Limitations of the Meta-Reduction Technique: The Case of Schnorr Signatures

Marc Fischlin[1] and Nils Fleischhacker[2]

[1] Technische Universität Darmstadt
www.cryptoplexity.de
[2] Saarland University
ca.cs.uni-saarland.de

Abstract. We revisit the security of Fiat-Shamir signatures in the non-programmable random oracle model. The well-known proof by Pointcheval and Stern for such signature schemes (Journal of Cryptology, 2000) relies on the ability to re-program the random oracle, and it has been unknown if this property is inherent. Pailler and Vergnaud (Asiacrypt 2005) gave some first evidence of the hardness by showing via meta-reduction techniques that algebraic reductions cannot succeed in reducing key-only attacks against unforgeability to the discrete-log assumptions. We also use meta-reductions to show that the security of Schnorr signatures cannot be proven equivalent to the discrete logarithm problem without programming the random oracle. Our result also holds under the one-more discrete logarithm assumption but applies to a large class of reductions, we call *single-instance* reductions, subsuming those used in previous proofs of security in the (programmable) random oracle model. In contrast to algebraic reductions, our class allows arbitrary operations, but can only invoke a single resettable adversary instance, making our class incomparable to algebraic reductions.

Our main result, however, is about meta-reductions and the question if this technique can be used to further strengthen the separations above. Our answer is negative. We present, to the best of our knowledge for the first time, limitations of the meta-reduction technique in the sense that finding a meta-reduction for general reductions is most likely infeasible. In fact, we prove that finding a meta-reduction against a potential reduction is equivalent to finding a "meta-meta-reduction" against the strong existential unforgeability of the signature scheme. This means that the existence of a meta-reduction implies that the scheme must be insecure (against a slightly stronger attack) in the first place.

1 Introduction

On a technical level, we investigate the security of Fiat-Shamir (FS) signatures [10] in the non-programmable random oracle model (NPROM), i.e., where programming the hash function is prohibited. Such programming has been exploited in the security proof for common FS signatures by Pointcheval and Stern [23], bringing forward the question if the security result remains valid in the more

stringent model of non-programmable random oracles. Conceptually, though, the more interesting result in the paper refers to limitations of so-called meta-reductions. Such meta-reductions are also called "reductions against the reductions" as they basically treat the reduction as an adversary itself and reduce the existence of such a reduction to a presumably hard problem, ruling out reductions and therefore security proofs for the underlying scheme. This proof technique recently gained quite some attention as it rules out certain reductions, especially those which only treat the adversary but not the underlying primitive as a black box (e.g., [9,20,15,14,21,16,2,27]). We show, via a "meta-meta-reduction", that one cannot use the meta-reduction technique to show impossibility results for FS signatures in the NPROM.

1.1 Fiat-Shamir Signatures in the NPROM

The class of FS signatures comprises all transformed three-move identification schemes in which the challenge CH, sent by the verifier in return of the prover's initial commitment COM, is replaced by the hash value $H(\text{COM}, m)$ for message m. The prover's response RESP, together with COM, then yields the signature (COM, RESP) for m. For some cases, like the Schnorr signature scheme [26], the signature can be shortened by using (CH, RESP) instead.

The common security proof for FS signature schemes in the random oracle model [6], given in [23], basically works as follows. The reduction to the underlying problem, such as the discrete logarithm problem for the Schnorr scheme, runs the adversary twice. In the first runs the reduction gets a signature forgery (COM, RESP) for message m and challenge $\text{CH} = H(\text{COM}, m)$. In the second run it re-programs H to yield a distinct challenge $\text{CH}' = H(\text{COM}, m)$ and response RESP'. From both signatures the reduction can then compute a solution to the underlying problem. Clearly, this technique relies on the programmability of the hash function.[3]

Fischlin et al. [13] later defined reductions in the non-programmable random oracle model (NPROM) by externalizing the hash function to both the adversary *and* the reduction.[4] In the NPROM the reduction may still observe the adversary's queries to the hash function, but cannot change the reply. Obviously, this non-programming property matches much closer our understanding of "real" hash functions and instantiations through, say, SHA-3. Interestingly, though, Fischlin et al. [13] do not investigate this arguably most prominent application of the random oracle methodology. Instead, they separate programming and non-programming reductions (and an intermediate notion called weakly programming reductions) through the case of OAEP encryption, FDH signatures,

---

[3] Note that this proof reduces the security of the signature scheme to the underlying number-theoretic problem via special soundness. Abdalla et al. [1] more generally consider FS schemes with reductions to the identification schemes. We do not cover the latter type of reductions and schemes here.

[4] The role of programmability was first investigated by Nielsen [18], even though not for reductions as in the proofs of Fiat-Shamir schemes.

and trapdoor permutation based KEMs. Weakly programming reductions are allowed to reset the random oracle and redirect the value to some (external) random answer. Our first result is to formally confirm the intuition that FS signatures should still be secure in the weakly programmable random oracle model.

1.2 Limitations Through Meta-Reductions

The more interesting question is if FS signatures can be shown to be secure in the NPROM. Our first result in this regard is negative and applies to discrete log schemes like the Schnorr signature scheme [26] or the RSA-based Guillou-Quisquater scheme [17]. Namely, we first consider any reduction $\mathcal{R}$ which initiates only a single (black-box) instance of the adversary $\mathcal{A}$ for some public key pk, but such that it can reset $\mathcal{A}$ arbitrarily to the point after having handed over the public key (from the fixed group). Note that the reduction in the *programmable* ROM in [23] is of this kind, only that it can also change the behavior of the random oracle, unlike our reductions here in the NPROM. We show that this type of *single-instance* reduction to the discrete log problem cannot succeed in the NPROM under the one-more discrete log assumption [5].

Our impossibility result follows from presenting a meta-reduction $\mathcal{M}$ against $\mathcal{R}$. That is, we show that if one can find a reduction $\mathcal{R}$ which successfully solves the DL problem given black-box access to any successful adversary $\mathcal{A}$ against the signature scheme, then there is a meta-reduction $\mathcal{M}$ breaking the one-more DL problem directly. Since we also present a successful (unbounded) adversary $\mathcal{A}$ which $\mathcal{M}$ can simulate towards $\mathcal{R}$ efficiently, we conclude that the existence of reduction $\mathcal{R}$ would already contradict the one-more DL problem.

We observe that our meta-reduction, too, works in the NPROM and thus cannot program the random oracle for $\mathcal{R}$; else the meta-reduction would violate the idea of modeling hash functions as non-programmable. It is also easy to show that, if the meta-reduction, unlike the reduction, was allowed to program the random oracle, this "unfair" situation would straightforwardly dismiss the possibility of such reductions. However, such an approach seems to violate the idea behind non-programmable oracles as a mean to capture real-world hash functions over which no party, not even the meta-reduction, has control.

