})$ for $i \in \{1, \ldots, \ell\}$ and $b \in \{0,1\}$. The public key is $(A, s, \{y_{i,b}\})$ and the secret key is $\{x_{i,b}\}$.

Signing: To sign a message $m \in \{0,1\}^k$, first compute $\bar{m} = m \cdot A \in \{0,1\}^\ell$. The signature then consists of the ℓ values $x_{1,\bar{m}_1}, \ldots, x_{\ell,\bar{m}_\ell}$.

Verification: Given a signature $x_1, \ldots, x_\ell$ on the message m with respect to the public key $(A, s, \{y_{i,b}\})$, first compute $\bar{m} = m \cdot A$ and then output 1 iff $y_{i,\bar{m}_i} \stackrel{?}{=} H_s(x_i)$ for all i.

Theorem 2 *If H is a UOWHF then the scheme above is a one-time signature scheme that is fully $(\frac{1}{4} - \epsilon) \cdot n$-leakage resilient, where $n = 2\ell \cdot \ell_{in}$ denotes the length of the secret key.*

Proof Let Π denote the scheme given above, and let $\mathcal{A}$ be a PPT adversary with $\delta = \delta(k) \stackrel{\text{def}}{=} \Pr[\mathsf{Succ}^{\lambda\text{-leakage}^*}_{\mathcal{A},\Pi}(k)]$. We construct an adversary $\mathcal{B}$ breaking the security of H with probability at least $(\delta - \mathsf{negl}(k))/4\ell$, implying that δ must be negligible.

$\mathcal{B}$ chooses random $A \in \{0,1\}^{k \times \ell}$ and $x_{i,0}, x_{i,1} \leftarrow \{0,1\}^{\ell_{in}}$ for $i = 1, \ldots, \ell$; we let $\mathcal{X} = \{x_{i,b}\}$ denote the set of secret key values $\mathcal{B}$ chooses and observe that $H_\infty(\mathcal{X}) = 2\ell \cdot \ell_{in}$. Next, $\mathcal{B}$ selects a random $b^* \in \{0,1\}$ and a random index $i^* \in \{1, \ldots, \ell\}$, and outputs x_{i^*,b^*}; it is given in return a hash key s. Then $\mathcal{B}$ computes $y_{i,b} := H_s(x_{i,b})$ for all i, b and gives the public key $(A, s, \{y_{i,b}\})$ to $\mathcal{A}$.

$\mathcal{B}$ answers the signing and leakage queries of $\mathcal{A}$ using the secret key $\{x_{i,b}\}$ that it knows. Since this secret key is distributed identically to the secret key of an honest signer, the simulation for $\mathcal{A}$ is perfect and $\mathcal{A}$ outputs a forgery with probability δ.

Let $\bar{m}$ denote the encoding of the message m whose signature was requested by $\mathcal{A}$. The information $\mathcal{A}$ has about $\mathcal{X}$ consists of: (1) the signature $(x_{1,\bar{m}_1}, \ldots, x_{\ell,\bar{m}_\ell})$ it obtained; (2) the values $\{y_{i,1-\bar{m}_i}\}_{i=1}^{\ell}$ from the public key and (3) the answers to the leakage queries asked by $\mathcal{A}$. Together, these total $\ell \cdot \ell_{in} + \ell k + (\frac{1}{4} - \epsilon) \cdot 2\ell \cdot \ell_{in}$ bits. By Lemma 1, it follows that $H_\infty(\mathcal{X} \mid \mathcal{A}\text{'s view}) > (\frac{1}{2} + \epsilon) \cdot \ell \cdot \ell_{in}$ except with probability at most

$$2^{\left(\ell \cdot \ell_{in} + \ell k + (\frac{1}{2} - 2\epsilon)\ell \cdot \ell_{in}\right) - 2\ell \cdot \ell_{in} + (\frac{1}{2} + \epsilon) \cdot \ell \cdot \ell_{in}} = 2^{\ell k - \epsilon \ell \cdot \ell_{in}},$$

which is negligible.

