

Efficient Public-Key Cryptography in the Presence of Key Leakage

Yevgeniy Dodis, Kristiyan Haralambiev, Adriana López-Alt, and Daniel Wichs

New York University

dodis@cs.nyu.edu, kkh@cs.nyu.edu, lopez@cs.nyu.edu, wicks@cs.nyu.edu

Abstract. We study the design of cryptographic primitives resistant to a large class of side-channel attacks, called “memory attacks”, where an attacker can repeatedly and adaptively learn information about the secret key, subject *only* to the constraint that the *overall amount* of such information is bounded by some parameter ℓ . Although the study of such primitives was initiated only recently by Akavia et al. [2], subsequent work already produced many such “leakage-resilient” primitives [48,4,42], including signature, encryption, identification (ID) and authenticated key agreement (AKA) schemes. Unfortunately, every existing scheme, — for any of the four fundamental primitives above, — fails to satisfy at least one of the following desirable properties:

- **Efficiency.** While the construction may be generic, it should have some *efficient* instantiations, based on standard cryptographic assumptions, and without relying on random oracles.
- **Strong Security.** The construction should satisfy the strongest possible definition of security (even in the presence of leakage). For example, encryption schemes should be secure against chosen *ciphertext* attack (CCA), while signatures should be *existentially* unforgeable.
- **Leakage Flexibility.** It should be possible to set the scheme parameters so that the leakage bound ℓ can come arbitrarily close to the secret-key size.

In this work we design the first signature, encryption, ID and AKA schemes which overcome these limitations, and satisfy all the properties above. Moreover, all our constructions are generic, in several cases elegantly simplifying and generalizing the prior constructions (which did not have any efficient instantiations). We also introduce several tools of independent interest, such as the abstraction (and constructions) of *true-simulation extractable* NIZK arguments, and a new *deniable* DH-based AKA protocol based on any CCA-secure encryption.

1 Introduction

Traditionally, the security of cryptographic schemes has been analyzed in an idealized setting, where an adversary only sees the specified “input/output behavior” of a scheme, but has no other access to its internal secret state. Unfortunately, in the real world, an adversary may often learn some partial information about secret state via various *key leakage* attacks. Such attacks come in a large variety and include *side-channel attacks* [43,10,7,44,54,27], where the physical realization of a cryptographic primitive can leak additional information, such as the computation-time, power-consumption, radiation/noise/heat emission etc. The cold-boot attack of Halderman et al. [34] is another example of a key-leakage attack, where an adversary can learn (imperfect) information

about memory contents of a machine, even after the machine is powered down. Schemes that are proven secure in an idealized setting, without key leakage, may become completely insecure if the adversary learns even a small amount of information about the secret key. Indeed, even very limited leakage attacks have been shown to have devastating consequences for the security of many natural schemes.

Unfortunately, it is unrealistic to assume that we can foresee, let alone block, all of the possible means through which key leakage can occur in real-world implementations of cryptographic schemes. Therefore, the cryptographic community has recently initiated the investigation of increasingly general (formally modeled) classes of leakage attacks, with the aim of constructing *leakage-resilient* cryptographic schemes that remain provably secure even in the presence of such attacks. Of course, if an adversary can get unrestricted information about the secret key, then she can learn the key in its entirety and the security of the system is necessarily compromised. Therefore, we must first place some “upper bound” on the type or amount of information that the adversary can learn. The nature of such bounds varies in the literature, as we survey later. For this work, we only restrict the *amount*, but not the *type*, of information that an adversary can learn through a key-leakage attack. In particular, we will assume that the attacker can learn *any efficiently computable function of the secret key sk* , subject only to the constraint that the total amount of information learned (i.e. the output size of the leakage function) is bounded by ℓ bits, where ℓ is called the “leakage parameter” of the system.¹ Clearly, at this level of generality, the secret-key size s must be strictly greater than the leakage-parameter ℓ .² Therefore, the quantity ℓ/s can be thought as the *relative leakage* of the system, with the obvious goal to make it as close to 1 as possible.

Our model of leakage-resilience was recently introduced by Akavia et al. [2], but already attracted a lot of attention from the cryptographic community [48,4,42,3]. In particular, as we survey later, we already know many “leakage-resilient” primitives, including such fundamental primitives as signature schemes, encryption schemes, identification (ID) schemes and authenticated key agreement (AKA) protocols. Unfortunately, we observe that every existing scheme, — for any of the four fundamental primitives above, — fails to satisfy at least one of the following desirable properties:

- **Efficiency.** While the proposed construction may be based on some generic cryptographic primitives, — which is in fact preferable for modular design, — it should have some *efficient* instantiations, based on standard cryptographic assumptions, and without relying on random oracles. We view this property as the main property we will strive to achieve.
- **Strong Security.** The construction should satisfy the strongest possible definition of security (even in the presence of leakage). For example, encryption schemes should be secure against chosen *ciphertext* attack (CCA), while signatures should be *existentially* unforgeable, etc.

¹ More formally, we allow adaptive measurements, as long as the sum of leaked outputs is bounded by ℓ .

² In fact, our actual constructions easily extend to the more general “noisy leakage” model of Naor and Segev [48], where the outputs can be longer than s , as long as the “average min-entropy” of sk drops by at most ℓ bits. However, we do not pursue this generalization, in order to keep our notation simple.

- **Leakage Flexibility.** It should be possible to set the parameters of the schemes so that the relative leakage ℓ/s is arbitrarily close to 1. We call such schemes *leakage-flexible*.

1.1 Our Results

In this work we design the first signature, encryption, ID and AKA schemes which simultaneously satisfy the efficiency, strong security and leakage flexibility properties mentioned above. Moreover, all our constructions are generic. This means that the actual construction is modularly defined and explained using natural simpler blocks, and its security against key leakage is also proven no matter how these simpler blocks are (securely) implemented. However, unlike the prior generic constructions, which did not have any known efficient instantiations (at least, with the desired security and flexibility we seek), ours are yet more general, which will allow us to obtain several efficient instantiations. Given this fact, it is not surprising that our contributions can be roughly split into two categories: “conceptual” contributions, allowing us to obtain more general (and, yet, conceptually simpler) leakage-resilient constructions, and “concrete” contributions, allowing us to instantiate our general schemes efficiently.

CONCEPTUAL CONTRIBUTIONS. As we will see, existing schemes (e.g., signature and CCA-encryption) could be largely divided into two categories: potentially efficient schemes, with some *inherent* limitation not allowing them to achieve relative leakage approaching 1 (which also prevents us from using these ideas for our purposes), and more theoretical schemes [48,42], achieving good relative leakage, but relying on the notion of *simulation-sound* non-interactive zero-knowledge (ss-NIZK) [56]. Informally, ss-NIZK proofs remain sound even if the attacker can see simulated proofs of arbitrary (even false) statements. Unfortunately, it appears that the existing cryptographic machinery does not allow us to instantiate non-trivial ss-NIZK proofs efficiently.³ On the other hand, a recent breakthrough result of Groth-Sahai [33] showed that one can obtain efficient *non-simulation-sound* NIZK proofs for a non-trivial class of languages. While the techniques of [31] could be applied to Groth-Sahai proofs to achieve ss-NIZKs, it is a non-trivial “exercise” and the resulting proofs are *significantly* less efficient, as the construction involves OR-proofs for Groth-Sahai languages. Therefore, our first idea was to try to generalize the existing constructions sufficiently, making them rely only on regular NIZKs, in the hope that such NIZKs can then be instantiated using the powerful Groth-Sahai techniques.

In the end, this is indeed what we realized. However, in the process we also abstracted away an elegant notion of independent interest: *true-simulation extractable* (tSE) NIZKs. While similar to the notion of simulation-sound extractable NIZKs [31], it involves a subtle but rather important difference: whether the adversary has oracle access to simulated proofs for arbitrary (even false) statements or only true ones. Intuitively, both the Naor-Segev’s leakage-resilient CCA encryption [48] and Katz-Vaikuntanathan’s leakage-resilient signature scheme [42] used the technique of encrypting a witness x for some relation R , and then providing a ss-NIZK proof φ that the

³ The work of [31] constructs ss-NIZK proofs for practical languages and uses them to construct group signatures, but the resulting scheme has signature size of “thousands or perhaps even millions of group elements” [32] despite being constant.