The noteworthy property of our meta-reduction $\mathcal{M}$ is that, unlike most of the previous proposals (cf. [11]), it does not work by resetting the reduction $\mathcal{R}$. The reset strategy is usually used to rewind the reduction and, in case of signature schemes, get an additional signature through a signing query in an execution branch, and display this signature back to $\mathcal{R}$ as a forgery in the main branch. However, this means that one needs to take care of correlations between the additional signature and the reduction's state. Instead of using such resets, our meta-reduction will essentially run two independent copies of the reduction and use the signatures of one execution in the other one. The independence of the executions thus "decorrelates" the additional signature from the reduction's state, avoiding many complications from the resetting strategy.

1.3 Limitations of Meta-Reductions

Does our meta-reduction impossibility result for non-programming reductions extend to other cases like the discrete logarithm problem? We show that this is unlikely, thus showing limitations of the meta-reduction technique. The idea is to consider the meta-reduction itself as a reduction, and to use the meta-reduction technique against this reduction. Hence, we obtain a "meta-meta-reduction" $\mathcal{N}$ which now simulates the reduction $\mathcal{R}$ for $\mathcal{M}$, just as the meta-reduction simulates the adversary for $\mathcal{R}$.

More concretely, assume that we consider reductions $\mathcal{R}$ transforming an adversary $\mathcal{A}$ against the signature scheme in a black-box way into a solver for some cryptographic problem $\Pi_{\mathcal{R}}$. Then, a meta-reduction $\mathcal{M}$ should turn $\mathcal{R}$ (to which it has black-box access) into a successful solver for some problem $\Pi_{\mathcal{M}}$. For technical reasons, in our case this problem $\Pi_{\mathcal{M}}$ has to be non-interactive, e.g., correspond to the discrete logarithm problem; this also circumvents the case of our previous meta-reduction for the interactive one-more DL problem. Then we show that such a meta-reduction can be used to build a meta-meta-reduction $\mathcal{N}$ against the strong unforgeabilty of the signature scheme.

In other words, the meta-reduction technique cannot help to rule out black-box reductions to arbitrary problems, unless the signature is insecure in the first place. Here, insecurity refers to the notion of strong unforgeability where the adversary also succeeds by outputting a new signature to a previously signed message. In fact, in the programmable ROM the security proof in [23] actually shows that the FS schemes achieve this stronger notion.

1.4 Related Work

As mentioned before, meta-reductions have been used in several recent results to rule out black-box reductions for Fiat-Shamir schemes, and especially for Schnorr signatures. Paillier and Vergnaud [20] analyzed the security of Schnorr Signatures in the standard model. They showed with the help of meta-reductions that, if the one-more discrete logarithm assumption holds, the security of Schnorr signatures cannot be reduced to the (one-more) discrete logarithm problem, at least using algebraic reductions. While algebraic reductions where first defined by Boneh and Venkatesan [7], Paillier and Vergnaud [20], however, use a slightly more liberal definition of algebraicity. Their notion basically states that, given the discrete logarithm of all of the reduction's inputs and access to the reduction, it is possible to compute the discrete logarithm of any group element output by the reduction. We note that the ability to trace the discrete logarithms of the group elements produced by the reduction is important to their result and allows them to prove impossibility even for key-only attacks.

Paillier and Vergnaud [20] also extended their result to other signature schemes, including the Guillou-Quisquater scheme [17] and the one-more RSA assumption [5]. They also considered the tightness loss in the Pointcheval-Stern proof for the Schnorr signature scheme in the programmable random oracle model. They showed, again for algebraic reductions, that the security loss of a factor

$\sqrt{q_H}$ is inevitable, where q_H is the maximum number of random oracle queries by the adversary. This bound was later raised to $q_H^{2/3}$ by Garg et al. [15] in the same setting. Seurin [27] recently improved this bound further to $\mathcal{O}(q_H)$. Using meta-reduction techniques and considering algebraic reductions, too, he proved it is unlikely that a tighter reduction exists.

In a recent work, Baldimtsi and Lysyanskaya [4] showed, via meta-reductions, that one cannot prove blind Schnorr signatures secure via black-box reductions. Their meta-reduction, like ours here, has the interesting feature of being non-resetting. Remarkably, though, they seem to rule out the more liberal *programming* reductions, whereas our result is against the "more confined" *non-programming* reductions. However, their result considers a special type of programming reduction, called naive. This roughly means that one can predict the reduction's programmed random oracle answers by reading the reduction's random tape. This property is inherently tied to the programmability and is clearly not fulfilled by non-programmable, external random oracles; for such oracles even the reduction does not know the answers in advance. This, unfortunately, also means that their meta-reduction technique may not apply to non-programmable hash functions. In other words, one may be able to bypass their impossibility result and may still be able to find a cryptographic security proof for such schemes, by switching to the non-programmable random oracle model, or even to standard-model hash functions.

1.5   Organization

In Section 2 we first recall some basic facts about signatures and (general and discrete-log specific) cryptographic problems. Then we show that FS signatures are secure in the weakly programmable random oracle model, and prove our meta-reduction impossiblity result for single-instance reductions in the NPROM in Section 3. Our main result about meta-meta-reductions appears in Section 4.

2   Preliminaries

We use standard notions for digital signature schemes $\mathcal{S} = (\mathsf{KGen}, \mathsf{Sign}, \mathsf{Vrfy})$ such as existential unforgeability and strong existential unforgeability. We usually assume (non-trivially) randomized signature schemes, where the signature algorithm has super-logarithmic min-entropy for the security parameter κ, i.e., $H_\infty(\mathsf{Sign}(\mathsf{sk}, m)) \in \omega(\log(\kappa))$ for all keys sk, all messages m, and given the random oracle. For formal definitions refer to the full version of this paper.

2.1   Cryptographic Problems

We define a cryptographic problem as a game between a challenger and an adversary. The challenger uses an instance generator to generate a fresh instance of the problem. The adversary is then supposed to find a solution for said instance. The challenger may assist the adversary by providing access to some oracle, like

a decryption oracle in a chosen-ciphertext attack against indistinguishability. Eventually the adversary outputs a solution for the problem instance and the challenger uses a verification algorithm to check whether the solution is correct.

For many problems there exist trivial adversaries, e.g., succeeding in an indistinguishability game by pure guessing. One is usually interested in the advantage of adversaries beyond such trivial strategies. We therefore introduce a so-called *threshold algorithm* to cover such trivial attacks and measure any adversary against this threshold adversary.

