

Optimal Security Proofs for Signatures from Identification Schemes

Eike Kiltz¹ Daniel Masny¹ Jiaxin Pan^{1,2}

¹ Ruhr-University Bochum, Germany
{eike.kiltz, daniel.masny, jiaxin.pan}@rub.de
² Karlsruhe Institute of Technology, Germany

Abstract

We perform a concrete security treatment of digital signature schemes obtained from canonical identification schemes via the Fiat-Shamir transform. If the identification scheme is random self-reducible and satisfies the weakest possible security notion (hardness of key-recoverability), then the signature scheme obtained via Fiat-Shamir is unforgeable against chosen-message attacks in the multi-user setting. Our security reduction is in the random oracle model and loses a factor of roughly Q_h , the number of hash queries. Previous reductions incorporated an additional multiplicative loss of N , the number of users in the system. Our analysis is done in small steps via intermediate security notions, and all our implications have relatively simple proofs. Furthermore, for each step, we show the optimality of the given reduction in terms of model assumptions and tightness.

As an important application of our framework, we obtain a concrete security treatment for Schnorr signatures in the multi-user setting.

Keywords: Signatures, Identification, Schnorr, tightness

1 Introduction

CANONICAL IDENTIFICATION SCHEMES AND THE FIAT-SHAMIR TRANSFORM. A canonical identification scheme ID as formalized by Abdalla et al. [AABN02] is a three-move public-key authentication protocol of a specific form. The prover (holding the secret-key) sends a commitment R to the verifier. The verifier (holding the public-key) returns a random challenge h , uniformly chosen from a set ChSet (of exponential size). The prover sends a response s . Finally, using the verification algorithm, the verifier publicly checks correctness of the transcript (R, h, s) . There is a large number of canonical identification schemes known (e.g. [FS87, GQ90, Bet88, MS90, Sch91, BM91, Gir91, OS91, Oka93, KW03, GJKW07]), the most popular among them being the scheme by Schnorr [Sch91]. The Fiat-Shamir method [FS87] transforms any such canonical identification scheme into a digital signature scheme SIG[ID] using a hash function.

DIGITAL SIGNATURES IN THE MULTI-USER SETTING. When it comes to security of digital signature schemes, in the literature almost exclusively the standard security notion of unforgeability against chosen message attacks (UF-CMA) [GMR88] is considered. This is a *single-user setting*, where an adversary obtains one single public-key and it is said to break the scheme's security if he can produce (after obtaining Q_s many signatures on messages of his choice) a valid forgery, i.e. a message-signature pair that verifies on the given public-key. However, in the real world the attacker is usually confronted with many public-keys and presumably he is happy if he can produce a valid forgery under any of the given public-keys. This scenario is captured in the *multi-user setting* for signatures schemes. Concretely, in multi-user unforgeability against chosen message attacks (MU-UF-CMA) the attacker obtains N independent public-keys and is said to break the scheme's security if he can produce (after obtaining Q_s many signatures on public-keys of his choice) a valid forgery that verifies under any of the public-keys.

There are essentially two reasons why one typically only analyzes signatures in the single-user setting. First, the single-user security notion and consequently their analysis are simpler. Second, there exists a simple generic security reduction [GMS02] between multi-user security and standard single-user security. Namely, for any signature system, attacking the scheme in the multi-user setting with N public-keys cannot increase the attacker's success ratio (i.e., the quotient of its success probability and its running time) by a factor more than N compared to attacking the scheme in the single-user setting. As the number of public-keys N is bounded by a polynomial, asymptotically, the single-user and the multi-user

setting are equivalent. However, the security reduction is not tight: it has a loss of a non-constant factor N . This is clearly not satisfactory as in complex environments one can easily assume the existence of at least $N = 2^{30}$ (≈ 1 billion) public-keys, thereby increasing the upper bound on the attacker’s success ratio by a factor of 2^{30} . For example, if we assume the best algorithm breaking the single-user security having success ratio $\rho = 2^{-80}$, then it can only be argued that the best algorithm breaking the multi-user security has success ratio $\rho' = 2^{-80} \cdot 2^{30} = 2^{-50}$, which is not a safe security margin that defends against today’s attackers.