Assuming $H_\infty(\mathcal{X} \mid \mathcal{A}\text{'s view}) > (\frac{1}{2} + \epsilon) \cdot \ell \cdot \ell_{in}$, there is *no* set $I \subseteq [\ell]$ with $|I| \geq (\frac{1}{2} - \epsilon) \cdot \ell$ for which the values $\{x_{i,1-\bar{m}_i}\}_{i \in I}$ are all fixed given $\mathcal{A}$'s view. To see this, assume the contrary. Then

$$H_\infty(\mathcal{X} \mid \mathcal{A}\text{'s view}) \leq \sum_{i \notin I} H_\infty(x_{i,1-\bar{m}_i} \mid \mathcal{A}\text{'s view}) \leq \left(\frac{1}{2} + \epsilon\right) \ell \cdot \ell_{in},$$

in contradiction to the assumed bound on the conditional min-entropy of $\mathcal{X}$.

Let $(m^*, (x_1^*, \ldots, x_\ell^*))$ denote the forgery output by $\mathcal{A}$, and let $\bar{m}^* = m^* \cdot A$ denote the encoding of m^*. Let I be the set of indices where $\bar{m}$ and $\bar{m}^*$ differ; with all but negligible probability over choice of the matrix A it holds that $|I| \geq (\frac{1}{2} - \epsilon) \cdot \ell$ and so we assume this to be the case. By the argument of the previous paragraph, it cannot be the case that the $\{x_{i,1-\bar{m}_i}\}_{i \in I}$ are all fixed given $\mathcal{A}$'s view. But then with probability at least half we have $x_i^* \neq x_{i,\bar{m}_i^*}$ for at least one index $i \in I$. Assuming this to be the case, with probability at least $1/2\ell$ this difference occurs at the index (i^*, b^*) guessed at the outset by $\mathcal{B}$; when this happens $\mathcal{B}$ has found a collision in H for the given hash key s. Putting everything together, we see that $\mathcal{B}$ finds a collision in H with probability at least $(\delta - \mathsf{negl}(k)) \cdot \frac{1}{2} \cdot \frac{1}{2\ell}$, as claimed. ■

A t-time signature scheme. The idea above can be further extended to give a fully leakage resilient t-time signature scheme using *cover-free families.* We follow the definition of [20].

Definition 3 A family of non-empty sets $\mathcal{S} = \{S_1, \ldots, S_N\}$, where $S_i \subset U$, is $(t, \frac{1}{2})$-*cover-free* if for all $S, S_1, \ldots, S_t \in \mathcal{S}$ it holds that $\left|S \setminus \cup_{i=1}^{t} S_i\right| \geq |S|/2$. ◇

Kumar et al. [20] show an explicit construction that, for any t and k, yields a $(t, \frac{1}{2})$-cover free family $\mathcal{S} = \{S_1, \ldots, S_N\}$ where the number of sets is $N = \Omega(2^k)$, the size of each set is $|S_i| = O(kt)$, and the universe size is $|U| = O(kt^3)$. If we let $f : \{0,1\}^k \to \mathcal{S}$ denote an injective map, we obtain the following scheme:

Key generation: Set $\ell = O(kt^3)$ and $\ell_{in} = 8t^2k$. Choose $x_i \leftarrow \{0,1\}^{\ell_{in}}$ for $i = 1, \ldots, \ell$. Generate $s \leftarrow \mathsf{Gen}_H(1^k)$, and compute $y_i := H_s(x_i)$ for $i \in \{1, \ldots, \ell\}$. The public key is $(s, \{y_i\}_{i=1}^{\ell})$ and the secret key is $\{x_i\}_{i=1}^{\ell}$.

Signing: To sign a message $m \in \{0,1\}^k$, first compute $f(m) = S_m \in \mathcal{S}$. The signature then consists of $\{x_i\}_{i \in S_m}$.

Verification: Given a signature $\{x_i\}$ on the message m with respect to the public key $(s, \{y_{i,b}\})$, first compute $S_m = f(m)$ and then output 1 iff $y_i \stackrel{?}{=} H_s(x_i)$ for all $i \in S_m$.

A proof of the following proceeds along exactly the same lines as the proof of Theorem 2:

Theorem 3 *If H is a UOWHF then the scheme above is a t-time signature scheme that is fully $\Theta(n/t^2)$-leakage resilient, where $n = \ell \cdot \ell_{in}$ denotes the length of the secret key.*

Acknowledgments

I am grateful to Yael Tauman Kalai and Vinod Vaikuntanathan for stimulating discussions about leakage resilience and their earlier work [1, 10, 32]. I also thank Krzysztof Pietrzak for suggesting the extension to the one-time signature scheme that tolerates leakage linear in the secret key length.