Definition 1 (Cryptographic Problem). *A cryptographic problem $\Pi = $ (IGen, Orcl, Vrfy, Thresh) consists of four algorithms:*

- *The instance generator IGen takes as input the security parameter 1^κ and outputs a problem instance z. The set of all possible instances output by IGen is called **Inst**.*
- *The computationally unbounded and stateful oracle algorithm Orcl takes as input a query $q \in \{0,1\}^*$ and outputs a response $r \in \{0,1\}^*$ or a special symbol $\bot$ indicating that q was not a valid query.*
- *The deterministic verification algorithm Vrfy takes as input a problem instance $z \in \mathbf{Inst}$ and a candidate solution $x \in \mathbf{Sol}$. The algorithm outputs $b \in \{0,1\}$. We say x is a valid solution to instance z if and only if $b \stackrel{?}{=} 1$.*
- *The efficient threshold algorithm Thresh takes as input a problem instance z and outputs some x. The threshold algorithm is a special adversary and as such also has access to Orcl.*

We note that the algorithms IGen, Orcl, Vrfy potentially have access to shared state that persists for the duration of an experiment.

Definition 2 (Hard Cryptographic Problem). *For a cryptographic problem $\Pi = $ (IGen, Orcl, Vrfy, Thresh) and an adversary $\mathcal{A}$ we define the following experiment:*

$$\mathsf{Exp}_\Pi^\mathcal{A}(\kappa) : [z \leftarrow \mathsf{IGen}(1^\kappa); x \leftarrow \mathcal{A}^\mathsf{Orcl}(z); b \leftarrow \mathsf{Vrfy}(z,x); \text{output } b].$$

The problem Π is said to be hard *if and only if for all probabilistic polynomial-time algorithms $\mathcal{A}$ the following advantage function is negligible in the security parameter κ:*

$$\mathsf{Adv}_\Pi^\mathcal{A}(\kappa) = \Pr\left[\mathsf{Exp}_\Pi^\mathcal{A}(\kappa) \stackrel{?}{=} 1\right] - \Pr\left[\mathsf{Exp}_\Pi^\mathsf{Thresh}(\kappa) \stackrel{?}{=} 1\right],$$

where the probability is taken over the random tapes of IGen and $\mathcal{A}$.

We sometimes require some additional properties of cryptographic problems, summarized in the following definition:

Definition 3 (Specific Cryptographic Problems). *Let $\Pi = $ (IGen, Orcl, Vrfy, Thresh) be a cryptographic problem as defined in Definition 1.*

- *The problem Π is said to be* non-interactive *if and only if Π.Orcl is the algorithm that always outputs $\bot$ and never changes the shared state.*
- *The problem Π is said to be* efficiently generatable *if and only if Π.IGen is a polynomial-time algorithm.*
- *The problem Π is said to be* solvable *if and only if Π.**Sol** is recursively enumerable, and the following holds:*

$$\forall z \leftarrow \Pi.\mathsf{IGen}(1^\kappa) : (\exists x \in \Pi.\mathbf{Sol} : \Pi.\mathsf{Vrfy}(z,x) \stackrel{?}{=} 1).$$

- *The problem Π is said to be* monotone *if and only if for all instances $z \leftarrow \Pi.\mathsf{IGen}(1^\kappa)$, all solutions $x \in \Pi.\mathbf{Sol}$, all $n \in \mathbb{N}$, and all sequences of queries $(q_1, \ldots, q_n)$ the following holds: If $\Pi.\mathsf{Vrfy}(z,x) \stackrel{?}{=} 1$ holds after executing the queries $\Pi.\mathsf{Orcl}(q_1); \ldots; \Pi.\mathsf{Orcl}(q_n)$, then this already held before $\Pi.\mathsf{Orcl}(q_n)$ was executed.*

Intuitively, an algorithm solving a monotone problem is not punished for issuing fewer queries. In particular, if a solution is valid after some sequence of queries, it is also valid if no queries were executed at all.

2.2 Discrete Logarithm Assumptions

The discrete logarithm problem with its corresponding hardness assumption is a specific instance of a non-interactive, efficiently generatable, and solvable problem. The assumption about the computational infeasibility of computing logarithms in certain groups is formally defined in Definition 4.

Definition 4 (Discrete Logarithm Assumption). *Let $\mathbb{G} = \langle g \rangle$ be a group of prime order q with $|q| = \kappa$. The* discrete logarithm (DL) *problem over $\mathbb{G}$ —written $\mathsf{DL}_\mathbb{G}$— is defined as follows:*

Instance and Solution space: *The instance space **Inst** is $\mathbb{G}$ and the solution space **Sol** is $\mathbb{Z}_q$.*

Instance Generation: *The instance generator $\mathsf{IGen}(1^\kappa)$ chooses $z \stackrel{\$}{\leftarrow} \mathbb{G}$ and outputs z. Note that this sampling of z may require to pick a random $w \stackrel{\$}{\leftarrow} \mathbb{Z}_q$ and compute $z = g^w$.*

Verification: *The verification algorithm $\mathsf{Vrfy}(z,x)$ computes $z' = g^x$. If $z' \stackrel{?}{=} z$, then it outputs 1, otherwise it outputs 0.*

Threshold: *The threshold algorithm $\mathsf{Thresh}(z)$ chooses $x \stackrel{\$}{\leftarrow} \mathbb{Z}_q$ and outputs x.*

The discrete logarithm *assumption is said to hold over $\mathbb{G}$ if $\mathsf{DL}_\mathbb{G}$ is hard.*

A natural extension of the discrete logarithm problem are the interactive, efficiently generatable, monotone, and solvable *one-more discrete logarithm* problems first introduced by Bellare et al. [5]. They are interactive, as the adversary is given access to an oracle capable of solving the $\mathsf{DL}_\mathbb{G}$ problem. However, an adversary computing $n+1$ discrete logarithms can only request at most n discrete logarithms from the DL oracle, hence, the name *one-more* discrete-log problem. The problems with their corresponding hardness assumptions are formally described in Definition 5. The assumptions are believed to be stronger than the regular DL assumption [8].

Definition 5 (n-One-More Discrete Logarithm Assumption [5]). *Let $\mathbb{G} = \langle g \rangle$ be a group of prime order q with $|q| = \kappa$. The n-one-more discrete logarithm (n-DL) problem over $\mathbb{G}$ –written n-$\mathsf{DL}_{\mathbb{G}}$– is defined as follows:*

Instance and Solution space: *The instance space* **Inst** *is $\mathbb{G}^{n+1}$ and the solution space* **Sol** *is $\mathbb{Z}_q^{n+1}$.*

Shared State: *The shared state consists only of a single counter variable i.*

Instance Generation: *The instance generator $\mathsf{IGen}(1^\kappa)$ initializes $i := 0$ in the shared state, chooses $z_0, \dots, z_n \xleftarrow{\$} \mathbb{G}$, and outputs $(z_0, \dots, z_n)$.*

Oracles: *The oracle algorithm $\mathsf{Orcl}(z)$, on input $z \in \mathbb{G}$, increments $i := i + 1$. It then exhaustively searches $\mathbb{Z}_q$ for an x such that $g^x \stackrel{?}{=} z$ and outputs x. On input some $z \notin \mathbb{G}$, Orcl outputs $\perp$.*

Verification: *The verification algorithm $\mathsf{Vrfy}((z_0, \dots, z_n), (x_0, \dots, x_n))$ computes $z_j' = g^{x_j}$. If $z_j' \stackrel{?}{=} z_j$ for all j and if $i \leq n$, then it outputs 1, otherwise it outputs 0.*

Threshold: *The threshold algorithm $\mathsf{Thresh}(z)$ chooses $x_0, \dots, x_n \xleftarrow{\$} \mathbb{Z}_q$ and outputs $(x_0, \dots, x_n)$.*

The n-one-more discrete logarithm (n-DL) assumption is said to hold over $\mathbb{G}$, if and only if the n-$\mathsf{DL}_{\mathbb{G}}$ problem is hard.