TIGHTNESS. Generally, we call a security implication between two problems *tight* [BR96], if the success ratio ρ of any adversary attacking the first problem cannot decrease by more than a small constant factor compared to the success ratio ρ' of any adversary attacking the second problem [Gal04, BR09]. Here the success ratio ρ is defined as the quotient between the adversary’s success probability and its running time. We note that this notion of tightness is slightly weaker than requiring that both, success probability and running time, cannot decrease by more than a small constant factor (called strong tightness in [Gal04]). However, the main goal of a concrete security analysis is to derive parameters provably guaranteeing *k-bit security*. As the term *k-bit security* is commonly defined as the non-existence of any adversary that breaks the scheme with a success ratio better than 2^{-k} (see, e.g., [BR09, CKMS16]), our definition of tightness is sufficient for this purpose.

1.1 Our Contributions

This work contains a concrete and modular security analysis of signatures SIG[ID] obtained via the Fiat-Shamir transform. Throughout this paper we assume that our identification schemes ID are Σ -protocols, i.e. they are honest-verifier zero-knowledge (HVZK), have special soundness (SS), and commitments R are sampled at random from a sufficiently large set. For some of our tight implications we furthermore require ID to be random self-reducible (RSR), a property we formally define in Definition 2.5. Most known canonical identification schemes satisfy the above properties.

SECURITY NOTIONS. For identification schemes we consider XXX-YYY security, where $\text{XXX} \in \{\text{KR}, \text{IMP}, \text{PIMP}\}$ denotes the attacker’s goal and $\text{YYY} \in \{\text{KOA}, \text{PA}\}$ the attacker’s capabilities. If the attacker’s goals defined as follows: in key-recovery (KR), it tries to compute a valid secret-key; in impersonation (IMP), it tries to impersonate a prover by convincing an honest verifier; parallel impersonation (PIMP) is a parallel version of IMP, where the adversary tries to convince a verifier in one of Q_{Ch} many parallel sessions. The attacker’s capabilities are defined as follows: in a key-only attack (KOA), the adversary is only given the public-key; in a passive attack (PA), the adversary is provided with valid transcripts between an honest prover and verifier. In total, we obtain $3 \times 2 = 6$ different security notions that that were all previously considered in the literature [PS00, OO98, AABN02], except PIMP-YYY security.

OVERVIEW. We show via a chain of implications that KR-KOA-security (the weakest possible security notion for ID where the adversary has to compute a secret-key from a given public-key without any further oracle access) implies multi-user unforgeability against chosen message attacks (MU-UF-CMA) of SIG[ID]. The diagram in Figure 1 summarizes our results. All implications are optimal in terms of tightness and model requirements in the following sense. If one implication makes use of a special model requirement, we prove its impossibility without this requirement. For example, our implication PIMP-KOA \rightarrow UF-KOA requires the random oracle model [BR93] (with its well-known deficiencies [CGH98]) and we show that the non-programmable random oracle model [FLR⁺10] is not sufficient to prove the same implication. Exactly one of our implications, namely IMP-KOA \rightarrow PIMP-KOA is non-tight, and we prove the impossibility of such a tight implication. We now discuss the implications from Figure 1 in more detail.

FROM IDENTIFICATION TO SINGLE USER SECURITY FOR SIGNATURES. Our first main theorem can be informally stated as follows.

Theorem 1.1. *If the identification scheme is KR-KOA-secure against any adversary having success ratio ρ , then SIG[ID] is UF-CMA-secure in the random oracle model against any adversary having success ratio $\rho' \approx \rho/Q_h$, where Q_h is the maximal number of the adversary’s random oracle queries.*

The proof of this theorem is obtained by combining four independent Lemmas 3.5, 3.6, 3.7, and 3.8 via intermediate security notions IMP-KOA, PIMP-KOA, and UF-KOA¹ security, see Figure 1. We certainly

¹Unforgeability against key-only attack (UF-KOA security) is the same as standard UF-CMA security, but the adversary