Reference	Unforgeability	Model	Leakage	Efficient?
[4]	Existential	<i>Random Oracle</i>	<i>1/2</i>	Yes
[4]	<i>Entropic</i>	<i>Random Oracle</i>	1	Yes
[42]	Existential	Standard	1	<i>No</i>
This Work	Existential	Standard	1	Yes

Table 1. Previous work on leakage-resilient signatures and results of this work

ciphertext c indeed contains the encryption of a valid witness x . The main reason for using this technique is to allow the reduction to extract a valid witness from any “new” valid pair (c^*, φ^*) produced by the attacker \mathcal{A} (who saw many such valid pairs earlier). In this paper, we will abstract this property into the tSE notion mentioned above (of which the above mentioned technique is a specific example, where the pair (c, φ) together makes up a single tSE-NIZK proof). Moreover, we show that true-simulation extractability, as we abstract it, is *precisely* the right notion for generalizing and proving the security of the previous constructions. This has two positive effects. First, it makes the generic constructions of CCA-encryption and signatures somewhat more intuitive, both for proving and understanding. For example, the traditional “double-encryption” paradigm of Naor-Yung [49] for designing CCA-secure schemes from chosen-plaintext secure (CPA-secure) schemes, also used by [48] in the context of key leakage, can be stated as “CPA-encrypting message m under two keys and proving plaintext equality”. Using our more general “simulation-extractability view”, it is now stated as “CPA-encrypting m and proving that one knows the plaintext”. We believe that the latter view is not only more general, but also more intuitive as a way of explaining “CPA-to-CCA” transformation. It also follows the original intuition of Rackoff and Simon [55], who combine CPA-encryption with NIZK-POK to achieve CCA-encryption, but in the model where the sender also has a secret key. A similar discussion is true for our signature constructions.

Second, we show a generic way to build tSE-NIZKs which *avoids using (expensive) ss-NIZKs*. Instead, our method uses *regular* NIZKs and *any* CCA-secure encryption scheme.⁴ Perhaps surprisingly, given the current state-of-the-art NIZK and CCA schemes, the combination “CCA + NIZK” appears to be much more efficient in practice than the combination “CPA + ss-NIZK”.⁵ As a result, we were able to provide a general framework for building leakage-flexible signature and CCA-encryption schemes, eventually allowing us to efficiently instantiate our schemes (by avoiding using ss-NIZKs). We summarize our results for signature and CCA-encryption schemes in Tables 1 and 2, also comparing them to the best prior constructions. In all the tables, the “sub-optimal” entries (for efficiency, security, model or relative leakage of prior constructions) are written in italics, and most prior rows are also explained in the related work Section 1.2. For signatures, we stress that no efficient construction in the standard model was known prior to our work, for any non-trivial relative leakage fraction (let alone 1).

Once we have efficient leakage-flexible signature schemes, we can obtain ID and AKA schemes with the same properties. The signature-based AKA protocol is not de-

⁴ This is OK for the signature application, but might appear strange for our CCA-encryption application, as we need “CCA to get CCA”. However, as a building block for tSE-NIZKs, we only need *standard* CCA schemes and as a result obtain *leakage-resilient* CCA schemes.

⁵ Indirectly, the same realization was made by Groth [32] and Camenisch et al. [11].

Reference	Attack	Model	Leakage	Efficient?
[2,48]	CPA	Standard	1	Yes
[48]	CCA	Standard	$1/6$	Yes
[48]	CCA	Standard	1	No
This Work	CCA	Standard	1	Yes

Table 2. Previous work on leakage-resilient encryption and results of this work

Reference	Security	Model	Leakage	Efficient?
[4]	<i>Pre-Impersonation</i>	Standard	1	Yes
[4]	Anytime	Standard	$1/2$	Yes
[42] (implicit)	Anytime	Standard	1	No
This Work	Anytime	Standard	1	Yes

Table 3. Previous work on leakage-resilient identification schemes and results of this work

niable. However, we also construct a *deniable* AKA protocol based on our construction of leakage-flexible CCA-secure encryption. We summarize our results for ID schemes in Table 3 and for AKA protocols in Table 4. See Section 6 for details.

CONCRETE CONTRIBUTIONS. As we explained above, we generically reduce the question of building efficient leakage-flexible ID schemes and AKA protocol to the question of efficiently instantiating our leakage-flexible signature and/or encryption schemes. Such instantiations are given in Section 5. We also explained how the latter instantiations became possible in our work, since we gave generic constructions of both primitives based on the new notion of tSE-NIZK, and then showed that satisfying this notion may be possible using *ordinary* NIZKs for appropriate languages, without relying on the expensive simulation-sound NIZKs. Unfortunately, efficient construction of (even ordinary) NIZKs, due to Groth and Sahai [33], are only known for a pretty restrictive class or languages in bilinear groups. Thus, obtaining a *concrete* efficient instantiation still requires quite a substantial effort.

Specifically, all the building blocks have to be instantiated efficiently, and expressed in a form such that the resulting NP relation satisfies the severe limitations imposed by the Groth-Sahai NIZKs. For example, to build leakage-resilient CCA-encryption, we need to have an efficient leakage-flexible CPA scheme, a CCA scheme supporting labels and a one-time signature scheme, all connected together by an efficient NIZK for a complicated “plaintext equality” relation. Similarly, for leakage-resilient signature schemes, we need an efficient second-preimage resistant (SPR; see Definition 1) relation and a CCA scheme supporting labels, once again connected by an efficient NIZK for a complex relation. Not surprisingly, such tasks cannot typically be done by simply combining “off-the-shelf” schemes from the literature. At best, it requires very careful selection of parameters to make everything “match”, followed by a round of further efficiency optimizations. Usually, though, it requires the design of new primitives, which work well with other known primitives, to enable efficient NIZK. For example, in this work, we designed two new SPR relations (see Section 5), since prior SPR relations did not appear to mesh well with our CCA encryption scheme. To emphasize the importance of the new SPR relations, we point out that combining previous constructions with Groth-Sahai proofs would require committing to the witness bit-by-bit in order to achieve full extractability.

Reference	Model	Leakage	Deniable?	Efficient?
[4]	<i>Random Oracle</i>	1	<i>No</i>	Yes
[4,42]	Standard	1	<i>No</i>	<i>No</i>
This Work	Standard	1	<i>No/Yes*</i>	Yes

* Our first AKA protocol is not deniable; our second — is.

Table 4. Previous work on leakage-resilient AKA and results of this work.

Overall, we get two different efficient instantiations of both leakage-resilient signature and CCA encryption schemes in the standard model, based on standard (static and “fixed-length”) assumptions in bilinear groups, called external Diffie-Hellman (SXDH) and Decision Linear (DLIN). The high-level idea of these schemes, as well as their efficiency, is described in Section 5. The actual low-level details of how to put “everything together” in the most efficient manner, is described in the full version [18].

1.2 Related Work

LEAKAGE-RESILIENCE AND MEMORY ATTACKS. Our model of leakage, sometimes called memory-attacks, was first proposed by Akavia et al. [2], who also constructed CPA secure PKE and IBE schemes in this model under the *learning with errors (LWE)* assumption. Later Naor and Segev [48] generalized the main ideas behind these constructions to show that all schemes based on *hash proof systems* (see [15]) are leakage-resilient. In particular, this resulted in efficient constructions based on the DDH and K -Linear assumptions, where the relative leakage on the secret key could be made to approach 1. Moreover, [48] showed how to also achieve CCA security in this model by either: (1) relying on the generic (and inefficient) Naor-Yung paradigm where the leakage-rate can be made to approach 1 or (2) using efficient hash proof systems with leakage-rate only approaching $1/6$. Unfortunately, it seems that the hash proof system approach to building CCA encryption is inherently limited to leakage-rates below $1/2$: this is because the secret-key consists of two components (one for verifying that the ciphertext is well-formed and one for decrypting it) and the proofs break down if either of the components is individually leaked in its entirety. The work of [3] generalizes [48] still further by showing how to construct leakage-resilient IBE schemes generically based on *identity-based hash proof systems*, with several instantiations.

