

On the Impossibility of Tight Cryptographic Reductions^{*}

Christoph Bader, Tibor Jäger^{**}, Yong Li, and Sven Schäge^{***}

Horst Görtz Institute for IT Security, Ruhr-University Bochum

Abstract. The existence of *tight* reductions in cryptographic security proofs is an important question, motivated by the theoretical search for cryptosystems whose security guarantees are truly independent of adversarial behavior and the practical necessity of concrete security bounds for the theoretically-sound selection of cryptographic parameters. At Eurocrypt 2002, Coron described a *meta-reduction* technique that allows to prove the *impossibility* of tight reductions for certain digital signature schemes. This seminal result has found many further interesting applications. However, due to a technical subtlety in the argument, the applicability of this technique beyond digital signatures in the *single-user* setting has turned out to be rather limited. We describe a new meta-reduction technique for proving such impossibility results, which improves on known ones in several ways. It enables interesting novel applications, including a formal proof that for certain cryptographic primitives (including public-key encryption/key encapsulation mechanisms and digital signatures), the security loss incurred when the primitive is transferred from an idealized single-user setting to the more realistic multi-user setting is *impossible* to avoid, and a lower tightness bound for non-interactive key exchange protocols. Moreover, the technique allows to rule out tight reductions from a very general class of non-interactive complexity assumptions. Furthermore, the proofs and bounds are simpler than in Coron’s technique and its extensions.

1 Introduction

Provable security. In modern cryptography, new cryptosystems are usually constructed together with a *proof of security*. Usually this security proof consists of a reduction Λ (in a complexity-theoretic sense), which turns an efficient adversary \mathcal{A} into a machine $\Lambda^{\mathcal{A}}$ solving a well-studied, assumed-to-be-hard computational problem. Under the assumption that this computational problem is not efficiently solvable, this implies that the cryptosystem is secure. This approach is usually called “provable security”, it is inspired by the analysis of relations between computational problems in complexity theory, and allows to show that breaking the security of a cryptosystem is at least as hard as solving a certain well-defined hard computational problem.

^{*} This is the full version of an article with the same title that was accepted for publication at Eurocrypt’16.

^{**} Supported by DFG grant JA 2445/1-1.

^{***} Supported by UbiCrypt, DFG grant GRK 1817/1.

The security loss in reduction-based security proofs. The “quality” of a reduction can be measured by comparing the running time and success probability of $\Lambda^{\mathcal{A}}$ to the running time and success probability of attacker \mathcal{A} . Ideally, $\Lambda^{\mathcal{A}}$ has about the same running time and success probability as \mathcal{A} . However, most security proofs describe reductions where $\Lambda^{\mathcal{A}}$ has either a significantly larger running time or a significantly smaller success probability than \mathcal{A} (or both). Thus, the reduction “loses” efficiency and/or efficacy.

Since provable security is inspired by classical complexity theory, security proofs have traditionally been formulated asymptotically. The running time and success probability of Turing machines are modeled as functions in a *security parameter* $k \in \mathbb{N}$. Let $t_{\Lambda^{\mathcal{A}}}(k)$ denote the running time and $\epsilon_{\Lambda^{\mathcal{A}}}(k)$ denote the success probability of $\Lambda^{\mathcal{A}}$. Likewise, let $t_{\mathcal{A}}(k)$ and $\epsilon_{\mathcal{A}}(k)$ denote the running time and success probability of \mathcal{A} . Then it holds that

$$t_{\Lambda^{\mathcal{A}}}(k)/\epsilon_{\Lambda^{\mathcal{A}}}(k) = \ell(k) \cdot t_{\mathcal{A}}(k)/\epsilon_{\mathcal{A}}(k)$$

for some “loss” $\ell(k)$. A reduction Λ is considered *efficient*, if its loss $\ell(k)$ is bounded by a polynomial. Note that in this approach the concrete size of polynomial ℓ (i.e., its degree and the size of its coefficients) does not matter. As common in classical complexity theory, it was considered sufficient to show that ℓ is polynomially-bounded.

Concrete security proofs, the notion of tightness, and its relevance. In order to deploy a cryptosystem in practice, the size of cryptographic parameters (like for instance the length of moduli or the size of underlying algebraic groups) has to be selected. However, the asymptotic approach described above does not allow to derive concrete recommendations for such parameters, as it only shows that sufficiently large parameters *exist*. This is because the size of parameters depends on the concrete value of ℓ , the loss of the reduction. A larger loss requires larger parameters.

The more recent approach, termed *concrete security*, makes the concrete security loss of a reduction explicit. This allows to derive concrete recommendations for parameters in a theoretically sound way (see e.g. [7] for a detailed treatment). Ideally, $\ell(k)$ is constant. In this case the reduction is said to be *tight*.¹ The existence of cryptosystems whose security is independent of deployment parameters is of course an interesting theoretical question in its own right. Moreover, it has a strong practical motivation, because the tightness of a reduction directly influences the selection of the size of cryptographic parameters, and thus has a direct impact to the efficiency of cryptosystems. An example is given in Appendix A.

Coron’s result and its refinements. Coron [19] considered the existence of tight reductions for *unique*² signature schemes in the single user setting, and described a “rewinding argument” (cf. Goldwasser *et al.* [28]), which allowed to prove lower tightness

¹ When speaking of tight reductions in this paper, we mean tight reductions from non-interactive computational problems, like integer factorization, the discrete logarithm problem, etc., rather than (often trivial) tight reductions from interactive or contrived non-standard computational problems, which sometimes are very similar to the assumption that the cryptosystem is secure.

² For a unique signature scheme there exists exactly one unique valid signature for each message. For instance, important instantiations of the famous Full-Domain Hash construction are unique signature schemes, see [32].

bounds for such signature schemes. In particular, Coron considered “simple”³ reductions, which convert a forger F breaking the security⁴ of a unique signature scheme into a machine solving a computationally hard problem Π . He showed that any such reduction yields an algorithm \mathcal{B} solving Π *directly* with probability $\epsilon_{\mathcal{B}}$, where

$$\epsilon_{\mathcal{B}} \geq \epsilon_{\Lambda} - \frac{\epsilon_F}{\exp(1) \cdot n} \cdot \left(1 - \frac{n}{|\mathcal{M}|}\right)^{-1}. \quad (1)$$

Here ϵ_{Λ} is the success probability of Λ , ϵ_F is the success probability of the signature forger F used by Λ , n is the number of signatures queried by F in the EUF-CMA security experiment, and $|\mathcal{M}|$ is the size of the message space. Note that if $|\mathcal{M}| \gg n$, which is a reasonable for signature schemes, then the bound in (1) essentially implies that the success probability of ϵ_{Λ} of the reduction can not substantially exceed $\epsilon_F / (\exp(1) \cdot n)$, unless there exists an algorithm \mathcal{B} solving Π efficiently. The latter, however, contradicts the hardness assumption on Π . This result was later revisited by Kakvi and Kiltz [32], and generalized by Hofheinz *et al.* [31] to (non-unique) signature schemes with efficiently re-randomizable signatures, see also Appendix B.

Limitations of known meta-reductions. Unfortunately, Coron’s result has found only limited applications beyond digital signatures in the single-user setting. Most previous works [19,32,31] consider this setting, the (to our best knowledge) only exception is due to Lewko and Waters [34], which considers hierarchical identity-based encryption. Why isn’t it possible to apply it to other primitives? One reason is that the bound in Equation (1) ceases to be useful for reasonable values of ϵ_{Λ} and ϵ_F if $n \approx |\mathcal{M}|$. This can be easily seen by setting $n = |\mathcal{M}| - 1$. The assumption that $|\mathcal{M}| \gg n$ is a prerequisite for the arguments in [19,32,31] to work, thus, it is not possible to apply this technique to settings, where the assumption $|\mathcal{M}| \gg n$ is *not* reasonable.

Therefore Coron’s technique is not applicable when $|\mathcal{M}|$ is polynomially-bounded. However, such a situation appears often when considering cryptographic primitives beyond digital signatures in the single-user setting. Consider, for instance, a security model where the adversary is provided with $\mathcal{M} = \{pk_1, \dots, pk_n\}$, where pk_1, \dots, pk_n is a list of public keys. The adversary may learn all but one of the corresponding secret keys, and is considered successful if it “breaks security” with respect to an uncorrupted key. This is a quite common setting, which occurs for instance in security models for signatures or public-key encryption in the multi-user setting with corruptions [3,4], all common security models for authenticated key exchange [9,16,4], and non-interactive key exchange [26] protocols. *How can we analyze the existence of inherent tightness bounds in these settings?*

Our contributions. We develop a new meta-reduction technique, which is also applicable in settings where $|\mathcal{M}|$ is polynomially bounded. In comparison to [19,32,31], we

³ Intuitively, a “simple” reduction is a reduction which has black-box access to the adversary, and runs the adversary only *sequentially*. Most reductions in cryptographic security proofs are of this type. A more precise definition is given in the body of the paper.

⁴ In the sense of existential unforgeability under chosen-message attacks (EUF-CMA, cf. Definition 21).

achieve the simpler bound

$$\epsilon_{\mathcal{B}} \geq \epsilon_{\Lambda} - 1/n.$$

which is independent of $|\mathcal{M}|$.

Our new technique allows to rule out tight reductions from any non-interactive complexity assumption (cf. Definition 5). This includes also “decisional” assumptions (like decisional Diffie-Hellman). It avoids the combinatorial lemma of Coron [19, Lemma 1], which has a relatively technical proof. Our approach does not require such a combinatorial argument, but is more “direct”.

This simplicity allows us to describe a generalized experiment with an abstract computable relation that captures the necessary properties for our tightness bounds. Then we explain that the standard security experiments for many cryptographic primitives are specific instances of this abstract experiment.

Technical idea. To describe our technical idea, let us consider the example of digital signatures in the single-user settings, as considered in [19,32,31], for this introduction. As sketched above, the result will later be generalized and applied to other settings as well. We consider a weakened signature security definition, where the security experiment proceeds as follows.

1. The adversary receives as input a verification key vk along with n random but pairwise distinct messages m_1, \dots, m_n .
2. The adversary selects an index j^* , and receives in response $n - 1$ signatures σ_i for all messages m_i with $i \neq j^*$.
3. Finally, the adversary wins the experiment if it outputs σ^* that is a valid signature for m_{j^*} with respect to j^* .

Note that this is a very weak security definition, because the adversary is only able to observe signatures of random messages. However, note also that any lower tightness bound for such a weaker security definition implies a corresponding bound for any stronger definition. In particular, the above definition is weaker than the standard security definition *existential unforgeability under chosen message attacks* considered in [19,32,31], where messages may be adaptively chosen by the adversary.

Essentially, we argue that once a reduction has started the adversary in Step 1 of the above experiment, and thus has “committed” to a verification key vk and messages m_1, \dots, m_n , there can only be a single choice of j^* for which this reduction is able to output valid signatures σ_i for all $i \neq j^*$. Thus, for any adversary which chooses j^* uniformly at random the reduction has probability at most $1/n$ to succeed. We prove this by contradiction, by showing essentially that any reduction which is successful for two distinct choices of j^* , say j_0, j_1 , can be used to construct a machine that breaks the underlying security assumption directly.

Technically, we proceed in two steps: first we describe an *inefficient* adversary against the reduction which chooses j^* uniformly random, and computes the signature σ^* for m_{j^*} by exhaustive search. Next, we show that this adversary can efficiently be simulated by our meta-reduction, if the reduction could succeed for two different choices j_0 and j_1 after committing to (vk, m_1, \dots, m_n) . The meta-reduction simulates the inefficient adversary by rewinding the reduction. Essentially, if the reduction could

succeed for two different values j_0, j_1 , then it must also be able output the signatures for *all* n messages. Therefore we start the reduction and let it run until it reaches a “break point” where it outputs (vk, m_1, \dots, m_n) . Next, we run the reduction n -times, each time starting from the break point and using a different index j , to search for two values j_0, j_1 such that $j_0 \neq j_1$ such that the reduction outputs valid signatures for all-but-one messages. If indeed there exist two such indices j_0, j_1 , then we now have learned signatures for all messages (m_1, \dots, m_n) which are valid w.r.t. vk . Thus, we can run the reduction one last time from the break point, this time to the end, using index j_0 (or equivalently j_1), and we simulate the inefficient adversary using the fact that we know a valid signature for m_{j_0} (or m_{j_1}). Importantly, in the last execution of the reduction we are able to simulate the inefficient adversary perfectly, so the reduction will help us to break the non-interactive complexity assumption.

We caution that the rigorous proof of the above is more complex than the intuition provided in this introduction, and we have to put restrictions on the signature scheme, which depend on the considered application. For instance, when considering signatures in the single-user setting as above, we have to require that signatures are efficiently re-randomizable. In the generalized setting we will consider other applications, which require different but usually simple-to-check properties, like for instance that for each public key vk there exists a *unique* secret key. In this way, our result provides simple criteria to check whether a cryptographic construction can have a tight proof at all. At the same time it implicitly provides guidelines for the construction of tightly secure cryptographic schemes, since all tightly secure constructions must circumvent our result in one way or the other.

The fact that we consider a weakened security experiment has several nice features. We think that the approach and its analysis described above are much simpler than previous works, which enables more involved impossibility results. We will show that it achieves a simpler bound and yields a qualitatively stronger result, as it even rules out tight reductions for such weak security experiments. Like previous works, we only consider reductions that execute the adversary sequentially and in a black-box fashion. We stress that most reductions in cryptography have this property.

We generalize the above idea from signature schemes in a single-user setting to abstract relations, which capture the relevant properties required for our impossibility argument to go through. We show that this abstraction allows to apply the result relatively easily to other cryptographic primitives, by describing applications to public-key encryption and signatures in the multi-user setting, and non-interactive key exchange.

Overview of Applications. A first, immediate application of our new technique are strengthened versions of the results of [19,32,31], but with significantly simpler proofs and tightness bounds even for weaker security notions (which is a stronger result). In contrast to previous works [19,32,31], the impossibility results hold also for “decisional” complexity assumptions.

Additionally, the fact that our meta-reduction does not require the combinatorial lemma of Coron enables further, novel applications in settings with polynomially-bounded spaces (where Coron’s result worked only for exponential-sized spaces). As a first novel application of our generalized theorem, we analyze the tightness loss that occurs when security proofs in idealized single-user settings are transferred to the more realistic

multi-user setting. Classical security models for standard cryptographic primitives often consider an idealized setting. For instance, the standard IND-CPA and IND-CCA security experiments for public-key encryption consider a setting with only one challenge public key and only a single challenge ciphertext. This is of course unrealistic for many practical applications. Public-key encryption is typically used in settings where an attacker sees many public keys and ciphertexts, and is (potentially) able to corrupt secret keys adaptively. Even though there is a reduction from breaking security in the multi-user setting to breaking security in the idealized setting, this reduction comes with a security loss which is linear in the number of users and ciphertexts. We show that under certain conditions (e.g., for schemes where there exists a *unique* secret key for each public key) this loss is impossible to avoid. This gives an insight into which properties a cryptosystem must or must not meet in order to allow a tight reduction in the multi-user setting.

Another novel application is the analysis of the existence of *non-interactive* key exchange (NIKE). In non-interactive key exchange (NIKE) two parties are able to derive a common shared secret. However, in contrast to traditional key exchange protocols, they do not need to exchange any messages. Besides the secret key of one party the key derivation algorithm only requires the availability of the public key of the communication partner. Security is defined solely by requiring indistinguishability of the derived shared secret from a random value. We show how to apply our main result to rule out tight reductions for a large class of NIKE protocols from a standard assumption in any sufficiently strong security model (such as the CKS-heavy model from [26]).

On certified public keys and the results of Kakvi and Kiltz. Several years after the publication of the paper of Coron [19] it has turned out that this paper contains a subtle technical flaw. Essentially, it is implicitly assumed that the value output by the reduction to the adversary is a *correct* signature public key (recall that Coron considered only digital signature schemes in the single-user setting). This misses the fact that a reduction may possibly output *incorrect* keys which are computationally indistinguishable from correct ones. Indeed, such keys lead to the technical problem that a meta-reduction may *not* be able to simulate the adversary constructed in the meta-reduction of Coron correctly.

This flaw was identified and corrected by Kakvi and Kiltz [32]. Essentially, Kakvi and Kiltz enforce that the reduction outputs only public keys which can be efficiently recognized as correct, by introducing the notion of *certified* public keys. A different (but similar in spirit), slightly more general approach is due to Hofheinz *et al.* [31], who require that signatures are *efficiently re-randomizable* with respect to the public key output by from the reduction (regardless of whether this key is correct or not). Both these approaches [32,31] essentially overcome the subtle issue from Coron’s paper by ensuring that the adversaries simulated by the meta-reductions are always able to output *correctly distributed* signatures.