3 Security of Schnorr Signatures

We first recall the definition of the Schnorr signature scheme (SSS) [25,26] as derived from the Schnorr identification scheme via the Fiat-Shamir transform [10]. Afterwards, we analyze the security of the resulting signature scheme in two variants of the random oracle model, in which reductions are limited in the way they can program the random oracle.

Definition 6 (Schnorr Signature Scheme). *Let $\mathbb{G}$ be a cyclic group of prime order q with generator g and let $\mathcal{H} : \{0, 1\}^* \to \mathbb{Z}_q$ be a hash function modeled as a random oracle. The Schnorr signature scheme, working over $\mathbb{G}$, is defined as follows:*

Key Generation: *The key generation algorithm $\mathsf{KGen}(1^\kappa)$ proceeds as follows: Pick $\mathsf{sk} \xleftarrow{\$} \mathbb{Z}_q$, compute $\mathsf{pk} := g^{\mathsf{sk}}$, and output $(\mathsf{sk}, \mathsf{pk})$.*

Signature Generation: *The signing algorithm $\mathsf{Sign}(\mathsf{sk}, m; r)$ proceeds as follows: Use $r \in \mathbb{Z}_q$ and compute $R := g^r$. Compute $c := \mathcal{H}(R, m)$ and $y := r + \mathsf{sk} \cdot c \mod q$. Output $\sigma := (c, y)$.*

Signature Verification *The verification algorithm $\mathsf{Vrfy}(\mathsf{pk}, m, \sigma)$ proceeds as follows: Parse σ as (c, y). If $c \stackrel{?}{=} \mathcal{H}(\mathsf{pk}^{-c} g^y, m)$, then output 1, otherwise output 0.*

3.1 Unforgeability of Schnorr Signatures Under Randomly Programming Reductions

We begin by showing that the original proof by Pointcheval and Stern [22,23] still holds for randomly programming reductions. Randomly programming reductions as defined in [13] do not simulate the random oracle themselves. Instead, they can re-set the random oracle to another hash value. As shown in [13] such randomly programming reductions are equivalent to the weakly-programmable random oracle model (WPROM) which is in between the programmable and non-programmable ROM. Whereas a conventional random oracle has only a single interface implementing a random mapping from domain **Dom** to range **Rng**, a weakly programmable random oracle has three interfaces, which allow for programming but only in a weak sense: one cannot freely re-program the hash values but only re-set them to another random value:

Definition 7 (Weakly Programmable Random Oracle). *A* weakly programmable random oracle *(WPRO) exposes three interfaces to the caller:*

Evaluation: *The evaluation interface* RO_{eval} *behaves as a conventional random oracle, mapping* **Dom** $\to$ **Rng**.
Random: *The random interface* RO_{rand} *takes as input bit strings of arbitrary length and implements a random mapping* $\{0,1\}^* \to$ **Rng**.
Programming: *The programming interface* RO_{prog} *takes as input a pair* $(a,b) \in$ **Dom** $\times \{0,1\}^*$ *and programs* $\mathsf{RO}_{eval}(a)$ *to evaluate to* $\mathsf{RO}_{rand}(b)$.

The randomly programming reduction gets oracle access to all three interfaces, whereas the adversary only gets access to the RO_{eval} interface. We now show that randomly programming reductions are sufficient to prove SSS secure in the ROM.

Theorem 1 (EUF-CMA Security of SSS Under Randomly Programming Reductions). *The* EUF-CMA *security of* SSS *is reducible to the* discrete logarithm *problem over* $\mathbb{G}$ *using a randomly programming reduction* $\mathcal{R}$.

The proof is close to the one in the programmable case and omitted here; the reader may refer to the full version for a sketch. We thus show that the limited programmability of a randomly programming random oracle is sufficient to obtain a (loose) proof of security for Schnorr signatures. In particular, choosing range points of the random oracle at will is not required for the proof. We note that the above result transfers to other FS schemes such as [19,17,12].

3.2 Schnorr Signatures are not Provably Secure Under Non-Programming Single-Instance Reductions

We now show that the Schnorr Signature Scheme cannot be proven existentially unforgeable under chosen message attacks without programming the random oracle —at least with respect to a slightly restricted type of reduction. We actually prove that, if such a reduction exists, the 1-one-more discrete logarithm assumption does not hold over $\mathbb{G}$.

We term the restricted class of reductions as *single-instance reductions*. Such single-instance reductions only invoke a single instance of the adversary and, while they may rewind the adversary, they may not rewind it to a point before it received the public key for the first time. This class of reductions is especially relevant, because both the original security reduction by Pointcheval and Stern [23] as well as the one in Theorem 1 are of this type.

Instead of simulating the random oracle itself, a non-programming reduction works relative to an external fixed random function and it is required to honestly answer all random oracle queries. That is, the black-box reduction can observe the adversary's queries to the random oracle, but cannot change the answers. We omit a formal approach (see [13]) because the definition reflects the intuition straightforwardly. We remark that the approach assumes fully-black-box reductions [24] (or, in terms of the CAP taxonomy of [3], the BBB-type of reduction) which need to work for any (unbounded) adversary oracle. In particular, and we will in fact exploit this below, the adversary can thus depend on the reduction. We may therefore think of the adversary as a family $\mathbb{A}$ of adversaries $\mathcal{A}_{\mathcal{R},a}$, depending on the reduction $\mathcal{R}$ and using some randomness a. We believe it is conceptually easier in this case here to make the randomness a explicit, as opposed to having a single adversary that internally chooses a at the beginning of the execution. It is nonetheless sometimes convenient to omit these subindices and to simply write $\mathcal{A}$.