Leakage-resilient signature schemes in the model of memory attacks were constructed in the random-oracle model by [4,42], and in the standard model by [42]. The random-oracle schemes are highly-efficient but suffer from two limitations. Firstly they rely on the Fiat-Shamir [25] transform which is only known to be secure in the Random Oracle model and is not sound in general [30]. Secondly, the schemes can only tolerate leakage which approaches $1/2$ of the secret key. On the other hand, the standard-model schemes allow for relative-leakage approaching 1, but are based on generic simulation-sound NIZKs and do not come with an efficient instantiation.

The work of [4] also constructs ID schemes and AKA protocols. For ID schemes, two notions of security were considered: a weaker notion called pre-impersonation leakage-resilience and a stronger notion called anytime leakage-resilience. Although efficient schemes in the standard model were given for both notions, the leakage resilience could be made to approach 1 only for pre-impersonation leakage while, for anytime leakage, the given schemes can only tolerate a leakage-rate below $1/2$. For AKA schemes, a construction was given based on leakage-resilient signatures (only re-

quiring a weakened notion of security called entropic-unforgeability). Using the appropriate signature schemes, this yielded two types of constructions: efficient constructions in the random-oracle model and generic but inefficient constructions in the standard model (both of which have leakage-rates approaching 1).

OTHER MODELS OF LEAKAGE-RESILIENCE. Several other models of leakage-resilience have appeared in the literature. They differ from the model we described in that they restrict the *type*, as well as *amount*, of information that the adversary can learn. For example, *exposure resilient cryptography* [12,20,41] studies the case where an adversary can only learn some small *subset of the physical bits of the secret key*. Similarly, [38] studies how to implement arbitrary computation in the setting where an adversary can observe a small *subset of the physical wires of a circuit*. Most recently, [24] study a similar problem, where the adversary can observe a low-complexity (e.g. AC^0) function of the wires. Unfortunately, these models fail to capture many meaningful side-channel attacks, such as learning the hamming-weight of the bits or their parity.

In their seminal work, Micali and Reyzin [46] initiated the formal modeling of side-channel attacks under the axiom that “*only computation leaks information*” (OCLI), where each invocation of a cryptographic primitive leaks a function of *only* the bits accessed during that invocation. Several primitives have been constructed in this setting including stream ciphers [22,53] and signatures [23]. More recently, [40] construct a general compiler that can secure *all primitives* in this setting assuming the use of some limited leak-free components and the existence of fully homomorphic encryption. On the positive side, the OCLI model only imposes a bound on the amount of information learned during each invocation of a primitive, but not on the overall amount of information that the attacker can get throughout the lifetime of the system. On the negative side, this model fails to capture many leakage-attacks, such as the cold-boot attack of [34], where *all* memory contents leak information, even if they were never accessed.

Lastly, we mention models of leakage-resilience which are strictly stronger than the memory-attacks model. Firstly, the Bounded-Retrieval Model [16,21,4,3] imposes an additional requirement on leakage-resilient schemes, by insisting that they provide a way to “grow” the secret-key (possibly to many Gigabytes) so as to proportionally increase the amount of tolerated leakage, but without increasing the size of the public-key, the computational or communication efficiency of the scheme, or the lengths of the ciphertexts or signatures. The work of [4] constructs “entropic” signatures, ID schemes and AKA protocols in this setting, while the work of [3] constructs PKE and IBE schemes in this model. A different strengthening is the auxiliary input model [19,17] where the leakage is not necessarily bounded in length, but it is (only) assumed to be computationally hard to recover the secret-key from the leakage. The work of [19] constructs symmetric-key encryption in this model, under a strengthening of the learning parity with noise (LPN) assumption, while [17] constructs public-key encryption under the DDH and LWE assumptions. Yet another strengthening of the memory-attacks model, proposed by [29], is to require that there is a single scheme (parameterized only by the security parameter) which can tolerate essentially any amount of relative-leakage where the exact-security of the scheme degrades smoothly as the relative-leakage increases. In this model, [29] construct a symmetric-key encryption scheme.

2 Definitions of Leakage-Resilient Primitives

We model leakage attacks by giving the adversary access to a *leakage oracle*, which he can adaptively access to learn leakage on the secret key. A leakage oracle $\mathcal{O}_{sk}^{\lambda, \ell}(\cdot)$ is parametrized by a secret key sk , a leakage parameter ℓ , and a security parameter λ . A query to the leakage oracle consists of a function $h_i : \{0, 1\}^* \rightarrow \{0, 1\}^{\alpha_i}$, to which the oracle answers with $y_i = h_i(sk)$. We only require that the functions h_i be efficiently computable, and the total number of bits leaked is $\sum_i \alpha_i \leq \ell$.

Definition 1 (Leakage Resilient Hard Relation). A relation R with a randomized PPT sampling algorithm KeyGen is an ℓ -leakage resilient hard relation if:

- For any $(sk, pk) \leftarrow \text{KeyGen}(1^\lambda)$, we have $(sk, pk) \in R$.
- There is a poly-time algorithm that decides if $(sk, pk) \in R$.
- For all PPT adversaries $\mathcal{A}^{\mathcal{O}_{sk}^{\lambda, \ell}(\cdot)}$ with access to the leakage oracle $\mathcal{O}_{sk}^{\lambda, \ell}(\cdot)$:

$$\Pr \left[R(sk^*, pk) = 1 \mid (pk, sk) \leftarrow \text{KeyGen}(1^\lambda), sk^* \leftarrow \mathcal{A}^{\mathcal{O}_{sk}^{\lambda, \ell}(\cdot)}(pk) \right] \leq \text{negl}(\lambda)$$

Notice that without loss of generality, we can assume that \mathcal{A} queries $\mathcal{O}_{sk}^{\lambda, \ell}(\cdot)$ only once with a function h whose output is ℓ bits.

Definition 2 (Leakage Resilient Signatures). A signature scheme $\mathcal{S} = (\text{KeyGen}, \text{Sign}, \text{SigVer})$ is ℓ -leakage resilient if \forall PPT \mathcal{A} we have $\Pr[\mathcal{A} \text{ wins}] \leq \text{negl}(\lambda)$ in the following game:

1. **Key Generation:** The challenger runs $(vk, sk) \leftarrow \text{KeyGen}(1^\lambda)$ and gives vk to \mathcal{A} .
2. **Signing and leakage queries:** $\mathcal{A}^{\mathcal{O}_{sk}^{\lambda, \ell}(\cdot), \mathcal{S}_{sk}(\cdot)}$ is given access to the leakage oracle $\mathcal{O}_{sk}^{\lambda, \ell}(\cdot)$ and the signing oracle $\mathcal{S}_{sk}(\cdot)$. A query to the signing oracle $\mathcal{S}_{sk}(\cdot)$ consists of a message m , to which the oracle responds with $\sigma = \text{Sign}_{sk}(m)$.
3. \mathcal{A} outputs (m^*, σ^*) and wins if $\text{SigVer}_{vk}(m^*, \sigma^*) = 1$ and m^* was not given to $\mathcal{S}_{sk}(\cdot)$ as a signing query.

Definition 3 (Leakage Resilient CCA-Secure Encryption). We say that an encryption scheme $\mathcal{E} = (\text{KeyGen}, \text{Enc}, \text{Dec})$ is ℓ -leakage resilient CCA-secure if \forall PPT \mathcal{A} we have $\Pr[\mathcal{A} \text{ wins}] \leq \frac{1}{2} + \text{negl}(\lambda)$ in the following game:

1. **Key Generation:** The challenger runs $(pk, sk) \leftarrow \text{KeyGen}(1^\lambda)$ and gives pk to \mathcal{A} .
2. **Decryption and leakage queries:** $\mathcal{A}^{\mathcal{O}_{sk}^{\lambda, \ell}(\cdot), \mathcal{D}_{sk}(\cdot)}$ is given access to the leakage oracle $\mathcal{O}_{sk}^{\lambda, \ell}(\cdot)$ and the decryption oracle $\mathcal{D}_{sk}(\cdot)$. A query to the decryption oracle $\mathcal{D}_{sk}(\cdot)$ consists of a ciphertext c , to which the oracle responds with $m = \text{Dec}_{sk}(c)$.
3. **Challenge generation:** \mathcal{A} sends plaintexts m_0, m_1 to the challenger. The challenger chooses $b \xleftarrow{\$} \{0, 1\}$, and sends $c^* \leftarrow \text{Enc}_{pk}(m_b)$ to \mathcal{A} .
4. **Decryption queries:** $\mathcal{A}^{\mathcal{D}_{sk}(\cdot)}$ is given access to the decryption oracle $\mathcal{D}_{sk}(\cdot)$ with the restriction that \mathcal{A} cannot send c^* as a decryption query. Notice also that $\mathcal{A}^{\mathcal{D}_{sk}(\cdot)}$ is not given access to the leakage oracle $\mathcal{O}_{sk}^{\lambda, \ell}(\cdot)$.
5. \mathcal{A} outputs b' , and wins if $b = b'$.