In this paper, we introduce the notion of *efficiently re-randomizable relations* to overcome the subtle issue pointed out by Kakvi and Kiltz [32]. This notion further generalizes the approach of [31] in a way that suits our more general setting.

Relation to tightly-secure constructions. There exist various constructions of tightly-secure cryptosystems, which have to avoid our impossibility results in one way or another. The signature schemes constructed in [20,33,37,1,30,10], for example, are tightly-secure in a single-user setting. They avoid our impossibility result because they do not have unique signatures or no efficient re-randomization algorithm is known. The same holds for the signature schemes derived from the IBE schemes of [18,11]. Bader *et al.* [4] constructed signature schemes with tight security even in the multi-user setting with adaptive secret-key corruptions. Again, our impossibility results are avoided here because signatures are not efficiently re-randomizable. The encryption schemes of Bellare, Boldyreva and Micali [6] are tightly-secure in a multi-user setting, but only without corruptions. We consider impossibility results for the multi-user setting with corruptions. The key encapsulation mechanism presented in [4] is tightly-secure even in a multi-user setting with corruptions. It avoids our impossibility result because it does not have unique secret keys.

More related work. Since their introduction by Boneh and Venkatesan in 1998 [13] meta-reductions have proven to be a versatile tool in many areas of provably security. Previous works have mainly used meta-reductions to derive impossibility results and efficiency/security bounds on signatures schemes [22,21,35,27,38,5,23,25], blind-signature schemes [24] and encryption systems [36]. In particular, among these results there exist several works that consider the existence of (tight) security proofs for the Schnorr signature scheme [35,27,38,5,25]. The results in [14,15] use meta-reductions to derive relationships among cryptographic one-more type problems. Lewko and Waters [34], building on [31], showed that under certain conditions it is impossible to prove security of hierarchical IBE (HIBE) schemes. To this end, Lewko and Waters extend the approach of [31] from signatures to hierarchical IBE to show that for certain HIBE schemes an *exponential* tightness loss is impossible to avoid. Finally, the inexistence of certain meta-reductions was considered in [23].

Outline. We begin with considering essentially the same setting as Coron and follow-up works [19,32,31], namely digital signatures in the single-user setting, as an instructive example. We prove a strengthened variant of the results of [19,32,31]. This allows us to explain how our new technique works in a known setting, which may be helpful for readers already familiar with these works. A generalized, much more abstract version will be presented in Section 4, Section 5 gives many further interesting applications, which seem not achievable using the previous approach of [19,32,31].

2 The New Meta-Reduction Technique

2.1 Preliminaries

Notation. We write $[n]$ to denote the set $[n] := \{1, 2, \dots, n\}$, and for $j \in [n]$ we write $[n \setminus j]$ to denote the set $[n] \setminus \{j\}$. If A is a set then $a \leftarrow^{\$} A$ denotes the action of sampling a uniformly from A . Given a set A we denote by U_A the uniform distribution on A . If A is a Turing machine (TM) then $a \leftarrow A(x; r)$ denotes that A outputs a when run with input x and random coins r . By $A(x)$ we denote the distribution of $a \leftarrow A(x; r)$ over

the uniform choice of r . If x is a binary string, then $|x|$ denotes its length. If M is a Turing machine, we denote by \widehat{M} its description as a bitstring.

If $t : \mathbb{N} \rightarrow \mathbb{N}$ and there exists a constant c such that $t(k) \leq k^c$ for all but finitely many $k \in \mathbb{N}$, then we say that $t \in \text{poly}(k)$. We denote by $\text{poly}^{-1}(k)$ the set $\text{poly}^{-1}(k) := \{\delta : \frac{1}{\delta} \in \text{poly}(k)\}$. We say that $\epsilon : \mathbb{N} \rightarrow [0, 1]$ is *negligible* if for all $c \in \mathbb{N}$ it holds that $\epsilon(k) > k^{-c}$ is true only for at most finitely many $k \in \mathbb{N}$. We write $\epsilon \in \text{negl}(k)$ to denote that ϵ is negligible.

Digital Signatures. A digital signature scheme $\text{SIG} = (\text{Setup}, \text{Gen}, \text{Sign}, \text{Vfy})$ is a four-tuple of PPT-TMs:

Public Parameters. The public parameter generation machine $\Pi \xleftarrow{\$} \text{Setup}(1^k)$ takes the security parameter k as input and returns public parameters Π .

Key Generation. The key generation machine takes as input public parameters Π and outputs a key pair, $(vk, sk) \xleftarrow{\$} \text{Gen}(\Pi)$.

Signing. The signing machine takes as input a secret key sk and a message m and returns a signature $\sigma \xleftarrow{\$} \text{Sign}(sk, m)$.

Verification. The verification machine, on input a public key vk , a signature σ and a message m , outputs 0 or 1, $\text{Vfy}(vk, m, \sigma) \in \{0, 1\}$.

Unique and re-randomizable signatures. Let $\Sigma(vk, m) := \{\sigma : \text{Vfy}(vk, m, \sigma) = 1\}$ denote the set of all valid signatures σ w.r.t. a given message m and verification key vk .

Definition 1 (Unique signatures). We say that SIG is a unique signature scheme, if $|\Sigma(vk, m)| = 1$ for all vk and m .

Definition 2 (Re-randomizable signatures). We say that SIG is t_{ReRand} -re-randomizable, if there exists a TM SIG.ReRand which takes as input (vk, m, σ) and outputs a signature $\sigma' \xleftarrow{\$} \text{SIG.ReRand}(vk, m, \sigma)$ with the following properties.

1. SIG.ReRand runs in time at most t_{ReRand}
2. If $\text{Vfy}(vk, m, \sigma) = 1$, then σ' is distributed uniformly over $\Sigma(vk, m)$.

Remark 1. Note that we do not put any bounds on t_{ReRand} . Thus, any signature scheme is t_{ReRand} -re-randomizable for sufficiently large t_{ReRand} . However, there are many examples of signature schemes which are *efficiently* re-randomizable, like the class of schemes considered in [31]. In particular, all unique signature schemes are efficiently re-randomizable by the Turing machine $\sigma \xleftarrow{\$} \text{SIG.ReRand}(vk, m, \sigma)$ which simply outputs its input σ .

Unforgeability under static message attacks. The UF-SMA security experiment is depicted in Figure 1.

Definition 3. Let $\text{UF-SMA}_{\text{SIG}}^{n, \mathcal{A}}(1^k)$ denote the UF-SMA security experiment depicted in Figure 1, executed with signature scheme SIG and attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$. We say that $\mathcal{A}(\text{t}_{\mathcal{A}}, n, \epsilon_{\mathcal{A}})$ -breaks the UF-SMA-security of SIG , if it runs in time $\text{t}_{\mathcal{A}}$ and

$$\Pr \left[\text{UF-SMA}_{\text{SIG}}^{n, \mathcal{A}}(1^k) \Rightarrow 1 \right] \geq \epsilon_{\mathcal{A}} .$$

<p>Game $\text{UF-SMA}_{\text{SIG}}^{n, \mathcal{A}}(1^k)$</p> <p>$\Pi \xleftarrow{\\$} \text{SIG.Setup}(1^k); \rho_{\mathcal{A}} \xleftarrow{\\$} \{0, 1\}^k$</p> <p>$(vk, sk) \xleftarrow{\\$} \text{SIG.Gen}(\Pi)$</p> <p>$m_1, \dots, m_n \xleftarrow{\\$} \mathcal{M}$ s.t. $m_i \neq m_j$ for all $i \neq j$</p> <p>$\sigma_i \xleftarrow{\\$} \text{SIG.Sign}(sk, m_i)$ for all $i \in [n]$</p> <p>$(j, st) \leftarrow \mathcal{A}_1(vk, (m_i)_{i \in [n]}; \rho_{\mathcal{A}})$</p> <p>$\sigma_j \leftarrow \mathcal{A}_2(st, (\sigma_i)_{i \in [n] \setminus \{j\}})$</p> <p>return $\text{SIG.Vfy}(vk, m_j, \sigma_j)$</p>
--

Fig. 1. The UF-SMA-security Game with attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$.

Remark 2. Observe that the messages in the UF-SMA security experiment from Figure 1 are chosen at random (but pairwise distinct). We do this for simplicity, but stress that for our tightness bound we actually do not have to make any assumption about the distribution of messages, apart from being pairwise distinct. For instance, the messages could alternatively be the lexicographically first n messages of the message space, for instance.

Non-interactive complexity assumptions. The following very general definition of non-interactive complexity assumptions is due to Abe et al. [2].

Definition 4. A non-interactive complexity assumption $N = (\mathsf{T}, \mathsf{V}, \mathsf{U})$ consists of three TMs. The instance generation machine $(c, w) \xleftarrow{\$} \mathsf{T}(1^k)$ takes the security parameter as input, and outputs a problem instance c and a witness w . U is a probabilistic polynomial-time machine, which takes as input c and outputs a candidate solution s . The verification TM V takes as input (c, w) and a candidate solution s . If $\mathsf{V}(c, w, s) = 1$, then we say that s is a correct solution to the challenge c .

Intuitively, U is a probabilistic polynomial-time machine which implements a suitable “trivial” attack strategy for N . This algorithm is used to define what “breaking” N with non-trivial success probability means, cf. Definition 5 below and [2].

Consider the following experiment $\text{NICA}_N^B(1^k)$.

1. The experiment runs the instance generator of N to generate a problem instance $(c, w) \xleftarrow{\$} \mathsf{T}(1^k)$. Then it samples uniformly random coins $\rho_B \xleftarrow{\$} \{0, 1\}^k$ for B .
2. B is executed on input (c, ρ_B) , it outputs a candidate solution s .
3. The experiment returns whatever $\mathsf{V}(c, w, s)$ returns.

Definition 5. We say that B (t, ϵ) -breaks assumption N , if A runs in time $t(k)$ and it holds that

$$\left| \Pr \left[\text{NICA}_N^B(1^k) \Rightarrow 1 \right] - \Pr \left[\text{NICA}_N^{\mathsf{U}}(1^k) \Rightarrow 1 \right] \right| \geq \epsilon(k)$$

where the probability is taken over the random coins consumed by T and the uniformly random choices of ρ_B and ρ_N respectively.

Simple reductions from non-interactive complexity assumptions to breaking UF-SMA-security. A reduction from breaking the UF-SMA-security of a signature scheme SIG to breaking the security of a non-interactive complexity assumption $N = (T, V, U)$ is a TM, which turns an attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ according to Definition 3 into a TM $\Lambda^{\mathcal{A}}$ according to Definition 5.

Following [19,32,31,34], we will consider a specific class of reductions in the sequel. We consider reductions having *black-box access* to the attacker, and which execute the attacker *only once* and *without rewinding*. We will generalize this later to reductions that may execute the attacker several times sequentially. Following [34], we call such reductions *simple*. At first sight we heavily constrain the class of reductions to that our result applies. However, as explained in [34], we include reductions that perform hybrid steps. Moreover, most reductions in cryptography are simple.

For preciseness and clarity, we define such a reduction as a triplet of Turing machines $\Lambda = (\Lambda_1, \Lambda_2, \Lambda_3)$. From these TMs and an attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, we construct a Turing machine $\Lambda^{\mathcal{A}}$ for a non-interactive complexity assumption as follows.

1. Machine $\Lambda^{\mathcal{A}}$ receives as input a challenge c of the considered non-interactive complexity assumption, as well as random coins $\rho_{\Lambda} \leftarrow^{\$} \{0, 1\}^k$. It first runs $\Lambda_1(c, \rho_{\Lambda})$, which returns the input to \mathcal{A}_1 , consisting of a verification key vk , a sequence of messages $(m_i)_{i \in [n]}$, and random coins $\rho_{\mathcal{A}}$, as well as some state st_{Λ_2} .
2. Then $\Lambda^{\mathcal{A}}$ executes the attacker \mathcal{A}_1 on input $(vk, (m_i)_{i \in [n]}, \rho_{\mathcal{A}})$, which returns an index $j^* \in [n]$ and some state $st_{\mathcal{A}}$.
3. TM Λ_2 receives as input j^* and state st_{Λ_2} , and returns a list of signatures $(\sigma_i)_{i \in [n \setminus j^*]}$ and an updated state st_{Λ_3} .
4. The attacker \mathcal{A}_2 is executed on $(\sigma_i)_{i \in [n \setminus j^*]}$ and state $st_{\mathcal{A}}$, it returns a signature σ^* .
5. Finally, $\Lambda^{\mathcal{A}}$ runs $\Lambda_3(\sigma^*, j^*, st_{\Lambda_3})$, which produces a candidate solution s , and outputs s .

Definition 6. We say that a Turing machine $\Lambda = (\Lambda_1, \Lambda_2, \Lambda_3)$ is a simple $(t_{\Lambda}, n, \epsilon_{\Lambda}, \epsilon_{\mathcal{A}})$ -reduction from breaking $N = (T, V, U)$ to breaking the UF-SMA-security of SIG, if for any TM \mathcal{A} that $(t_{\mathcal{A}}, n, \epsilon_{\mathcal{A}})$ -breaks the UF-SMA security of SIG, TM $\Lambda^{\mathcal{A}}$ $(t_{\Lambda} + t_{\mathcal{A}}, \epsilon_{\Lambda})$ -breaks N .

Definition 7. Let $\ell : \mathbb{N} \rightarrow \mathbb{N}$. We say that reduction Λ loses ℓ , if there exists an adversary \mathcal{A} that $(t_{\mathcal{A}}, n, \epsilon_{\mathcal{A}})$ -breaks the UF-SMA security of SIG, such that $\Lambda^{\mathcal{A}}$ $(t_{\Lambda} + t_{\mathcal{A}}, \epsilon_{\Lambda})$ -breaks N with

$$\frac{t_{\Lambda}(k) + t_{\mathcal{A}}(k)}{\epsilon_{\Lambda}(k)} \geq \ell(k) \cdot \frac{t_{\mathcal{A}}(k)}{\epsilon_{\mathcal{A}}(k)}.$$

Remark 3. The quotient $t_{\mathcal{A}}(k)/\epsilon_{\mathcal{A}}(k)$ of the running time $t_{\mathcal{A}}(k)$ and the success probability $\epsilon_{\mathcal{A}}(k)$ of a Turing machine \mathcal{A} is called the *work factor* of \mathcal{A} [8]. Thus, the factor ℓ in Definition 6 relates the work factor of attacker \mathcal{A} to the work factor of TM $\Lambda^{\mathcal{A}}$, which allows us to measure the *tightness* of a cryptographic reduction. The smaller ℓ , the tighter is the reduction.

2.2 Bound for Simple Reductions without Rewinding

For simplicity, we will consider reductions that have access to a “perfect” adversary \mathcal{A} , which $(t_{\mathcal{A}}, \epsilon_{\mathcal{A}})$ -breaks the signature scheme with $\epsilon_{\mathcal{A}} = 1$. We explain in Section 2.4 why the extension to adversaries with $\epsilon_{\mathcal{A}} < 1$ is straightforward.

Theorem 1. *Let $N = (T, V, U)$ be a non-interactive complexity assumption, $n \in \text{poly}(k)$ and let SIG be a signature scheme. For any simple $(t_{\mathcal{A}}, n, \epsilon_{\mathcal{A}}, 1)$ -reduction from breaking N to breaking the UF-SMA-security of SIG, there exists a Turing machine \mathcal{B} that $(t_{\mathcal{B}}, \epsilon_{\mathcal{B}})$ -breaks N where*

$$t_{\mathcal{B}} \leq n \cdot t_{\mathcal{A}} + n \cdot (n - 1) \cdot t_{\text{Vfy}} + t_{\text{ReRand}} \quad \text{and} \quad \epsilon_{\mathcal{B}} \geq \epsilon_{\mathcal{A}} - 1/n.$$

Here, t_{ReRand} is the time required to re-randomize a signature, and t_{Vfy} is the running time of the verification machine of SIG.

Proof. Our proof structure follows the structure of [31] (also used in [34]). That is, we first describe a hypothetical, inefficient adversary, then we show how to simulate it efficiently for certain reductions.

The Hypothetical Adversary. The hypothetical adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ consists of two procedures that work as follows.

