

Signatures Resilient to Continual Leakage on Memory and Computation

Tal Malkin¹, Isamu Teranishi^{1,2}, Yevgeniy Vahlis¹, and Moti Yung^{1,3}

¹ Columbia University.

{tal, evahlis}@cs.columbia.edu

² NEC Japan.

teranisi@ah.jp.nec.com

³ Google Inc.

moti@cs.columbia.edu

Abstract. Recent breakthrough results by Brakerski *et al* and Dodis *et al* have shown that signature schemes can be made secure even if the adversary *continually* obtains information leakage from the secret key of the scheme. However, the schemes currently do not allow leakage on the secret key and randomness *during signing*, except in the random oracle model. Further, the random oracle based schemes require updates to the secret key in order to maintain security, even when no leakage during computation is present.

We present the first signature scheme that is resilient to full continual leakage: memory leakage as well as leakage from processing during signing (both from the secret key and the randomness), in key generation, and in update. Our scheme can tolerate leakage of a $1 - o(1)$ fraction of the secret key between updates, and is proven secure in the standard model based on the symmetric external DDH (SXDH) assumption in bilinear groups. The time periods between updates are a function of the amount of leakage in the period (and nothing more).

As an additional technical contribution, we introduce a new tool: independent pre-image resistant hash functions, which may be of independent interest.

1 Introduction

Cryptographic schemes have traditionally been modeled under the assumption that the secret key is hidden completely within a device and attacks are of “black box” nature. However, cryptographic engineering has taught us that devices are not perfect and can leak information about the key, primarily via what is known as “side channels.” These channels give the adversary some partial information by observing physical leakage correlated with the secret key, such as timing or radiation or power consumption associated with computations (either directly by probing the circuit or via antennas probing unshielded devices), as well as memory leakage that reveals information about the secret key (cf., [BB05, QS01, KJJ99, HSH+08].)

The threat of partial leakage of keys has not escaped cryptographers; perhaps the earliest works on the subject are Shamir’s secret sharing [S79], and later Rivest’s all or nothing transform [R97]. In the last few years a large body of work on leakage-resilient cryptography has emerged (cf., [ISW03,MR04,SMY09,AGV09,DKL09,P09,NS09,ADW09,KV09,FKPR10,DGK+10,FRR+10,BG10,DP10,JV10,GR10,DHLW10,BKKV10]), which provide different leakage models requiring that cryptographic schemes be secure even if partial information about the secret key is leaked to an adversary. In order to have an abstract model that is not directly dependent upon the hardware architecture itself, but rather can produce some general guidelines to systems and hardware designers regarding how much leakage can be tolerated overall (within a given setting), leakage was formally modeled as functions associated with algorithmic steps and states. These leakage functions $f_1(sk), \dots, f_q(sk)$ of the secret key sk , are from some given family of functions, where the specific f_i is selected by the adversary arbitrarily. The leakage (i.e., the function result which contains only partial information about the key) is applied to the relevant state or computation and is given to the adversary (in addition to the black-box access it gets).

Wiki-Leakage: The variety of Leakage Models. There are many types of leakage sub-models based on various parameters of leakage, which we briefly review here. Adaptive leakage, where the function is chosen dynamically by the adversary, obviously gives it more power than non-adaptive. Such adversaries that rather than observing the device once execute side channel attacks repetitively, are stronger and require the schemes to be *continuous leakage resilient*. Repeating (continual) adaptive leakage requires the scheme to have an update procedure for the secret key. Otherwise, the adversary can eventually leak enough information about the key to break the scheme, even if the amount of leakage per time period is small. However, the public key should *not* be changed by the secret key update. For example, a signature verification key should remain valid throughout the lifetime of the scheme.

Next we note that leakage can be *processing leakage* (i.e., only computation leaks) so that physical computations on the device are the only source of leakage. Procedures such as key generation, key updates, and other kinds of processing (i.e., signing, decryption, etc.) which involve random values may be allowed different amounts of leakage, but are all related to some physics of computations (e.g., timing of operations). The influential works of Ishai, Sahai, and Wagner [ISW03] and Micali and Reyzin [MR04] enable us to construct schemes under the “any computation, and only computation, leak information,” model, which has led to many recent achievements. In contrast, *memory leakage* [AGV09] (which, in some sense, can be traced to the original works of Shamir and Rivest mentioned above) are produced as a function of the memory state itself. This type of leakage is orthogonal to computational leakage: an adversary can get memory leakage by probing memories even if the memories are not currently used in any computation (e.g., the cold-boot attacks [HSH+08]). For example, the scheme of [DHLW10] is secure against memory attacks (even continual), but

assumes that the signing process leaks no information. The most general model allows *full leakage* which includes leakage both from processing and memory.

The most demanding case for designing digital signature schemes seems to be the case of adaptive and continual full leakage that is available to the adversary from both computational and memory sources (without protection of sub-steps of computations). However, to date, there are no known schemes which achieve a digital signature scheme in this adversarial setting in the standard model. All known schemes with full (memory and processing) leakage either do not have a key update algorithm and thus are not continual (cf., [KV09]), have a key update algorithm but require some restrictions (e.g., [ADW09] which requires an additional leakage-free master key), or are based on the random oracle model (with a relaxation of the definition of a “time period”) [BKKV10].

1.1 Our Contributions

We propose the first signature scheme satisfying all of the above requirements, whose security can be proven in the standard model and without further relaxation. Specifically, our scheme protects against (1) continual memory leakages combined with (2) all types of continual processing (computational) leakages, namely leakages from key generation, key updates, and signing.

Moreover, the amount of information that our scheme allows to leak in each time period is optimal, in the sense that our scheme remains secure even if $1 - o(1)$ fraction of the secret key of each time period is leaked. Here “time period” is the period between two consecutive updates of the secret key (the time period is a function of the accumulated leakage itself and not a relaxed notion which depends on other parameters). We stress that our scheme remains secure even when the leakage during signing is a function $f(sk, r)$ of both the secret key and the randomness. Further, the function f can be adaptively chosen by the adversary. Using standard techniques, our scheme also allows $O(\log \kappa)$ leakage in the key generation and in each of the key updates, where κ is a security parameter. (The secret key has $O(\kappa)$ bit length. The fraction of leakage is therefore $O(\log \kappa / \kappa)$.)

Comparison with Recent Schemes. Let us compare our scheme with the recent breakthrough results of Dodis, Haralambiev, Lopez-Alt, and Wichs [DHLW10] and Brakerski, Kalai, Katz, and Vaikuntanathan [BKKV10]. The signature scheme of [DHLW10], as noted above, has to assume that there is no leakage from the signing process. The work in [BKKV10] proposed two signature schemes. Their first scheme is secure on if there is no leakage from the signing process. The second scheme relies on random oracle to protect against leakage during signing, and further requires signatures to be treated as leakage. That is, even if there is no actual side-channel leakage during a certain time period, the signing key must be refreshed to preserve security. In contrast, our signature scheme is proved secure in the standard model and the key needs to be refreshed only if leakage occurs (i.e. signatures do not constitute leakage).