We refer to a 0-leakage-resilient CCA-secure as simply CCA-secure.

Recall that we can define labeled CCA encryption in which a message is encrypted and decrypted according to a public label L . If an encryption scheme $\mathcal{E} = (\text{KeyGen}, \text{Enc}, \text{Dec})$ supports labels, we use the syntax $\text{Enc}^L(m)$ to denote the encryption of message m under label L . Similarly, we use $\text{Dec}^L(c)$ to denote the decryption of ciphertext c under the label L . In this case, we extend the correctness of encryption/decryption to requiring that $\text{Dec}^L(\text{Enc}^L(m)) = m$. The security definition described in Definition 3 can also be easily modified as follows. A query to the decryption oracle now consists of a ciphertext c and a label L , to which the oracle responds with $m = \text{Dec}_{sk}^L(c)$. In the challenge generation stage, \mathcal{A} submits a label L^* as well as messages m_0, m_1 and the challenger computes $c^* \leftarrow \text{Enc}_{pk}^{L^*}(m_b)$ for $b \xleftarrow{\$} \{0, 1\}$. Finally, in the second stage of decryption queries we require that the adversary is allowed to ask for decryptions of any ciphertext c under label L only subject to $(L, c) \neq (L^*, c^*)$.

Definition 4 (Leakage Resilient CPA-Secure Encryption). *We say that an encryption scheme $\mathcal{E} = (\text{KeyGen}, \text{Enc}, \text{Dec})$ is ℓ -leakage resilient CPA-secure if \forall PPT \mathcal{A} we have $\Pr[\mathcal{A} \text{ wins}] \leq \frac{1}{2} + \text{negl}(\lambda)$ in the game described above with the modification that \mathcal{A} does not have access to the decryption oracle $\mathcal{D}_{sk}(\cdot)$. If an encryption scheme is 0-leakage-resilient CPA-secure we simply refer to it as being CPA secure.*

3 Simulation Extractability

We start by briefly recalling the notion of *non-interactive zero-knowledge (NIZK)* [8]. For our purposes, it will be slightly more convenient to use the notion of (*same-string*) *NIZK argument* from [57]. Note, however, that the definitions and constructions given in this section can be extended to the case of NIZK proofs.

Let R be an NP relation on pairs (x, y) with corresponding language $L_R = \{y \mid \exists x \text{ s.t. } (x, y) \in R\}$. A *non-interactive zero-knowledge (NIZK) argument* for a relation R consists of three algorithms ($\text{Setup}, \text{Prove}, \text{Verify}$) with syntax:

- $(\text{CRS}, \text{TK}) \leftarrow \text{Setup}(1^\lambda)$: Creates a common reference string (CRS) and a trapdoor key to the CRS.
- $\pi \leftarrow \text{Prove}_{\text{CRS}}(x, y)$: Creates an argument that $R(x, y) = 1$.
- $0/1 \leftarrow \text{Verify}_{\text{CRS}}(y, \pi)$: Verifies whether or not the argument π is correct.

For the sake of clarity, we write $\text{Prove}, \text{Verify}$ without the CRS in the subscript when the CRS can be inferred from the context. We require that the following properties hold:

Completeness: For any $(x, y) \in R$, if $(\text{CRS}, \text{TK}) \leftarrow \text{Setup}(1^\lambda)$, $\pi \leftarrow \text{Prove}(x, y)$, then $\text{Verify}(y, \pi) = 1$.

Soundness: For any PPT adversary \mathcal{A} , $\Pr[\text{Verify}(y, \pi^*) = 1 \wedge y \notin L_R] \leq \text{negl}(\lambda)$, where the probability is taken over $(\text{CRS}, \text{TK}) \leftarrow \text{Setup}(1^\lambda), (y, \pi^*) \leftarrow \mathcal{A}(\text{CRS})$.

Composable Zero-Knowledge: There exists PPT simulator Sim such that, for any PPT adversary \mathcal{A} we have $|\Pr[\mathcal{A} \text{ wins}] - \frac{1}{2}| \leq \text{negl}(\lambda)$ in the following game:

- The challenger samples $(\text{CRS}, \text{TK}) \leftarrow \text{Setup}(1^\lambda)$ and gives (CRS, TK) to \mathcal{A} .
- \mathcal{A} chooses $(x, y) \in R$ and gives these to the challenger.
- The challenger samples $\pi_0 \leftarrow \text{Prove}(x, y), \pi_1 \leftarrow \text{Sim}(y, \text{TK}), b \leftarrow \{0, 1\}$ and gives π_b to \mathcal{A} .
- \mathcal{A} outputs a bit \tilde{b} and wins if $\tilde{b} = b$.

We revisit the notion of simulation extractable NIZK arguments [58,13,50,51,31], and define a new primitive called *true-simulation extractable* NIZK arguments. Apart from satisfying the properties described above, an NIZK argument is simulation extractable if there exists a PPT *extractor* Ext which given an additional trapdoor to the CRS, extracts a witness x' from any proof π produced by a malicious prover P^* , even if P^* has previously seen some *simulated proofs* for other statements. We make an important distinction between our new definition of *true-simulation extractability*, where all simulated proofs seen by P^* are only of *true* statements, and the stronger notion of *any-simulation extractability*, where P^* can also see proofs of *false* statements. As we will see, the former notion is often simpler to construct and sufficient in our applications.

We extend our definition to *f-extractability*, where Ext only needs to output some function $f(x')$ of a valid witness x' . We further extend this definition to support *labels*, so that the Prove , Verify , Sim , and Ext algorithms also take a public label L as input, and the correctness, soundness, and zero-knowledge properties are updated accordingly. If $\Pi = (\text{Setup}, \text{Prove}, \text{Verify})$ is an NIZK argument with simulator Sim and extractor Ext , we write $\text{Prove}^L, \text{Verify}^L, \text{Sim}^L, \text{Ext}^L$ to denote proof, verification, simulation, and extraction under label L , respectively.

We start by defining a simulation oracle $\mathcal{SIM}_{\text{TK}}(\cdot)$. A query to the simulation oracle consists of a pair (x, y) and a label L . The oracle checks if $(x, y) \in R$. If true, it ignores x and outputs a simulated argument $\text{Sim}^L(\text{TK}, y)$, and otherwise outputs \perp . We now give a formal definition of true-simulation extractability.

Definition 5 (True-Simulation f -Extractability). *Let f be a fixed efficiently computable function and let $\Pi = (\text{Setup}, \text{Prove}, \text{Verify})$ be an NIZK argument for a relation R , satisfying the completeness, soundness and zero-knowledge properties above. We say that Π is true-simulation f -extractable (f -tSE) with labels if:*

- Apart from outputting a CRS and a trapdoor key, Setup also outputs an extraction key: $(\text{CRS}, \text{TK}, \text{EK}) \leftarrow \text{Setup}(1^\lambda)$.
- There exists a PPT algorithm $\text{Ext}(y, \varphi, \text{EK})$ such that for all P^* , $\Pr[P^* \text{ wins}] \leq \text{negl}(\lambda)$ in the following game:
 1. **Key Generation:** The challenger runs $(\text{CRS}, \text{TK}, \text{EK}) \leftarrow \text{Setup}(1^\lambda)$ and gives CRS to P^* .
 2. **Simulation queries:** $P^* \mathcal{SIM}_{\text{TK}}(\cdot)$ is given access to the simulation oracle $\mathcal{SIM}_{\text{TK}}(\cdot)$, which it can adaptively access.
 3. **Adversary Output:** P^* outputs a tuple (y^*, L^*, φ^*) .
 4. **Extraction:** The challenger runs $z^* \leftarrow \text{Ext}^{L^*}(y^*, \varphi^*, \text{EK})$.
 5. P^* wins if (a) the pair (y^*, L^*) was not part of a simulator query, (b) $\text{Verify}^{L^*}(y^*, \varphi^*) = 1$, and (c) for all x' such that $f(x') = z^*$ we have $R(x', y^*) = 0$.⁶

In the case when f is the identity function, we simply say that Π is true-simulation extractable (tSE).