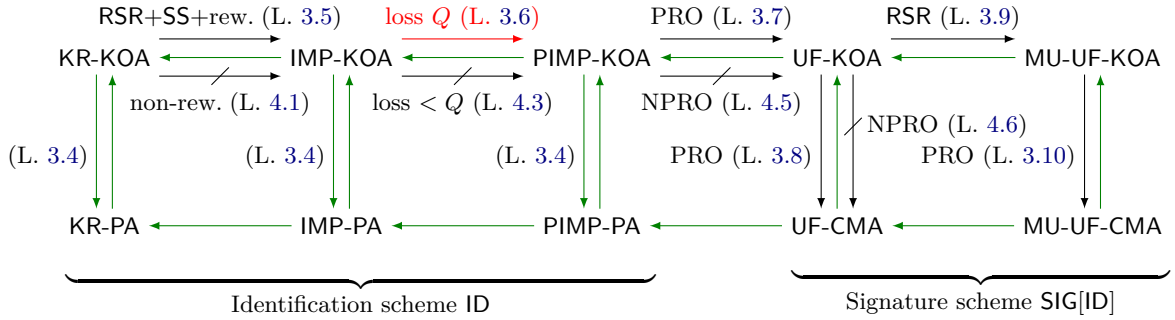


Figure 1: Overview of our notions and results for canonical identification schemes ID and their implied signature schemes SIG[ID]. $X \xrightarrow{Z} Y$ means that X-security implies Y-security under condition Z. Trivial implications are denoted with green arrows. All implications are tight except the one marked with red. The conditions are: rew. (reduction rewinds), loss Q (reduction loses a factor of Q), PRO (reduction is in the programmable random oracle model), SS (reduction uses special soundness), and RSR (reduction uses random self-reducibility for tightness). All implications from top to bottom require HVZK. $X \not\xrightarrow{Z} Y$ means that X-security does not imply Y-security if only condition Z is fulfilled. The conditions are: non-rew. (reduction does not rewind), loss $< Q$ (reduction loses a factor smaller than Q), and NPRO (reduction is in the non-programmable random oracle model).

do not claim any novelty of the above lemmas, nor a new proof technique. For example, the implication $\text{IMP-KOA} \rightarrow \text{UF-CMA}$ is already explicitly contained in [OO98] (and implicitly in the seminal paper by Pointcheval and Stern [PS00]). However, by our specific choice of the intermediate security notions, all four proofs are simple and intuitive. In particular, unlike previous proofs, none of our proofs requires the full power of the Forking Lemma [PS00, BN06]. At the core of Lemma 3.5 ($\text{KR-KOA} \rightarrow \text{IMP-KOA}$) we use a new Multi-Instance Reset Lemma (Lemma 3.3) which is a generalization of Bellare and Palacio’s (Single-Instance) Reset Lemma [BP02] and may be of independent interest. The key to simplicity is the fact that IMP-KOA security only deals with one single impersonation session, which greatly simplifies the probability analysis. Even though the reduction uses rewinding, the RSR property makes the implication $\text{KR-KOA} \rightarrow \text{IMP-KOA}$ tight. We view identifying the intermediate security notions that allow for simple proofs as a conceptual contribution. Our result show that IMP-KOA and PIMP-KOA security can be seen as the *tightness barrier* for identification schemes in the sense that PIMP-KOA is the weakest of our notions for ID that is tightly equivalent to (multi-user) UF-CMA security of SIG[ID] in the random oracle model, whereas IMP-KOA is tightly equivalent to KR-KOA.

One particular advantage of our modular approach is that we are able to prove optimality of all four implications via meta-reductions (Lemmas 4.1, 4.3, 4.5, and 4.6). Lemma 4.3 proving the impossibility of a tight reduction between PIMP-KOA and IMP-KOA security is a generalization of Seurin’s impossibility result to canonical identification schemes [Seu12]; Lemmas 4.5 and 4.6 proving the impossibility of a reduction in the non-programmable random oracle model between PIMP-KOA, UF-KOA, and UF-CMA can be considered as a fine-grained version of a general impossibility result by Fukumitsu and Hasegawa [FH15] who only consider the implication $\text{IMP-PA} \rightarrow \text{UF-CMA}$; All our impossibility results assume the reductions to be key-preserving [PV05] and are conditional in the sense that the existence of a reduction would imply that ID does not satisfy some other natural security property (that is believed to hold).

FROM SINGLE-USER TO MULTI-USER SECURITY FOR SIGNATURES. Our second main theorem can be informally stated as follows.