Concurrent work. In a concurrent work, Boyle, Segev, and Wichs [BSW10] construct a fully leakage resilient signature scheme using different techniques. [BSW10] take a more generic approach than we do. On the other hand, our scheme is somewhat more efficient.

Other related work. In the recent years a very rich body of research on leakage resilient cryptography has been developed. Dziembowski and Pietrzak [DP08], and Pietrzak [P09] describe the first stream ciphers resilient to continual leakage in the only computation leaks model. Faust *et al* [FKPR10] construct signature schemes resilient to continual leakage in the only computation leaks model. Faust *et al* [FRR+10] give a general compiler using secure hardware that protects an arbitrary circuit against continual leakage that can be modeled as a shallow (AC^0) boolean circuit. Juma and Vahlis [JV10], and separately Goldwasser and Rothblum [GR10], give compilers that protect any algorithm against continual leakage (without complexity restrictions), using secure hardware. Recently, Dodis and Pietrzak [DP10] show how to build continual leakage resilient pseudorandom functions that are secure against non-adaptive leakage.

Discussion on processing leakage. Continual memory attacks are a very powerful leakage model that allows the adversary to continuously obtain leakage from the entire secret key. A natural question to ask is whether adding processing leakage into the mix adds anything to adversarial power. Indeed, the only additional information available to the adversary during processing leakage is ephemeral randomness that is used in the computation. In many cases, such as in the case of public key encryption or decryption (using the corresponding private key), leakage on ephemeral randomness does not provide any useful information to the adversary about the secret key. In fact, in public key encryption and decryption, the adversary can simulate the leakage from the randomness of these algorithms on her own. However, this is not the case for signature schemes.

In a signature scheme, a signature is computed using the secret key, and made public. Consequently, signatures can be viewed as a very restricted type of leakage on the secret key. A signature scheme is considered secure if such “signature leakage” is useless to any efficient adversary. When the adversary is given leakage from the randomness of the signing process, she may be able to obtain information that will allow her to extract useful information from the accompanying signature. For example, each signature may contain an encryption of the secret key under a random key that is generated during signing, and then forgotten. If the adversary is able to leak this random key, she trivially breaks the security of the scheme.

1.2 The Idea Behind Our Construction

We are motivated in our construction by the (non-continual) memory-attack resilient signature schemes of Alwen, Dodis, and Wichs [ADW09], and of Katz and Vaikuntanathan [KV09]. Both [ADW09] and [KV09] use the following high level approach, which is based on the Fiat-Shamir heuristic [FS86]: a signature of a message M relative to a public key pk is an extractable proof of knowledge

(in the random oracle model) of a value sk for which $H(sk) = pk$. Here is H is a hash function.

The security proof of these signature schemes relies on the *second preimage resistance* of the hash function H , and the witness extractability of the proofs that are used. That is, they use the property that it is infeasible to find $sk^* \neq sk$ satisfying $H(sk^*) = H(sk)$.

To prove security, they construct a simulator that generates a secret key sk randomly and computes $pk = H(sk)$. The simulator then can answer leakage queries and signing queries using the secret key that it itself has generated. If an adversary can forge a message/signature pair, the simulator extracts the witness sk' . Now, if the fraction of leakage of sk is less than $1 - o(1)$, the exact key sk that is used by the simulator is information theoretically undetermined in the view of the adversary (specifically, there are at least two possible keys, given the leakage). Therefore, with probability at least $1/2$, the witness sk' is different from sk , which breaks the second pre-image resistance of H .

We start with this basic approach, and enhance it along three dimensions. Specifically, we:

1. Remove the requirement for a random oracle, and get a scheme secure in the standard model.
2. Add a key update procedure that refreshes the secret key, while keeping the public key fixed. This yields a signature scheme resilient against *continual* memory attacks [BKKV10].
3. Develop a proof method that allows leakage of randomness used in signing within a period (allowing optimal leakage).

Removing the Random Oracle The simplest idea to remove the random oracle from the scheme of [ADW09,KV09] is to prove the knowledge of the secret key not by using Fiat-Shamir heuristic, but by using other known non-interactive proof systems in the standard model.

This initial attempt fails for the following reason: the argument of [ADW09,KV09] showing $sk^* \neq sk$ is purely information theoretic. Hence, if we want to use the argument of [ADW09,KV09], the proof systems should hide sk not in a computational sense, but in an information theoretic sense. But if the proof system hides sk information theoretically, the secret key sk^* used by an adversary is also hidden information theoretically (since no random oracle is available). Hence, it is impossible for the simulator to get the second pre-image sk^* from the adversary's forgery, and so we cannot rely on second pre-image resistance.

To overcome the above problem, we use the Waters' function

$$h(\mathcal{H}, M) = H_0 + \sum_k M_k H_k,$$

where M_k is the k -th bit of a message M and $\mathcal{H} = (H_0, \dots, H_m)$ is a tuple of group elements¹. The function is introduced in [W05] in order to construct

¹ Here we use *additive notation* to describe the function.

an adaptively secure ID-based encryption scheme. The Waters' function has the property that a simulator can choose the parameters H_0, \dots, H_m in a special way that defines a subset \mathcal{M} of the range of h . It can then be shown that with non-negligible probability, all signing queries M of the adversary satisfy $h(\mathcal{H}, M) \in \mathcal{M}$, and the forgery M^* satisfies $h(\mathcal{H}, M^*) \notin \mathcal{M}$.

We construct a Waters-like function h' such that \mathcal{M} is a set of all non-DDH tuples in the range of h' . Consequently, we get that with non-negligible probability all signing queries of the adversary map to non-DDH tuples, and the forgery maps to a DDH tuple.

We then combine the above hash function h' with the Groth-Sahai [GS08] proof system. Groth-Sahai is a proof system which uses a common reference string (CRS). The proof system hides the witness information theoretically if the CRS is a non-DDH tuple and hides it only computationally, and the witness therefore is extractable, if the CRS is a DDH tuple.

Hence, by using Groth-Sahai proofs as signatures, and $h'(M)$ as the CRS of the Groth-Sahai proof, we get a proof system that hides the witness sk information theoretically when the simulator generates proofs as answers to signing queries, and allows the witness sk^* to be extracted from the proof generated by an adversary as a forged signature.

The scheme before adding key update. Our scheme therefore has the following basic structure. The private key consists of a vector of group elements \mathbf{W} , and the public key consists of group elements \mathbf{A} and T such that $e(\mathbf{A}, \mathbf{W}) = T$. The public key also contains the description of a Waters-like hash function h' . To sign a message M , we first use h' to compute a CRS $h'(\mathcal{H}, M)$ for the Groth-Sahai proof system. Then, the signature is a proof, under the CRS $h'(\mathcal{H}, M)$, that $e(\mathbf{A}, \mathbf{W}) = T$.