Theorem 2 (Non-Programming Irreducibility for SSS). *Assume that the 1-one-more discrete logarithm assumption holds over $\mathbb{G}$. There exists no non-programming single-instance fully-black-box reduction that reduces the EUF-CMA security of SSS over $\mathbb{G}$ to the discrete logarithm problem over $\mathbb{G}$.*

More precisely, assume there exists a non-programming single-instance fully-black-box reduction $\mathcal{R}$ that converts any adversary $\mathcal{A}$ against the EUF-CMA security of SSS working in group $\mathbb{G}$ into an adversary against the DL problem over $\mathbb{G}$. Assume further that the reduction has success probability $\mathsf{Succ}_{\mathsf{DL},\mathbb{G}}^{\mathcal{R}^{\mathcal{A}}}(\kappa)$ for given $\mathcal{A}$ and runtime $\mathsf{Time}_{\mathcal{R}}(\kappa)$. Then, there exists a family $\mathbb{A}$ of successful (but possibly inefficient) adversaries $\mathcal{A}_{\mathcal{R},a}$ against the EUF-CMA security of SSS and a meta-reduction $\mathcal{M}$ that breaks the 1-DL assumption over $\mathbb{G}$ with non-negligible success probability $\mathsf{Succ}_{\mathsf{1\text{-}DL},\mathbb{G}}^{\mathcal{M}}(\kappa) \geq (\mathsf{Succ}_{\mathsf{DL},\mathbb{G}}^{\mathcal{R}^{\mathcal{A}_{\mathcal{R},a}}}(\kappa))^2$ for a random $\mathcal{A}_{\mathcal{R},a} \in \mathbb{A}$ and runtime $\mathsf{Time}_{\mathcal{M}}(\kappa) = 2 \cdot \mathsf{Time}_{\mathcal{R}}(\kappa) + \mathsf{poly}(\kappa)$.

Note that the fact that $\mathcal{A}$ breaks SSS working over $\mathbb{G}$ implies that for any public key pk output by $\mathcal{R}$ it holds that $\mathsf{pk} \in \mathbb{G}$.

Proof. (Sketch) Roughly, the meta-reduction $\mathcal{M}$ – depicted in Figure 1 – with inputs z_0, z_1 works as follows: It invokes two instances $\mathcal{R}_0$ and $\mathcal{R}_1$ of the reduction in a black-box way, on inputs z_0 and z_1, respectively, and independent random tapes. When the instances of $\mathcal{R}$ invoke the forger with public key pk_0 and pk_1 respectively, $\mathcal{M}$ simulates a specific (inefficient) forger. To do so, the meta-reduction queries random messages to the sign oracles and obtains signatures on them. It then queries the quotient of the two public keys, i.e., $\mathsf{pk}_0\mathsf{pk}_1^{-1}$, to the
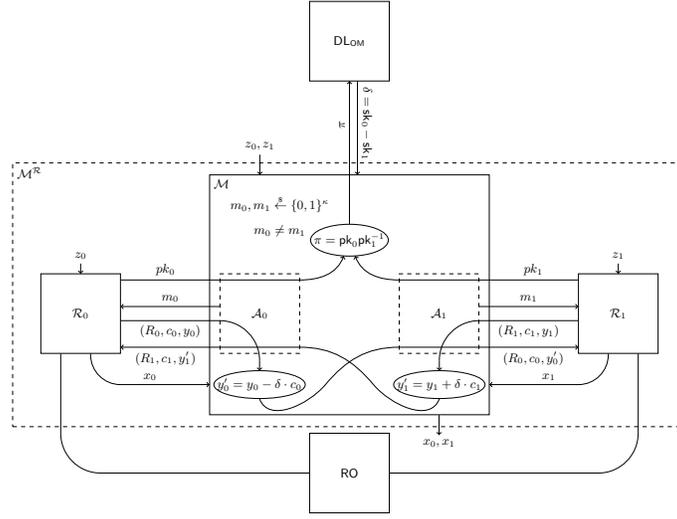
Fig. 1. The meta-reduction uses two instances of $\mathcal{R}$ and simulates the adversary $\mathcal{A}$ by obtaining the difference between the secret keys and adapting the signatures output by $\mathcal{R}$ to the other key, respectively.

$\mathsf{DL_{OM}}$ oracle, thus obtaining the difference between the secret keys. The difference between the secret keys can then be used to adapt the obtained signatures to the other public key, respectively. These adapted signatures are then returned to the reductions as forgeries. As we are working in the (non-programmable) random oracle model, the instances of $\mathcal{R}$ expect to see all the random oracle queries, the (simulated) adversary would issue. The meta-reduction $\mathcal{M}$ therefore makes sure to issue exactly those queries. Then the meta-reduction mimics the behavior of the adversary closely, and succeeds in solving the 1-one-more DL problem. □

Remark 1. Note that the restriction to single-instance reductions is crucial at this point. Consider a reduction that would output a second public key, either by invoking another instance of $\mathcal{A}$ or by rewinding the adversary to a point before it received the public key. The meta-reduction would then need to issue another query to the $\mathsf{DL_{OM}}$ oracle to simulate the signing oracle. Obviously, $\mathcal{M}$ would then have made 2 queries to the $\mathsf{DL_{OM}}$ oracle and could, thus, no longer win in the 1-DL experiment.

Remark 2. It should be noted, that the meta-reduction employed in the proof of Theorem 2 only works because SSS is defined relative to a single fixed random oracle. If one uses a common variant of the Fiat-Shamir transform, in which the random oracle is "personalized" by including the public key in the hash query, $c = H(\mathsf{pk}, R, m)$, the meta-reduction no longer works. This is due to the fact that in this case signatures can no longer be simply adapted to another public key, using only the secret keys' difference.

Remark 3. The idea immediately applies to other FS signature schemes with unique keys, where there is a related one-more problem, such as the RSA-GQ scheme [17].

4 Limitations of the Meta-Reduction Technique

Paillier and Vergnaud [20], as well as we here, have used meta-reductions to provide evidence that, once we drop programmability, the security of Schnorr signatures might not be equivalent to the discrete log problem after all. However, it is interesting to note that in both cases the meta-reduction-based proofs rely on the one-more discrete log assumption. As the discrete log assumption does not seem to imply its one-more variants [8] the results are, thus, conditional and not as strong as they could be. The obvious question is therefore: "Can we do better?" Unfortunately, the answer turns out to be "Not without finding an actual adversary."

Our results actually holds for *any* randomized signature scheme $\mathcal{S}$ (where, as explained in Section 2, the signing algorithm has super-logarithmic min entropy) for which the signing algorithm's hash queries in any signature generation can always be reconstructed from the signature alone, in the right order. We call them randomized signature schemes *with reconstructible hash queries*, for a formal definition refer to the full version. These schemes include Fiat-Shamir transformed schemes such as Schnorr but also cover (randomized versions of) FDH-RSA signatures. We show that finding a meta-reduction to a non-interactive problem such as the discrete log problem is at least as hard as finding an adversary against the strong existential unforgeability of $\mathcal{S}$. For this, we first describe an inefficient reduction $\mathcal{R}$ that is capable of detecting when the forgery it receives is actually one of the signatures it produced itself as an answer to a signing query. For example, a meta-reduction may make several signing requests to a reduction and then reset these requests in order to use one additional message-signature pairs as the forgery. Our reduction will be able to spot such attempts.