$\mathcal{A}_1(vk, (m_i)_{i \in [n]}; \rho_{\mathcal{A}})$. On input a public key vk and messages m_1, \dots, m_n , \mathcal{A}_1 samples $j \leftarrow^{\$} [n]$ uniformly random and outputs (j, st) , where $st = (vk, (m_i)_{i \in [n]}, j)$. $\mathcal{A}_2((\sigma_i)_{i \in [n \setminus j]}, st)$. \mathcal{A}_2 checks whether $\text{SIG.Vfy}(vk, m_i, \sigma_i) = 1$ for all $i \in [n \setminus j]$.

If this holds, then it samples a uniformly random signature $\sigma_j \leftarrow^{\$} \Sigma(vk, m_j)$ for m_j . Finally, it outputs σ_j .

Note that $\mathcal{A}(t_{\mathcal{A}}, 1)$ -breaks the UF-SMA-security of SIG. Note also that the second step of this adversary may not be efficiently computable, which is why we call this adversary *hypothetical*.

Simulating \mathcal{A} . Consider the following TM \mathcal{B} , which runs reduction $\Lambda = (\Lambda_1, \Lambda_2, \Lambda_3)$ as a subroutine and attempts to break N . \mathcal{B} receives as input $c \leftarrow^{\$} T(1^k)$. It maintains an array A with n entries, which are all initialized to \emptyset , and proceeds as follows.

1. \mathcal{B} first runs $(vk, (m_i)_{i \in [n]}, \rho_{\mathcal{A}}, st_{\Lambda_2}) \leftarrow^{\$} \Lambda_1(c; \rho_{\mathcal{A}})$ for uniformly random $\rho_{\mathcal{A}} \leftarrow^{\$} \{0, 1\}^k$.
2. Next, \mathcal{B} runs $\Lambda_2(j, st_{\Lambda_2})$ for each $j \in [n]$. Let $((\sigma_{i,j})_{i \in [n \setminus j]}, st_{\Lambda_3,j})$ denote the output of the j -th execution of Λ_2 . Whenever Λ_2 outputs $(\sigma_{i,j})_{i \in [n \setminus j]}$ such that

$$\text{SIG.Vfy}(vk, m_i, \sigma_{i,j}) = 1 \text{ for all } i \in [n \setminus j]$$

then it sets $A[i] \leftarrow \sigma_{i,j}$ for all $i \in [n \setminus j]$.

3. \mathcal{B} samples $j^* \leftarrow^{\$} [n]$. Then it proceeds as follows.
 - If there exists an index $i \in [n \setminus j^*]$ such that $\text{SIG.Vfy}(vk, m_i, \sigma_{i,j^*}) \neq 1$, then \mathcal{B} sets $\sigma^* := \perp$.
 - Otherwise, if $\text{SIG.Vfy}(vk, m_i, \sigma_{i,j^*}) = 1$ for all $i \in [n \setminus j^*]$, then \mathcal{B} computes

$$\sigma^* \leftarrow^{\$} \text{SIG.ReRand}(vk, m_{j^*}, A[j^*]).$$

4. Finally, \mathcal{B} runs $s \leftarrow \Lambda_3(\sigma^*, j^*, st_{\Lambda_3,j^*})$ and outputs s . Note that the state st_{Λ_3,j^*} used to execute Λ_3 corresponds to the state returned by Λ_2 on its j^* -th execution.

Running time of \mathcal{B} . \mathcal{B} essentially runs each part of Turing machine $\Lambda = (\Lambda_1, \Lambda_2, \Lambda_3)$ once, plus $n - 1$ additional executions of Λ_2 . Moreover, it executes SIG.Vfy $n(n - 1)$ times, and the re-randomization TM SIG.ReRand once. Thus, the total running time of \mathcal{B} is at most

$$t_{\mathcal{B}} \leq n \cdot t_{\Lambda} + n \cdot (n - 1) \cdot t_{\text{Vfy}} + t_{\text{ReRand}}.$$

Success probability of \mathcal{B} . To analyze the success probability of \mathcal{B} , let us define an event bad. Intuitively, this event occurs, if j^* is the *only* (with respect to state st_{Λ_2}) value such that $\Lambda_2(st_{\Lambda_2}, j)$ outputs signatures which are all valid. More formally, for both experiments $\text{NICA}_N^{\mathcal{B}}(1^k)$ and $\text{NICA}_N^{\Lambda^{\mathcal{A}}}(1^k)$, let st_{Λ_2} denote the (in both experiments unique) value computed by $\Lambda_1(c; \rho_{\Lambda})$, and let j^* denote the (in both experiments unique) value given as input to $\Lambda_3(\sigma^*, j^*, st_{\Lambda_3, j^*})$. We say that bad occurs (in either $\text{NICA}_N^{\mathcal{B}}(1^k)$ or $\text{NICA}_N^{\Lambda^{\mathcal{A}}}(1^k)$), if $\text{pred}(st_{\Lambda_2}, j^*) = 1 \wedge \text{pred}(st_{\Lambda_2}, j) = 0 \forall j \in [n \setminus j^*]$, where predicate pred is defined as

$$\begin{aligned} \text{pred}(st_{\Lambda_2}, j) &= 1 \\ \iff \bigwedge_{i \in [n \setminus j]} \text{SIG.Vfy}(vk, m_i, \sigma_i) &= 1, \text{ where } ((\sigma_i)_{i \in [n \setminus j]}, st_{\Lambda_3}) \leftarrow \Lambda_2(st_{\Lambda_2}, j). \end{aligned}$$

Note that pred is well-defined, because Λ_2 is a deterministic TM.

Let us write $S(\mathcal{F})$ shorthand for the event $\text{NICA}_N^{\mathcal{F}}(1^k) \Rightarrow 1$ to abbreviate our notation. Then, it holds that

$$|\Pr[S(\mathcal{B})] - \Pr[S(\Lambda^{\mathcal{A}})]| \leq |\Pr[S(\mathcal{B}) \cap \neg \text{bad}] - \Pr[S(\Lambda^{\mathcal{A}}) \cap \neg \text{bad}]| + \Pr[\text{bad}]. \quad (2)$$

Bounding $\Pr[\text{bad}]$. Recall that event bad occurs only if

$$\text{pred}(st_{\Lambda_2}, j^*) = 1 \wedge \text{pred}(st_{\Lambda_2}, j) = 0 \forall j \in [n \setminus j^*] \quad (3)$$

where st_{Λ_2} is the value computed by $\Lambda_1(c; \rho_{\Lambda})$, and j^* is the value given as input to $\Lambda_3(\sigma^*, j^*, st_{\Lambda_3, j^*})$. Suppose that indeed st_{Λ_2} is such that there exist at least one $j^* \in [n]$ such that (3) holds. We claim that even then we have

$$\Pr[\text{bad}] \leq 1/n. \quad (4)$$

To see this, note first that for each st_{Λ_2} there can be at most one value j^* that satisfies (3). Moreover, both the hypothetical adversary \mathcal{A} and the adversary simulated by \mathcal{B} choose $j^* \leftarrow^{\$} [n]$ independently and uniformly random, which yields (4).

Proving $\Pr[S(\mathcal{B}) \cap \neg \text{bad}] = \Pr[S(\Lambda^{\mathcal{A}}) \cap \neg \text{bad}]$. Note that \mathcal{B} executes in particular

1. $(vk, (m_i)_{i \in [n]}, st_{\Lambda_2}) \leftarrow^{\$} \Lambda_1(c; \rho_{\Lambda})$
2. $((\sigma_{i, j^*})_{i \in [n \setminus j^*]}, st_{\Lambda_3}) \leftarrow^{\$} \Lambda_2(j^*, st_{\Lambda_2})$
3. $s \leftarrow \Lambda_3(\sigma^*, j^*, st_{\Lambda_3})$.

We show that if $\neg \text{bad}$ occurs, then \mathcal{B} simulates the hypothetical adversary \mathcal{A} perfectly. To this end, consider the distribution of σ^* computed by \mathcal{B} in following two cases.

1. Machine $\mathcal{A}_2(j^*, st_{\mathcal{A}_2})$ outputs $((\sigma_{i,j^*})_{i \in [n \setminus j^*]}, st_{\mathcal{A}_3,j^*})$ such that there exists an index $i \in [n \setminus j^*]$ with $\text{SIG.Vfy}(vk, m_i, \sigma_{i,j^*}) \neq 1$.
In this case, \mathcal{A} would compute $\sigma^* := \perp$. \mathcal{B} also sets $\sigma^* := \perp$ in this case.
2. TM $\mathcal{A}_2(j^*, st_{\mathcal{A}_2})$ outputs $((\sigma_{i,j^*})_{i \in [n \setminus j^*]}, st_{\mathcal{A}_3,j^*})$ such that for all $i \in [n \setminus j^*]$ it holds that

$$\text{SIG.Vfy}(vk, m_i, \sigma_{i,j^*}) = 1 .$$

In this case, \mathcal{A} would output a uniformly random signature $\sigma^* \xleftarrow{\$} \Sigma(vk, m_{j^*})$. Note that in this case \mathcal{B} outputs a re-randomized signature $\sigma^* \xleftarrow{\$} \text{SIG.ReRand}(vk, m_{j^*}, A[j^*])$, which is a uniformly distributed valid signature for m_{j^*} provided that $A[j^*] \neq \emptyset$. The latter happens whenever bad does not occur.

Thus, \mathcal{B} simulates \mathcal{A} perfectly in either case, provided that $\neg \text{bad}$. This implies $S(\mathcal{B}) \cap \neg \text{bad} \iff S(\mathcal{A}) \cap \neg \text{bad}$, which yields

$$\Pr[S(\mathcal{B}) \cap \neg \text{bad}] = \Pr[S(\mathcal{A}) \cap \neg \text{bad}] . \quad (5)$$

Finishing the proof of Theorem 1. By plugging (4) and (5) into Inequality (2), we obtain

$$|\Pr[S(\mathcal{B})] - \Pr[S(\mathcal{A})]| \leq 1/n$$

which implies

$$\epsilon_{\mathcal{B}} = |\Pr[S(\mathcal{B})] - \Pr[S(\mathcal{U})]| \geq |\Pr[S(\mathcal{A})] - \Pr[S(\mathcal{U})]| - 1/n = \epsilon_{\mathcal{A}} - 1/n .$$

2.3 Interpretation

Assuming that no adversary \mathcal{B} is able to (t_N, ϵ_N) -break the security of NICA with $t_N = t_{\mathcal{B}} = n \cdot t_{\mathcal{A}} + n \cdot (n-1) \cdot t_{\text{Vfy}} + t_{\text{ReRand}}$, we must have $\epsilon_{\mathcal{B}} \leq \epsilon_N$. By Theorem 1, we thus must have

$$\epsilon_{\mathcal{A}} \leq \epsilon_{\mathcal{B}} + 1/n \leq \epsilon_N + 1/n$$

for all reductions \mathcal{A} . In particular, the hypothetical adversary \mathcal{A} constructed in the proof of Theorem 1 is an example of an adversary such that

$$\frac{t_{\mathcal{A}} + t_{\mathcal{A}}}{\epsilon_{\mathcal{A}}} \geq \frac{t_{\mathcal{A}}}{\epsilon_N + 1/n} = (\epsilon_N + 1/n)^{-1} \cdot \frac{t_{\mathcal{A}}}{1} = (\epsilon_N + 1/n)^{-1} \cdot \frac{t_{\mathcal{A}}}{\epsilon_{\mathcal{A}}} .$$

Thus, any reduction \mathcal{A} from breaking the security of NICA N to breaking the UF-SMA-security of signature scheme SIG loses (in the sense of Definition 7) at least a factor of $\ell \geq 1/(\epsilon_N + 1/n)$. In particular, note that $\ell \approx n$ if ϵ_N is very small. This yields the following informal theorem.

Theorem 2 (informal). *Any simple reduction from breaking the security of NICA N to breaking the UF-SMA-security (or any stronger security notion, like EUF-CMA-security, cf. Definition 22) of signature scheme SIG that provides efficient signature re-randomization loses a factor that is at least linear in the number n of sign queries issued by the attacker, or N is easy to solve.*

Remark 4. Since a unique signature scheme is trivially efficiently re-randomizable, Theorem 2 applies also to unique signature schemes.

```

TM  $r$ - $\Lambda^A(c; \rho_A)$ 
 $st_{\Lambda_{1,1}} \leftarrow \Lambda_0(c, \rho_A)$ 
for  $1 \leq l \leq r$  do:
     $(vk^l, (m_i^l)_{i \in [n]}, \rho_A, st_{\Lambda_{l,2}}) \leftarrow \Lambda_{l,1}(st_{\Lambda_{l,1}})$ 
     $(j^{l*}, st_A) \leftarrow \mathcal{A}_1(vk^l, (m_i^l)_{i \in [n]}; \rho_A)$ 
     $((\sigma_i^l)_{i \in [n \setminus j^{l*}]}, st_{\Lambda_{l,3}}) \leftarrow \Lambda_{l,2}(j^{l*}, st_{\Lambda_{l,2}})$ 
     $\sigma_{j^{l*}}^l \leftarrow \mathcal{A}_2((\sigma_i^l)_{i \in [n \setminus j^{l*}]}, st_A)$ 
     $st_{\Lambda_{l+1,1}} \leftarrow \Lambda_{l,3}(\sigma_{j^{l*}}^l, j^{l*}, st_{\Lambda_{l,3}})$ 
 $s \leftarrow \Lambda_3(st_{\Lambda_{r+1,1}})$ 
return  $s$ 

```

Fig. 2. TM r - Λ^A that solves a non-interactive complexity assumption according to Definition 5, constructed from a r -simple reduction r - $\Lambda = (\Lambda_0, (\Lambda_{l,1}, \Lambda_{l,2}, \Lambda_{l,3})_{l \in [r]}, \Lambda_3)$ and an attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$.

2.4 Extension to “non-perfect” adversaries

Note that the proof of Theorem 1 trivially generalizes to $(t_\Lambda, n, \epsilon_\Lambda, \epsilon_{\mathcal{A}})$ -reductions with $\epsilon_{\mathcal{A}} < 1$, that is, reductions that have access to an adversary which has success probability $\epsilon_{\mathcal{A}} < 1$. To this end, we first would have to describe a hypothetical adversary, which has success probability $\epsilon_{\mathcal{A}}$. This is simple, because we can simply let the hypothetical adversary constructed above toss a biased coin χ with $\Pr[\chi = 1] = \epsilon_{\mathcal{A}}$, such that \mathcal{A} outputs σ^* only if $\chi = 1$. Note that in the proof of Theorem 1 we are even able to simulate a perfect adversary \mathcal{A} . Therefore we would also be able to simulate the non-perfect adversary sketched above, by tossing a biased coin χ and outputting σ^* only if $\chi = 1$. This yields the following theorem.

Theorem 3. *Let $N = (T, V, U)$ be a non-interactive complexity assumption, $n \in \text{poly}(k)$ and let SIG be a signature scheme. For any simple $(t_\Lambda, n, \epsilon_\Lambda, \epsilon_{\mathcal{A}})$ -reduction from breaking the UF-SMA-security of SIG to breaking N , there exists a Turing machine \mathcal{B} that (t_B, ϵ_B) -breaks N where*

$$t_B \leq n \cdot t_\Lambda + n \cdot (n - 1) \cdot t_{\text{Vfy}} + t_{\text{ReRand}} \quad \text{and} \quad \epsilon_B \geq \epsilon_\Lambda - 1/n.$$

Here, t_{ReRand} is the time to re-randomize a given valid signature over a message and t_{Vfy} is the time needed to execute the verification machine of SIG.

3 Bound for Reductions with Sequential Rewinding

Theorem 1 applies only to reductions that run the forger only once. Here we show that under assumptions similar to that in Theorem 1 the work factor of any reduction that is allowed to run or rewind the adversary r times *sequentially* cannot decrease significantly below $\frac{n}{r}$ if N is hard.

Let r be an upper bound on the number of times that the adversary can be rewound by the reduction. We then consider a reduction r - Λ as a $3 \cdot r + 2$ -tuple of Turing machines

$r\text{-}\Lambda = (\Lambda_0, (\Lambda_{l,1}, \Lambda_{l,2}, \Lambda_{l,3})_{l \in [r]}, \Lambda_3)$. Let now $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be an attacker against the UF-SMA-security of SIG. From these TMs we construct a Turing machine $r\text{-}\Lambda^{\mathcal{A}}$ that solves a NICA N as depicted in Figure 2. We shortly explain Figure 2 here.

Λ_0 . $r\text{-}\Lambda$ inputs a challenge c of the considered non-interactive complexity assumption and random coins ρ_{Λ} . It processes these inputs by running Λ_0 which outputs a state st_{Λ} .