We give several variations of this new primitive. First, we define *one-time* simulation extractability, in which the adversary P^* is only given a *single* query to the simulation oracle $\mathcal{SIM}_{\text{TK}}(\cdot)$. Second, we define the notion of *strong* simulation extractability

⁶ In other words, the adversary wins if the extractor fails to extract a good value z^* which corresponds to at least one valid witness x' ; i.e. $f(x') = z^*$. For the identity function, $f(x) = x$, this corresponds to the statement: $R(z^*, y) = 0$.

by changing the winning condition so that P^* is now required to output a new statement/argument pair instead of a new statement. More formally, condition 5a becomes: the tuple (y^*, L^*, φ^*) is new, that is, either (y^*, L^*) was not part of a simulator query, or if it was, the argument φ^* is different from the one(s) given to P^* by $\widetilde{\text{SIM}}_{\text{TK}}(\cdot)$. We observe that we can generically construct strong f -tSE NIZK arguments from (standard) f -tSE NIZK arguments if we additionally use a strongly-secure one-time signature. In particular, the prover now computes the standard f -tSE argument, signs it, and attaches the verification key vk to the public label. To verify, we first check that the signature is valid and then verify the f -tSE argument.

Finally, we say that an NIZK argument Π is *any-simulation f -extractable* (f -aSE) (similar to the notion of simulation-sound extractability of [31]) if the adversary P^* instead has access to a modified simulation oracle $\widetilde{\text{SIM}}_{\text{TK}}(\cdot)$ that responds to all simulation queries without checking that $R(x, y) = 1$ (and hence might also give simulated arguments of false statements). In this work we do not make use of this variation, but state it here because as we will see, this notion has been implicitly used in prior works. However, f -aSE is a stronger notion than f -tSE and is *not needed*, as we will show that f -tSE is sufficient in constructing leakage-resilient signatures and CCA-encryption.

4 Generic Constructions

In this section we give generic constructions of leakage-resilient hard relations, signatures, and CCA-secure encryption. In the latter two we use the f -tSE NIZK primitive that we defined in Section 3. Finally, we give a construction of f -tSE NIZK arguments.

LEAKAGE-RESILIENT HARD RELATIONS. We begin by showing how to generically construct leakage-resilient hard relations from SPR relations. Informally, we say that a relation R is *second-preimage resistant* (SPR) if given a random $(x, y) \in R$ it is difficult to find $x' \neq x$ such that $(x', y) \in R$. We formalize this in the following definition.

Definition 6 (Second-Preimage Resistant (SPR) Relation). *A relation R with a randomized PPT sampling algorithm KeyGen is second-preimage resistant if:*

- For any $(x, y) \leftarrow \text{KeyGen}(1^\lambda)$, we have $(x, y) \in R$.
- There is a poly-time algorithm that decides if $(x, y) \in R$.
- For any PPT algorithm \mathcal{A} , $\Pr[(x', y) \in R \wedge x' \neq x] \leq \text{negl}(\lambda)$, where the probability is taken over $(x, y) \leftarrow \text{KeyGen}(1^\lambda)$, $x' \leftarrow \mathcal{A}(x, y)$.

We define the average-case pre-image entropy of the SPR relation to be $\mathbf{H}_{\text{avg}}(R) = \widetilde{\mathbf{H}}_\infty(X | Y)$, where random variables (X, Y) are distributed according to $\text{KeyGen}(1^\lambda)$. (We refer the reader to the full version [18] for the definition of $\widetilde{\mathbf{H}}_\infty(X | Y)$.)

Theorem 1. *If $R(x, y)$ is an SPR relation, then it is also an ℓ -leakage resilient hard relation for $\ell = \mathbf{H}_{\text{avg}}(R) - \omega(\log \lambda)$, where λ is the security parameter.*

LEAKAGE-RESILIENT SIGNATURES. We give a generic construction of leakage-resilient signatures based on leakage-resilient hard relations and tSE-NIZK arguments. Let $R(x, y)$ be an ℓ -leakage resilient hard relation with sampling algorithm $\text{KeyGen}_R(1^\lambda)$. Let $\Pi = (\text{Setup}, \text{Prove}, \text{Verify})$ be a tSE-NIZK argument for relation R supporting labels. Consider the following signature scheme:

- $\text{KeyGen}(1^\lambda)$: Output $sk = x$ and $vk = (\text{CRS}, y)$ where $(x, y) \leftarrow \text{KeyGen}_R(1^\lambda)$, $(\text{CRS}, \text{TK}, \text{EK}) \leftarrow \text{Setup}(1^\lambda)$.

- $\text{Sign}_{sk}(m)$: Output $\sigma = \varphi$ where $\varphi \leftarrow \text{Prove}^m(x, y)$. (Note that m is the *label*.)
- $\text{SigVer}_{vk}(m, \sigma)$: Output $\text{Verify}^m(y, \sigma)$.

Theorem 2. *If $R(x, y)$ is an ℓ -leakage resilient hard relation and Π is a labeled tSE-NIZK argument for R , then the above scheme is an ℓ -leakage resilient signature scheme.*

LEAKAGE-RESILIENT CCA-SECURE ENCRYPTION. We give a generic construction of leakage-resilient CCA-secure encryption from leakage-resilient CPA-secure encryption and strong f -tSE NIZK arguments. Let $\mathcal{E} = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be an ℓ -LR-CPA secure encryption scheme and let $\Pi = (\text{Setup}, \text{Prove}, \text{Verify})$ be a one-time strong f -tSE NIZK argument for the relation $R_{\text{enc}} = \{(m, r), (pk, c) \mid c = \text{Enc}_{pk}(m; r)\}$, where $f(m, r) = m$ (i.e. the extractor only needs to extract the message m , but not the randomness r of encryption). We show how to use \mathcal{E}, Π to construct an ℓ -LR-CCA encryption scheme \mathcal{E}^* . Define $\mathcal{E}^* = (\text{KeyGen}^*, \text{Enc}^*, \text{Dec}^*)$ by:

- $\text{KeyGen}^*(1^\lambda)$: Output $pk = (pk_0, \text{CRS}), sk = sk_0$ where $(pk_0, sk_0) \leftarrow \text{KeyGen}(1^\lambda), (\text{CRS}, \text{TK}, \text{EK}) \leftarrow \text{Setup}(1^\lambda)$.
- $\text{Enc}_{pk}^*(m; r)$: Output $C = (c, \pi)$ where $c \leftarrow \text{Enc}_{pk_0}(m; r), \pi \leftarrow \text{Prove}_{\text{CRS}}((pk_0, c), (m, r))$.
- $\text{Dec}_{sk}^*(C)$: Parse $C = (c, \pi)$. If π verifies output $\text{Dec}_{sk}(c)$, else output \perp .

Theorem 3. *Assume that \mathcal{E} is ℓ -LR-CPA secure, and Π is a strong one-time f -tSE NIZK argument for the relation R_{enc} where, for any witness (m, r) , we define $f(m, r) = m$. Then the scheme \mathcal{E}^* defined above is ℓ -LR-CCA secure.*

We note that if the tSE NIZK construction allows labels, then we can naturally extend our construction above to yield a ℓ -LR-CCA encryption *with labels*, by simply putting the encryption labels into the NIZK proofs (and using them to verify the proofs).