References

[1] A. Akavia, S. Goldwasser, and V. Vaikuntanathan. Simultaneous hardcore bits and cryptography against memory attacks. In *6th Theory of Cryptography Conference — TCC 2009*, volume 5444 of *LNCS*, pages 474–495. Springer, 2009.

[2] J. Alwen, Y. Dodis, and D. Wichs. Public key cryptography in the bounded retrieval model and security against side-channel attacks. In *Advances in Cryptology — Crypto 2009*, LNCS, pages ???–??? Springer, 2009. Available at `http://eprint.iacr.org`.

[3] M. Bellare, O. Goldreich, and S. Goldwasser. Incremental cryptography: The case of hashing and signing. In *Advances in Cryptology — Crypto '94*, volume 839 of *LNCS*, pages 216–233. Springer, 1994.

[4] E. Biham, Y. Carmeli, and A. Shamir. Bug attacks. In *Advances in Cryptology — Crypto 2008*, volume 5157 of *LNCS*, pages 221–240. Springer, 2008.

[5] D. Boneh and D. Brumley. Remote timing attacks are practical. *Computer Networks*, 48(5):701–716, 2005.

[6] D. Boneh, R. A. DeMillo, and R. J. Lipton. On the importance of checking cryptographic protocols for faults. In *Advances in Cryptology — Eurocrypt '97*, volume 1233 of *LNCS*, pages 37–51. Springer, 1997.

[7] S. Brands. An efficient off-line electronic cash system based on the representation problem. Technical Report CS-R9323, CWI, 1993.

[8] R. Canetti, Y. Dodis, S. Halevi, E. Kushilevitz, and A. Sahai. Exposure-resilient functions and all-or-nothing transforms. In *Advances in Cryptology — Eurocrypt 2000*, volume 1807 of *LNCS*, pages 453–469. Springer, 2000.

[9] A. De Santis, G. Di Crescenzo, R. Ostrovsky, G. Persiano, and A. Sahai. Robust non-interactive zero knowledge. In *Advances in Cryptology — Crypto 2001*, volume 2139 of *LNCS*, pages 566–598. Springer, 2001.

[10] Y. Dodis, Y. Kalai, and S. Lovett. On cryptography with auxiliary input. In *41st Annual ACM Symposium on Theory of Computing (STOC)*, pages ???–??? ACM Press, 2009.

[11] S. Dziembowski and K. Pietrzak. Leakage-resilient cryptography. In *49th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 293–302. IEEE Computer Society Press, 2008. Full version available at `http://eprint.iacr.org/2008/240`.

[12] A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *Advances in Cryptology — Crypto '86*, volume 263 of *LNCS*, pages 186–194. Springer, 1987.

[13] M. Fischlin and R. Fischlin. The representation problem based on factoring. In *Cryptographers' Track — RSA 2002*, volume 2271 of *LNCS*, pages 96–113. Springer, 2002.

[14] S. Goldwasser, S. Micali, and R. L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, 1988.

[15] L. C. Guillou and J.-J. Quisquater. A "paradoxical" indentity-based signature scheme resulting from zero-knowledge. In *Advances in Cryptology — Crypto '88*, volume 403 of *LNCS*, pages 216–231. Springer, 1990.

[16] A. Halderman, S. Schoen, N. Heninger, W. Clarkson, W. Paul, J. Calandrino, A. Feldman, J. Applebaum, and E. Felten. Lest we remember: Cold boot attacks on encryption keys. In *Proc. 17th USENIX Security Symposium*, pages 45–60. USENIX Association, 2008.

[17] Y. Ishai, A. Sahai, and D. Wagner. Private circuits: Securing hardware against probing attacks. In *Advances in Cryptology — Crypto 2003*, volume 2729 of *LNCS*, pages 463–481. Springer, 2003.

[18] P. C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In *Advances in Cryptology — Crypto '96*, volume 1109 of *LNCS*, pages 104–113. Springer, 1996.

[19] P. C. Kocher, J. Jaffe, and B. Jun. Differential power analysis. In *Advances in Cryptology — Crypto '99*, volume 1666 of *LNCS*, pages 388–397. Springer, 1999.

[20] R. Kumar, S. Rajagopalan, and A. Sahai. Coding constructions for blacklisting problems without computational assumptions. In *Advances in Cryptology — Crypto '99*, volume 1666 of *LNCS*, pages 609–623. Springer, 1999.