Before proceeding to describe our key update procedure, we believe it is instructive to see why the above scheme, without update, is secure. Intuitively, this follows from the second pre-image resistance of the hash function $H_{\mathbf{A}}(\mathbf{X}) := e(\mathbf{A}, \mathbf{X})$. From the perfect witness indistinguishability of Groth-Sahai proofs we know that the adversary learns about the specific witness \mathbf{W} only from leakage. However, since the amount of leakage is bounded, the actual witness \mathbf{W} in the private key remains information theoretically undetermined. Finally, when the adversary submits a successful forgery, we use the indistinguishability of common reference strings technique discussed above to show that, with non-negligible probability, a witness \mathbf{W}' can be extracted from the forgery. This witness is likely to be different from \mathbf{W} , which would give the simulator two inputs \mathbf{W} and \mathbf{W}' such that $H_{\mathbf{A}}(\mathbf{W}) = H_{\mathbf{A}}(\mathbf{W}')$. This in turn violates the second pre-image resistance of $H_{\mathbf{A}}$.

We remark that the above technique is somewhat reminiscent of the Feige-Lapidot-Shamir [FLS90] method for using witness indistinguishability to achieve witness hiding.

Adding Key Updates The approach of Section 1.2 allows us to move from the random oracle to the standard model. However, the above scheme can still

tolerate only a bounded amount of leakage. We now describe a method for choosing the private keys of the scheme that allows us to randomize the key without having to issue a new public key.

We start by observing that if our scheme relies for security on the second pre-image resistance of the hash function H , then no key update algorithm can exist. This is because otherwise we could use the key update algorithm itself to break second pre-image resistance as follows:

If we can get new key $sk^{[i+1]}$ efficiently by updating $sk^{[i]}$, this means that one can get two keys $sk^{[i]}$ and $sk^{[i+1]}$ satisfying of $T = H(sk^{[i]})$ and $T = H(sk^{[i+1]})$, where T is the public key. The function H therefore cannot be second pre-image resistant.

Note that our model does not allow to update the public key for the convenience of verifiers. The above collision of the function H is therefore unavoidable.

(n, k) -independent pre-image resistant (IPIR) hash functions. We overcome the above problem by introducing the following new notion: (n, k) -independent pre-image resistant hash function H , where n is an integer and $k \leq n - 2$. This is a linear function H from an n -dimensional vector space \mathbb{H}^n to a 1-dimensional vector space \mathbb{T} , over \mathbb{Z}_p . We require that, given certain trapdoor information about H , one can find a tuple $(\mathbf{Y}_1, \dots, \mathbf{Y}_{k+1}, T)$ satisfying $T = H(\mathbf{Y}_j)$ for all $j \in [k + 1]$. However, it must be infeasible to find $\mathbf{Y}_* \in \mathbb{H}^n$ satisfying $T = H(\mathbf{Y}_*)$ and $\mathbf{Y}_* \notin \text{Aff}(\mathbf{Y}_1, \dots, \mathbf{Y}_{k+1})$, where $\text{Aff}(\mathbf{Y}_1, \dots, \mathbf{Y}_{k+1})$ is the smallest affine space spanned by $\mathbf{Y}_1, \dots, \mathbf{Y}_{k+1}$. We call the hardness property (n, k) -independent preimage resistance.

We note that although in this paper we give a construction of an (n, k) -IPIR hash function, we do not use it as a black box in our signature scheme. Indeed, the Groth-Sahai proofs that are generated use the parameters of the hash function, which are of a very specific form. However, the notion of IPIR seems useful in itself, and we hope that other applications will be found for it. Furthermore, abstracting the properties that we need from H allows us to present a modular and structured proof of security for our signature scheme.

Generating and updating keys. Using the linearity of H , we can generate any linear sum of $\mathbf{Y}_1, \dots, \mathbf{Y}_{k+1}$ in polynomial time. We use this property to perform key update. And we use the IPIR (instead of the second pre-image resistance) when we show that no adversary can forge a signature.

In slightly greater detail, we construct the key generation algorithm and the update algorithm of our scheme as follows. In the key generation, by using the trapdoor of the hash function H , we generate $(\mathbf{Y}_1, \mathbf{Y}_2, T)$ satisfying $T = H(\mathbf{Y}_1) = H(\mathbf{Y}_2)$. We then compute $\mathbf{Q} \leftarrow \mathbf{Y}_1 - \mathbf{Y}_2$ and publish \mathbf{Q} as a part of the public key. The secret key is $\mathbf{W}^{[0]} \leftarrow \mathbf{Y}_2$. Note that $H(\mathbf{Q}) = H(\mathbf{Y}_1) - H(\mathbf{Y}_2) = 0$ holds.

Key update then works as follows: in the $(i + 1)$ -st round, we select $s^{[i]} \xleftarrow{\$} \mathbb{Z}_p$ randomly and compute $\mathbf{W}^{[i+1]} \leftarrow \mathbf{W}^{[i]} + s^{[i]} \mathbf{Q}$. Based on the linearity of H , and the equality $H(\mathbf{Q}) = 0$, one can easily show that $H(\mathbf{W}^{[i+1]}) =$

$H(\mathbf{W}^{[i]}) = \dots H(\mathbf{W}^{[0]}) = H(\mathbf{Y}_2) = T$ holds, and that $\mathbf{W}^{[i]}$ is an element of $\text{Aff}(\mathbf{Y}_1, \mathbf{Y}_2)$. The latter holds because $\mathbf{W}^{[i]}$ are linear sums of $\mathbf{W}^{[0]} = \mathbf{Y}_2$ and $\mathbf{Q} = \mathbf{Y}_1 - \mathbf{Y}_2$. We then use an adaptation of the “leakage resilient subspace” technique from [BKKV10] to show that the affine subspace $\text{Aff}(\mathbf{Y}_1, \mathbf{Y}_2)$ (or even $\text{Aff}(\mathbf{Y}_1, \dots, \mathbf{Y}_{k+1})$) is hidden from the adversary, even given continual leakage (assuming the refresh procedure above is periodically performed). Given the hiding property of the affine space, it follows that if the adversary forges a signature $\sigma^* = \text{Prf}_{M^*}(\mathbf{W}^*)$ for some message M^* , she likely uses a witness $\mathbf{W}^* \notin \text{Aff}(\mathbf{Y}_1, \mathbf{Y}_2)$. However, this violates the IPIR of H . The security of the scheme follows.

Security Against Leakage in Signing The main challenge in achieving security under leakage from the signing process is that the signature and the secret key are correlated through the randomness that was used to produce the signature. When the adversary obtains leakage on the randomness, the signature may become much more valuable, and potentially allow the adversary to break the scheme (as we discussed in the introduction).

In the proof based signature schemes we described above, there is no guarantee that leakage on the randomness of the prover does not break the zero-knowledge or witness indistinguishability property of the proof system. We solve this problem through a combination of several tools: first, as described above, we rely on Groth-Sahai proofs, which have a dual mode – witness hiding and witness binding. When the proof is (perfectly) hiding, it is information theoretically independent from the witness, which is the secret key, if there is no leakage on the randomness of the prover.

We use the above fact to “invert” the order in which components of the signature are generated: first the GS proof σ in the signature is generated using some globally fixed witness \mathcal{Y} (note that this is only done by the simulator in the analysis, and so there is no leakage on the witness \mathcal{Y}). Then, given an actual witness \mathbf{W} for the proof, we “reverse engineer” the randomness R that would yield the same proof σ , and compute leakage on (\mathbf{W}, R) . We use an additional property of Groth-Sahai that for every pair of proof and witness (σ, \mathbf{W}) there exists a unique randomness $R_{\sigma, \mathbf{W}}$ that causes the prover to output σ given witness \mathbf{W} . Moreover, since the proof is perfectly witness hiding, for all witnesses \mathbf{W} , the distribution on the tuple $(\sigma, R_{\sigma, \mathbf{W}})$ are identical whether we first generate the proof using witness \mathbf{V} and then determine the randomness, or choose the randomness uniformly, and compute the proof directly.