Theorem 1.2. *If ID is UF-KOA-secure against any adversary having success ratio ρ , then it is MU-UF-CMA-secure in the random oracle model against any adversary having success ratio $\rho' \approx \rho/4$, independent of the number of users N in the multi-user scenario.*

This theorem improves the bound implied by previous generic reductions [GMS02] by a factor of N . Following our modular approach, the theorem is proved in two steps via Lemmas 3.9 and 3.10. Lemma

is not allowed to ask any signing query.

3.9 proves that UF-KOA tightly implies MU-UF-KOA. Tightness stems from the RSR property, meaning that from a given public key pk we can derive properly distributed pk_1, \dots, pk_N such that any signature σ which is valid under pk can be transformed into a signature σ_i which is valid under pk_i and vice-versa.

Lemma 3.10 is our main technical contribution and proves MU-UF-KOA \rightarrow MU-UF-CMA in the programmable random oracle model, again with a tight reduction. One is tempted to believe that it can be proved the same way as in the single user setting (i.e., the same way as UF-KOA \rightarrow UF-CMA). In the single user setting, the reduction simulates signatures on m_j using the HVZK property to obtain a valid transcript (R_j, h_j, s_j) and programs the random oracle as $H(R_j, m_j) := h_j$. However, in the MU-UF-KOA experiment an adversary can ask for a signature under pk_1 on message m which makes the reduction program the random oracle $H(R_1, m) := h_1$. Now, if the adversary submits a forgery (R_1, s_2) under pk_2 on the same message m , the reduction cannot use this forgery to break the MU-UF-KOA experiment because the random oracle $H(R_1, m)$ was externally defined by the reduction. Hence, for the MU-UF-KOA experiment, $m, (R_1, s_2)$ does not constitute a valid forgery. In order to circumvent the above problem we make a simple probabilistic argument. In our reduction, about one half of the multi-user public-keys are coming from the MU-UF-KOA experiment, for the other half the reduction knows the corresponding secret-keys. Which secret-keys are known is hidden from the adversary's view. Now, if the multi-user adversary first obtains a signature on message m under pk_1 and then submits a forgery on the same message m under pk_2 , the reduction hopes for the good case that one of the public-keys comes from the MU-UF-KOA experiment and the other one is known. This happens with probability $1/4$ which is precisely the loss of our new reduction.

1.2 Example Instantiations

SCHNORR SIGNATURES. One of the most important and signature schemes in the discrete logarithm setting is the Schnorr signature scheme [Sch91]. It is obtained via the Fiat-Shamir transform applied to the Schnorr identification protocol. The recent expiry of the patent in 2008 has triggered a number of initiatives to obtain standardized versions of it.

Theorems 1.1 and 1.2 can be used to derive a concrete security bound for strong multi-user MU-UF-CMA-security of Schnorr signatures in the random oracle model from the DLOG problem.² Our reduction loses a factor of roughly Q_h , the number of random oracle queries. This improves previous bounds by a factor of N , the number of users in the system. We derive concrete example parameters for a provably secure instantiation. Figure 1 shows that DLOG is tightly equivalent to IMP-KOA-security and PIMP-KOA-security is tightly equivalent to MU-UF-CMA-security, meaning the tightness barrier for Schnorr lies precisely between IMP-KOA and PIMP-KOA security.

KATZ-WANG SIGNATURES. The Chaum-Pedersen identification scheme [CP93] is a double-generator version of Schnorr. It is at least as secure as Schnorr which means one cannot hope for a tight proof under the DLOG assumption. However, we can use a simple argument from [KW03, GJKW07] for a tight security proof of its PIMP-KOA security under the (stronger) Decision Diffie-Hellman Assumption. The resulting signature scheme is known as the Katz-Wang signature scheme [KW03] and our framework yields a tight proof of its strong MU-UF-CMA-security. Again, this improves previous bounds by a factor of N , the number of users in the system.

GUILLOU-QUISQUATER SIGNATURES. Another canonical identification scheme of interest with the required properties is the one by Guillou-Quisquater [GQ90]. Similar to Katz-Wang, for the Guillou-Quisquater scheme, we can use an argument from [ABP13] for a tight proof of PIMP-KOA security under the Phi-hiding assumption. Alternatively, we can give a proof with loss Q_h under the Factoring assumption. Our framework also shows that this loss is unavoidable.