$\Lambda_l = (\Lambda_{l,1}, \Lambda_{l,2}, \Lambda_{l,3})$. Now, for each $l \in [r]$, we have a triplet of TMs $\Lambda_l = (\Lambda_{l,1}, \Lambda_{l,2}, \Lambda_{l,3})$ that has black box access to attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$. Note that the state st_{Λ} may be passed over from $\Lambda_{l,3}$ to $\Lambda_{l+1,1}$ (and Λ_3) while the state $st_{\mathcal{A}}$ of \mathcal{A}_2 may not be passed over to the next execution of \mathcal{A}_1 .

$\Lambda_{l,1}$. $\Lambda_{l,1}$ inputs the current state $st_{\Lambda_{l,1}}$ and outputs a public key vk^l , distinct messages $m_i^l, i \in [n]$, a random tape $\rho_{\mathcal{A}}$ for \mathcal{A}_1 and a state $st_{\Lambda_{l,2}}$. Next, \mathcal{A}_1 is run on input $(vk^l, (m_i)_{i \in [n]}; \rho_{\mathcal{A}})$ and returns a state $st_{\mathcal{A}}$ and an index j^l .

$\Lambda_{l,2}$. On input index j^l and state $st_{\Lambda_{l,2}}$, $\Lambda_{l,2}$ returns signatures $(\sigma_i^l)_{i \in [n \setminus j]}$ and state $st_{\Lambda_{l,2}}$. Now, \mathcal{A}_2 is run on $((\sigma_i^l)_{i \in [n \setminus j^l]}, st_{\mathcal{A}})$ and returns $\sigma_{j^l}^l$.

$\Lambda_{l,3}$. $\Lambda_{l,3}$ inputs the signature output by $\Lambda_{l,2}$ and the current state $st_{\Lambda_{l,2}}$. It returns the state $st_{\Lambda_{l+1,1}}$.

Λ_3 . Finally, Λ_3 inputs the current state of $r\text{-}\Lambda$ and returns s . $r\text{-}\Lambda$ is considered successful if $V(c, w, s) = 1$.

Definition 8. We say that a Turing machine $r\text{-}\Lambda = (\Lambda_0, (\Lambda_{l,1}, \Lambda_{l,2}, \Lambda_{l,3})_{l \in [r]}, \Lambda_3)$ is an r -simple $(t_{\Lambda}, n, \epsilon_{\Lambda}, \epsilon_{\mathcal{A}})$ -reduction from breaking $N = (T, V, U)$ to breaking the UF-SMA-security of SIG, if for any TM \mathcal{A} that $(t_{\mathcal{A}}, n, \epsilon_{\mathcal{A}})$ -breaks the UF-SMA security of SIG, TM $r\text{-}\Lambda^{\mathcal{A}}$ (as constructed above) $(t_{\Lambda} + r \cdot t_{\mathcal{A}}, \epsilon_{\Lambda})$ -breaks N .

Definition 9. Let $\ell : \mathbb{N} \rightarrow \mathbb{N}$. We say that an r -simple reduction Λ from breaking a non-interactive complexity assumption N to breaking the UF-SMA security of a signature scheme SIG loses ℓ if there exists an adversary \mathcal{A} that $(t_{\mathcal{A}}, n, \epsilon_{\mathcal{A}})$ -breaks such that $\Lambda^{\mathcal{A}}$ $(t_{\Lambda} + r \cdot t_{\mathcal{A}}, \epsilon_{\Lambda})$ -breaks N where

$$\frac{t_{\Lambda}(k) + r \cdot t_{\mathcal{A}}(k)}{\epsilon_{\Lambda}} \geq \ell(k) \cdot \frac{t_{\mathcal{A}}(k)}{\epsilon_{\mathcal{A}}(k)}.$$

Theorem 4. Let $N = (T, V, U)$ be a non-interactive complexity assumption, $n, r \in \text{poly}(k)$ and let SIG be a signature scheme. Then for any r -simple $(t_{\Lambda}, n, \epsilon_{\Lambda}, 1)$ -reduction Λ from breaking N to breaking the UF-SMA-security of SIG there exists a TM \mathcal{B} that $(t_{\mathcal{B}}, \epsilon_{\mathcal{B}})$ -breaks N where

$$\begin{aligned} t_{\mathcal{B}} &\leq r \cdot n \cdot t_{\Lambda} + r \cdot n \cdot (n-1) \cdot t_{\text{Vfy}} + r \cdot t_{\text{ReRand}} \\ \epsilon_{\mathcal{B}} &\geq \epsilon_{\Lambda} - \frac{r}{n}. \end{aligned}$$

Here, t_{ReRand} is the time to re-randomize a given valid signature over a message and t_{Vfy} is the time needed to run the verification machine of SIG.

Proof. The proof is structured as the proof of Theorem 1. Namely, we consider again our hypothetical attacker \mathcal{A} (cf. Page 11) that breaks the UF-SMA-security of SIG. Since \mathcal{A} works exactly as in the proof of Theorem 1, we do not describe it here again but we start directly with our analysis. We stress however, once again, that \mathcal{A} $(t_{\mathcal{A}}, 1)$ -breaks N .

Next, we show how to simulate \mathcal{A} . Roughly speaking, we will apply the technique from the proof of Theorem 1 r times.

Simulating \mathcal{A} r times sequentially. As in the proof of Theorem 1, we describe a TM \mathcal{B} that runs TM r - $\mathcal{A} = (\mathcal{A}_0, (\mathcal{A}_{l,1}, \mathcal{A}_{l,2}, \mathcal{A}_{l,3}), \mathcal{A}_3)$ as a subroutine. Recall that the goal of \mathcal{B} is to break N and that \mathcal{B} is called on input $c \leftarrow^{\$} \mathcal{T}(1^k)$ and random tape ρ . \mathcal{B} proceeds as follows:

- i. Run $st_{\mathcal{A}_{l,1}} \leftarrow \mathcal{A}_0(c, \rho_{\mathcal{A}})$ for uniformly $\rho_{\mathcal{A}}$. If \mathcal{A}_0 does not output $st_{\mathcal{A}_{l,1}}$, \mathcal{B} aborts. Note that, since the input to \mathcal{A}_0 is fixed (including random coins $\rho_{\mathcal{A}}$), we may view \mathcal{A}_0 (and all following subroutines of \mathcal{A}) as deterministic.
- ii. For $1 \leq l \leq r$, \mathcal{B} performs the following steps (that bare similarity with the simulation in the proof of Theorem 1).

Round l . Initialize an array A^l that has n entries, all initialized to \emptyset .

- a) Run $\mathcal{A}_{l,1}(st_{\mathcal{A}_{l,1}})$. If \mathcal{A} outputs $(vk^l, (m_i^l)_{i \in [n]}, \rho_{\mathcal{A}}, st_{\mathcal{A}_{l,2}})$, continue. Otherwise stop.
- b) Next, run $((\sigma_{i,j}^l)_{i \in [n \setminus j]}, st_{\mathcal{A}_{l,3}}) \leftarrow \mathcal{A}_{l,2}(j, st_{\mathcal{A}_{l,2}})$ for each $j \in [n]$. In case all signatures returned by $\mathcal{A}_{l,2}$ are valid, i.e., if

$$\bigwedge_{i \in [n \setminus j]} \text{SIG.Vfy}(vk^l, m_i^l, \sigma_{i,j}^l) = 1$$

set $A^l[i] \leftarrow \sigma_{i,j}^l$ for all $i \in [n \setminus j]$.

- c) Sample $j^{l*} \leftarrow^{\$} [n]$. We distinguish between two cases:

- 1) If for any $i \in [n \setminus j^{l*}]$ it holds that $\text{SIG.Vfy}(vk^l, m_i^l, \sigma_{i,j^{l*}}^l) \neq 1$, set $\sigma^* \leftarrow \perp$.
- 2) Otherwise, i.e., if $\text{SIG.Vfy}(vk^l, m_i^l, \sigma_{i,j^{l*}}^l) = 1$ for all $i \in [n \setminus j^{l*}]$ set $\sigma^* \leftarrow^{\$} \text{SIG.ReRand}(vk^l, m_{j^{l*}}^l, A^l[j^{l*}])$.

- d) Run $st_{\mathcal{A}_{l+1,1}} \leftarrow \mathcal{A}_{l,3}(\sigma^*, st_{\mathcal{A}_{l,3}})$.

- iii. Finally, \mathcal{B} runs $s \leftarrow \mathcal{A}_3(st_{\mathcal{A}_{r+1,1}})$ and returns s .

Analysis. Let us first note that the running time of \mathcal{B} is essentially bounded by the running time of r - \mathcal{A} plus the running time to execute $\mathcal{A}_{l,2}$ $n - 1$ times, the time to run SIG.Vfy for $n \cdot (n - 1)$ times and the time to run SIG.ReRand for all $1 \leq l \leq r$. We thus obtain an upper bound on the total running time t of \mathcal{A} by

$$t \leq r \cdot n \cdot t_{\mathcal{A}} + r \cdot n \cdot (n - 1) \cdot t_{\text{Vfy}} + r \cdot t_{\text{ReRand}}$$

Success probability of \mathcal{B} . Similar to the proof of Theorem 1 we define events $\text{bad}(l)$, $l \in [r]$ to analyze the success probability of \mathcal{B} . Informally, $\text{bad}(l)$ occurs if, given $st_{\Lambda_{l,2}}, j^{l*}$ is the *only* value such that $\Lambda_{l,2}(j, st_{\Lambda_{l,2}})$ outputs signatures that are all valid. As before, let us denote by $st_{\Lambda_{l,2}}$ the unique state computed by $\Lambda_{l,1}$ and let j^{l*} denote the unique value input to $\Lambda_{l,1}$ ($\sigma^*, j^{l*}, st_{\Lambda_{l,3}}$). We note that $st_{\Lambda_{l,2}}$ as well as j^{l*} are well-defined in both experiments $\text{NICA}_N^{\mathcal{B}}(1^k)$ and $\text{NICA}_N^{\mathcal{A}}(1^k)$.

We say that $\text{bad}(l)$ occurs if $\text{pred}(st_{\Lambda_{l,2}}, j^{l*}) = 1 \wedge \text{pred}(st_{\Lambda_{l,2}}, j) = 0 \forall j \in [n \setminus j^{l*}]$ where

$$\begin{aligned} \text{pred}(st_{\Lambda_{l,2}}, j) &= 1 \\ \Leftrightarrow \bigwedge \text{SIG.Vfy}(vk^l, m_i^l, \sigma_i^l) &= 1 : \left((\sigma_i^l)_{i \in [n \setminus j]}, st_{\Lambda_{l,3}} \right) \leftarrow \Lambda_{l,2}(st_{\Lambda_{l,2}}, j) \end{aligned}$$

Now, we define event $\text{bad} = \bigvee_{l \in [r]} \text{bad}(l)$.

Let us denote by $S(\mathcal{F})$ denote the event $\text{NICA}_N^{\mathcal{F}}(1^k) \Rightarrow 1$. As in the proof of Theorem 1 (cf. page12) we obtain:

$$\Pr[S(\mathcal{A})] - \Pr[S(\mathcal{B})] \leq \Pr[S(\mathcal{A}) \cap \neg \text{bad}] - \Pr[S(\mathcal{B}) \cap \neg \text{bad}] + \Pr[\text{bad}] \quad (6)$$

Bounding $\Pr[\text{bad}]$. Recall that event bad occurs only if there exists any $l \in [r]$ such that $\text{bad}(l)$ occurs, i.e.,

$$\text{pred}(st_{\Lambda_{l,2}}, j^{l*}) = 1 \wedge \text{pred}(st_{\Lambda_{l,2}}, j) = 0 \forall j \in [n \setminus j^{l*}] \quad (7)$$

where $st_{\Lambda_{l,2}}$ is the value computed by $\Lambda_{l,1}(c; st_{\Lambda_{l,1}})$, and j^{l*} is the value given as input to $\Lambda_{l,3}(\sigma^*, j^{l*}, st_{\Lambda_{l,3}})$. We claim that then we have $\Pr[\text{bad}(l)] \leq 1/n$. To see this, note first that for $\text{bad}(l)$ to occur it is necessary that there is only one value j^{l*} that satisfies $\text{pred}(st_{\Lambda_{l,2}}, j^{l*})$. Moreover, bad does only occur if this index j^{l*} is input to $\Lambda_{l,3}$. Note that both the hypothetical adversary \mathcal{A} and the adversary simulated by \mathcal{B} choose j^{l*} independently and uniformly random, which yields the claim. Now, by the union bound, we obtain

$$\Pr[\text{bad}] \leq r/n \quad (8)$$

$\Pr[S(\mathcal{B}) \cap \neg \text{bad}] = \Pr[S(\mathcal{A}) \cap \neg \text{bad}]$. Note that \mathcal{B} executes in particular:

- 1) $st_{\Lambda_{1,1}} \leftarrow \Lambda_0(c; \rho_{\mathcal{A}})$
- 2a) $(vk^l, (m_i^l)_{i \in [n]}, st_{\Lambda_{l,2}}) \xleftarrow{\$} \Lambda_{l,1}(c, st_{\Lambda_{l,1}})$
- 2b) $((\sigma_{i,j^{l*}})_{i \in [n \setminus j^{l*}]}, st_{\Lambda_{l,3}}) \xleftarrow{\$} \Lambda_{l,2}(j^{l*}, st_{\Lambda_{l,3}})$
- 2c) $st_{\Lambda_{l+1,1}} \leftarrow \Lambda_{l,3}(\sigma^*, j^{l*}, st_{\Lambda_{l,3}})$
- 3) $s \leftarrow \Lambda_3(st_{\Lambda_{r+1,1}})$

where steps 2a) - 2c) are carried out for each $l \in [r]$. We show that if $\neg \text{bad}(l)$ occurs, then \mathcal{B} simulates the hypothetical adversary \mathcal{A} perfectly for $\Lambda_l = (\Lambda_{l,1}, \Lambda_{l,2}, \Lambda_{l,3})$. To this end, consider the distribution of σ^* that is input to $\Lambda_{l,3}$ in following two cases.

1. TM $\Lambda_{l,2}(j^{l*}, st_{\Lambda_{l,2}})$ outputs $((\sigma_{i,j^{l*}})_{i \in [n \setminus j^{l*}]}, st_{\Lambda_{l,3}})$ such that there exists an index $i \in [n \setminus j^{l*}]$ with $\text{SIG.Vfy}(vk^l, m_i^l, \sigma_{i,j^{l*}}) \neq 1$. In this case, \mathcal{A} would compute $\sigma^* := \perp$ which is also output by \mathcal{B} .

2. Turing machine $A_{l,2}(j^{l*}, st_{A_{l,2}})$ outputs $((\sigma_{i,j^{l*}})_{i \in [n \setminus j^{l*}]} st_{A_{l,3}})$ such that

$$\text{SIG.Vfy}(vk^l, m_i^l, \sigma_{i,j^{l*}}^l) = 1$$

for all $i \in [n \setminus j^{l*}]$. In this case, \mathcal{A} would output a uniformly random signature $\sigma^* \leftarrow^{\$} \Sigma(vk^l, m_{j^{l*}}^l)$. Note that in this case \mathcal{B} outputs a re-randomized signature $\sigma^* \leftarrow^{\$} \text{SIG.ReRand}(vk^l, m_{j^{l*}}^l, A^l[j^{l*}])$, which is a uniformly distributed valid signature for $m_{j^{l*}}^l$ provided that $A^l[j^{l*}] \neq \emptyset$. The latter happens whenever $\text{bad}(l)$ does not occur.