The above approach, however, does not work as is, since the process of finding R may not be efficiently computable! We therefore rely on an information theoretic leakage resilience property of random subspaces that was shown in [BKKV10] (in fact, we prove a slightly stronger version that suits our construction). We combine both techniques together, simultaneously using the randomness reverse engineering technique described above to handle leakage from signing, and information theoretic indistinguishability of random subspaces under leakage. Using these techniques together, we show that the adversary is unable

to gain information about the subspace from which we generate private keys during update, even if leakage on the signing process is available. We then use independent pre-image resistance to conclude that our scheme is secure.

2 Preliminaries

Notations. Let $[n]$ denote $\{1, \dots, n\}$ and $[k..n]$ denote $\{k, \dots, n\}$. For two random variables X and Y , $\text{dist}(X, Y)$ denote the statistical distance between X and Y .

Linear Algebra. Unless otherwise stated, vectors in this paper are column vectors. We represent a row vector as a transpose of a column vector in this paper. For natural numbers n and m , let $\mathbb{Z}_p^{n \times m}$ denote the set of $n \times m$ matrices over \mathbb{Z}_p . Let A^T denote the transposed of a matrix $A = (a_{i,j})_{i,j}$, that is, $A^T = (a_{j,i})_{i,j}$. For two vectors $\mathbf{u} = (u_1, \dots, u_n)$, $\mathbf{v} = (v_1, \dots, v_n) \in \mathbb{Z}_p^n$, we let $\langle \mathbf{u}, \mathbf{v} \rangle$ denote the inner product of them in \mathbb{Z}_p , that is, $\langle \mathbf{u}, \mathbf{v} \rangle = \sum_i u_i v_i \pmod p$.

For a vector $\mathbf{v} = (v_1, \dots, v_n) \in \mathbb{Z}_p^n$ and an element W of the group \mathbb{G} or \mathbb{H} , let $\mathbf{v}W$ denote the vector (v_1W, \dots, v_nW) .

For a vector space \mathcal{V} and vectors $\mathbf{Y}_1, \dots, \mathbf{Y}_{k+1} \in \mathcal{V}$, let $\text{Span}(\mathbf{Y}_1, \dots, \mathbf{Y}_k)$ denote the smallest vector subspace of \mathcal{V}^n which contains all of $(\mathbf{Y}_j)_{j \in [k]}$, that is,

$$\text{Span}(\mathbf{Y}_1, \dots, \mathbf{Y}_k) = \{\mathbf{Y} \in \mathcal{V}^n \mid \exists s_1, \dots, s_k \in \mathbb{Z}_p : \mathbf{Y} = \sum_{j \in [k]} s_j \mathbf{Y}_j\}.$$

Similarly, let $\text{Aff}(\mathbf{Y}_1, \dots, \mathbf{Y}_{k+1})$ is the smallest affine subspace of \mathcal{V}^n which contains all of $(\mathbf{Y}_j)_{j \in [k+1]}$, that is,

$$\text{Aff}(\mathbf{Y}_1, \dots, \mathbf{Y}_{k+1}) = \{\mathbf{Y} \in \mathcal{V}^n \mid \exists s_1, \dots, s_{k+1} \in \mathbb{Z}_p : \mathbf{Y} = \sum_{j \in [k+1]} s_j \mathbf{Y}_j, \sum_{j \in [k+1]} s_j = 1\}.$$

Note that the space $\text{Aff}(\mathbf{Y}_1, \dots, \mathbf{Y}_{k+1})$ becomes k dimensional, when $k+1$ vectors $\mathbf{Y}_1, \dots, \mathbf{Y}_{k+1}$ are linear independent.

2.1 Signature Schemes Resilient Against Continual Leakage

A *signature scheme with key update* \mathcal{SGN} consists of four algorithms Kg , Sig , Ver , and Update . The inputs and outputs of Kg , Sig , and Ver are the same as in standard signature schemes. Update takes as input a secret key and a public key and outputs a new element of the secret key space. $\mathcal{SGN} = (\text{Kg}, \text{Sig}, \text{Ver}, \text{Update})$ has to satisfy the following property:

- (**Correctness**) For any integer $n \geq 0$ and any message M , if we compute $(pk, sk_0) \leftarrow \text{Gen}(1^\kappa)$, $sk_1 \leftarrow \text{Update}_{pk}(sk_0)$, \dots , $sk_n \leftarrow \text{Update}_{pk}(sk_{n-1})$, and $\sigma \leftarrow \text{Sig}(sk_n, M)$, $\text{Ver}(pk, M, \sigma) = 1$ always holds.

We follow the definition of [BKKV10] of leakage resilient signatures.

Definition 1 ([BKKV10]) Let $\rho_G, \rho_U, \rho_M,$ and ρ_S be elements of the real range $[0, 1]$. We say that $\mathcal{SGN} = (\text{Gen}, \text{Sig}, \text{Ver}, \text{Update})$ is $(\rho_G, \rho_U, \rho_M, \rho_S)$ -*EU-CMA-CML secure* (stand for existentially unforgeable under chosen message attack in the CML model) if no PPT adversary \mathcal{A} succeeds in the game of Fig.1 with non-negligible probability. Here $\text{Rnd}[\text{Algo}]$ denote the set of randomnesses for algorithm Algo .

Setup $\mathcal{A}(1^\kappa)$ sends to the challenger a function f satisfying $|f(R)| \leq \rho_G |R|$ for all R . The challenger then selects $R \xleftarrow{\$} \text{Rnd}[\text{Gen}]$ computes $(pk, sk_0) \leftarrow \text{Gen}(1^\kappa; R)$ sends $(pk, f(R))$ to \mathcal{A} , and initializes $i \leftarrow 0$ and $L_i \leftarrow |f(R)|$. Here i represents the number of updates and L_i denote the bit length of all leakages about the i -th secret key.

Queries \mathcal{A} makes queries of the following three types polynomial number of times:

- Update queries (update, f) where f is a circuit satisfying $|f(sk, R)| \leq \rho_U(|sk| + |R|)$ for any (sk, R) . If $L_i + |f(sk_i, R)| \leq \rho_M |sk_i|$ holds, the challenger chooses $R \xleftarrow{\$} \text{Rnd}[\text{Update}]$ randomly, computes $sk_{i+1} \leftarrow \text{Update}_{pk}(sk_i)$ and sends $f(sk_i, R)$ back to \mathcal{A} and resets $i \leftarrow i + 1$ and $L_i \leftarrow |f(sk_i, R)|$. Otherwise, the challenger aborts.
- (Memory) leak queries (leak, f), where f is a circuit. If $L_i + |f(sk_i)| \leq \rho_M |sk_i|$ holds, the challenger sends $f(sk_i)$ to adversary and resets $L_i \leftarrow L_i + |f(sk_i)|$. Otherwise, the challenger aborts.
- Signing queries (sig, M, f) where f is a circuit with $|f(sk, R)| \leq \rho_S(|sk| + |R|)$ for any (sk, R) . The challenger chooses $R \leftarrow \text{Rnd}[\text{Sig}]$ randomly, computes $\sigma \leftarrow \text{Sig}(sk_i, M; R)$ and sends $(\sigma, f(sk_i, R))$ back to \mathcal{A} .