TRUE-SIMULATION f -EXTRACTABLE (f -TSE) NIZK. Let f be any efficiently computable function, and let $R(x, y)$ be an NP relation. We show how to construct an f -tSE NIZK argument Ψ from any labeled CCA-secure encryption scheme, and (standard) NIZK arguments. Let $\mathcal{E} = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be a CCA-secure encryption scheme supporting labels, and let $\Pi = (\text{Setup}_\Pi, \text{Prove}_\Pi, \text{Verify}_\Pi)$ be an NIZK argument for the relation $R_\Pi = \{(x, r), (y, c, pk, L) \mid R(x, y) = 1 \wedge c = \text{Enc}_{pk}^L(f(x); r)\}$. We define f -tSE NIZK argument Ψ (supporting labels) as follows:

- $\text{Setup}(1^\lambda)$: Output $\text{CRS} = (\text{CRS}_\Pi, pk), \text{TK} = \text{TK}_\Pi, \text{EK} = sk$ where $(pk, sk) \leftarrow \text{KeyGen}(1^\lambda), (\text{CRS}_\Pi, \text{TK}_\Pi) \leftarrow \text{Setup}_\Pi(1^\lambda)$.
- $\text{Prove}^L(x, y; r)$: Output $\varphi = (c, \pi)$ where $c \leftarrow \text{Enc}_{pk}^L(f(x); r), \pi \leftarrow \text{Prove}_\Pi((x, r), (y, c, pk, L))$.
- $\text{Verify}^L(y, \varphi)$: Parse $\varphi = (c, \pi)$ and run $\text{Verify}_\Pi((y, c, pk, L), \pi)$.

Theorem 4. *If \mathcal{E} is a labeled CCA-secure encryption scheme and Π is an NIZK argument for relation R_Π , then Ψ is a f -tSE NIZK argument for relation R .*

COMPARISON OF OUR GENERIC CONSTRUCTIONS TO PRIOR WORK. The idea of using an SPR relation to construct leakage-resilient hard relations was implicit in [4,42], and explicitly described in [5] for the case of leakage-resilient one-way functions.

Our constructions of leakage-resilient CCA encryption and signatures from tSE NIZKs bear significant resemblance to prior constructions. In particular, we observe that

an alternate construction of tSE NIZK could be achieved by using a CPA-encryption scheme instead of a CCA one, and a ss-NIZK argument system [56] instead of a standard one. In fact, the resulting construction would yield an *any*-simulation extractable (aSE) NIZK argument. This instantiation of aSE NIZKs is implicitly used by [42] in their construction of leakage-resilient signatures. It is also used implicitly in the Naor-Yung “double-decryption” paradigm [49,55,56,45] for CCA security, which was later used in [48] to construct leakage-resilient CCA-encryption. However, as we have seen, tSE is sufficient for constructing *both* leakage-resilient signatures and CCA-encryption and thus, the stronger notion of aSE is not needed. Furthermore, given the current state of efficient encryption schemes and NIZK, the difference in efficiency between ss-NIZK and standard NIZK is *significantly* greater than the difference between CCA and CPA-secure encryption⁷, thus making tSE superior in both simplicity and efficiency.

We note that our construction of tSE NIZKs (based on CCA-encryption and standard NIZKs) was implicitly used by [31] to construct signatures of group elements, and by [11] to construct efficient CCA-encryption with key-dependent message (KDM) security from KDM-secure CPA-encryption. Still, the abstraction of tSE has not been explicitly defined in prior work despite its apparent usefulness.

5 Instantiations

ASSUMPTIONS. We review several standard hardness assumptions on which we will base our constructions.

Decisional Diffie-Hellman (DDH). Let \mathbb{G} be a group of prime order q . Let $g_1, g_2 \xleftarrow{\$} \mathbb{G}$ and $r, r_1, r_2 \xleftarrow{\$} \mathbb{Z}_q$. The decisional Diffie-Hellman (DDH) assumption states that the following two distributions are computationally indistinguishable: $(\mathbb{G}, g_1, g_2, g_1^{r_1}, g_2^{r_2})$ and $(\mathbb{G}, g_1, g_2, g_1^r, g_2^r)$.

Let $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T$ be groups of prime order q and let $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ be a non-degenerate efficiently computable bilinear map.

Symmetric External Diffie-Hellman (SXDH) [59,9,6,26,61]. The symmetric external Diffie-Hellman assumption (SXDH) is that the DDH problem is hard in *both* \mathbb{G}_1 and \mathbb{G}_2 . The assumption is clearly invalid for symmetric pairings (when $\mathbb{G}_1 = \mathbb{G}_2$), but is believed to hold when there is no efficiently computable mapping between \mathbb{G}_1 and \mathbb{G}_2 .

K -Linear [37,60] and DLIN [9]. Let \mathbb{G} be a group of prime order q and let $K \geq 1$ be constant. Let $g_0, g_1, \dots, g_K \xleftarrow{\$} \mathbb{G}$ and $x_0, x_1, \dots, x_K \xleftarrow{\$} \mathbb{Z}_q$. The K -Linear assumption states that the following two distributions are computationally indistinguishable: $(\mathbb{G}, g_0, g_1, \dots, g_K, g_1^{x_1}, \dots, g_K^{x_K}, g_0^{x_0})$, and $(\mathbb{G}, g_0, g_1, \dots, g_K, g_1^{x_1}, \dots, g_K^{x_K}, g_0^X)$, with $X = \sum_{i=1}^K x_i$. Note that for $K = 1$, the K -Linear is the same as DDH, and that it does not hold when working with symmetric pairings. In that setting, the 2-Linear assumption is usually assumed to hold, and is often referred to as the Decisional Linear (DLIN) assumption. *Throughout this paper we assume the K -Linear assumption holds in both \mathbb{G}_1 and \mathbb{G}_2 , which is the case when working with symmetric pairings, and slightly abuse notation when $K = 1$ and assume SXDH holds in that case.*

⁷ Informally, the difference between CCA and CPA-secure encryption is only 2 group elements, whereas the size of a ss-NIZK proof is *more than twice* the size of a standard NIZK proof.

OUR INSTANTIATIONS. We show efficient instantiations of the leakage-resilient signature and CCA-secure encryption constructions described in Section 4. For each scheme, we give two instantiations based on bilinear maps: one secure under the SXDH assumption, and a second, secure under the DLIN assumption. The first can be used with asymmetric pairings, while the second applies to the case of symmetric pairings. We give details of all instantiations in the full version [18] but give a high-level idea below.

Signatures. Recall that to instantiate the signature scheme from Section 4, we need a leakage-resilient hard relation R (which we will derive from an SPR relation) and a true-simulation extractable (tSE) NIZK argument, which we build from CCA-secure encryption and a standard NIZK argument for the relation $\{(x, r), (y, c, pk, L) \mid R(x, y) = 1 \wedge c = \text{Enc}_{pk}^L(f(x); r)\}$. We show our choice of instantiations for these components:

- *CCA-Secure Encryption:* Under both the SXDH and DLIN assumptions, we use efficient encryption schemes in the style of Cramer-Shoup [14,60].
- *NIZK Argument:* We use the Groth-Sahai proof system [33], which can be instantiated both under SXDH and DLIN.
- *SPR Relation:* Previous constructions of leakage-resilient primitives use the SPR function $g_1^{x_1} g_2^{x_2} \dots g_n^{x_n}$. However, this function has the problem that the witness lies in the exponent. This means that we cannot combine it with an encryption scheme for elements in \mathbb{G} (unless each witness component is committed bit by bit which, among other things, results in proofs growing linearly with the security parameter), and unfortunately encryption schemes for messages in \mathbb{Z}_q cannot be combined with the Groth-Sahai system. We therefore construct two new SPR relations based on pairing-product equations. For our SXDH instantiation, we use the relation $e(h_1, x_1) e(h_2, x_2) \dots e(h_n, x_n) = e(y, \tilde{g})$, where \tilde{g} is a generator of \mathbb{G}_2 . We prove that this relation is SPR under the SXDH assumption. In the DLIN case, we use the relation: $e(h_1, x_1) e(h_2, x_2) \dots e(h_n, x_n) = e(y_1, g), e(h_1, x_1) e(h_2, x_2) \dots e(h_n, x_n) = e(y_2, g)$, where g is a generator of \mathbb{G} . We prove that this relation is SPR under the DLIN assumption. To achieve a $(1 - \epsilon)$ leakage ratio, we let n (the number of witness components) in the SPR relation be inversely proportional to ϵ .