Thus, \mathcal{B} simulates \mathcal{A} perfectly for Λ provided that $\neg \text{bad}$. This implies $S(\mathcal{B}) \cap \neg \text{bad} \iff S(\Lambda^A) \cap \neg \text{bad}$, which yields

$$\Pr[S(\mathcal{B}) \cap \neg \text{bad}] = \Pr[S(\Lambda^A) \cap \neg \text{bad}] \quad (9)$$

Finishing the proof. By plugging (8) and (9) into Inequality (6), we obtain

$$\Pr[S(\Lambda^A)] - \Pr[S(\mathcal{B})] \leq r/n \implies \epsilon_{\mathcal{B}} \geq \epsilon_{\Lambda} - r/n$$

3.1 Interpretation

Assuming that no adversary \mathcal{B} is able to (t_N, ϵ_N) -break the security of NICA with $t_N = t_{\mathcal{B}} = r \cdot n \cdot t_{\Lambda} + r \cdot n \cdot (n-1) \cdot t_{\text{Vfy}} + r \cdot t_{\text{ReRand}}$, we must have $\epsilon_{\mathcal{B}} \leq \epsilon_N$. By Theorem 4, we thus must have

$$\epsilon_{\Lambda} \leq \epsilon_{\mathcal{B}} + r/n \leq \epsilon_N + r/n$$

for all reductions Λ . In particular, the hypothetical adversary \mathcal{A} constructed in the proof of Theorem 1 is an example of an adversary such that

$$\frac{t_{\Lambda} + r \cdot t_{\mathcal{A}}}{\epsilon_{\Lambda}} \geq \frac{r \cdot t_{\mathcal{A}}}{\epsilon_N + r/n} = (\epsilon_N + r/n)^{-1} \cdot r \cdot \frac{t_{\mathcal{A}}}{1} = (\epsilon_N + r/n)^{-1} \cdot r \cdot \frac{t_{\mathcal{A}}}{\epsilon_{\mathcal{A}}}.$$

Thus, any reduction Λ from breaking the security of NICA N to breaking the UF-SMA-security of signature scheme SIG loses (in the sense of Definition 7) at least a factor of $\ell \geq r/(\epsilon_N + r/n)$. In particular, note that $\ell \approx n$ if ϵ_N is very small.

4 A Generalized Meta-Reduction

In this section we state and prove our main result, which generalizes the results from Section 2. Essentially, we observe that for the proof to work we do not need all structural elements a signature scheme possesses. In particular we do not require dedicated parameter generation-, key generation- and sign-algorithms. Instead, we consider an abstract security experiment with the following properties:

1. The values that are publicly available “induce a relation” $R(x, y)$ that is efficiently verifiable for the adversary during the security experiment.

2. The adversary is provided with statements y_1, \dots, y_n at the beginning of the security experiment and has access to an oracle that when queried y_i returns x_i such that $R(x_i, y_i), i \in [n]$.
3. If the adversary is able to output x_j such that $R(x_j, y_j)$ and it did not query its oracle on y_j , this is sufficient to win the security game.

Remark 5. To show the usefulness of such an abstract experiment, we note that for instance the security experiments for public key encryption or key encapsulation mechanisms in the multi-user setting with corruptions [4], or digital signature schemes in the multi-user (MU) setting with corruptions [3,4], naturally satisfy these properties as follows. Essentially, we define a relation $R(sk, pk)$ over pairs of public keys and secret keys such that $R(sk, pk) = 1$ whenever sk “matches” pk . The adversary is provided with public keys at the beginning of the experiment, and is able to obtain secret keys corresponding to public keys of its choice. Finally, if the adversary is able to output an uncorrupted secret key, it is clearly able to compute a signature over a message that was not signed before (i.e., winning the signature security game) or decrypt the challenge ciphertext (i.e., winning the PKE/KEM security game). Thus, all three requirements are satisfied. Details on how to apply the result to, e.g., digital signatures and PKE/KEMs in the multi user setting with corruptions we refer to Section 5.

4.1 Definitions

Re-randomizable relations. Let $R \subseteq X \times Y$ be a relation. For (x, y) with $R(x, y) = 1$ we call x the *witness* and y the *statement*. We use $X(R, y)$ to denote the set

$$X(R, y) := \{x : R(x, y) = 1\}$$

of all witnesses x for statement y with respect to R . We denote by $L(R) := \{y : \exists x \text{ s.t. } R(x, y) = 1\} \subseteq Y$ the language consisting of statements in R .

In the sequel we will consider *computable* relations. We will therefore identify a relation R with a machine \hat{R} that computes R . We say that a relation R is t_{Vfy} -computable, if there is a deterministic Turing machine \hat{R} that runs in time at most $t_{\text{Vfy}}(|x| + |y|)$ such that $\hat{R}(x, y) = R(x, y)$.

Definition 10. Let $\mathcal{R} := \{R_i\}_{i \in I}$ be a family of computable relations. We say that \mathcal{R} is t_{ReRand} -re-randomizable if there is a probabilistic Turing machine $\mathcal{R}.\text{ReRand}$ that inputs (R_i, y, x) , runs in time at most t_{ReRand} , and outputs x' which is uniformly distributed over $X(R, y_i)$ whenever $R_i(x, y) = 1$, with probability 1.

Example 1. Digital signatures in the single user setting, as considered in Section 2, may be described in terms of families of relations. We set $R_{\Pi, vk}$ to the relation over signatures and messages that is defined by a verification key vk . In this case, we have that $X(R, y) = \Sigma(vk, y)$ is the set of all valid signatures over message y with respect to public key vk . Note that the family of relations $(R_{\Pi, vk})_{\Pi, vk}$ is t_{ReRand} -re-randomizable, if the signature scheme is t_{ReRand} -re-randomizable (cf. Definition 2).

<p>Game $\text{UF-SSA}_{\mathcal{R}}^{n,\mathcal{A}}(1^k)$</p> <p>$R = R_i \xleftarrow{\\$} \mathcal{R}$</p> <p>$y_1, \dots, y_n \xleftarrow{\\$} L(R)$ s.t. $y_i \neq y_j$ for all $i \neq j$</p> <p>$x_i \xleftarrow{\\$} X(R, y_i)$ for all $i \in [n]$</p> <p>$(j, st) \leftarrow \mathcal{A}_1(\bar{R}, (y_i)_{i \in [n]}; \rho_{\mathcal{A}})$</p> <p>$x_j \leftarrow \mathcal{A}_2(st, (x_i)_{i \in [n] \setminus \{j\}})$</p> <p>return $R(x_j, y_j)$</p>
--

Fig. 3. The UF-SSA-security Game with attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$.

Witness unforgeability under static statement attacks. We will consider a weak security experiment for computable relations, which is inspired by the UF-SMA-security experiment considered in Section 2, but abstract and general enough to be applicable in other useful settings. Jumping slightly ahead, we will show in Section 5 that this includes applications to signatures, public-key encryption, key encapsulation mechanisms in the multi-user setting, and non-interactive key exchange.

The security experiment is described in Figure 3. It is parametrized by a family \mathcal{R} of computable relations, $\mathcal{R} = \{R_i\}_{i \in I}$, and the number n of statements the adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ is provided with. These statements need to be pairwise distinct. \mathcal{A} may *non-adaptively* ask for witnesses for *all but one* statement, and is considered successful if it manages to output a “valid” witness for the remaining statement.

Definition 11. Let $\mathcal{R} = \{R_i\}_{i \in I}$ be a family of computable relations. We say that an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ (t, n, ϵ) -breaks the witness unforgeability under static statement attacks of \mathcal{R} if it runs in time t and

$$\Pr[\text{UF-SSA}_{\mathcal{R}}^n(\mathcal{A}) \Rightarrow 1] \geq \epsilon$$

where $\text{UF-SSA}_{\mathcal{R}}^n(\mathcal{A})$ is the security game depicted in Figure 3.

Simple reductions from non-interactive complexity assumptions to breaking UF-SSA-security. Informally, a reduction from breaking the UF-SSA-security of a family of relations \mathcal{R} to breaking the security of a non-interactive complexity assumption $N = (T, U, V)$ is a Turing machine Γ , which turns an attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ against \mathcal{R} according to Definition 11 into a TM $\Gamma^{\mathcal{A}}$ that breaks N according to Definition 5. As in Section 2, we will only consider simple reductions, i.e., reductions that have *black-box* access to the attacker and that may run the attacker at most r times *sequentially*.

We define a reduction from breaking the security of \mathcal{R} to breaking N as an $(3r+2)$ -tuple of TMs $\Gamma = (\Gamma_0, (\Gamma_{l,1}, \Gamma_{l,2}, \Gamma_{l,3})_{l \in [r]}, \Gamma_3)$, which turn a TM \mathcal{A} breaking the security of \mathcal{R} into a TM $\Gamma^{\mathcal{A}}$ breaking N , as described in Figure 4. Note that this Turing machine works almost identical to that considered in Section 3, except that we consider a more general class of relations.

```

TM  $r\text{-}\Gamma^{\mathcal{A}}(c; \rho_{\mathcal{A}})$ 
 $st_{\Gamma} \leftarrow \Gamma_0(c, \rho_{\Gamma})$ 
for  $1 \leq l \leq r$  do:
     $(\widehat{R}^l, (y_i^l)_{i \in [n]}, \rho_{\mathcal{A}}, st_{\Gamma}) \leftarrow \Gamma_{l,1}(st_{\Gamma})$ 
     $(j^l, st_{\mathcal{A}}) \leftarrow \mathcal{A}_1(\widehat{R}^l, (y_i^l)_{i \in [n]}; \rho_{\mathcal{A}})$ 
     $((x_i^l)_{i \in [n \setminus j^l]}, st_{\Gamma}) \leftarrow \Gamma_{l,2}(j^l, st_{\Gamma})$ 
     $x_j^l \leftarrow \mathcal{A}_2((x_i^l)_{i \in [n \setminus j^l]}, st_{\mathcal{A}})$ 
     $st_{\Gamma} \leftarrow \Gamma_{l,3}(x_j^l, st_{\Gamma})$ 
 $s \leftarrow \Gamma_3(st_{\Gamma})$ 
return  $s$ 

```

Fig. 4. TM $r\text{-}\Gamma^{\mathcal{A}}$ that solves a non-interactive complexity assumption according to Definition 5, constructed from a r -simple reduction $r\text{-}\Gamma = (\Gamma_0, (\Gamma_{l,1}, \Gamma_{l,2}, \Gamma_{l,3})_{l \in [r]}, \Gamma_3)$ and an attacker $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$.

Definition 12. We say that a TM $r\text{-}\Gamma = (\Gamma_0, (\Gamma_{l,1}, \Gamma_{l,2}, \Gamma_{l,3})_{l \in [r]}, \Gamma_3)$ is an r -simple $(t_{\Gamma}, n, \epsilon_{\Gamma}, \epsilon_{\mathcal{A}})$ -reduction from breaking $N = (\mathsf{T}, \mathsf{V}, \mathsf{U})$ to breaking the UF-SSA-security of a family of relations \mathcal{R} , if for any TM \mathcal{A} that $(t_{\mathcal{A}}, n, \epsilon_{\mathcal{A}})$ -breaks the UF-SSA security of \mathcal{R} , TM $r\text{-}\Gamma^{\mathcal{A}}$ (cf. Figure 4) $(t_{\mathcal{A}} + r \cdot t_{\mathcal{A}}, \epsilon_{\mathcal{A}})$ -breaks N .

We define the loss of an r -simple reduction $r\text{-}\Gamma$ from breaking N to breaking the UF-SSA-security of a family of computable relations \mathcal{R} similar to Definition 9.

4.2 Main Result

In this Section we establish the following result that generalizes Theorem 4.

Theorem 5. Let $N = (\mathsf{T}, \mathsf{V}, \mathsf{U})$ be a non-interactive complexity assumption, $n, r \in \text{poly}(k)$ and let \mathcal{R} be a family of computable relations. Then for any r -simple $(t_{\Gamma}, n, \epsilon_{\Gamma}, 1)$ -reduction Γ from breaking N to breaking the UF-SSA-security of \mathcal{R} there exists a TM \mathcal{B} that $(t_{\mathcal{B}}, \epsilon_{\mathcal{B}})$ -breaks N where

$$t_{\mathcal{B}} \leq r \cdot n \cdot t_{\Gamma} + r \cdot n \cdot (n - 1) \cdot t_{\text{Vfy}} + r \cdot t_{\text{ReRand}}$$

$$\epsilon_{\mathcal{B}} \geq \epsilon_{\Gamma} - \frac{r}{n}.$$

Here, t_{ReRand} is the time to re-randomize a given valid witness and t_{Vfy} is the maximum time needed to compute $R \in \mathcal{R}$.

The proof of Theorem 5 is nearly identical to the proof of Theorem 4, and therefore omitted. Also the interpretation of Theorem 5 is nearly identical to the interpretation described in Section 2.3. Assuming that no adversary \mathcal{B} is able to $(t_{\mathcal{N}}, \epsilon_{\mathcal{N}})$ -break the security of NICA with $t_{\mathcal{N}} = t_{\mathcal{B}} = r \cdot n \cdot t_{\mathcal{A}} + r \cdot n \cdot (n - 1) \cdot t_{\text{Vfy}} + r \cdot t_{\text{ReRand}}$, we must have $\epsilon_{\mathcal{B}} \leq \epsilon_{\mathcal{N}}$. Thus, if \mathcal{R} is efficiently computable and re-randomizable, the loss of any simple reduction from breaking N to breaking the UF-SSA-security of \mathcal{R} is at least linear in n .

5 New Applications

5.1 Signatures in the Multi-User Setting

Definitions. The syntax of digital signature schemes is defined in Section 2. Here, we define additional properties of signature schemes that are required to establish our result. Let $\text{SIG} = (\text{Setup}, \text{Gen}, \text{Sign}, \text{Vfy})$ be a signature scheme. In the sequel we require *perfect correctness*, i.e., that for all $k \in \mathbb{N}$, all $\Pi \leftarrow^{\$} \text{Setup}(1^k)$, all $(vk, sk) \leftarrow^{\$} \text{Gen}(\Pi)$ and all m it holds that:

$$\Pr \left[\text{SIG.Vfy}(vk, m, \sigma) = 1 : \sigma \leftarrow^{\$} \text{SIG.Sign}(sk, m) \right] = 1.$$

Moreover, let $\Pi \leftarrow^{\$} \text{Setup}(1^k)$ and let us recall that Π is contained in vk . We require an additional deterministic TM SKCheck_{Π} that takes as input strings sk and pk and outputs 0 or 1 such that:

$$\begin{aligned} \text{SKCheck}_{\Pi}(pk, sk) &= 1 \\ &\iff \\ \Pr \left[\text{Vfy}(pk, m, \sigma) = 1 : m \leftarrow^{\$} |\mathcal{M}| \wedge \sigma \leftarrow^{\$} \text{Sign}(sk, m) \right] &= 1. \end{aligned}$$

That is, SKCheck takes inputs sk and pk and returns 1 if and only if pk is a valid public key and sk is a corresponding secret key. Since we require perfect correctness for signature schemes, we have $\text{SKCheck}(vk, sk) = 1$ whenever $(vk, sk) \leftarrow^{\$} \text{Gen}(\Pi)$.

Definition 13 (Key re-randomization). We say that a signature encryption scheme SIG is t_{ReRand} -key re-randomizable if there exists a Turing machine SIG.ReRand that runs in time at most t_{ReRand} , takes as input $\Pi(vk, sk)$ and returns sk uniformly distributed over $\{sk : \text{SKCheck}_{\Pi}(vk, sk) = 1\}$ whenever $\text{SKCheck}_{\Pi}(vk, sk) = 1$.

Example 2. If we consider, for example, the Waters signature scheme [40], a public key consists among others of elements $g, g_1, g_2 \in \mathcal{G}$ where $g_1 = g^{\alpha}$. The key generation algorithm outputs a corresponding secret key as $sk = g_2^{\alpha}$. However, there may be other secret keys that might be accepted by SKCheck .

To investigate this issue we shortly recall the signing and verification algorithms of [40]. The signing algorithm, when given as input a secret key and a message returns $\sigma = (\sigma_1, \sigma_2) = (g^r, sk \cdot (H(m))^r)$ where r is uniformly random chosen from \mathbb{Z}_p . Verification returns $e(g_1, g_2) \stackrel{?}{=} e(g, \sigma_2) \cdot e(\sigma_1, H(m))^{-1} = e(g, sk) \cdot e(g, H(m))^r \cdot e(g, H(m))^{-r}$.

We observe that by definition of SKCheck we must have $\text{SKCheck}(vk, sk) = 1 \iff e(g_1, g_2) = e(g, sk)$. Thus there is an efficient SKCheck procedure. Moreover, since there is only one value that satisfies this equation in prime order groups we have an efficient secret key re-randomization algorithm, namely, the identity map. This is all that is to verify before applying our result.