1.3 Related Work

SINGLE-USER SECURITY. There have been many different works addressing the single-user security of Fiat-Shamir based signature schemes SIG[ID]. In pioneering work, Pointcheval and Stern [PS00] introduced the Forking Lemma as a tool to prove UF-CMA security of SIG[ID] from HVZK, SS and KR-KOA-security.

²We can even prove *strong* MU-UF-CMA security of Schnorr signatures in the sense that a new signature on a previously signed message already counts as a valid forgery.

Ohta and Okamoto [OO98] gave an alternative proof from IMP-KOA security and HVZK. Abdalla et al. [AABN02] prove the equivalence of IMP-PA-security of ID and UF-CMA security of SIG[ID] in the random oracle model. All above results incorporate a security loss of at least Q_h and can be seen as a special case of our framework. Furthermore, [BP02] consider stronger security notions (e.g., IMP-AA and man-in-the-middle security) for the Schnorr and GQ identification schemes. Abdalla et al. [AFLT12] show that lossy identification schemes tightly imply UF-CMA-secure signatures in the random oracle model from decisional assumptions. Our Multi-Instance Reset Lemma (Lemma 3.3) is a generalization to the Reset Lemma of Bellare and Palacio [BP02].

MULTI-USER SECURITY. To mitigate the generic security loss problem in the multi-user setting for the special case of Schnorr’s signature scheme, Galbraith, Malone-Lee, and Smart (GMLS) proved [GMS02] a tight reduction, namely that attacking the Schnorr signatures in the multi-user setting with N public-keys provably cannot decrease (by more than a small constant factor) the attacker’s success ratio compared to attacking the scheme in the single-user setting. Unfortunately, Bernstein [Ber15b] recently pointed out an error in the GMLS proof leaving a tight security reduction for Schnorr signatures as an open problem. Even worse, Bernstein identifies an “apparently insurmountable obstacle to the claimed [GMLS] theorem”. Section 4.3 of [Ber15b] further expands on the insurmountable obstacle. Our Theorem 1.2 shows there is such a tight security reduction for Schnorr signatures if one is willing to rely on the random oracle model. Additionally, in Theorem B.1 we also prove an alternative tight reduction in the standard model which assumes *strong* UF-CMA security. (Schnorr is generally believed to be strongly UF-CMA secure and this is provably equivalent to UF-CMA security in the random oracle model.) Proving the original GMLS theorem (i.e., without random oracles and from standard UF-CMA security) remains an open problem.

IMPOSSIBILITY RESULTS. In terms of impossibility results, Seurin [Seu12], building on earlier work of [PV05, GBL08], proves that there is no tight reduction from the (one-more) discrete logarithm assumption to UF-KOA-security of Schnorr signatures. A more recent result by [FJS14] even excludes a reduction from any non-interactive assumption.³ Fukumitsu and Hasegawa [FH15], generalizing earlier work on Schnorr signatures [FF13, PV05], prove that SIG[ID] cannot be proved secure in the non-programmable random oracle model only assuming IMP-PA security of ID.

SCHNORR SIGNATURES VS. KEY-PREFIXED SCHNORR SIGNATURES. After identifying the error in the GMLS proof, Bernstein [Ber15b] uses the lack of a tight security reduction for Schnorr’s signature scheme as a motivation to promote a “key-prefixed” modification to Schnorr’s signature scheme which includes the verifier’s public-key in the hash function. The EdDSA signature scheme by Bernstein, Duif, Lange, Schwabe, and Yang [BDL⁺11] is essentially a key-prefixing variant of Schnorr’s signature scheme. (In the context of security in a multi-user setting, key-prefixing was considered before, e.g., in [BGLS03].) In [BDL⁺11] key-prefixing is advertized as “an inexpensive way to alleviate concerns that several public keys could be attacked simultaneously.” Indeed, Bernstein [Ber15b] proves that single-user security of the original Schnorr signatures scheme tightly implies multi-user security of the key-prefixed variant of the scheme. That is, the key-prefixed variant has the advantage of a standard model proof of its tight multi-user security, whereas for standard Schnorr signatures one has to assume strong security or rely on the random oracle model.