Challenge Assuming the challenger did not aborts, \mathcal{A} outputs (M_*, σ_*) . It succeeds if $\text{Ver}(pk, M_*, \sigma_*) = 1$ holds and \mathcal{A} never made query (sig, M_*).

Fig. 1. Game of $(\rho_G, \rho_U, \rho_M, \rho_S)$ -EU-CMA-CML secure

2.2 Bilinear Pairings

In our paper, we are working on a bilinear pairing, $e : \mathbb{G} \times \mathbb{H} \rightarrow \mathbb{T}$ with prime order p , where $\mathbb{G} \neq \mathbb{H}$ holds and there is no efficiently computable homomorphism between two groups \mathbb{G} and \mathbb{H} (Such a pairing is called Type III [GPS08]). We denote by gk bilinear map parameters of the form $(p, \mathbb{G}, \mathbb{H}, \mathbb{T}, e)$. We will occasionally refer to gk as *group description*.

Our proofs rely on basic properties of linear algebra. We therefore find it convenient to use *additive notation* for pairings. For example, we write $e((a + b)A, W) = a \cdot e(A, W) + b \cdot e(A, W)$. For two (column) vectors $\mathbf{A} = (A_1, \dots, A_n)^T \in \mathbb{G}$ and $\mathbf{W} = (W_1, \dots, W_n)^T \in \mathbb{H}$, we denote

$$e(\mathbf{A}^T, \mathbf{W}) = \sum_{i \in [n]} e(A_i, W_i)$$

Assumption 2 (SXDH assumption [GS08]) *We say that $gk = (p, \mathbb{G}, \mathbb{H}, \mathbb{T}, \mathbf{e})$ satisfies the Symmetric external Diffie-Hellman (SXDH) assumption, if the DDH assumption holds both over \mathbb{G} and \mathbb{H} (note that this is possible in type III pairings).*

2.3 Groth-Sahai Proofs

Groth and Sahai [GS08] proposed efficient non-interactive witness indistinguishable proof systems for settings where a bilinear pairing is available. Their system allows efficient proofs of various statements about the groups of the pairing, and the pairing relation.

In this work we prove statements of the form $\mathbf{e}(\mathbf{A}^\top, \mathbf{W}) = T$, where an instance $(\mathbf{A}, T) \in \mathbb{G}^n \times \mathbb{T}$ is the input to the verifier, and \mathbf{W} is the witness used by the prover.

Let $gk = (p, \mathbb{G}, \mathbb{H}, \mathbb{T}, \mathbf{e})$ be a group description and $crs = (\mathbf{G}, \mathbf{H}) \in \mathbb{H}^2$. The Groth-Sahai proof using crs as CRS (Common Reference String) is as follows.

- $\text{Prf}(gk, crs, (T, \mathbf{A}), \mathbf{W})$: Parse gk and crs as $(p, \mathbb{G}, \mathbb{H}, \mathbb{T}, \mathbf{e})$ and (\mathbf{G}, \mathbf{H}) respectively. Select $R \xleftarrow{\$} \mathbb{Z}_p^{n \times 2}$ randomly, and compute

$$\begin{aligned} (\mathbf{C}, \mathbf{D}) &\leftarrow (R \cdot \mathbf{G}, \mathbf{W} + R \cdot \mathbf{H}) \\ \mathbf{\Pi} &\leftarrow R^\top \mathbf{A} \end{aligned}$$

and output $\sigma = (\mathbf{C}, \mathbf{D}, \mathbf{\Pi})$.

- $\text{Vrf}(gk, crs, (\mathbf{A}, T), \sigma)$: Parse gk , crs , and σ as $(p, \mathbb{G}, \mathbb{H}, \mathbb{T}, \mathbf{e})$, (\mathbf{G}, \mathbf{H}) and $(\mathbf{C}, \mathbf{D}, \mathbf{\Pi})$ respectively.
Output 1 iff the following equality holds.

$$(\mathbf{e}(\mathbf{A}^\top, \mathbf{C}), \mathbf{e}(\mathbf{A}^\top, \mathbf{D})) \stackrel{?}{=} (\mathbf{e}(\mathbf{\Pi}^\top, \mathbf{G}), T + \mathbf{e}(\mathbf{\Pi}^\top, \mathbf{H}))$$

One can easily show that a correctly generated proof is always accepted by Vrf.

Groth and Sahai [GS08] gave the following two algorithms HideCRS and BindCRS to generate two different types of CRS: hiding and binding. When a hiding CRS is used for the proof, the witness is perfectly (information theoretically) hidden. When a binding CRS is used, BindCRS provides a trapdoor along with the CRS, and the trapdoor can be used to extract the witness from any proof. Finally, it is required that the two types of CRS are computationally indistinguishable.

For the above proof system, let $gk = (p, \mathbb{G}, \mathbb{H}, \mathbb{T}, \mathbf{e})$. The algorithms HideCRS and BindCRS are defined as follows:

- HideCRS(gk) : Select $\mathbf{G}, \mathbf{H} \xleftarrow{\$} \mathbb{H}^2$ randomly and output $crs \leftarrow (\mathbf{G}, \mathbf{H})$.
- BindCRS(gk) : Select $\mathbf{G} \xleftarrow{\$} \mathbb{H}^2$ and $\alpha \xleftarrow{\$} \mathbb{Z}_p$ randomly, compute $\mathbf{H} \leftarrow \alpha \mathbf{G}$ and output $crs \leftarrow (\mathbf{G}, \mathbf{H})$ and the trapdoor α .

Groth-Sahai show that in the above proof system, a hiding CRS and a binding CRS are indistinguishable under the SXDH assumption. Formally, the perfect witness indistinguishability, and the witness extractability properties are defined as follows [GS08]. Below, $\text{Setup}(1^\kappa)$ be an algorithm which generates a group description $gk = (p, \mathbb{G}, \mathbb{H}, \mathbb{T}, \mathbf{e})$.

- **(Composable Perfect Witness Indistinguishability)** For all (possibly unbounded) adversaries \mathcal{A}

$$\begin{aligned} & \Pr \left[gk \leftarrow \text{Setup}(1^\kappa), crs \leftarrow \text{HideCRS}(gk), \right. \\ & \left. (\mathbf{A}, T, \mathbf{W}_0, \mathbf{W}_1, st) \leftarrow \mathcal{A}(gk, crs), \sigma \leftarrow \text{Prf}(gk, crs, (\mathbf{A}, T), \mathbf{W}_0), b \leftarrow \mathcal{A}(\sigma, st) : b = 1 \right] \\ &= \Pr \left[gk \leftarrow \text{Setup}(1^\kappa), crs \leftarrow \text{HideCRS}(gk), \right. \\ & \left. (\mathbf{A}, T, \mathbf{W}_0, \mathbf{W}_1, st) \leftarrow \mathcal{A}(gk, crs), \sigma \leftarrow \text{Prf}(gk, crs, (\mathbf{A}, T), \mathbf{W}_1), b \leftarrow \mathcal{A}(\sigma, st) : b = 1 \right] \end{aligned}$$

where we require $\mathbf{e}(\mathbf{A}, \mathbf{W}^{[0]}) = \mathbf{e}(\mathbf{A}, \mathbf{W}^{[1]}) = T$.