Game MU-EUF-CMA-C- $\text{SIG}^{n,\mu}(\mathcal{A})$

$\Pi \leftarrow^{\$} \text{SIG.Setup}(1^k)$	$\mathcal{O}.\text{Sign}(m, i)$
$(vk_i, sk_i) \leftarrow^{\$} \text{SIG.Gen}(\Pi)$	if $ Q_i \geq \mu$ return \perp
$\rho_{\mathcal{A}} \leftarrow^{\$} \{0, 1\}^k$	$Q_i \leftarrow Q_i \cup \{m\}$
$Q^{\text{Corrupt}} = Q_1 = \dots = Q_n \leftarrow \emptyset$	return $\sigma \leftarrow^{\$} \text{SIG.Sign}(sk_i, m)$
$(i, m^*, \sigma^*) \leftarrow \mathcal{A}^{\mathcal{O}.\text{Sign}(\cdot, \cdot), \mathcal{O}.\text{Corrupt}(\cdot)}((vk_i)_{i \in [n]}; \rho_{\mathcal{A}})$	$\mathcal{O}.\text{Corrupt}(i)$
return $vk_{i^*} \notin Q^{\text{Corrupt}} \wedge m^* \notin Q_{i^*} \wedge \text{SIG.Vfy}(vk_{i^*}, m^*, \sigma^*)$	$Q^{\text{Corrupt}} \leftarrow Q^{\text{Corrupt}} \cup \{vk_i\}$
	return sk_i

Fig. 5. MU-EUF-CMA-C-security Game. The attacker has access to a signing oracle $\mathcal{O}.\text{Sign}$ and a corrupt oracle $\mathcal{O}.\text{Corrupt}$.

Security definition. The MU-EUF-CMA-C-security game is depicted in Figure 5. Here the adversary \mathcal{A} is provided with public keys vk_1, \dots, vk_n of the signature scheme. It may now adaptively issue *sign* and *corrupt*-queries. To issue a sign query it specifies a message m and a public key $vk_i, i \in [n]$ and obtains a valid signature σ over m that is valid with respect to vk_i . In order to issue a corrupt query, \mathcal{A} specifies an index $i \in [n]$ and obtains a secret key sk_i that “matches” vk_i . Finally, \mathcal{A} outputs a triplet (i, m, σ) and is considered successful if it did neither issue a corrupt query for i nor a sign query for (m, vk_i) and at the same time σ is valid over m with respect to vk_i .

Definition 14 (MU-EUF-CMA-C-security). We say that an adversary (t, n, μ, ϵ) -breaks the MU-EUF-CMA-C-security of a signature scheme SIG if it runs in time t and

$$\Pr [\text{MU-EUF-CMA-C-}\text{SIG}^{n,\mu}(\mathcal{A}) \Rightarrow 1] \geq \epsilon .$$

Definition 15. We say that a Turing machine $r\text{-}\Gamma$ is an r -simple $(t_{\Lambda}, n, \mu, \epsilon_{\Lambda}, \epsilon_{\mathcal{A}})$ -reduction from breaking $N = (T, V, U)$ to breaking the MU-EUF-CMA-C-security of SIG, if for any TM \mathcal{A} that $(t_{\mathcal{A}}, n, \mu, \epsilon_{\mathcal{A}})$ -breaks the MU-EUF-CMA-C security of SIG, TM $\Lambda^{\mathcal{A}}(t_{\Lambda} + r \cdot t_{\mathcal{A}}, \epsilon_{\Lambda})$ -breaks N .

The loss of an r -simple reduction Γ from breaking N to breaking the MU-EUF-CMA-C-security of SIG is defined similar to definition 7.

Defining a suitable relation. Let $\text{SIG} = (\text{Setup}, \text{Gen}, \text{Sign}, \text{Vfy})$ be a signature scheme and let I be the range of Setup. We set $\mathcal{R}_{\text{SIG}} = \{R_{\Pi}\}_{\Pi \in I}$ where $R_{\Pi}(x, y) := \text{SKCheck}_{\Pi}(y, x)$. Now, if SIG is t_{ReRand} -key re-randomizable then \mathcal{R}_{SIG} is t_{ReRand} re-randomizable.

UF-SSA security for \mathcal{R}_{SIG} is weaker than MU-EUF-CMA-C-security for SIG. Let now SIG be a perfectly correct signature scheme and let \mathcal{R}_{SIG} be derived from SIG as described in Section 5.1.

Claim. If there is an attacker \mathcal{A} that (t, n, ϵ) -breaks the UF-SSA-security for \mathcal{R}_{SIG} then there is an attacker \mathcal{B} that $(t', n, 0, \epsilon')$ -breaks the MU-EUF-CMA-C-security of SIG with $t' = \mathcal{O}(t)$ and $\epsilon' \geq \epsilon$.

Proof. We construct \mathcal{B} that $(t', n, 0, \epsilon')$ -breaks the MU-EUF-CMA-C-security of SIG, given black box access to \mathcal{A} as follows:

1. \mathcal{B} is called on input a set of public key $(vk)_{i \in [n]}$ and random tape ρ . Recall that Π are contained in vk . First, \mathcal{B} samples and $\rho_{\mathcal{A}}$, the random coins of \mathcal{A} . After that, it runs $(j, st_{\mathcal{A}}) \leftarrow \mathcal{A}_1(\Pi, (vk)_{i \in [n]}, \rho_{\mathcal{A}})$.
2. \mathcal{B} will issue a corrupt-query to oracle $\mathcal{O}.\text{Corrupt}$ for all $i \in [n \setminus j]$. It will obtain sk_i such that $\text{SKCheck}_{\Pi}(vk_i, sk_i)$. Next, \mathcal{B} runs $sk_j \leftarrow^{\$} \mathcal{A}_2((sk_i)_{i \in [n \setminus j]}, st_{\mathcal{A}})$. Note that $\text{SKCheck}_{\Pi}(vk_j, sk_j) = 1$ with probability ϵ .
3. \mathcal{B} samples $m \leftarrow^{\$} \mathcal{M}$ and computes $\sigma \leftarrow^{\$} \text{SIG.Sign}(sk_j, m)$ and outputs (j, m, σ) . Note that $vk_j \notin Q^{\text{Corrupt}}$ and $m \notin Q_j$. Moreover, by the property of SKCheck we have $\text{SIG.Vfy}(vk_j, m, \sigma) = 1$.

Tightness Bound.

Theorem 6 (informal). *Any simple reduction from breaking the security of a NICA N to breaking the MU-EUF-CMA-C-security of a perfectly correct signature scheme SIG (cf. Definition 15) that provides efficient key re-randomization and that supports an efficient SKCheck loses a factor that is linear in the number of public keys the attacker is provided with and that it may corrupt, or N is easy to solve.*

We prove the Theorem via the following technical Theorem, which follows immediately from Theorem 5.

Theorem 7. *Let $N = (\mathsf{T}, \mathsf{V}, \mathsf{U})$ be a non-interactive complexity assumption, $n, r \in \text{poly}(k)$ and let \mathcal{R}_{SIG} be a family of computable relations as described above. Then for any r -simple $(t_{\Gamma}, n, \epsilon_{\Gamma}, 1)$ -reduction Γ from breaking N to breaking the UF-SSA-security of \mathcal{R}_{SIG} there exists a TM \mathcal{B} that $(t_{\mathcal{B}}, \epsilon_{\mathcal{B}})$ -breaks N where*

$$\begin{aligned} t_{\mathcal{B}} &\leq r \cdot n \cdot t_{\Gamma} + r \cdot n \cdot (n - 1) \cdot t_{\text{Vfy}} + r \cdot t_{\text{ReRand}} \\ \epsilon_{\mathcal{B}} &\geq \epsilon_{\Gamma} - \frac{r}{n}. \end{aligned}$$

Here, t_{ReRand} is the time to re-randomize a given valid witness and t_{Vfy} is the maximum time needed to compute $R \in \mathcal{R}_{\text{SIG}}$.

5.2 Public-Key Encryption in the Multi-User Setting

Our main result also applies to public key encryption in the multi-user setting with corruptions (and a similar result for key encapsulation mechanisms is straightforward).

Definitions. A public key encryption scheme PKE is a four tuple of PPT-TMs $\text{PKE} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{Dec})$:

Public Parameters. The parameter generation machine $\Pi \leftarrow^{\$} \text{Setup}(1^k)$ returns public parameters when input the security parameter. We assume Π to contain at least 1^k .

Game MU-IND-CPA-C- $\mathcal{C}_{\text{PKE}}^{n,\mu}(\mathcal{A})$	
$\Pi \leftarrow^{\$} \text{PKE.Setup}(1^k)$	$\mathcal{O}.\text{Encrypt}(m_0, m_1, i^*)$
$(pk_i, sk_i) \leftarrow^{\$} \text{PKE.Gen}(\Pi)$	if $ m_0 \neq m_1 $ return \perp
$\rho_{\mathcal{A}} \leftarrow^{\$} \{0, 1\}^k$	$b \leftarrow^{\$} \{0, 1\}$
$Q^{\text{Corrupt}} \leftarrow \emptyset$	return $c \leftarrow^{\$} \text{Enc}(pk_i, m_b)$
$b' \leftarrow \mathcal{A}^{\mathcal{O}.\text{Encrypt}(\cdot, \cdot, \cdot), \mathcal{O}.\text{Corrupt}(\cdot)} \left((pk_i)_{i \in [n]}; \rho_{\mathcal{A}} \right)$	$\mathcal{O}.\text{Corrupt}(i)$
return $b = b' \wedge pk_{i^*} \notin Q^{\text{Corrupt}}$	$Q^{\text{Corrupt}} \leftarrow Q^{\text{Corrupt}} \cup \{pk_i\}$
	return sk_i

Fig. 6. MU-IND-CPA-C-security Game. The attacker has access to an encryption oracle $\mathcal{O}.\text{Encrypt}$ which may be queried only once and a corrupt oracle $\mathcal{O}.\text{Corrupt}$.

Key Generation. The key generation machine inputs public parameters and returns a key pair as $(pk, sk) \leftarrow^{\$} \text{Gen}(\Pi)$. We assume pk to contain Π .

Encryption. The encryption machine takes as input a public key pk and a message m and returns a ciphertext, i.e., $c \leftarrow^{\$} \text{Enc}(pk, m)$.

Decryption. The decryption machine returns a message m when input a ciphertext c and a secret key sk . We require perfect correctness, i.e., that for all $k \in \mathbb{N}$, all $\Pi \leftarrow^{\$} \text{Setup}(1^k)$ and all $(pk, sk) \leftarrow^{\$} \text{Gen}(\Pi)$ it holds that

$$\Pr [\text{Dec}(sk, \text{Enc}(pk, m)) = m : m \leftarrow^{\$} \mathcal{M}] = 1$$

Security definition. As in the MU-EUF-CMA-C-security game for digital signatures, the adversary is provided with n public keys pk_1, \dots, pk_n at the beginning of the experiment. It may now adaptively issue *corrupt* and one *test*-query. To issue a corrupt query, it specifies an index $i \in [n]$ and obtains sk_i that works for pk_i . A test-query may be issued only once throughout the security game. Here, the adversary specifies an index $i^* \in [n]$ and two messages m_0, m_1 of equal length and obtains an encryption of m_b , where $b \leftarrow^{\$} \{0, 1\}$. The adversary terminates outputting a bit b' and is successful if it did not issue a corrupt query for i^* and $b = b'$. The security experiment is depicted in Figure 6.

Definition 16 (MU-IND-CPA-C-security). We say that an adversary (t, n, μ, ϵ) -breaks the MU-IND-CPA-C-security of a public key encryption scheme PKE if it runs in time t and

$$\Pr [\text{MU-IND-CPA-C-}\mathcal{C}_{\text{PKE}}^{n,\mu}(\mathcal{A}) \Rightarrow 1] \geq \epsilon.$$

Definition 17. We say that a Turing machine $r\text{-}\Gamma$ is an r -simple $(t_{\Lambda}, n, \mu, \epsilon_{\Lambda}, \epsilon_{\mathcal{A}})$ -reduction from breaking $N = (\text{T}, \text{V}, \text{U})$ to breaking the MU-IND-CPA-C-security of PKE, if for any TM \mathcal{A} that $(t_{\mathcal{A}}, n, \mu, \epsilon_{\mathcal{A}})$ -breaks the MU-IND-CPA-C security of PKE, TM $\Lambda^{\mathcal{A}}(t_{\Lambda} + r \cdot t_{\mathcal{A}}, \epsilon_{\Lambda})$ -breaks N .

The loss of an r -simple reduction Γ from breaking N to breaking the MU-IND-CPA-C-security of PKE is defined similar to definition 7.

Remark 6. For simplicity, we consider *chosen-plaintext* security in the multi-user setting. The result generalizes easily to other security notions for public-key encryption, like *chosen-ciphertext* security, for example.

Additional properties. Similar to the case of digital signatures in the multi user setting we require an additional TM SKCheck_Π where $\Pi \leftarrow^{\$} \text{Setup}(1^k)$.

$$\text{SKCheck}_\Pi(pk, sk) = 1 \iff \Pr \left[\text{Dec}(sk, \text{Enc}(pk, m)) = m : m \leftarrow^{\$} \mathcal{M} \right] = 1.$$

That is, SKCheck takes inputs sk and pk and returns 1 if and only if pk is a PKE public key and sk is a secret key corresponding to public key pk .

Definition 18 (Key re-randomization). We say that a public key encryption scheme PKE is t_{ReRand} -key re-randomizable if there exists a Turing machine PKE.ReRand that runs in time at most t_{ReRand} , takes as input (Π, pk, sk) and returns sk uniformly distributed over $\{sk : \text{SKCheck}_\Pi(pk, sk) = 1\}$ if $\text{SKCheck}_\Pi(pk, sk) = 1$.

Example 3. Let us consider the linear encryption scheme from [12]. Here, a public key is a tuple of group elements $g, h, k \in \mathcal{G}$ where \mathcal{G} is a group of prime order p . A corresponding secret key, as generated by the key generation machine, consists of $\alpha, \beta \in \mathbb{Z}_p$ such that $h = g^\alpha$ and $k = g^\beta$. For further investigation we recall the encryption and decryption procedures. The encryption machine outputs $c = (c_2, c_1, c_0) = (h^s, k^t, g^{s+t} \cdot m)$ where s, t are chosen uniformly random from \mathbb{Z}_p . Decryption returns $m \leftarrow c_0 \cdot c_1^{-\beta^{-1}} \cdot c_2^{-\alpha^{-1}}$. Now suppose some tuple (α', β') that will be accepted by SKCheck . By definition of SKCheck the following holds for all m and $s, t \leftarrow^{\$} \mathbb{Z}_p$:

$$\begin{aligned} 0 &= s + t - \log(k) \cdot t \cdot \beta'^{-1} - \log(h) \cdot s \cdot \alpha'^{-1} \\ &= t \cdot (1 - \log(k) \cdot \beta'^{-1}) + s \cdot (1 - \log(h) \cdot \alpha'^{-1}). \end{aligned}$$

We conclude that for SKCheck to accept a proposed secret key, (α', β') must satisfy $\alpha' = \log(h)$ and $\beta' = \log(k)$. Thus, secret keys are unique and in particular efficiently re-randomizable. Moreover, there is an efficient SKCheck -procedure that, given a purported key (α', β') returns $h = ? g^{\alpha'} \wedge k = ? g^{\beta'}$.

Defining a suitable relation. Let $\text{PKE} = (\text{Setup}, \text{Gen}, \text{Enc}, \text{Dec})$ be a public key encryption scheme and let I be the range of Setup . To define a suitable relation, we set $\mathcal{R}_{\text{PKE}} = \{R_\Pi\}_{\Pi \in I}$ where $R_\Pi(x, y) := \text{SKCheck}_\Pi(y, x)$ and I is the range of Setup .

Next, we show that any adversary that breaks the UF-SSA-security of \mathcal{R}_{PKE} then there is an attacker that breaks the MU-IND-CPA-C-security of PKE.

Claim. If there is an attacker \mathcal{A} that (t, n, ϵ) -breaks the UF-SSA-security of \mathcal{R}_{PKE} then there is an attacker \mathcal{B} that $(t', n, 0, \epsilon')$ -breaks the MU-IND-CPA-C-security of PKE with $t' = \mathcal{O}(t)$ and $\epsilon' \geq \epsilon$.