Theorem 5. *Let $\mathbb{G}_1, \mathbb{G}_2$ be groups of prime order q . For any $\epsilon > 0$, there exists a $(1 - \epsilon)|sk|$ -leakage resilient signature scheme, secure under the SXDH assumption, using signatures consisting of $(9/\epsilon)(1 + \omega(\log \lambda)/\log q) + 24$ group elements and 2 elements in \mathbb{Z}_q . Similarly, for any $\epsilon > 0$, there exists a $(1 - \epsilon)|sk|$ -leakage resilient signature scheme, secure under the DLIN assumption, using signatures consisting of $(19/\epsilon)(2 + \omega(\log \lambda)/\log q) + 70$ group elements and 6 elements in \mathbb{Z}_q .*

CCA-Secure Encryption. Recall that for leakage-resilient encryption, we need leakage-resilient CPA-secure encryption, standard CCA-secure encryption and strong tSE NIZK, which we can get from combining regular tSE NIZK with a strong one-time signature. We build regular tSE NIZK from CCA-secure encryption and regular NIZK. We describe our choices for each of these below.

- *LR-CPA-Secure Encryption:* We construct a new leakage-resilient CPA-secure encryption scheme for our purpose in the style of ElGamal (similar to ones used in

[48,11] but making it more efficient). The leakage that our new CCA-secure encryption tolerates is the same as the leakage tolerated by the CPA-secure scheme. Informally, we achieve a $(1 - \epsilon)$ leakage ratio in the CPA-secure scheme by increasing the number of generators used in the public key and ciphertext. This number will be inversely proportional to ϵ .

- *CCA-Secure Encryption*: Under both the SXDH and DLIN assumptions, we use efficient encryption schemes in the style of Cramer-Shoup [14,60].
- *NIZK Argument*: We use the Groth-Sahai proof system [33], which can be instantiated both under SXDH and DLIN.
- *One-Time Signature*: We observe that *any* strong one-time signature secure under these assumptions can be used. Here, we opt for the scheme of [31], secure under the Discrete Log assumption (implied by both SDXH and DLIN), because its signature size is small, namely 2 elements in \mathbb{Z}_q .

Theorem 6. *Let $\mathbb{G}_1, \mathbb{G}_2$ be groups of primer order q . For any $\epsilon > 0$, there exists a $(1 - \epsilon)|sk|$ -leakage resilient encryption scheme, secure under the SXDH assumption, using ciphertexts consisting of $(2/\epsilon)(2 + \lambda/\log q) + 15$ group elements and 2 elements in \mathbb{Z}_q . Similarly, for any $\epsilon > 0$, there exists a $(1 - \epsilon)|sk|$ -leakage resilient encryption scheme, secure under the DLIN assumption, using ciphertexts consisting of $(3/\epsilon)(3 + \lambda/\log q) + 34$ group elements and 2 elements in \mathbb{Z}_q .*

6 Other Applications

Once we have efficient leakage-flexible signature schemes, we observe that the standard signature-based ID scheme, where the verifier asks the prover to sign a random message, easily extends to the leakage setting. Moreover, the resulting actively secure ID scheme inherits its relative leakage from the corresponding signature scheme, and satisfies the strongest notion of “anytime-leakage” [4], where the leakage can occur even during the impersonation attack. Although our method is pretty simple, we notice that the other two popular methods of building ID schemes — the use of Σ -protocols for hard relations analyzed in [4] (see first two rows of Tables 3), and the use of CCA-secure encryption (where the prover decrypts a random challenge ciphertext) — inherently do not allow us to obtain optimal results, even when instantiated with leakage-flexible hard relations or CCA-encryption schemes.

Finally, we obtain two efficient leakage-flexible AKA protocols. First, similarly to the case of ID schemes, we can obtain leakage-resilient AKA schemes from any leakage-resilient signature scheme, as formally explained in [4]. The idea is to essentially sign every flow of a standard Diffie-Hellman-based protocol, but with a leakage-resilient signature scheme. We notice, though, that the resulting protocol is not *deniable*. Namely, the transcript of the protocol leaves irrefutable evidence that the protocol took place. Motivated by this deficiency, we design another general AKA protocol based on CCA-encryption. The details are given in the full version [18], but, intuitively, the parties encrypt the flows of the standard Diffie-Hellman-based protocol, effectively proving their identities by successfully re-encrypting the appropriate flows. Although we do not formalize this, this protocols is “deniable”, because the transcript of the protocol can be simulated without the knowledge of parties’ secret keys. To the best of our knowledge, this protocol was not suggested and analyzed even in the leakage-free setting, where it

appears interesting already. Here we actually show that our (new) deniable AKA protocol works even in the presence of leakage.

Acknowledgments

We would like to thank Victor Shoup for helpful discussions, as well as the anonymous reviewers for their useful suggestions.

References

1. *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings*. IEEE Computer Society, 2003.
2. A. Akavia, S. Goldwasser, and V. Vaikuntanathan. Simultaneous hardcore bits and cryptography against memory attacks. In O. Reingold, editor, *TCC*, volume 5444 of *LNCS*, pages 474–495. Springer, 2009.
3. J. Alwen, Y. Dodis, M. Naor, G. Segev, S. Walfish, and D. Wichs. Public-key encryption in the bounded-retrieval model. In Gilbert [28], pages 113–134.
4. J. Alwen, Y. Dodis, and D. Wichs. Leakage-resilient public-key cryptography in the bounded-retrieval model. In Halevi [35], pages 36–54.
5. J. Alwen, Y. Dodis, and D. Wichs. Survey: Leakage resilience and the bounded retrieval model. In *ICITS*, 2009.
6. L. Ballard, M. Green, B. de Medeiros, and F. Monrose. Correlation-resistant storage via keyword-searchable encryption. *Cryptology ePrint Archive*, Report 2005/417, 2005.
7. E. Biham and A. Shamir. Differential fault analysis of secret key cryptosystems. In B. S. K. Jr., editor, *CRYPTO*, volume 1294 of *LNCS*, pages 513–525. Springer, 1997.
8. M. Blum, P. Feldman, and S. Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *STOC*, pages 103–112. ACM, 1988.
9. D. Boneh, X. Boyen, and H. Shacham. Short group signatures. In M. K. Franklin, editor, *CRYPTO*, volume 3152 of *LNCS*, pages 41–55. Springer, 2004.
10. D. Boneh, R. A. DeMillo, and R. J. Lipton. On the importance of checking cryptographic protocols for faults (extended abstract). In W. Fumy, editor, *EUROCRYPT*, volume 1233 of *LNCS*, pages 37–51. Springer, 1997.
11. J. Camenisch, N. Chandran, and V. Shoup. A public key encryption scheme secure against key dependent chosen plaintext and adaptive chosen ciphertext attacks. In Joux [39], pages 351–368.
12. R. Canetti, Y. Dodis, S. Halevi, E. Kushilevitz, and A. Sahai. Exposure-resilient functions and all-or-nothing transforms. In B. Preneel, editor, *EUROCRYPT*, volume 1807 of *LNCS*, pages 453–469. Springer, 2000.
13. R. Canetti, Y. Lindell, R. Ostrovsky, and A. Sahai. Universally composable two-party and multi-party secure computation. In *STOC*, pages 494–503, 2002.
14. R. Cramer and V. Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In H. Krawczyk, editor, *CRYPTO*, volume 1462 of *LNCS*, pages 13–25. Springer, 1998.
15. R. Cramer and V. Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In L. R. Knudsen, editor, *EUROCRYPT*, volume 2332 of *LNCS*, pages 45–64. Springer, 2002.
16. G. D. Crescenzo, R. J. Lipton, and S. Walfish. Perfectly secure password protocols in the bounded retrieval model. In Halevi and Rabin [36], pages 225–244.
17. Y. Dodis, S. Goldwasser, Y. T. Kalai, C. Peikert, and V. Vaikuntanathan. Public-key encryption schemes with auxiliary inputs. In Micciancio [47], pages 361–381.
18. Y. Dodis, K. Haralambiev, A. Lopez-Alt, and D. Wichs. Efficient public-key cryptography in the presence of key leakage. *Cryptology ePrint Archive*, Report 2010/154, 2010.