The meta-reduction result of the previous section does not apply here, even though our reduction here will be of the single-instance type. The reason is that the meta-reduction there assumed an (interactive) one-more DL problem –and made use of the DL oracle– whereas the meta-reduction here should work for non-interactive problems such as the discrete log problem.

4.1 An Inefficient Reduction for Randomized Signature Schemes with Reconstructible Hash Queries

Let $\mathcal{S}$ be a randomized signature scheme and let $\Pi_{\mathcal{R}}$ be a monotone solvable problem. Let $\mathbb{Q}$ be the set of message-signature pairs (m_i, σ_i) resulting from queries to $\mathcal{R}$'s signing oracle. Furthermore, let p be the maximum number of signature queries issued by a forger $\mathcal{A}$ and assume that $\mathcal{R}$ knows the polynomially bounded p. We note that for the adversary in our single-instance reductions in the previous section, the reduction could have been given $p = 1$, too. For

the moment, the reader may think of the meta-reduction as running a single instance of the reduction; we will later reduce the multi-instance case to the single-instance case via standard "guess-and-insert" techniques.

The reduction $\mathcal{R}$ on input an instance z of $\Pi_{\mathcal{R}}$ first generates a key pair $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathcal{S}.\mathsf{KGen}(1^\kappa)$, then initializes the counter variable $i := 0$, and chooses a random function $\mathcal{O} : \{0,1\}^{2^\kappa} \times \mathbb{Z}_p \rightarrow \mathcal{C}oins_{\mathsf{pk}}$. The public key pk is then output as the key under which the forger is supposed to forge a signature. When the forger queries a message m to the signing oracle, $\mathcal{R}$ determines random coins $\omega \leftarrow \mathcal{O}(m, i)$, computes the signature as $\sigma \leftarrow \mathcal{S}.\mathsf{Sign}(\mathsf{sk}, m; \omega)$, and returns σ to the forger. The counter i is then incremented by one. If the counter is ever incremented to $p + 1$, then $\mathcal{R}$ aborts, as it is obviously not interacting with the real adversary.

Eventually, the forger outputs a forgery (m^*, σ^*). If the signature does not verify, i.e., $\mathcal{S}.\mathsf{Vrfy}(\mathsf{pk}, m^*, \sigma^*) \overset{?}{=} 0$, then $\mathcal{R}$ immediately aborts. Otherwise, the reduction computes $\sigma_j \leftarrow \mathcal{S}.\mathsf{Sign}(\mathsf{sk}, m^*; \omega_j)$ with $\omega_j \leftarrow \mathcal{O}(m^*, j)$ for all $j \in \mathbb{Z}_p$ and checks whether $\sigma_j \overset{?}{=} \sigma^*$. If the check holds for any σ_j, then $\mathcal{R}$ also immediately aborts. Otherwise, $\mathcal{R}$ enumerates all possible solutions $x \in \Pi_{\mathcal{R}}.\mathbf{Sol}$ and checks whether $\Pi_{\mathcal{R}}.\mathsf{Vrfy}(x, z) \overset{?}{=} 1$. Once such an x is found, it is output by $\mathcal{R}$ as the solution. Because $\Pi_{\mathcal{R}}$ is monotone and solvable, it is guaranteed that there exists a valid solution even though $\mathcal{R}$ never issues a single oracle query and that the enumeration of possible solutions will terminate in finite time.

Observe that the adversary $\mathcal{A}$ used by $\mathcal{R}$ is an $\mathsf{EUF\text{-}CMA}$ adversary, therefore, whenever $\mathcal{A}$ forges successfully, it forges a signature for a message m^* that has not been queried before. The probability that $\mathcal{R}$ will reject such a forgery is the probability that at least one of the σ_i collides with σ^*. As $\mathcal{O}$ is a random function, all values to which $\mathcal{O}$ evaluates on input m^* and some number i are uniformly and independently distributed. For each σ_i, the probability that it matches σ^* is thus bounded through the min-entropy of the random variable describing $\mathcal{S}.\mathsf{Sign}(\mathsf{sk}, m^*)$, i.e., $\forall i \in \mathbb{Z}_p : (\Pr[\sigma_i \overset{?}{=} \sigma^*] \le 2^{-H_\infty(\mathcal{S}.\mathsf{Sign}(\mathsf{sk}, m^*))})$.

Therefore, the probability that $\mathcal{R}$ will accept a forgery is at least $1 - p \cdot 2^{-H_\infty(\mathcal{S}.\mathsf{Sign}(\mathsf{sk}, m^*))}$. As $\mathcal{S}$ is randomized, the probability for each σ_i to match is negligible and thus

$$\mathsf{Succ}^{\mathcal{R}^{\mathcal{A}}}_{\Pi_{\mathcal{R}}}(\kappa) \ge (1 - p \cdot \epsilon(\kappa)) \cdot \mathsf{Succ}^{\mathcal{S},\mathcal{A}}_{\mathsf{EUF\text{-}CMA}}(\kappa) = \mathsf{Succ}^{\mathcal{S},\mathcal{A}}_{\mathsf{EUF\text{-}CMA}}(\kappa) - \epsilon'(\kappa)$$

for negligible functions ϵ, ϵ'. Therefore, we conclude that $\mathsf{Succ}^{\mathcal{R}^{\mathcal{A}}}_{\Pi}(\kappa)$ is non-negligible for any successful adversary $\mathcal{A}$ and that $\mathcal{R}$ is, thus, a successful –albeit inefficient– reduction from problem $\Pi_{\mathcal{R}}$ to the $\mathsf{EUF\text{-}CMA}$ security of $\mathcal{S}$.

We next show that the checks of our reduction prevent the meta-reduction to replay signatures to the reduction. This step relies on the fact that the meta-reduction can only use the reduction in a black-box way, $\mathcal{M}^{\mathcal{R}}$, and has for example no control over the coin tosses of $\mathcal{R}$. First, we show that we can restrict ourselves to meta-reductions which actually take advantage of the reduction, at least if the meta-reduction's problem $\Pi_{\mathcal{M}}$ is hard:

Lemma 1 (Meta-Reductions Rely on the Reduction). *Let $\mathcal{M}$ be a non-programming meta-reduction that converts any $(\text{EUF-CMA}_{\mathcal{S}} \rightsquigarrow \Pi_{\mathcal{R}})$ reduction in a black-box way into an adversary against some hard problem $\Pi_{\mathcal{M}}$. Further, let the reduction used by $\mathcal{M}$ be $\mathcal{R}$ as described above. Then it holds that $\mathcal{M}$ provides $\mathcal{R}$ with a forgery (m^*, σ^*) with non-negligible advantage.*

Proof. Assume that this was not the case. Then one could easily simulate $\mathcal{R}$ and the meta-reduction interacting with this reduction would solve $\Pi_{\mathcal{M}}$ efficiently with non-negligible advantage. This contradicts the hardness of the problem. □