Proof. We construct \mathcal{B} that $(t', n, 0, \epsilon')$ -breaks the MU-IND-CPA-C-security of PKE, given black box access to \mathcal{A} as follows:

1. \mathcal{B} is called on input a public key $(pk)_{i \in [n]}$ and random tape ρ . Recall that Π is contained in pk . First, \mathcal{B} samples and $\rho_{\mathcal{A}}$, the random coins of \mathcal{A} . After that, it runs $(j, st_{\mathcal{A}}) \leftarrow \mathcal{A}_1 \left(\Pi, (pk)_{i \in [n]}, \rho_{\mathcal{A}} \right)$.
2. \mathcal{B} will issue a corrupt-query to oracle $\mathcal{O}.\text{Corrupt}$ for all $i \in [n \setminus j]$. It will obtain sk_i such that $\text{SKCheck}_{\Pi}(vk_i, sk_i)$. Next, \mathcal{B} runs $sk_j \leftarrow^{\$} \mathcal{A}_2 \left((sk_i)_{i \in [n \setminus j]}, st_{\mathcal{A}} \right)$. Note that $\text{SKCheck}_{\Pi}(vk_j, sk_j) = 1$ with probability ϵ .
3. Next, \mathcal{B} samples $m_0 \neq m_1 \leftarrow^{\$} \mathcal{M}$ such that $|m_0| = |m_1|$ and issues query (m_0, m_1, j) to oracle $\mathcal{O}.\text{Encrypt}$ which will respond with $c \leftarrow^{\$} \text{Enc}(pk_j, m_b)$. \mathcal{B} runs $m \leftarrow \text{Dec}(sk_j, c)$ and returns $b' \leftarrow m \stackrel{?}{=} m_1$. Note that $pk_j \notin Q^{\text{Corrupt}}$. Moreover, by the perfect correctness of PKE and the property of SKCheck, we have that $b' = b$ whenever \mathcal{A} is successful.

Tightness Bound.

Theorem 8 (informal). *Any simple reduction from breaking the security of a NICA N to breaking the MU-IND-CPA-C-security of a perfectly correct public key encryption scheme PKE (cf. Definition 16) that provides efficient key re-randomization and that supports an efficient SKCheck loses a factor that is linear in the number of public keys the attacker is provided with and that it may corrupt, or N is easy to solve.*

We prove the above Theorem via the following technical Theorem, which is a direct application of Theorem 5.

Theorem 9. *Let $N = (\mathsf{T}, \mathsf{V}, \mathsf{U})$ be a non-interactive complexity assumption, $n, r \in \text{poly}(k)$ and let \mathcal{R}_{PKE} be a family of computable relations as described above. Then for any r -simple $(t_{\Gamma}, n, \epsilon_{\Gamma}, 1)$ -reduction Γ from breaking N to breaking the UF-SSA-security of \mathcal{R}_{PKE} there exists a TM \mathcal{B} that $(t_{\mathcal{B}}, \epsilon_{\mathcal{B}})$ -breaks N where*

$$t_{\mathcal{B}} \leq r \cdot n \cdot t_{\Gamma} + r \cdot n \cdot (n - 1) \cdot t_{\text{Vfy}} + r \cdot t_{\text{ReRand}}$$

$$\epsilon_{\mathcal{B}} \geq \epsilon_{\Gamma} - \frac{r}{n}.$$

Here, t_{ReRand} is the time to re-randomize a given valid witness and t_{Vfy} is the maximum time needed to compute $R \in \mathcal{R}_{\text{PKE}}$.

5.3 Non-Interactive Key Exchange

In this section we will show how to apply our main result to non-interactive key exchange (NIKE) [26]. This case differs from the cases considered before in that we will have to define a relation $R(x, y)$, which is not efficiently verifiable, given just x and y . Instead, we will need additional information, which will be available in the NIKE security experiment. Formally, we consider again UF-SSA-security for some relation R but model \mathcal{A}_2 as an oracle machine. The responses of the oracle may depend on the output of \mathcal{A}_1 . We explain that this makes it possible to extend the range of covered cryptographic primitives to NIKE.

Definitions. Following [17,26], a NIKE protocol consists of three PPT-TMs with the following syntax:

Public Parameters. On input 1^k , the public parameter generation machine $\Pi \xleftarrow{\$}$ NIKE.Setup(1^k) outputs a set Π of system parameters.

Key Generation. The key generation machine takes as input Π and outputs a random key pair (sk_i, pk_i) for party i , i.e. $(sk_i, pk_i) \xleftarrow{\$}$ NIKE.Gen(Π). We assume that pk contains Π and 1^k .

Shared Key Generation. The deterministic shared key machine SharedKey takes as input (sk_i, pk_j) and outputs a shared key $K_{i,j}$ in time t_{vfy} , where $K_{i,j} = \perp$ if $i = j$.

We require perfect correctness, that is,

$$\Pr [\text{SharedKey}(sk_i, pk_j) = \text{SharedKey}(sk_j, pk_i)] = 1$$

for all $\Pi \xleftarrow{\$}$ NIKE.Setup(1^k) and $(pk_i, sk_i), (pk_j, sk_j) \xleftarrow{\$}$ NIKE.Gen(Π).

We require an additional Turing machine PKCheck that inputs strings Π and pk and evaluates to true if pk is in the range of NIKE.Gen(Π). Moreover, whenever two public keys pk and pk' are accepted by PKCheck, we require that the respective shared key is uniquely determined, given only pk and pk' . In the sequel we will denote this key by $K(pk, pk')$ and call NIKE *unique*. The pairing-based NIKE scheme from [26] satisfies uniqueness.

NIKE security. There exists several different, but polynomial-time equivalent [26] security models for NIKE. Of course the tightness of a reduction depends on the choice of the security model. Indeed, the weakest security model considered in [26] is the *CKS-light* model. However, this model is strongly idealized. The reduction from breaking security in a stronger and more realistic security model (called the CKS model in [26]) to breaking security in this idealized model loses a factor of n^2 , where n is the number of users. We show that this loss is inherent for NIKE schemes with the properties defined above.

CKS-security for NIKE. The *CKS-security* experiment is depicted in Figure 7.

Definition 19. We say that an adversary \mathcal{A} (t, n, ϵ) -breaks the CKS-security of a non-interactive key exchange protocol NIKE if it runs in time at most t and

$$\Pr \left[\text{CKS}_{\text{NIKE}}^{n, \mathcal{A}}(1^k) \Rightarrow 1 \right] \geq \epsilon.$$

Definition 20. We say that a Turing machine $r\text{-}\Gamma$ is an r -simple $(t_\Lambda, n, \epsilon_\Lambda, \epsilon_\mathcal{A})$ -reduction from breaking $N = (\mathsf{T}, \mathsf{V}, \mathsf{U})$ to breaking the CKS-security of NIKE, if for any TM \mathcal{A} that $(t_\mathcal{A}, n, \epsilon_\mathcal{A})$ -breaks the CKS security of NIKE, TM $\Lambda^\mathcal{A}$ $(t_\Lambda + r \cdot t_\mathcal{A}, \epsilon_\Lambda)$ -breaks N .

The loss of an r -simple reduction Γ from breaking the security of N to breaking the CKS-security of NIKE is defined similar to Definition 7.

Game $\text{CKS}_{\text{NIKE}}^{n, \mathcal{A}}(1^k)$	
$\Pi \xleftarrow{\$} \text{Setup}(1^k)$ $(pk_i, sk_i) \xleftarrow{\$} \text{NIKE.Gen}(\Pi)$ $\rho_{\mathcal{A}} \xleftarrow{\$} \{0, 1\}^k$ $Q^{\text{Corrupt}} = Q^{\text{Reveal}} \leftarrow \emptyset$ $b' \xleftarrow{\$} \mathcal{A}^{\mathcal{O}.\text{Corrupt}(\cdot), \mathcal{O}.\text{Reveal}(\cdot, \cdot), \mathcal{O}.\text{Test}(\cdot, \cdot)}(\Pi, (pk_i)_{i \in [n]}; \rho_{\mathcal{A}})$ return $b' = b \wedge pk_{i^*}, pk_{j^*} \notin Q^{\text{Corrupt}} \wedge (i^*, j^*) \notin Q^{\text{Reveal}}$	$\mathcal{O}.\text{Corrupt}(i)$ $Q^{\text{Corrupt}} \leftarrow Q^{\text{Corrupt}} \cup \{pk_i\}$ return sk_i
$\mathcal{O}.\text{Test}(i^*, j^*)$ $K_0 \leftarrow \text{SharedKey}(sk_{i^*}, pk_{j^*}); K_1 \xleftarrow{\$} \text{SharedKey}(\cdot, \cdot)$ $b \xleftarrow{\$} \{0, 1\}$ return K_b	$\mathcal{O}.\text{Reveal}(i, j)$ $Q^{\text{Reveal}} \leftarrow Q^{\text{Reveal}} \cup \{(i, j)\}$ return $\text{SharedKey}(sk_i, pk_j)$

Fig. 7. CKS-Security Game for NIKE. Oracle $\mathcal{O}.\text{Test}$ may be queried only once. K_1 is sampled uniform from the range of SharedKey .

Defining a suitable relation. Let $\text{NIKE} = (\text{Setup}, \text{Gen}, \text{SharedKey})$ be a unique NIKE scheme and let I be the range of Setup . We set $\mathcal{R}_{\text{NIKE}} = \{R_{\Pi}\}_{\Pi \in I}$ where

$$R_{\Pi}(x, (y_1, y_2)) = 1 \Leftrightarrow x = K(y_1, y_2) .$$

Let us fix Π for the moment. Note that the attacker is provided with $\tilde{n} = (n - 1) \cdot n$ R_{Π} statements if it is provided with n NIKE-public keys.

Let now $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ denote an attacker against the UF-SSA-security of $\mathcal{R}_{\text{NIKE}}$. Because R may not be efficiently verifiable, we let \mathcal{A}_2 have oracle access to Oracle $\text{Corrupt}_{i^*, j^*}$ that returns secret key sk_i when queried on input $i \in [n \setminus \{i^*, j^*\}]$. Here $K(pk_{i^*}, pk_{j^*})$ is the shared key that \mathcal{A} needs to compute to break the UF-SSA security of \mathcal{R} and n is the number of public keys that \mathcal{A} is provided with (note that this leads to \tilde{n} NIKE shared keys).

UF-SSA-security for $\mathcal{R}_{\text{NIKE}}$ is weaker than CKS-security for NIKE. Next, we show that any adversary that breaks the UF-SSA-security of $\mathcal{R}_{\text{NIKE}}$ then there is an attacker that breaks the CKS-security of NIKE.

Claim. If there is an attacker \mathcal{A} that (t, \tilde{n}, ϵ) -breaks the UF-SSA-security of $\mathcal{R}_{\text{NIKE}}$ then there is an attacker \mathcal{B} that (t', n, ϵ') -breaks the CKS-security of NIKE with $t' = \mathcal{O}(t)$ and $\epsilon' \geq \epsilon$.

Proof. We construct \mathcal{B} that (t', n, ϵ') -breaks the CKS-security of NIKE, given black box access to \mathcal{A} as follows:

1. \mathcal{B} is called on input a set of public keys $(pk)_{i \in [n]}$ and random tape ρ . Recall that Π is contained in pk . First, \mathcal{B} samples and $\rho_{\mathcal{A}}$, the random coins of \mathcal{A} . Next, it runs $((i^*, j^*), st_{\mathcal{A}}) \leftarrow \mathcal{A}_1(\Pi, (pk)_{i \in [n]}, \rho_{\mathcal{A}})$. Note that n public keys define $n \cdot (n - 1)$ statements for R_{Π} . The one that \mathcal{A} will compute is determined by i^* and j^* .

2. \mathcal{B} will issue a reveal-query to oracle $\mathcal{O}.\text{Reveal}$ for all $(i, j) \in [n]^2 \setminus \{(i^*, j^*)\}, i \neq j$. It will obtain $K_{i,j} = \text{SharedKey}(sk_i, pk_j)$. Next, \mathcal{B} runs

$$K^* \leftarrow \mathcal{A}_2^{\mathcal{O}.\text{Corrupt}_{i^*, j^*}(\cdot)} \left((K_{i,j})_{(i,j) \in [n]^2 \setminus \{(i^*, j^*)\}, i \neq j}, st_{\mathcal{A}} \right).$$

\mathcal{B} provides \mathcal{A} with oracle $\text{Corrupt}_{i^*, j^*}$ by forwarding all queries to oracle $\mathcal{O}.\text{Corrupt}()$ and forwarding the response back to \mathcal{A} . Note that, using sk_i , \mathcal{A} may efficiently check whether $K_{i,j} = \text{SharedKey}(sk_i, pk_j)$ for all $j \in [n]$. By assumption it holds that $K^* = \text{SharedKey}(sk_{i^*}, pk_{j^*})$ with probability at least ϵ .

3. Next, \mathcal{B} issues (i^*, j^*) to oracle $\mathcal{O}.\text{Test}()$ which will respond with K . \mathcal{B} returns 0 if $K = K^*$ and 1 otherwise. Note that by construction of oracle $\text{Corrupt}_{i^*, j^*}$ it holds that $i^*, j^* \notin Q^{\text{Corrupt}}$. Moreover, by the perfect correctness of NIKE and the uniqueness of shared keys \mathcal{B} is successful whenever \mathcal{A} is successful.

Tightness Bounds

Theorem 10 (informal). *Any simple reduction from breaking the security of a NICA N to breaking the CKS-security of a perfectly correct, unique NIKE scheme NIKE (cf. Definition 19) that supports an efficient PKCheck loses a factor that is quadratic in the number of public keys the attacker is provided with and that it may corrupt, or N is easy to solve.*

We prove the Theorem via the following technical Theorem.

Theorem 11. *Let $N = (\mathsf{T}, \mathsf{V}, \mathsf{U})$ be a non-interactive complexity assumption, $\tilde{n}, r \in \text{poly}(k)$ and let $\mathcal{R}_{\text{NIKE}}$ be a family of computable relations as described above. Then for any r -simple $(t_{\Gamma}, \tilde{n}, \epsilon_{\Gamma}, 1)$ -reduction Γ from breaking N to breaking the UF-SSA-security of $\mathcal{R}_{\text{NIKE}}$ there exists a TM \mathcal{B} that $(t_{\mathcal{B}}, \epsilon_{\mathcal{B}})$ -breaks N where*

$$t_{\mathcal{B}} \leq r \cdot \tilde{n} \cdot t_{\Gamma} + r \cdot \tilde{n} \cdot (\tilde{n} - 1) \cdot t_{\mathsf{Vfy}} \text{ and } \epsilon_{\mathcal{B}} \geq \epsilon_{\Gamma} - \frac{r}{\tilde{n}}.$$

Here, t_{Vfy} is the maximum time needed to compute $R \in \mathcal{R}_{\text{NIKE}}$ with access to $\text{Corrupt}_{i^*, j^*}$.

Interpretation. As mentioned before, if the attacker is provided with \tilde{n} statements, it is provided only with $\approx \sqrt{\tilde{n}}$ public keys. Thus, the loss of any r -simple reduction is quadratic in the number of public keys if the underlying problem is assumed to be hard.

Our lower bound for NIKE can easily be generalized to systems where keys are derived from $\ell = O(\log(k))$ parties for security parameter k . Syntactically, the difference is that SharedKey now takes as input $\ell - 1$ public keys and a single secret key. Now, the attacker obtains \tilde{n} statements and $\approx \tilde{n}^{1/\ell}$ public keys. Thus, the loss of any r -simple reduction grows with an exponent of ℓ in the number of public keys.

Extending the result to interactive key exchange. On the one hand, our NIKE bounds do not carry over directly to arbitrary interactive key exchange protocols, because these do not necessarily meet the properties of NIKE schemes that we need to put up. In

particular, we have to require that any pair of NIKE public keys uniquely determines the corresponding shared key (which limits the generality of the result, but appears very reasonable for natural (and possibly all) NIKE constructions, in particular it holds for the NIKE schemes of [26]). This requirement does not hold for interactive AKE protocols, where the shared key may additionally depend on ephemeral random values (nonces or Diffie-Hellman shares, for example) exchanged between parties.

On the other hand, our tightness bounds for signatures and public-key encryption (with unique/re-randomizable secret keys, in the multi-user setting with corruptions) directly imply tightness bounds for AKE protocols that use these primitives, and where the attacker is able to adaptively corrupt the secret keys of these signature/PKE schemes. Note that this includes the vast majority of all known AKE constructions. The tightly-secure key exchange protocol of [4] overcomes this hurdle by using a signature scheme that does not have unique/re-randomizable secret keys, and this is used in a crucial way (cf. the “Naor-Yung trick for signatures” in [4]).