- **(Perfect Extractability)** For all possible output $gk = (p, \mathbb{G}, \mathbb{H}, \mathbb{T}, \mathbf{e})$ of $\text{Setup}(1^\kappa)$, all possible output (crs, α) of $\text{BindCRS}(gk)$, all $(\mathbf{A}, T) \in \mathbb{G}^n \times \mathbb{T}$ and all $\sigma = (\mathbf{C}, \mathbf{D}, \mathbf{II}) \in \mathbb{H}^2 \times \mathbb{H}^2 \times \mathbb{G}^2$ satisfying $\text{Vrf}(gk, crs, (\mathbf{A}, T), \sigma) = 1$, if we set

$$\mathbf{W}^* \leftarrow \mathbf{D} - \alpha \mathbf{C},$$

the equality $\mathbf{e}(\mathbf{A}, \mathbf{W}^*) = T$ always holds.

3 Independent Preimage Resistant (IPIR) Hash Functions

We introduce a new notion: independent pre-image resistance (IPIR), that we use in the construction and analysis of our scheme. As we have already mentioned in the introduction, our construction does not use the IPIR hash function described below in a black box way. Nevertheless, we believe it to be instructive to define this notion separately, both to conceptually isolate the properties of the hash function that we use, and for potential future use.

Definition 3 (Independent Preimage Resistant Hash Function) Let n be a positive number, and let \mathbb{H} and \mathbb{T} be cyclic groups of order p . Let $\text{Gen}, H, \text{GenSim}, \text{Check}$ be polynomial time algorithms whose inputs and outputs are as follows. Below, k is the parameter representing the dimension. Note that a k -dimensional *affine* (not linear) subspace contains $k + 1$ vectors and GenSim of the below therefore outputs $k + 1$ vectors \mathcal{Y}_j .

- $\text{Gen}(1^\kappa)$ outputs some public information P .
- For any possible output P of Gen , H_P is a deterministic linear function from \mathbb{H}^n to \mathbb{T} .
- GenSim takes an integer $k \in [n - 2]$ as an input and outputs a public information P , a trapdoor td , $(T, \mathbf{Y}_1, \dots, \mathbf{Y}_{k+1}) \in \mathbb{T} \times (\mathbb{H}^n)^{k+1}$ satisfying $T = H_P(\mathbf{Y}_i)$ for all $i \in [k + 1]$.

- Check takes $P, (T, \mathbf{Y}_1, \dots, \mathbf{Y}_{k+1})$, an element \mathbf{Y}' of \mathbb{H}^n , and a trapdoor td as inputs and outputs 1 or 0.

For integers n and $k \in [n - 2]$, we say that H is (n, k) -independent pre-image resistant with respect to $(\text{Gen}, \text{GenSim}, \text{Check})$ if it satisfies the following properties.

- **(Correctness)** For all outputs $(P, td, (T, \mathbf{Y}_1, \dots, \mathbf{Y}_{k+1}))$ of GenSim and all $\mathbf{Y}' \in \mathbb{H}^n$, $\text{Check}(P, td, (T, \mathbf{Y}_1, \dots, \mathbf{Y}_{k+1}), \mathbf{Y}') = 1$ holds iff $T = H_P(\mathbf{Y}')$ and $\mathbf{Y}' \in \text{Aff}(\mathbf{Y}_1, \dots, \mathbf{Y}_{k+1})$ holds. Moreover, for an output $(P, td, (T, \mathbf{Y}_1, \dots, \mathbf{Y}_{k+1}))$ of GenSim , P have the same distribution as an output of $\text{Gen}(1^\kappa)$ and $(\mathbf{Y}_1, \dots, \mathbf{Y}_{k+1})$ uniformly distributed on $\{\mathbf{Y} \in \mathbb{H}^n \mid H_P(\mathbf{Y}) = T\}^{k+1}$.
- **(n, k) -Independent Preimage Resistance** For any polytime adversary \mathcal{A} ,

$$\Pr \left[\begin{array}{l} (P, td, (T, \mathbf{Y}_1, \dots, \mathbf{Y}_{k+1})) \leftarrow \text{GenSim}(1^\kappa), \quad T = H_P(\mathbf{Y}_*) \\ \mathbf{Y}_* \leftarrow \mathcal{A}(P, T, \mathbf{Y}_1, \dots, \mathbf{Y}_{k+1}) \quad : \quad \mathbf{Y}_* \notin \text{Aff}(\mathbf{Y}_1, \dots, \mathbf{Y}_{k+1}) \end{array} \right]$$

is negligible.

Note that one can check whether $\mathbf{Y}_* \notin \text{Aff}(\mathbf{Y}_1, \dots, \mathbf{Y}_{k+1})$ holds or not in polytime by using Check and the a trapdoor td .

Construction of an independent pre-image resistant function. In our signature scheme we use the function $H_{\mathbf{A}}(\mathbf{Y}) = \mathbf{e}(\mathbf{A}, \mathbf{Y})$. The algorithms Gen , GenSim , and Check for the function H are as follows. Below, $\text{Setup}(1^\kappa)$ is an algorithm which generates a group description $gk = (p, \mathbb{G}, \mathbb{H}, \mathbb{T}, \mathbf{e})$.

- $\text{Gen}(1^\kappa)$: Compute $gk \leftarrow \text{Setup}(1^\kappa)$ and generates $\mathbf{A} \xleftarrow{\$} \mathbb{G}^n$ randomly, and output $P \leftarrow (gk, \mathbf{A})$.
- $\text{GenSim}(1^\kappa)$: Compute $gk \leftarrow \text{Setup}(1^\kappa)$, and choose randomly $A \xleftarrow{\$} \mathbb{G}$, and $\mathbf{a} \xleftarrow{\$} \mathbb{Z}_p^n$. Compute $\mathbf{A} \leftarrow \mathbf{a}A$. Choose randomly $\mathbf{Y} \xleftarrow{\$} \mathbb{H}$, $t \xleftarrow{\$} \mathbb{Z}_p$, and $\mathbf{y}_j \in \mathbb{Z}_p^n$ satisfying $\langle \mathbf{a}, \mathbf{y}_j \rangle = t$ for $j \in [k + 1]$. Choose $n - k$ linearly independent vectors $\mathbf{e}_1, \dots, \mathbf{e}_{n-k}$ satisfying $\langle \mathbf{e}_i, \mathbf{y}_j \rangle = 0$ for all $i \in [n - k]$, $j \in [k]$. Output $P = (gk, \mathbf{A})$, $td \leftarrow (\mathbf{e}_i)_{i \in [n-k]}$, and $\mathbf{Y}_j \leftarrow \mathbf{y}_j \mathbf{Y}$ for $j \in [k]$, $T \leftarrow t\mathbf{Y}$.
- $\text{Check}(P, td, (T, \mathbf{Y}_1, \dots, \mathbf{Y}_{k+1}), \mathbf{Y}')$: Parse P and td as (gk, \mathbf{A}) and $(\mathbf{e}_i)_{i \in [n-k]}$ respectively. Output 1 iff $\mathbf{e}(\mathbf{A}, \mathbf{Y}') = T$ and $\langle \mathbf{e}_i, \mathbf{Y}' \rangle = 0$ holds for all $i \in [n - k]$.