Hence, from now on we condition on the meta-reduction to always provide the reduction with a forgery, without losing more than a negligible advantage. In this case we have:

Lemma 2 (Meta-Reductions Cannot Replay Signatures). *Let $\mathcal{M}$ be a non-programming meta-reduction that converts any $(\text{EUF-CMA}_{\mathcal{S}} \rightsquigarrow \Pi_{\mathcal{R}})$ reduction in a black-box way into an adversary against some hard problem $\Pi_{\mathcal{M}}$. Let $\mathbb{Q}$ be the set of message-signature pairs (m_i, σ_i) resulting from $\mathcal{M}$'s queries to the reduction's signing oracle, and let (m^*, σ^*) be the message-signature pair output by $\mathcal{M}$ as a forgery on behalf of the adversary. Further, let the reduction used by $\mathcal{M}$ be $\mathcal{R}$ as described above. Then it holds that $(m^*, \sigma^*) \notin \mathbb{Q}$.*

Proof. The proof is rather straightforward. Observe that by construction of $\mathcal{R}$ the following holds: $\forall (m, \sigma) \in \mathbb{Q} : \exists i \in \mathbb{Z}_p : \omega \leftarrow \mathcal{O}(m, i) \wedge \sigma \stackrel{?}{=} \mathcal{S}.\mathsf{Sign}(m, i; \omega)$. Therefore, it follows directly that, for $(m^*, \sigma^*) \in \mathbb{Q}$, the reduction $\mathcal{R}$ will abort and $\mathsf{Succ}_{\Pi_{\mathcal{R}}}^{\mathcal{R}^{\mathcal{M}}}(\kappa) = 0$ for $\mathcal{M}$ if it replays an element of $\mathbb{Q}$ as a forgery. Note that here we rely on the previous Lemma which assumes that $\mathcal{M}$ always provides such a forgery. As it, thus, would not be a successful meta-reduction it must hold that $(m^*, \sigma^*) \notin \mathbb{Q}$. □

4.2 A Reduction against the Meta-Reduction

Using the reduction described in the previous section, we now prove that finding an efficient meta-reduction for a randomized signature scheme is at least as hard as finding a strong existential forger.

Theorem 3 (Meta-Reductions to Non-Interactive Problems Are Hard). *Let $\mathcal{S}$ be a randomized signature scheme with reconstructible hash queries, $\Pi_{\mathcal{R}}$ be a monotone solvable problem, and $\Pi_{\mathcal{M}}$ be a non-interactive, efficiently generatable problem. If $\Pi_{\mathcal{M}}$ is hard, then finding an efficient meta-reduction $\mathcal{M}$ that converts any successful $(\text{EUF-CMA}_{\mathcal{S}} \rightsquigarrow \Pi_{\mathcal{R}})$-reduction in a black-box way into an efficient successful adversary against $\Pi_{\mathcal{M}}$ is at least as hard as finding an sEUF-CMA adversary against $\mathcal{S}$.*

More precisely, assume there exists an efficient non-programming black-box meta-reduction $\mathcal{M}$ that converts any $(\text{EUF-CMA}_{\mathcal{S}} \rightsquigarrow \Pi_{\mathcal{R}})$-reduction into an adversary against $\Pi_{\mathcal{M}}$. Then, there exists a meta-meta-reduction $\mathcal{N}$ that converts $\mathcal{M}$ into an adversary against the sEUF-CMA security of $\mathcal{S}$ with non-negligible

success probability $\mathsf{Succ}^{\mathcal{S},\mathcal{N}^{\mathcal{M}}}_{\mathsf{sEUF\text{-}CMA}}(\kappa) \geq \frac{1}{r} \cdot \mathsf{Succ}^{\mathcal{M}^{\mathcal{R}}}_{\Pi_{\mathcal{M}}}(\kappa)$ *and runtime* $\mathsf{Time}_{\mathcal{N}^{\mathcal{M}}}(\kappa) = \mathsf{Time}_{\mathcal{M}^{\mathcal{R}}}(\kappa) + \mathsf{poly}(\kappa)$, *where* $\mathcal{R}$ *is the reduction described above and* r *is the maximal number of reduction instances invoked by* $\mathcal{M}$.

Note that, since $\mathcal{M}$ needs to work for any (black-box) $\mathcal{R}$, we may assume that $\mathcal{R}$ knows r. Indeed, we take advantage of this fact in the proof.
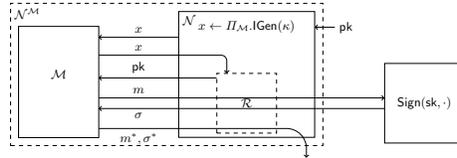
Fig. 2. The meta-meta-reduction relies on the fact that $\mathcal{M}$ cannot replay an old signature. It simply outputs the forgery provided by $\mathcal{M}$. The meta-meta-reduction can be generalized for multiple instances of $\mathcal{R}$ using a standard guess-and-insert approach.

The proof is omitted here, but can be found in the full version. The idea is outlined in Figure 2. The meta-meta-reduction picks one of the r reduction instances run by $\mathcal{M}$ at random and substitutes this instance with the help of its external signature oracle. The reconstruction property guarantees that $\mathcal{N}$ can still pretend towards $\mathcal{M}$ to have made the hash queries of externally provided signatures locally. All other reduction instances are simulated by $\mathcal{N}$ itself. In order to be successful, $\mathcal{M}$ needs to provide a forgery to some of the $\mathcal{R}$-instances, and with probability $1/r$ it will be the one which is "externalized" by $\mathcal{N}$. In this case, Lemma 2 ensures that the forgery is a strong forgery against the external signature oracle.

Acknowledgments

References

1. Abdalla, M., An, J.H., Bellare, M., Namprempre, C.: From identification to signatures via the Fiat-Shamir transform: Minimizing assumptions for security and forward-security. In Knudsen, L.R., ed.: Advances in Cryptology – EUROCRYPT 2002. Volume 2332 of Lecture Notes in Computer Science., Amsterdam, The Netherlands, Springer, Berlin, Germany (April 28 – May 2, 2002) 418–433

2. Abe, M., Groth, J., Ohkubo, M.: Separating short structure-preserving signatures from non-interactive assumptions. In Lee, D.H., Wang, X., eds.: Advances in Cryptology – ASIACRYPT 2011. Volume 7073 of Lecture Notes in Computer Science., Seoul, South Korea, Springer, Berlin, Germany (December 4–8, 2011) 628–646

3. Baecher, P., Brzuska, C., Fischlin, M.: Notions of black-box reductions, revisited. IACR Cryptology ePrint Archive (2013)

4. Baldimtsi, F., Lysyanskaya, A.: On the security of one-witness blind signature schemes. IACR Cryptology ePrint Archive (2012)