The TLS standard used to secure HTTPS connections is maintained by the Internet Engineering Task Force (IETF) which delegates research questions to the Internet Research Task Force (IRTF). Cryptographic research questions are usually discussed in the Crypto Forum Research Group (CFRG) mailing list. In the last months the CFRG discussed the issue of key-prefixing.

Key-prefixing comes with the disadvantage that the entire public-key has to be available at the time of signing. Specifically, in a CFRG message from September 2015 Hamburg [Ham15] argues “having to hold the public key along with the private key can be annoying” and “can matter for constrained devices”. Independent of efficiency, we believe that a cryptographic protocol should be as light as possible and prefixing (just as any other component) should only be included if its presence is justified. Naturally, in light of the GMLS proof, Hamburg [Ham15] and Struik [Str15] (among others) recommended against key prefixing for Schnorr. Shortly after, Bernstein [Ber15a] identifies the error in the GMLS theorem and posts a tight security proof for the key-prefixed variant of Schnorr signatures. In what happens next, the participant of the CFRG mailing list switched their minds and mutually agree that key-prefixing

³The main result of the published paper [FJS14] even excludes reduction from any *interactive* assumption (with special algebraic properties), but the proof turned out to be flawed.




Figure 2: A canonical identification scheme and its transcript (R, h, s) .

should be preferred, despite of its previously discussed disadvantages. Specifically, Brown writes about Schnorr signatures that “this justifies a MUST for inclusion of the public key in the message of the classic signature” [Bro15]. As a consequence, key-prefixing is contained in the current draft for EdDSA [JL]. In the light of our new results, we recommend to reconsider this decision.

2 Definitions

2.1 Preliminaries

For an integer p , define $[p] := \{1, \dots, p\}$ and \mathbb{Z}_p as the residual ring $\mathbb{Z}/p\mathbb{Z}$. If A is a set, then $a \stackrel{\boxtimes}{\leftarrow} A$ denotes picking a from A according to the uniform distribution. All our algorithms are probabilistic polynomial time unless stated otherwise. If \mathcal{A} is an algorithm, then $a \stackrel{\boxtimes}{\leftarrow} \mathcal{A}(b)$ denotes the random variable which is defined as the output of \mathcal{A} on input b . To make the randomness explicit, we use the notation $a := \mathcal{A}(b; \rho)$ meaning that the algorithm is executed on input b and randomness ρ . Note that \mathcal{A} ’s execution is now deterministic.

2.2 Canonical Identification Schemes

A canonical identification scheme ID is a three-move protocol of the form depicted in Figure 2. The prover’s first message R is called *commitment*, the verifier selects a uniform *challenge* h from set ChSet , and, upon receiving a *response* s from the prover, makes a deterministic decision.

Definition 2.1 (Canonical Identification Scheme). *A canonical identification scheme ID is defined as a tuple of algorithms $\text{ID} := (\text{IGen}, \mathsf{P}, \text{ChSet}, \mathsf{V})$.*

- The key generation algorithm IGen takes system parameters par as input and returns public and secret key (pk, sk) . We assume that pk defines ChSet , the set of challenges.
- The prover algorithm $\mathsf{P} = (\mathsf{P}_1, \mathsf{P}_2)$ is split into two algorithms. P_1 takes as input the secret key sk and returns a commitment R and a state St ; P_2 takes as input the secret key sk , a commitment R , a challenge h , and a state St and returns a response s .
- The verifier algorithm V takes the public key pk and the conversation transcript as input and outputs a deterministic decision, 1 (acceptance) or 0 (rejection).

We require that for all $(pk, sk) \in \text{IGen}(\text{par})$, all $(R, St) \in \mathsf{P}_1(sk)$, all $h \in \text{ChSet}$ and all $s \in \mathsf{P}_2(sk, R, h, St)$, we have $\mathsf{V}(pk, R, h, s) = 1$.