19. Y. Dodis, Y. T. Kalai, and S. Lovett. On cryptography with auxiliary input. In M. Mitzenmacher, editor, *STOC*, pages 621–630. ACM, 2009.
20. Y. Dodis, A. Sahai, and A. Smith. On perfect and adaptive security in exposure-resilient cryptography. In Pfitzmann [52], pages 301–324.
21. S. Dziembowski. Intrusion-resilience via the bounded-storage model. In Halevi and Rabin [36], pages 207–224.
22. S. Dziembowski and K. Pietrzak. Leakage-resilient cryptography. In *FOCS*, pages 293–302. IEEE Computer Society, 2008.
23. S. Faust, E. Kiltz, K. Pietrzak, and G. N. Rothblum. Leakage-resilient signatures. In Micciancio [47], pages 343–360.
24. S. Faust, T. Rabin, L. Reyzin, E. Tromer, and V. Vaikuntanathan. Protecting circuits from leakage: the computationally-bounded and noisy cases. In Gilbert [28], pages 135–156.
25. A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In A. M. Odlyzko, editor, *CRYPTO*, volume 263 of *LNCS*, pages 186–194. Springer, 1986.
26. S. D. Galbraith and V. Rotger. Easy decision-diffie-hellman groups. *LMS Journal of Computation and Mathematics*, 7:2004, 2004.
27. K. Gandolfi, C. Moutrel, and F. Olivier. Electromagnetic analysis: Concrete results. In Çetin Kaya Koç, D. Naccache, and C. Paar, editors, *CHES*, volume 2162 of *LNCS*, pages 251–261. Springer, 2001.
28. H. Gilbert, editor. *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, French Riviera, May 30 - June 3, 2010. Proceedings*, volume 6110 of *LNCS*. Springer, 2010.
29. S. Goldwasser, Y. Kalai, C. Peikert, and V. Vaikuntanathan. Robustness of the learning with errors assumption. In *Innovations in Computer Science (ICS)*, 2010.
30. S. Goldwasser and Y. T. Kalai. On the (in)security of the fiat-shamir paradigm. In *FOCS* [1], pages 102–.
31. J. Groth. Simulation-sound nizk proofs for a practical language and constant size group signatures. In X. Lai and K. Chen, editors, *ASIACRYPT*, volume 4284 of *LNCS*, pages 444–459. Springer, 2006.
32. J. Groth. Fully anonymous group signatures without random oracles. In K. Kurosawa, editor, *ASIACRYPT*, volume 4833 of *LNCS*, pages 164–180. Springer, 2007.
33. J. Groth and A. Sahai. Efficient non-interactive proof systems for bilinear groups. In N. P. Smart, editor, *EUROCRYPT*, volume 4965 of *LNCS*, pages 415–432. Springer, 2008.
34. J. A. Halderman, S. D. Schoen, N. Heninger, W. Clarkson, W. Paul, J. A. Calandrino, A. J. Feldman, J. Appelbaum, and E. W. Felten. Lest we remember: cold-boot attacks on encryption keys. *Commun. ACM*, 52(5):91–98, 2009.
35. S. Halevi, editor. *Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings*, volume 5677 of *LNCS*. Springer, 2009.
36. S. Halevi and T. Rabin, editors. *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, volume 3876 of *LNCS*. Springer, 2006.
37. D. Hofheinz and E. Kiltz. Secure hybrid encryption from weakened key encapsulation. In A. Menezes, editor, *CRYPTO*, volume 4622 of *LNCS*, pages 553–571. Springer, 2007.
38. Y. Ishai, A. Sahai, and D. Wagner. Private circuits: Securing hardware against probing attacks. In D. Boneh, editor, *CRYPTO*, volume 2729 of *LNCS*, pages 463–481. Springer, 2003.
39. A. Joux, editor. *Advances in Cryptology - EUROCRYPT 2009, 28th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cologne, Germany, April 26-30, 2009. Proceedings*, volume 5479 of *LNCS*. Springer, 2009.

40. A. Juma and Y. Vahlis. Protecting cryptographic keys against continual leakage. In T. Rabin, editor, *CRYPTO*, volume 6223 of *LNCS*, pages 41–58. Springer, 2010.
41. J. Kamp and D. Zuckerman. Deterministic extractors for bit-fixing sources and exposure-resilient cryptography. In *FOCS* [1], pages 92–101.
42. J. Katz and V. Vaikuntanathan. Signature schemes with bounded leakage resilience. In M. Matsui, editor, *ASIACRYPT*, volume 5912 of *LNCS*, pages 703–720. Springer, 2009.
43. P. C. Kocher. Timing attacks on implementations of diffie-hellman, rsa, dss, and other systems. In N. Kobitz, editor, *CRYPTO*, volume 1109 of *LNCS*, pages 104–113. Springer, 1996.
44. P. C. Kocher, J. Jaffe, and B. Jun. Differential power analysis. In M. J. Wiener, editor, *CRYPTO*, volume 1666 of *LNCS*, pages 388–397. Springer, 1999.
45. Y. Lindell. A simpler construction of cca2-secure public-key encryption under general assumptions. *J. Cryptology*, 19(3):359–377, 2006.
46. S. Micali and L. Reyzin. Physically observable cryptography (extended abstract). In M. Naor, editor, *TCC*, volume 2951 of *LNCS*, pages 278–296. Springer, 2004.
47. D. Micciancio, editor. *Theory of Cryptography, 7th Theory of Cryptography Conference, TCC 2010, Zurich, Switzerland, February 9-11, 2010. Proceedings*, volume 5978 of *LNCS*. Springer, 2010.
48. M. Naor and G. Segev. Public-key cryptosystems resilient to key leakage. In Halevi [35], pages 18–35.
49. M. Naor and M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC*, pages 427–437. ACM, 1990.
50. R. Pass and A. Rosen. Concurrent non-malleable commitments. In *FOCS*, pages 563–572. IEEE Computer Society, 2005.
51. R. Pass and A. Rosen. New and improved constructions of non-malleable cryptographic protocols. In H. N. Gabow and R. Fagin, editors, *STOC*, pages 533–542. ACM, 2005.
52. B. Pfitzmann, editor. *Advances in Cryptology - EUROCRYPT 2001, International Conference on the Theory and Application of Cryptographic Techniques, Innsbruck, Austria, May 6-10, 2001, Proceeding*, volume 2045 of *Lecture Notes in Computer Science*. Springer, 2001.
53. K. Pietrzak. A leakage-resilient mode of operation. In Joux [39], pages 462–482.
54. J.-J. Quisquater and D. Samyde. Electromagnetic analysis (ema): Measures and countermeasures for smart cards. In I. Attali and T. P. Jensen, editors, *E-smart*, volume 2140 of *LNCS*, pages 200–210. Springer, 2001.
55. C. Rackoff and D. R. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In J. Feigenbaum, editor, *CRYPTO*, volume 576 of *LNCS*, pages 433–444. Springer, 1991.
56. A. Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *FOCS*, pages 543–553, 1999.
57. A. D. Santis, G. D. Crescenzo, R. Ostrovsky, G. Persiano, and A. Sahai. Robust non-interactive zero knowledge. In J. Kilian, editor, *CRYPTO*, volume 2139 of *LNCS*, pages 566–598. Springer, 2001.
58. A. D. Santis and G. Persiano. Zero-knowledge proofs of knowledge without interaction (extended abstract). In *FOCS*, pages 427–436. IEEE, 1992.
59. M. Scott. Authenticated id-based key exchange and remote log-in with simple token and pin number. *Cryptology ePrint Archive*, Report 2002/164, 2002.
60. H. Shacham. A cramer-shoup encryption scheme from the linear assumption and from progressively weaker linear variants. *Cryptology ePrint Archive*, Report 2007/074, 2007.
61. E. R. Verheul. Evidence that xtr is more secure than supersingular elliptic curve cryptosystems. *J. Cryptology*, 17(4):277–296, 2004.