References

1. Michel Abdalla, Pierre-Alain Fouque, Vadim Lyubashevsky, and Mehdi Tibouchi. Tightly-secure signatures from lossy identification schemes. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 572–590. Springer, Heidelberg, April 2012.
2. Masayuki Abe, Jens Groth, and Miyako Ohkubo. Separating short structure-preserving signatures from non-interactive assumptions. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 628–646. Springer, Heidelberg, December 2011.
3. Christoph Bader. Efficient signatures with tight real world security in the random-oracle model. In Dimitris Gritzalis, Aggelos Kiayias, and Ioannis G. Askoxylakis, editors, *CANS 14*, volume 8813 of *LNCS*, pages 370–383. Springer, Heidelberg, October 2014.
4. Christoph Bader, Dennis Hofheinz, Tibor Jäger, Eike Kiltz, and Yong Li. Tightly-secure authenticated key exchange. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part I*, volume 9014 of *LNCS*, pages 629–658. Springer, Heidelberg, March 2015.
5. Foteini Baldimtsi and Anna Lysyanskaya. On the security of one-witness blind signature schemes. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 82–99. Springer, Heidelberg, December 2013.
6. Mihir Bellare, Alexandra Boldyreva, and Silvio Micali. Public-key encryption in a multi-user setting: Security proofs and improvements. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 259–274. Springer, Heidelberg, May 2000.
7. Mihir Bellare and Thomas Ristenpart. Simulation without the artificial abort: Simplified proof and improved concrete security for waters’ IBE scheme. Cryptology ePrint Archive, Report 2009/084, 2009. <http://eprint.iacr.org/2009/084>.
8. Mihir Bellare and Thomas Ristenpart. Simulation without the artificial abort: Simplified proof and improved concrete security for Waters’ IBE scheme. In Antoine Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 407–424. Springer, Heidelberg, April 2009.
9. Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In Douglas R. Stinson, editor, *CRYPTO’93*, volume 773 of *LNCS*, pages 232–249. Springer, Heidelberg, August 1994.

10. Olivier Blazy, Saqib A. Kakvi, Eike Kiltz, and Jiaxin Pan. Tightly-secure signatures from chameleon hash functions. In Jonathan Katz, editor, *PKC 2015*, volume 9020 of *LNCS*, pages 256–279. Springer, Heidelberg, March / April 2015.
11. Olivier Blazy, Eike Kiltz, and Jiaxin Pan. (Hierarchical) identity-based encryption from affine message authentication. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 408–425. Springer, Heidelberg, August 2014.
12. Dan Boneh, Xavier Boyen, and Hovav Shacham. Short group signatures. In Matthew Franklin, editor, *CRYPTO 2004*, volume 3152 of *LNCS*, pages 41–55. Springer, Heidelberg, August 2004.
13. Dan Boneh and Ramarathnam Venkatesan. Breaking RSA may not be equivalent to factoring. In Kaisa Nyberg, editor, *EUROCRYPT’98*, volume 1403 of *LNCS*, pages 59–71. Springer, Heidelberg, May / June 1998.
14. Emmanuel Bresson, Jean Monnerat, and Damien Vergnaud. Separation results on the “one-more” computational problems. In Tal Malkin, editor, *CT-RSA 2008*, volume 4964 of *LNCS*, pages 71–87. Springer, Heidelberg, April 2008.
15. Daniel R. L. Brown. Irreducibility to the one-more evaluation problems: More may be less. Cryptology ePrint Archive, Report 2007/435, 2007. <http://eprint.iacr.org/>.
16. Ran Canetti and Hugo Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 453–474. Springer, Heidelberg, May 2001.
17. David Cash, Eike Kiltz, and Victor Shoup. The twin Diffie-Hellman problem and applications. In Nigel P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 127–145. Springer, Heidelberg, April 2008.
18. Jie Chen and Hoeteck Wee. Fully, (almost) tightly secure IBE and dual system groups. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 435–460. Springer, Heidelberg, August 2013.
19. Jean-Sébastien Coron. Optimal security proofs for PSS and other signature schemes. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 272–287. Springer, Heidelberg, April / May 2002.
20. Ronald Cramer and Ivan Damgård. New generation of secure and practical RSA-based signatures. In Neal Koblitz, editor, *CRYPTO’96*, volume 1109 of *LNCS*, pages 173–185. Springer, Heidelberg, August 1996.
21. Yevgeniy Dodis, Roberto Oliveira, and Krzysztof Pietrzak. On the generic insecurity of the full domain hash. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 449–466. Springer, Heidelberg, August 2005.
22. Yevgeniy Dodis and Leonid Reyzin. On the power of claw-free permutations. In Stelvio Cimato, Clemente Galdi, and Giuseppe Persiano, editors, *SCN 02*, volume 2576 of *LNCS*, pages 55–73. Springer, Heidelberg, September 2003.
23. Marc Fischlin and Nils Fleischhacker. Limitations of the meta-reduction technique: The case of Schnorr signatures. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 444–460. Springer, Heidelberg, May 2013.
24. Marc Fischlin and Dominique Schröder. On the impossibility of three-move blind signature schemes. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 197–215. Springer, Heidelberg, May 2010.
25. Nils Fleischhacker, Tibor Jager, and Dominique Schröder. On tight security proofs for Schnorr signatures. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part I*, volume 8873 of *LNCS*, pages 512–531. Springer, Heidelberg, December 2014.
26. Eduarda S. V. Freire, Dennis Hofheinz, Eike Kiltz, and Kenneth G. Paterson. Non-interactive key exchange. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 254–271. Springer, Heidelberg, February / March 2013.

27. Sanjam Garg, Raghav Bhaskar, and Satyanarayana V. Lokam. Improved bounds on security reductions for discrete log based signatures. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 93–107. Springer, Heidelberg, August 2008.
28. Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.
29. Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, April 1988.
30. Dennis Hofheinz and Tibor Jager. Tightly secure signatures and public-key encryption. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 590–607. Springer, Heidelberg, August 2012.
31. Dennis Hofheinz, Tibor Jager, and Edward Knapp. Waters signatures with optimal security reduction. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 66–83. Springer, Heidelberg, May 2012.
32. Saqib A. Kakvi and Eike Kiltz. Optimal security proofs for full domain hash, revisited. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 537–553. Springer, Heidelberg, April 2012.
33. Jonathan Katz and Nan Wang. Efficiency improvements for signature schemes with tight security reductions. In Sushil Jajodia, Vijayalakshmi Atluri, and Trent Jaeger, editors, *ACM CCS 03*, pages 155–164. ACM Press, October 2003.
34. Allison B. Lewko and Brent Waters. Why proving HIBE systems secure is difficult. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 58–76. Springer, Heidelberg, May 2014.
35. Pascal Paillier and Damien Vergnaud. Discrete-log-based signatures may not be equivalent to discrete log. In Bimal K. Roy, editor, *ASIACRYPT 2005*, volume 3788 of *LNCS*, pages 1–20. Springer, Heidelberg, December 2005.
36. Pascal Paillier and Jorge L. Villar. Trading one-wayness against chosen-ciphertext security in factoring-based encryption. In Xuejia Lai and Kefei Chen, editors, *ASIACRYPT 2006*, volume 4284 of *LNCS*, pages 252–266. Springer, Heidelberg, December 2006.
37. Sven Schäge. Tight proofs for signature schemes without random oracles. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 189–206. Springer, Heidelberg, May 2011.
38. Yannick Seurin. On the exact security of Schnorr-type signatures in the random oracle model. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 554–571. Springer, Heidelberg, April 2012.
39. Nigel Smart. ECRYPT II yearly report on algorithms and key sizes (2011-2012). <http://www.ecrypt.eu.org/documents/D.SPA.20.pdf>, 2012.
40. Brent R. Waters. Efficient identity-based encryption without random oracles. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 114–127. Springer, Heidelberg, May 2005.

A Example for the Effect of Tight Reductions

Suppose we want to deploy a given cryptosystem in practice. For concreteness, let us consider the case of a digital signature scheme with a reduction from breaking the security of the scheme⁵ to solving the RSA problem. Suppose the reduction converts an

⁵ In some well-defined sense like, say, existential forgery under chosen-message attacks [29], for instance.

attacker that runs in time t and breaks the security of the signature scheme with success probability ϵ into an algorithm that solves the RSA problem in time t' and with success probability ϵ' such that

$$t' \approx t \quad \text{and} \quad \epsilon' \geq \frac{\epsilon}{n \cdot q_S},$$

where n denotes the number of public keys observed by the attacker, and q_S denotes the number of signatures issued per public key over the lifetime of the cryptosystem. Thus, the reduction has a multiplicative loss of $n \cdot q_S$.⁶

If we want to achieve a security level of k “bits of security” for the signature scheme, that is, $\epsilon \leq 2^{-k}$ for $t = 2^k$, then the security reduction shows that we can achieve this by selecting the instance of the RSA problem such that for $t' \approx 2^k$ holds that

$$\epsilon' \cdot n \cdot q_S \leq 2^{-k}.$$

The hardness of an instance of the RSA problem is (essentially) determined by the size of the modulus. Note that this choice depends on n and q_S . If we want to achieve a security level of $k = 100$ “bits of security” in a small-scale setting with $n = 32 = 2^5$ users and at most $q_S = 128 = 2^7$ signatures per user, we need to have

$$\epsilon' \leq 2^{-100-5-7} = 2^{-112}.$$

According to the ECRYPT II report on algorithms and key sizes [39], this can be achieved⁷ by instantiating the scheme with an 2432-bit RSA modulus. However, in a large-scale setting, where the numbers of users and signatures may be as large as $n = 2^{20}$ and $q_S = 2^{20}$ (such as in applications of the TLS protocol on the internet, for instance), we already need to have an RSA-modulus of size greater than 4000 bits in order to achieve the same level of security.

In contrast, if we knew a *tight* reduction for the same scheme, which achieves (say)

$$t' \approx t \quad \text{and} \quad \epsilon' \geq 2 \cdot \epsilon,$$

Then this would imply that even in the large-scale setting with $n = 2^{20}$ and $q_S = 2^{20}$ a modulus size of 2432 bits would provide sufficient security.

B Summary of Coron’s Meta-Reduction and its Generalizations

EUFCMA-security is commonly considered the standard security definition for digital signature schemes [29]. The security game is depicted in Figure 8.

Definition 21 (EUFCMA-security.). *We say that an attacker (t, n, ϵ) -breaks the EUFCMA-security of a signature scheme SIG if it runs in time t and*

$$\Pr \left[\text{EUFCMA}_{\text{SIG}}^{n, \mathcal{A}}(1^k) \Rightarrow 1 \right] \geq \epsilon.$$

⁶ This essentially equals the security loss of RSA-based security reductions for the well-known RSA Full-Domain Hash scheme [19,32].

⁷ Under standard heuristics, see [39] for details.

Game EUF-CMA _{SIG} ^{n, A} (1 ^k)	
$\Pi \leftarrow^{\$} \text{Setup}(1^k)$ $(vk, sk) \leftarrow^{\$} \text{Gen}(\Pi)$ $\rho_A \leftarrow^{\$} P_A$ $Q \leftarrow \emptyset$ $(m^*, \sigma^*) \leftarrow \mathcal{A}^{\mathcal{O}.\text{Sign}(\cdot)}(vk; \rho_A)$ return $m^* \notin Q \wedge \text{SIG}.\text{Vfy}(vk, m^*, \sigma^*)$	$\mathcal{O}.\text{Sign}(m)$ if $ Q \geq n$ return \perp $Q \leftarrow Q \cup \{m\}$ return $\sigma \leftarrow^{\$} \text{SIG}.\text{Sign}(sk, m)$

Fig. 8. EUF-CMA-Security Game. When called, the attacker has access to a signing oracle $\mathcal{O}.\text{Sign}$.

Definition 22. We say that a Turing machine $r\text{-}\Gamma$ is an r -simple $(t_\Lambda, n, \epsilon_\Lambda, \epsilon_A)$ -reduction from breaking $N = (\text{T}, \text{V}, \text{U})$ to breaking the EUF-CMA-security of SIG, if for any TM \mathcal{A} that (t_A, n, ϵ_A) -breaks the EUF-CMA security of SIG, TM $\Lambda^{\mathcal{A}}$ $(t_\Lambda + r \cdot t_A, \epsilon_\Lambda)$ -breaks N .

Definition 23. Let $\ell : \mathbb{N} \rightarrow \mathbb{N}$. We say that an r -simple reduction Γ from breaking N to breaking the EUF-CMA-security of SIG loses ℓ , if there exists an adversary \mathcal{A} that (t_A, n, ϵ_A) -breaks the EUF-CMA security of SIG, such that $\Lambda^{\mathcal{A}}$ $(t_\Lambda + t_A, \epsilon_\Lambda)$ -breaks N with

$$\frac{t_\Lambda(k) + t_A(k)}{\epsilon_\Lambda(k)} \geq \ell(k) \cdot \frac{t_A(k)}{\epsilon_A(k)}.$$

The following lemma is due to Hofheinz *et al.* [31] and generalizes a result from Coron [19].

Lemma 1 ([19,31]). Let N be a (t_N, ϵ_N) -secure non-interactive complexity assumption where $\epsilon_N \in \text{negl}(k)$ and let SIG be a unique signature scheme with message space of size 2^l . If Γ is a $(t_\Gamma, n, \epsilon_\Gamma)$ -reduction from breaking N to breaking the EUF-CMA-security of SIG and $t_N \geq 2 \cdot t_\Gamma + t_{\text{ReRand}}$ then

$$\epsilon_\Gamma \leq \epsilon_A \cdot \frac{\exp(-1)}{n} \cdot \left(1 - \frac{n}{2^l}\right)^{-1} + \text{negl}(k).$$

□

Coron [19] and Hofheinz *et al.* [31] conclude that we have $\epsilon_\Lambda = \mathcal{O}\left(\frac{\epsilon_A}{n}\right)$. The conclusion builds on the fact that $2^l \gg n$. This is reasonable for most digital signatures schemes.

C UF-SMA-security is strictly weaker than EUF-CMA-security

We show that any attacker \mathcal{A} that breaks the UF-SMA-security of a signature scheme SIG implies an attacker \mathcal{A}' that breaks the EUF-CMA-security (depicted in Figure 8) of SIG in roughly the same running time and with the same probability of success. Moreover UF-SMA-security and EUF-CMA-security are *not* polynomially equivalent.

Claim. Let SIG be a signature scheme. If there is an attacker \mathcal{A} that (t, n, ϵ) -breaks the UF-SMA-security of a signature scheme SIG then there is an attacker \mathcal{B} that (t', n, ϵ') -breaks the EUF-CMA-security of SIG where $t' = \mathcal{O}(t)$ and $\epsilon' \geq \epsilon$.

Proof. We construct \mathcal{B} that (t', n, ϵ') -breaks the EUF-CMA-security of SIG, given black box access to \mathcal{A} as follows:

1. \mathcal{B} is called on input a public key vk and random tape ρ . First, \mathcal{B} samples n distinct messages m_1, \dots, m_n from the message space and $\rho_{\mathcal{A}}$, the random coins of \mathcal{A} . After that, it runs $(j, st_{\mathcal{A}}) \leftarrow \mathcal{A}_1(vk, (m_i)_{i \in [n]}, \rho_{\mathcal{A}})$.
2. \mathcal{B} will issue a sign-query to oracle Sign for all messages $m_i, i \in [n \setminus j]$. It will obtain $\sigma_i \leftarrow^{\$} \text{SIG.Sign}(sk, m_i)$. Note that σ_i is a valid signature over m_i with respect to vk . Next, \mathcal{B} runs $\sigma_j \leftarrow^{\$} \mathcal{A}_2((\sigma_i)_{i \in [n \setminus j]}, st_{\mathcal{A}})$ which is valid with probability ϵ .
3. \mathcal{B} outputs (m_j, σ_j) . Note that due to the fact that $m_i \neq m_j$ for all $i \neq j$, this is a valid forgery which is valid with probability at least ϵ .

Let SIG be a signature scheme with exponential message space \mathcal{M} . Let $m \leftarrow^{\$} \mathcal{M}$. Then we define a signature scheme $\text{SIG}'(m)$ that works exactly like SIG except the $\text{SIG}'(m)$ -verification machine will accept 0 as a valid signature over m .

Claim. Suppose that no adversary (t, n, ϵ) -breaks the EUF-CMA-security of SIG. Then the following holds: 1. There is no adversary that (t, n, ϵ') -breaks the UF-SMA-security of $\text{SIG}'(m)$ with $\epsilon' \geq \epsilon + \frac{n}{|\mathcal{M}|}$. 2. There exists a trivial attack strategy that $(\mathcal{O}(1), 0, 1)$ -breaks the EUF-CMA-security of $\text{SIG}'(m)$.

Proof. 1. Recall that at the beginning of the UF-SMA security experiment, \mathcal{A} is called on input a verification key and n distinct messages that are sampled uniformly from \mathcal{M} . Now, the probability that $m_i = m$ for $i \in [n]$ is upper bounded by $\frac{n}{|\mathcal{M}|}$. However, if for all $i \in [n]$ we have $m_i \neq m$ then we can apply the previous claim. When called on vk , \mathcal{A} simply outputs $(m, 0)$ which is a valid forgery.