We make a couple of useful definitions. An identification scheme ID is called *unique* if for all $(pk, sk) \in \text{IGen}(\text{par})$, $(R, St) \in \mathsf{P}_1(sk)$, $h \in \text{ChSet}$, there exists at most one response $s \in \{0, 1\}^*$ such that $\mathsf{V}(pk, R, h, s) = 1$. A *transcript* is a three-tuple (R, h, s) . It is called *valid* (with respect to public-key pk) if $\mathsf{V}(pk, R, h, s) = 1$. Furthermore, it is called *real*, if it is the output of a real interaction between prover and verifier as depicted in Figure 2. A canonical identification schemes ID has α bits of min-entropy, if for all $(pk, sk) \in \text{IGen}(\text{par})$, the commitment generated by the prover algorithm is chosen from a distribution with at least α bits of min-entropy. That is, for all strings R' we have $\Pr[R = R'] \leq 2^{-\alpha}$, if $(R, St) \stackrel{\boxtimes}{\leftarrow} \mathsf{P}_1(sk)$ was honestly generated by the prover.

We now define (parallel) impersonation against key-only attack (KOA), passive attack (PA), and active attack (AA).

Definition 2.2 ((Parallel) Impersonation). Let $YYY \in \{\text{KOA}, \text{PA}, \text{AA}\}$. A canonical identification ID is said to be $(t, \varepsilon, Q_{\text{CH}}, Q_{\text{O}})$ -PIMP-YYY secure (parallel impersonation against YYY attacks) if for all adversaries \mathcal{A} running in time at most t and making at most Q_{CH} queries to the challenge oracle CH and Q_{O} queries to oracle O,

$$\Pr \left[\bigvee (pk, R_{i^*}, h_{i^*}, s_{i^*}) = 1 \wedge i^* \in [Q_{\text{CH}}] \left| \begin{array}{l} (pk, sk) \stackrel{\boxtimes}{\leftarrow} \text{IGen}(\text{par}) \\ St \stackrel{\boxtimes}{\leftarrow} \mathcal{A}^{\text{O}(\cdot)}(pk) \\ (i^*, s_{i^*}) \stackrel{\boxtimes}{\leftarrow} \mathcal{A}^{\text{CH}(\cdot)}(pk, St) \end{array} \right. \right] \leq \varepsilon,$$

where on the i -th query $\text{CH}(R_i)$ ($i \in [Q_{\text{CH}}]$), the challenge oracle returns $h_i \stackrel{\boxtimes}{\leftarrow} \text{ChSet}$ to \mathcal{A} .⁴ Depending on YYY, oracle O is defined as follows.

- If $YYY = \text{KOA}$ (key-only attack), then O always returns \perp .
- If $YYY = \text{PA}$ (passive attack), then $\text{O} := \text{TRAN}$, where on the j -th empty query $\text{TRAN}(\varepsilon)$ ($j \in Q_{\text{O}}$), the transcript oracle returns a real transcript (R'_j, h'_j, s'_j) to \mathcal{A} , where $(R'_j, St'_j) \stackrel{\boxtimes}{\leftarrow} \text{P}_1(sk)$, $h'_j \stackrel{\boxtimes}{\leftarrow} \text{ChSet}$; $s'_j \stackrel{\boxtimes}{\leftarrow} \text{P}_2(sk, R'_j, h'_j, St'_j)$.
- If $YYY = \text{AA}$ (active attack), then $\text{O} := \text{PROVER} = (\text{PROVER}_1, \text{PROVER}_2)$, where on the j -th query $\text{PROVER}_1(\varepsilon)$ ($j \in Q_{\text{O}}$), the prover oracle returns R'_j for $(R'_j, St'_j) \stackrel{\boxtimes}{\leftarrow} \text{P}_1(sk)$ to \mathcal{A} ; on query $\text{PROVER}_2(j, h'_j)$, the oracle returns $s'_j \stackrel{\boxtimes}{\leftarrow} \text{P}_2(sk, R'_j, h'_j, St'_j)$, if R'_j is already defined (and \perp otherwise).

If $YYY = \text{KOA}$, then the parameter Q_{O} is not used and we simply speak of $(t, \varepsilon, Q_{\text{CH}})$ -PIMP-KOA. Moreover, $(t, \varepsilon, Q_{\text{O}})$ -IMP-YYY (impersonation against YYY attack) security is defined as $(t, \varepsilon, 1, Q_{\text{O}})$ -PIMP-YYY security, i.e., the adversary is only allowed $Q_{\text{CH}} = 1$ query to the CH oracle.