Proposition 4 For any n and $k \leq n - 2$, the scheme $(\text{Setup}, H, \text{GenSim}, \text{Check})$ is (n, k) -independent pre-image resistant under the SXDH assumption.

Correctness is straightforward. We give the proof of independent pre-image resistance in the full paper.

4 Proposed Scheme

Let $n \geq 3$ and m be integers. Let Setup be a polytime algorithm that generates a group description $gk = (p, \mathbb{G}, \mathbb{H}, \mathbb{T}, \mathbf{e})$, as discussed above, where $\mathbf{e} : \mathbb{G} \times \mathbb{H} \rightarrow \mathbb{T}$. For $\mathcal{H} = (\mathbf{H}_0, \mathbf{H}_1, \dots, \mathbf{H}_m) \in (\mathbb{H}^2)^{m+1}$ and $M \in \{0, 1\}^m$, we define a Water's hash function h as

$$h_{gk}(\mathcal{H}, M) = \mathbf{H}_0 + \sum_{k \in [m]} M_k \mathbf{H}_k,$$

where M_k is the k -th bit of M . Let Prf and Vrf be the proof algorithm and the verification algorithm of the Groth-Sahai proof system reviewed in Section 2.3. Our signature scheme $\mathcal{SIG} = (\text{Kg}, \text{Update}, \text{Sig}, \text{Ver})$ works as follows.

Key Generation $\text{Kg}(1^\kappa)$: $gk \leftarrow (p, \mathbb{G}, \mathbb{H}, \mathbb{T}, \mathbf{e}) \leftarrow \text{Setup}(1^\kappa)$, $\mathbf{G} \leftarrow \mathbb{H}^2$, $\mathcal{H} \leftarrow (\mathbf{H}_0, \mathbf{H}_1, \dots, \mathbf{H}_m) \leftarrow (\mathbb{H}^2)^{m+1}$.

Randomly select $\mathbf{A} \xleftarrow{\$} \mathbb{G}$, $\mathbf{Q} \xleftarrow{\$} \mathbb{H}$, and $\mathbf{a}, \mathbf{q} \xleftarrow{\$} \mathbb{Z}_p^n$ satisfying $\langle \mathbf{a}, \mathbf{q} \rangle = 0$ and compute $\mathbf{A} \leftarrow \mathbf{a}\mathbf{A}$, $\mathbf{Q} \leftarrow \mathbf{q}\mathbf{Q}$. Select $\mathbf{W}^{[0]} \xleftarrow{\$} \mathbb{H}^n$ randomly, compute $T \leftarrow \mathbf{e}(\mathbf{A}, \mathbf{W}^{[0]})$, and outputs $pk \leftarrow (gk, \mathbf{G}, \mathcal{H}, \mathbf{A}, T, \mathbf{Q})$ and $sk^{[0]} \leftarrow \mathbf{W}^{[0]}$.

Key Update $\text{Update}_{pk}(sk^{[i]})$: Parse pk and $sk^{[i]}$ as $(gk, \mathbf{G}, \mathcal{H}, \mathbf{A}, T, \mathbf{Q})$ and $\mathbf{W}^{[i]}$ respectively, select $s \xleftarrow{\$} \mathbb{Z}_p$ randomly, and output $sk^{[i+1]} \leftarrow \mathbf{W}^{[i+1]} \leftarrow \mathbf{W}^{[i]} + s\mathbf{Q}$.

Signing $\text{Sig}(sk^{[i]}, M)$ for $M \in \{0, 1\}^m$: Parse pk and $sk^{[i]}$ as $(gk, \mathbf{G}, \mathcal{H}, \mathbf{A}, T, \mathbf{Q})$ and $\mathbf{W}^{[i]}$. Compute $\mathbf{H}_M \leftarrow h_{gk}(\mathcal{H}, M)$, set $crs_M \leftarrow (\mathbf{G}, \mathbf{H}_M)$, and $\sigma \leftarrow \text{Prf}(gk, crs_M, (\mathbf{A}, T), \mathbf{W}^{[i]})$ and output σ .

Verification $\text{Ver}(pk, M, \sigma)$: Parse pk as $(gk, \mathbf{G}, \mathcal{H}, \mathbf{A}, T, \mathbf{Q})$, compute $\mathbf{H}_M \leftarrow h_{gk}(\mathcal{H}, M)$, and set $crs_M \leftarrow (\mathbf{G}, \mathbf{H}_M)$. If $\text{Ver}(gk, crs_M, (\mathbf{A}, T), \sigma) = 1$, output 1. Otherwise, output 0.

Theorem 5. *For any constants $c > 0$ and any $\gamma = \Theta(1/\sqrt{\kappa})$, the proposed scheme \mathcal{SIG} is $(\rho_G, \rho_U, \rho_M, \rho_S)$ -EU-CMA-CML secure under the SXDH assumption. Here*

$$(\rho_G, \rho_U, \rho_M, \rho_S) = \left(\frac{c \cdot \log k}{n \log p}, \frac{c \cdot \log k}{n \log p}, 1 - \frac{2 + \gamma}{n}, 1 - \frac{2 + \gamma}{n} \right).$$

We can achieve the fraction $1 - o(1)$ of leakage in signing and in memory by setting $n = \kappa$.

4.1 Overview of Security Analysis

Our proof starts with a reduction that shows how to convert any adversary that obtains leakage on key generation and updates, to an adversary that does not require such leakage. This follows from the lemma of Brakerski *et al* [BKKV10]. (See our full paper for the proof.)

5 Conclusion

In this work, we propose a signature scheme that protects against (1) continual memory leakage combined with (2) all types of continual processing (computational) leakage, namely leakage from key generation, key updates, and signing. Our scheme remains secure even when the leakage during signing is a function $f(sk, r)$ of both the secret key and the randomness.

The security of our scheme is proven in the standard model. Moreover, the amount of information that our scheme is allowed to leak during each time period is optimal, in the sense that our scheme remains secure even if $1 - o(1)$ fraction of the secret key of each time period is leaked.