[21] L. Lamport. Constructing digital signatures from a one-way function. Technical Report SRI-CSL-98, SRI International Computer Science Laboratory, Oct. 1979.

[22] S. Micali and L. Reyzin. Physically observable cryptography. In *1st Theory of Cryptography Conference — TCC 2004*, volume 2951 of *LNCS*, pages 278–296. Springer, 2004.

[23] M. Naor and G. Segev. Public-key cryptosystems resilient to key leakage. In *Advances in Cryptology — Crypto 2009*, LNCS, pages ???–??? Springer, 2009. Available at `http://eprint.iacr.org/2009/105`.

[24] M. Naor and M. Yung. Universal one-way hash functions and their cryptographic applications. In *21st Annual ACM Symposium on Theory of Computing (STOC)*, pages 33–43. ACM Press, 1989.

[25] P. Q. Nguyen and I. Shparlinski. The insecurity of the digital signature algorithm with partially known nonces. *Journal of Cryptology*, 15(3):151–176, 2002.

[26] T. Okamoto. Provably secure and practical identification schemes and corresponding signature schemes. In *Advances in Cryptology — Crypto '92*, volume 740 of *LNCS*, pages 31–53. Springer, 1993.

[27] H. Ong and C.-P. Schnorr. Fast signature generation with a Fiat-Shamir-like scheme. In *Advances in Cryptology — Eurocrypt '90*, volume 473 of *LNCS*, pages 432–440. Springer, 1990.

[28] K. Pietrzak. A leakage-resilient mode of operation. In *Advances in Cryptology — Eurocrypt 2009*, LNCS, pages ???–??? Springer, 2009.

[29] A. Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *40th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 543–553. IEEE Computer Society Press, 1999.

[30] C.-P. Schnorr. Efficient identification and signatures for smart cards. In *Advances in Cryptology — Crypto '89*, volume 435 of *LNCS*, pages 239–252. Springer, 1990.

[31] F.-X. Standaert, T. Malkin, and M. Yung. A unified framework for the analysis of side-channel key recovery attacks. In *Advances in Cryptology — Eurocrypt 2009*, LNCS, pages 443–461. Springer, 2009.

[32] Y. Tauman and V. Vaikuntanathan. Public-key encryption schemes with auxiliary inputs and applications. Manuscript, 2009.

A Fully Leakage-Resilient Signatures in the Random Oracle Model

For completeness, we describe a construction of a fully leakage-resilient signature scheme in the random oracle model that was discovered independently by [2]. Although the schemes are identical in both works, we note two differences in the analysis: first, we make explicit the fact that the leakage can depend on all the state information of the signer; second, we allow leakage queries to depend on the random oracle.

For concreteness, we exemplify the approach using (a variant of) the Okamoto-Schnorr signature scheme [30, 26], thus basing security on the discrete logarithm assumption. The same ideas can also be applied to obtain schemes based on the RSA [15, 26] or factoring [27, 13] assumptions (but see the remark at the end of this section).

Let $\mathcal{G}$ be a group generation algorithm that on input 1^k outputs $(\mathbb{G}, q)$ where q is a k-bit prime and $\mathbb{G}$ is a cyclic group of order q. We assume group operations in $\mathbb{G}$ are efficient, and that a random element of $\mathbb{G}$ can be sampled "obliviously": i.e., there is an efficient algorithm $\mathsf{samp}_{\mathbb{G}}$ generating random elements of $\mathbb{G}$ with the property that, given $g \in \mathbb{G}$, one can sample uniformly from the set of coins ω for which $g := \mathsf{samp}_{\mathbb{G}}(\omega)$.[5] (We leave $\mathsf{samp}_{\mathbb{G}}$ implicit from now on.) This implies in particular that (assuming the discrete logarithm problem is hard in the first place) it is possible to generate a random element of $\mathbb{G}$ without learning its discrete logarithm.

Our signature scheme is parameterized by an integer ℓ, and is defined as follows:

Key generation: Compute $(\mathbb{G}, q) \leftarrow \mathbb{G}(1^k)$ and choose random $g_1, \ldots, g_\ell \leftarrow \mathbb{G}$ and random $x_1, \ldots, x_\ell \leftarrow \mathbb{Z}_q$. Set $h := \prod_i g_i^{x_i}$. The public key is $(\mathbb{G}, q, g_1, \ldots, g_\ell, h)$ and the secret key is $(x_1, \ldots, x_\ell)$. We also assume a random oracle $H : \{0,1\}^* \rightarrow \mathbb{Z}_q$.