5. Bellare, M., Namprempre, C., Pointcheval, D., Semanko, M.: The one-more-RSA-inversion problems and the security of Chaum's blind signature scheme. Journal of Cryptology **16**(3) (June 2003) 185–215

6. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In Ashby, V., ed.: ACM CCS 93: 1st Conference on Computer and Communications Security, Fairfax, Virginia, USA, ACM Press (November 3–5, 1993) 62–73

7. Boneh, D., Venkatesan, R.: Breaking RSA may not be equivalent to factoring. In Nyberg, K., ed.: Advances in Cryptology – EUROCRYPT'98. Volume 1403 of Lecture Notes in Computer Science., Espoo, Finland, Springer, Berlin, Germany (May 31 – June 4, 1998) 59–71

8. Bresson, E., Monnerat, J., Vergnaud, D.: Separation results on the "one-more" computational problems. In Malkin, T., ed.: Topics in Cryptology – CT-RSA 2008. Volume 4964 of Lecture Notes in Computer Science., San Francisco, CA, USA, Springer, Berlin, Germany (April 7–11, 2008) 71–87

9. Coron, J.S.: Optimal security proofs for PSS and other signature schemes. In Knudsen, L.R., ed.: Advances in Cryptology – EUROCRYPT 2002. Volume 2332 of Lecture Notes in Computer Science., Amsterdam, The Netherlands, Springer, Berlin, Germany (April 28 – May 2, 2002) 272–287

10. Fiat, A., Shamir, A.: How to prove yourself: Practical solutions to identification and signature problems. In Odlyzko, A.M., ed.: Advances in Cryptology – CRYPTO'86. Volume 263 of Lecture Notes in Computer Science., Santa Barbara, CA, USA, Springer, Berlin, Germany (August 1987) 186–194

11. Fischlin, M.: Black-box reductions and separations in cryptography. In: AFRICACRYPT. Volume 7374 of Lecture Notes in Computer Science., Springer (2012) 413–422

12. Fischlin, M., Fischlin, R.: The representation problem based on factoring. In Preneel, B., ed.: Topics in Cryptology – CT-RSA 2002. Volume 2271 of Lecture Notes in Computer Science., San Jose, CA, USA, Springer, Berlin, Germany (February 18–22, 2002) 96–113

13. Fischlin, M., Lehmann, A., Ristenpart, T., Shrimpton, T., Stam, M., Tessaro, S.: Random oracles with(out) programmability. In Abe, M., ed.: Advances in Cryptology – ASIACRYPT 2010. Volume 6477 of Lecture Notes in Computer Science., Singapore, Springer, Berlin, Germany (December 5–9, 2010) 303–320

14. Fischlin, M., Schröder, D.: On the impossibility of three-move blind signature schemes. In Gilbert, H., ed.: Advances in Cryptology – EUROCRYPT 2010. Volume 6110 of Lecture Notes in Computer Science., French Riviera, Springer, Berlin, Germany (May 30 – June 3, 2010) 197–215

15. Garg, S., Bhaskar, R., Lokam, S.V.: Improved bounds on security reductions for discrete log based signatures. In Wagner, D., ed.: Advances in Cryptology – CRYPTO 2008. Volume 5157 of Lecture Notes in Computer Science., Santa Barbara, CA, USA, Springer, Berlin, Germany (August 17–21, 2008) 93–107

16. Gentry, C., Wichs, D.: Separating succinct non-interactive arguments from all falsifiable assumptions. In Fortnow, L., Vadhan, S.P., eds.: 43rd Annual ACM Symposium on Theory of Computing, San Jose, California, USA, ACM Press (June 6–8, 2011) 99–108

17. Guillou, L.C., Quisquater, J.J.: A practical zero-knowledge protocol fitted to security microprocessor minimizing both trasmission and memory. In Günther, C.G., ed.: Advances in Cryptology – EUROCRYPT'88. Volume 330 of Lecture Notes in Computer Science., Davos, Switzerland, Springer, Berlin, Germany (May 25–27, 1988) 123–128

18. Nielsen, J.B.: Separating random oracle proofs from complexity theoretic proofs: The non-committing encryption case. In Yung, M., ed.: Advances in Cryptology – CRYPTO 2002. Volume 2442 of Lecture Notes in Computer Science., Santa Barbara, CA, USA, Springer, Berlin, Germany (August 18–22, 2002) 111–126

19. Okamoto, T.: Provably secure and practical identification schemes and corresponding signature schemes. In Brickell, E.F., ed.: Advances in Cryptology – CRYPTO'92. Volume 740 of Lecture Notes in Computer Science., Santa Barbara, CA, USA, Springer, Berlin, Germany (August 16–20, 1993) 31–53

20. Paillier, P., Vergnaud, D.: Discrete-log-based signatures may not be equivalent to discrete log. In Roy, B.K., ed.: Advances in Cryptology – ASIACRYPT 2005. Volume 3788 of Lecture Notes in Computer Science., Chennai, India, Springer, Berlin, Germany (December 4–8, 2005) 1–20

21. Pass, R.: Limits of provable security from standard assumptions. In Fortnow, L., Vadhan, S.P., eds.: 43rd Annual ACM Symposium on Theory of Computing, San Jose, California, USA, ACM Press (June 6–8, 2011) 109–118

22. Pointcheval, D., Stern, J.: Security proofs for signature schemes. In Maurer, U.M., ed.: Advances in Cryptology – EUROCRYPT'96. Volume 1070 of Lecture Notes in Computer Science., Saragossa, Spain, Springer, Berlin, Germany (May 12–16, 1996) 387–398

23. Pointcheval, D., Stern, J.: Security arguments for digital signatures and blind signatures. Journal of Cryptology **13**(3) (2000) 361–396

24. Reingold, O., Trevisan, L., Vadhan, S.P.: Notions of reducibility between cryptographic primitives. In Naor, M., ed.: TCC 2004: 1st Theory of Cryptography Conference. Volume 2951 of Lecture Notes in Computer Science., Cambridge, MA, USA, Springer, Berlin, Germany (February 19–21, 2004) 1–20

25. Schnorr, C.P.: Efficient identification and signatures for smart cards. In Brassard, G., ed.: Advances in Cryptology – CRYPTO'89. Volume 435 of Lecture Notes in Computer Science., Santa Barbara, CA, USA, Springer, Berlin, Germany (August 20–24, 1990) 239–252

26. Schnorr, C.P.: Efficient signature generation by smart cards. Journal of Cryptology **4**(3) (1991) 161–174

27. Seurin, Y.: On the exact security of schnorr-type signatures in the random oracle model. In Pointcheval, D., Johansson, T., eds.: Advances in Cryptology – EUROCRYPT 2012. Volume 7237 of Lecture Notes in Computer Science., Cambridge, UK, Springer, Berlin, Germany (April 15–19, 2012) 554–571