References

- [AFGKHO10] Masayuki Abe, Georg Fuchsbauer, Jens Groth, Kristiyan Haralambiev, Miyako Ohkubo: Structure-Preserving Signatures and Commitments to Group Elements. CRYPTO 2010: pp.209-236.
- [AGV09] Adi Akavia, Shafi Goldwasser, Vinod Vaikuntanathan: Simultaneous Hardcore Bits and Cryptography against Memory Attacks. TCC 2009: 474-495
- [ADW09] Joël Alwen, Yevgeniy Dodis, Daniel Wichs: Leakage-Resilient Public-Key Cryptography in the Bounded-Retrieval Model. CRYPTO 2009: pp.36-54
- [BSW10] Elett Boyle, Gil Segev, Daniel Wichs. Fully Leakage-Resilient Signatures. eprint archive. (2010/488), 2010.
- [BB05] David Brumley and Dan Boneh. Remote timing attacks are practical. Computer Networks, 48(5): 2005 pp.701–716,
- [BR09] Mihir Bellare, Thomas Ristenpart: Simulation: without the Artificial Abort: Simplified Proof and Improved Concrete Security for Waters' IBE Scheme. EUROCRYPT 2009, pp.407-424.
- [BFO08] Alexandra Boldyreva, Serge Fehr, Adam O'Neill: On Notions of Security for Deterministic Encryption, and Efficient Constructions without Random Oracles. CRYPTO 2008: pp.335-359
- [BSS99] Ian F. Blake, Gadiel Seroussi, Nigel Paul Smart: Elliptic Curves in Cryptography, London Mathematical Society 265, Cambridge University Press, 1999.
- [BB04] Dan Boneh, Xavier Boyen: Secure Identity Based Encryption Without Random Oracles. CRYPTO 2004: pp.443-459
- [BG10] Zvika Brakerski and Shafi Goldwasser. Circular and Leakage Resilient Public-Key Encryption under Subgroup Indistinguishability - (or: Quadratic Residuosity Strikes Back). CRYPTO 2010 pp.1-20.
- [BKKV10] Zvika Brakerski, Yael Tauman Kalai, Jonathan Katz, Vinod Vaikuntanathan. Overcoming the Hole in the Bucket: Public-Key Cryptography Resilient to Continual Memory Leakage. FOCS 2010.
- [DKL09] Yevgeniy Dodis and Yael Tauman Kalai and Shachar Lovett. On cryptography with auxiliary input. STOC 2009. pp.621-630
- [DHLW10] Yevgeniy Dodis, Kristiyan Haralambiev, Adriana Lopez-Alt, Daniel Wichs. Cryptography Against Continuous Memory Attacks. FOCS 2010.
- [DHLW10b] Yevgeniy Dodis, Kristiyan Haralambiev, Adriana Lopez-Alt, Daniel Wichs. Efficient Public-Key Cryptography in the Presence of Key Leakage. Cryptology ePrint Archive, Report 2010/154.

- [DGK+10] Yevgeniy Dodis and Shafi Goldwasser and Yael Tauman Kalai and Chris Peikert and Vinod Vaikuntanathan. Public-Key Encryption Schemes with Auxiliary Inputs. TCC 2010: pp.361-381
- [DP08] Stefan Dziembowski, Krzysztof Pietrzak: Leakage-Resilient Cryptography. FOCS 2008: 293-302
- [DS05] Yevgeniy Dodis, Adam Smith: Correcting errors without leaking partial information. STOC 2005: pp.654-663
- [DP10] Yevgeniy Dodis, Krzysztof Pietrzak: Leakage-Resilient Pseudorandom Functions and Side-Channel Attacks on Feistel Networks. CRYPTO 2010: pp.21-40
- [FKPR10] Sebastian Faust, Eike Kiltz, Krzysztof Pietrzak, Guy N. Rothblum: Leakage-Resilient Signatures. TCC 2010: pp.343-360
- [FRR+10] Sebastian Faust, Tal Rabin, Leonid Reyzin, Eran Tromer, Vinod Vaikuntanathan: Protecting Circuits from Leakage: the Computationally-Bounded and Noisy Cases. EUROCRYPT 2010: 135-156
- [FS86] Amos Fiat and Adi Shamir: How to Prove Yourself: Practical Solutions to Identification and Signature Problems. CRYPTO 1986: pp. 186-194
- [FLS90] Uriel Feige, Dror Lapidot, Adi Shamir: Multiple Non-Interactive Zero Knowledge Proofs Based on a Single Random String (Extended Abstract) FOCS 1990: pp.308-317
- [GPS08] Steven D. Galbraith, Kenneth G. Paterson, Nigel P. Smart: Pairings for cryptographers. Discrete Applied Mathematics 156(16): pp.3113-3121, 2008
- [GR10] Shafi Goldwasser, Guy N. Rothblum: Securing Computation against Continuous Leakage. CRYPTO 2010: pp.59-79
- [GSW10] Essam Ghadafi, Nigel P. Smart, Bogdan Warinschi: Groth-Sahai Proofs Revisited. Public Key Cryptography 2010: pp.177-192
- [GS08] Jens Groth, Amit Sahai: Efficient Non-interactive Proof Systems for Bilinear Groups. EUROCRYPT 2008: 415-432
- [HSH+08] J. Alex Halderman, Seth D. Schoen, Nadia Heninger, William Clarkson, William Paul, Joseph A. Calandrino, Ariel J. Feldman, Jacob Appelbaum, Edward W. Felten: Lest We Remember: Cold Boot Attacks on Encryption Keys. USENIX Security Symposium 2008: pp.45-60
- [JV10] Ali Juma, Yevgeniy Vahlis: Protecting Cryptographic Keys against Continual Leakage. CRYPTO 2010: pp.41-58
- [ISW03] Yuval Ishai, Amit Sahai, David Wagner: Private Circuits: Securing Hardware against Probing Attacks. CRYPTO 2003: 463-481
- [KJJ98] Paul Kocher and Joshua Jaffe and Benjamin Jun. Introduction to Differential Power Analysis and Related Attacks. 1998.
<http://www.cryptography.com/dpa/technical/>
- [KJJ99] Paul C. Kocher and Joshua Jaffe and Benjamin Jun. Differential Power Analysis. CRYPTO 1999: pp.388-397
- [KV09] Jonathan Katz, Vinod Vaikuntanathan: Signature Schemes with Bounded Leakage Resilience. ASIACRYPT 2009: pp.703-720
- [MR04] Silvio Micali, Leonid Reyzin: Physically Observable Cryptography (Extended Abstract). TCC 2004: 278-296
- [NS09] Moni Naor and Gil Segev. Public-Key Cryptosystems Resilient to Key Leakage. CRYPTO 2009. pp18-35.
- [P09] Krzysztof Pietrzak: A Leakage-Resilient Mode of Operation. EUROCRYPT 2009: pp.462-482
- [QS01] Jean-Jacques Quisquater, David Samyde: ElectroMagnetic Analysis (EMA): Measures and Counter-Measures for Smart Cards. E-SMART '01: Proceedings of the International Conference on Research in Smart Cards 2001: pp.200-210

- [R97] Ronald L. Rivest: All-or-Nothing Encryption and the Package Transform. FSE 1997: pp.210-218
- [S79] Adi Shamir: How to Share a Secret. Commun. ACM 22(11): 612-613 (1979)
- [SMY09] François-Xavier Standaert and Tal Malkin and Moti Yung. A Unified Framework for the Analysis of Side-Channel Key Recovery Attacks. EUROCRYPT 2009: pp.443-461
- [W05] Brent Waters: Efficient Identity-Based Encryption Without Random Oracles. EUROCRYPT 2005: pp.114-127