Signing: Signatures are perfectly witness-indistinguishable proofs of knowledge of a representation of h with respect to the basis $(g_1, \ldots, g_\ell)$, turned into signatures using the Fiat-Shamir transformation. In detail: to sign message m using secret key $(x_1, \ldots, x_\ell)$ and public key $(\mathbb{G}, q, g_1, \ldots, g_\ell, h)$, the signer chooses random $r_1, \ldots, r_\ell \leftarrow \mathbb{Z}_q$ and computes $A := \prod_i g_i^{r_i}$. It then sets $c := H(A, m)$, and outputs the signature $(A, cx_1 + r_1, \ldots, cx_\ell + r_\ell)$.

Verification: Given a signature $(A, \alpha_1, \ldots, \alpha_\ell)$ on the message m with respect to the public key $(\mathbb{G}, q, g_1, \ldots, g_\ell, h)$, compute $c := H(A, m)$ and output 1 iff $\prod_i g_i^{\alpha_i} \overset{?}{=} h^c \cdot A$.

Before proving security, we formally state the discrete logarithm assumption as well as an equivalent assumption we will use in our proof.

Definition 4 Let $\mathcal{G}$ be as above. We say the *discrete logarithm problem is hard for* $\mathcal{G}$ if the following is negligible for all PPT algorithms $\mathcal{A}$:

$$\Pr\left[(\mathbb{G}, q) \leftarrow \mathcal{G}(1^k); g, h \leftarrow \mathbb{G}; x \leftarrow \mathcal{A}(\mathbb{G}, q, g, h, \omega) : g^x = h\right],$$

where ω denotes the randomness used to generate $\mathbb{G}, q, g, h$. (Note that by our assumption on $\mathsf{samp}_{\mathbb{G}}$, the coins used to generate g, h are extraneous.) $\diamondsuit$

We stress that the above definition requires hardness to hold even against adversaries given the randomness used to generate the problem instance. (For concrete $\mathcal{G}$ used in practice, this additional randomness does not seem to make the problem any easier.)

Definition 5 Let $\mathcal{G}$ be as above. We say the *ℓ-representation problem is hard for* $\mathcal{G}$ if the following is negligible for all PPT algorithms $\mathcal{A}$:

$$\Pr\left[\begin{array}{c}(\mathbb{G}, q) \leftarrow \mathcal{G}(1^k); g_1, \ldots, g_\ell \leftarrow \mathbb{G}; \\ (x_1, \ldots, x_\ell), (x'_1, \ldots, x'_\ell) \leftarrow \mathcal{A}(\mathbb{G}, q, g_1, \ldots, g_\ell, \omega)\end{array} : \prod_i g_i^{x_i} = \prod_i g_i^{x'_i} \bigwedge \vec{x} \neq \vec{x}'\right],$$

where ω denotes the randomness used to generate $\mathbb{G}, q, g_1, \ldots, g_\ell$. $\diamondsuit$

Hardness of the discrete logarithm problem for $\mathcal{G}$ implies hardness of the ℓ-representation problem [7, 3] (for any polynomial ℓ). This implication carries over to our setting (where the randomness

[5] This property holds for concrete examples $\mathcal{G}$ used in practice.

used to generate the instance is given to the adversary) given our assumption that elements in $\mathbb{G}$ can be sampled without knowledge of their.

Theorem 4 *If the discrete logarithm problem is hard for $\mathcal{G}$, then for any $\epsilon > 0$ the signature scheme above is $\left(\frac{1}{2} - \frac{1}{2\ell} - \epsilon\right) \cdot n$-leakage resilient, where n denotes the length of the secret key.*

Thus, for any desired $1 > \epsilon' > 0$, setting $\ell > 1/\epsilon'$ (and taking $\epsilon = \epsilon'/2$) gives a scheme that is $(\frac{1}{2} - \epsilon') \cdot n$-leakage resilient.

Proof Let $\mathcal{A}$ be a probabilistic polynomial-time adversary attacking the scheme in the sense of Definition 2. We let q_H be a bound on the total number of hash queries asked throughout the entire experiment; such queries can be made directly by $\mathcal{A}$ and can also occur in the course of answering signing or leakage queries. Denote the success probability of $\mathcal{A}$ by $\delta = \delta(k)$.

We make a number of assumptions about $\mathcal{A}$ without loss of generality. First, we assume that if $\mathcal{A}$ outputs $(m, \sigma = (A, \alpha_1, \ldots, \alpha_\ell))$ then (1) $\mathcal{A}$ at some point queried $H(A, m)$ and (2) $\mathcal{A}$ never requested a signature on m. Second, for any leakage query $\mathsf{Leak}(f_i)$ we assume $f_i(\mathsf{state})$ makes the same number of H-oracle calls regardless of the value of state (this can always be ensured by adding dummy queries, as needed).

We construct a probabilistic polynomial-time algorithm $\mathcal{B}$ solving the ℓ-representation problem. Algorithm $\mathcal{B}$ proceeds as follows: on input $(\mathbb{G}, q, g_1, \ldots, g_\ell, \omega)$, it chooses random $x_1, \ldots, x_\ell$ and computes $h := \prod_i g_i^{x_i}$. It gives the public key $pk = (\mathbb{G}, q, g_1, \ldots, g_\ell, h)$ to $\mathcal{A}$ and runs the entire experiment of Definition 2, simulating the random oracle for $\mathcal{A}$. Note that $\mathcal{B}$ can run the entire experiment easily since it knows a legitimate secret key $(x_1, \ldots, x_\ell)$ corresponding to pk.

When $\mathcal{A}$ terminates, $\mathcal{B}$ examines $\mathcal{A}$'s output $(m, \sigma = (A, \alpha_1, \ldots, \alpha_\ell))$; we call the execution of $\mathcal{A}$ to this point the *first run* of $\mathcal{A}$. If $\mathsf{Vrfy}_{pk}(m, \sigma) = 1$, then $\mathcal{B}$ rewinds to the point in the experiment where the hash query $H(A, m) = c$ was first made; note that this may occur either as a result of a direct H-query made by $\mathcal{A}$, or during the course of answering a Leak query. $\mathcal{B}$ then chooses a fresh random value c' for the result of $H(A, m)$, and re-runs the experiment from that point; we will refer to the execution of $\mathcal{A}$ from this point to termination as the *second run* of $\mathcal{A}$. During the second run of $\mathcal{A}$, signing queries are answered using fresh randomness but H-oracle queries are answered consistently with the first run of $\mathcal{A}$. (I.e., if a hash query is made during the second run of $\mathcal{A}$ that was also made during the first run of $\mathcal{A}$, the same answer chosen during the first run is used in the second run. Any new hash queries are answered using fresh randomness.) We say $\mathcal{B}$ *succeeds* if, in this second run, $\mathcal{A}$ terminates with output $(m, \sigma' = (A, \alpha'_1, \ldots, \alpha'_\ell))$ where σ' is also a valid signature on m.

If $\mathcal{B}$ succeeds and also $c' \neq c$, then $\mathcal{B}$ computes (in the standard way [26]) values $(x'_1, \ldots, x'_\ell)$ such that $\prod_i g_i^{x'_i} = h$. By definition, $\mathcal{B}$ solves the representation problem if $(x'_1, \ldots, x'_\ell) \neq (x_1, \ldots, x_\ell)$. The following two claims complete the proof of the theorem.

Claim 1 $\Pr[\mathcal{B} \text{ succeeds}] \geq \delta^2/q_H$.

Proof Consider all possible states during the execution of $\mathcal{A}$ in the experiment of Definition 2 where a new hash query is made, and for any such state i let h_i denote the hash query made at that state. If an execution of $\mathcal{A}$ terminates with a valid forgery $(m, \sigma = (A, \alpha_1, \ldots, \alpha_\ell))$, say that state i is *associated with* the forgery if $h_i = (m, A)$.

For any state i where a new hash query is made, let a_i be the probability that this state is reached in the experiment of Definition 2, where this probability is over the entire specification of the random oracle H (*except* for the value of $H(h_i)$, which does not affect the probability since the query h_i is being made for the first time in state i) as well as the randomness of $\mathcal{A}$, the randomness used to generate the public key, and the randomness used to answer any signing queries made up to that point. For the same state i, let b_i be the probability that, starting from state i, the execution of $\mathcal{A}$ terminates with a successful forgery associated with i. This probability is over the value of $H(h_i)$ and the randomness used to answer any signing queries made after this point. Since every successful forgery is associated with a unique state of $\mathcal{A}$, we have $\sum_i a_i \cdot b_i = \delta$, the overall success probability of $\mathcal{A}$. Furthermore, we have $\sum_i a_i = \mathbf{E}[\text{number of hash queries made by } \mathcal{A}] \leq q_h$.

By construction of $\mathcal{B}$, we have $\Pr[\mathcal{B} \text{ succeeds}] = \sum_i a_i \cdot (b_i)^2$. Using Jensen's inequality:

$$\begin{aligned} \sum_i a_i \cdot (b_i)^2 &\geq \sum_i a_i \cdot \left(\frac{\sum_i a_i \cdot b_i}{\sum_i a_i} \right)^2 \\ &\geq \frac{\delta^2}{\sum_i a_i} \geq \delta^2 / q_H \,, \end{aligned}$$

completing the proof of the claim. □

Claim 2 *The probability that $\mathcal{B}$ solves the representation problem is at least*

$$\frac{1}{2} \cdot \left(\Pr[\mathcal{B} \text{ succeeds}] - 1/q - q_H / q^{2\epsilon\ell} \right).$$

Proof When $\mathcal{B}$ succeeds, there are two bad events that can prevent $\mathcal{B}$ from solving the representation problem. First, it may be the case that $c' = c$; this happens with probability $1/q$. Second, it may be the case that the extracted representation $\vec{x}' = (x'_1, \ldots, x'_\ell)$ is equal to the original representation $\vec{x} = (x_1, \ldots, x_\ell)$. We show that except with probability at most $q_H / q^{2\epsilon\ell}$, the min-entropy of $\vec{x}$ conditioned on the view of $\mathcal{A}$ (in both its runs) is greater than 0; given this, the probability that $\vec{x}' \neq \vec{x}$ is at least $1/2$ and the claim follows.

Let $\lambda = \left(\frac{1}{2} - \frac{1}{2\ell} - \epsilon\right) \cdot \ell \cdot \log q$, an upper bound on the number of leaked bits in each run of $\mathcal{A}$. The public key pk constrains $\vec{x}$ to lie in an $(\ell - 1)$-dimensional vector space, and it is well-known [26] that signature queries do not further constrain $\vec{x}$. Thus, the min-entropy of $\vec{x}$ conditioned on the public key and the observed signatures is $(\ell - 1) \cdot \log q$ bits. The views of $\mathcal{A}$ in its two runs contain only the following additional information about $\vec{x}$: at most $2 \cdot \lambda$ bits from the leakage functions (i.e., λ bits in each view), and $\log q_H$ bits indicating the relevant state associated with the first forgery (cf. the proof of Claim 1).[6] Applying Lemma 1, we see that the conditional min-entropy of $\vec{x}$ is greater than 0 except with probability at most

$$2^{2\lambda + \log q_h - (\ell - 1) \cdot \log q} \leq q_H \cdot q^{-2\epsilon\ell}.$$

The claim follows. □

Taking the two claims together, we see that if $\mathcal{A}$ succeeds with probability δ then $\mathcal{B}$ solves the representation problem with probability at least

$$\frac{1}{2} \cdot \left(\delta^2 / q_H - 1/q - q_H / q^{2\epsilon\ell} \right).$$

[6] Note that this information may not be evident from $\mathcal{A}$'s view. For example, consider the leakage query f_1 defined as follows: "if $x_1 = 0$ then query $H(A, m)$ (otherwise do nothing); in any case, return the first bit of x_2." If $\mathcal{A}$ later queries $H(A, m)$ and eventually outputs a forgery $(A, \ldots)$ on m, then the state associated with this forgery depends on $\vec{x}$ and cannot be determined solely from the view of $\mathcal{A}$.

Since q_H is polynomial, $\epsilon > 0$ is a constant, and $1/q$ is negligible, the above is non-negligible whenever δ is. This completes the proof of the theorem. ■

It is an interesting open question to improve the tightness of the security reduction.

Remark: instantiating this approach using factoring-based assumptions. As noted at the beginning of this section, the same approach as above can also be applied to obtain leakage-resilient schemes based on the RSA [15, 26] or factoring [27, 13] assumptions. In these cases, however, we must either (1) restrict the leakage functions to apply only to those values used by the signer *after* the key generation phase (i.e., the secret key and all state variables used to sign, but not the random coins used to generate the secret/public key), or (2) generate the modulus N via oblivious sampling, leading to a significant loss in efficiency.