Definition 2.3 (Key-recovery). Let $YYY \in \{\text{KOA}, \text{PA}, \text{AA}\}$. A canonical identification ID is said to be (t, ε) -KR-YYY secure (key recovery under YYY attack) if for all adversaries \mathcal{A} running in time at most t ,

$$\Pr \left[(sk^*, pk) \in \text{IGen}(\text{par}) \left| \begin{array}{l} (pk, sk) \stackrel{\boxtimes}{\leftarrow} \text{IGen}(\text{par}) \\ sk^* \stackrel{\boxtimes}{\leftarrow} \mathcal{A}^{\text{O}(\cdot)}(pk) \end{array} \right. \right] \leq \varepsilon,$$

where depending on YYY oracle O is defined as in Definition 2.2. The winning condition $(sk^*, pk) \in \text{IGen}(\text{par})$ means that the tuple (sk^*, pk) is in the support of $\text{IGen}(\text{par})$, i.e., that \mathcal{A} outputs a valid secret-key sk^* with respect to pk .

Definition 2.4 (Special Soundness). A canonical identification ID is said to be SS (special sound) if there exists an extractor algorithm Ext such that, for all $(pk, sk) \in \text{IGen}(\text{par})$, given any two accepting transcripts (R, h, s) and (R, h', s') (where $h \neq h'$), we have $\Pr[(sk^*, pk) \in \text{IGen}(\text{par}) \mid sk^* \stackrel{\boxtimes}{\leftarrow} \text{Ext}(pk, R, h, s, h', s')] = 1$.

Definition 2.5 (Random Self-reducibility). A canonical identification ID is said to be RSR (random self-reducible) if there is an algorithm Rerand and two deterministic algorithms Tran and Derand such that, for all $(pk, sk) \in \text{IGen}(\text{par})$:

- pk' and pk'' have the same distribution, where $(pk', \tau') \stackrel{\boxtimes}{\leftarrow} \text{Rerand}(pk)$ is the rerandomized key-pair and $(pk'', sk'') \stackrel{\boxtimes}{\leftarrow} \text{IGen}(\text{par})$ is a freshly generated key-pair.
- For all $(pk', \tau') \in \text{Rerand}(pk)$, all $(pk', sk') \in \text{IGen}(\text{par})$, and $sk^* = \text{Derand}(pk, pk', sk', \tau')$, we have $(pk, sk^*) \in \text{IGen}(\text{par})$, i.e., Derand returns a valid secret-key sk^* with respect to pk , given any valid sk' for pk' .
- For all $(pk', \tau') \in \text{Rerand}(pk)$, all transcripts (R', h', s') that are valid with respect to pk' , the transcript $(R', h', s := \text{Tran}(pk, pk', \tau', (R', h', s')))$ is valid with respect to pk .

Definition 2.6 (Honest-verifier Zero-knowledge). A canonical identification ID is said to be (perfect) HVZK (honest-verifier zero-knowledge) if there exists an algorithm Sim that, given public key pk , outputs (R, h, s) such that (R, h, s) is a real (i.e., properly distributed) transcript with respect to pk .

2.3 Digital Signatures

We now define syntax and security of a digital signature scheme. Let par be common system parameters shared among all participants.

⁴On two queries $\text{CH}(R_i)$ and $\text{CH}(R'_i)$ with the same input $R_i = R'_i$ the oracle returns two independent random challenges $h_i \stackrel{\boxtimes}{\leftarrow} \text{ChSet}$ and $h'_i \stackrel{\boxtimes}{\leftarrow} \text{ChSet}$.

Definition 2.7 (Digital Signature). A digital signature scheme SIG is defined as a triple of algorithms $\text{SIG} = (\text{Gen}, \text{Sign}, \text{Ver})$.

- The key generation algorithm $\text{Gen}(\text{par})$ returns the public and secret keys (pk, sk) .
- The signing algorithm $\text{Sign}(sk, m)$ returns a signature σ .
- The deterministic verification algorithm $\text{Ver}(pk, m, \sigma)$ returns 1 (accept) or 0 (reject).

We require that for all $(pk, sk) \in \text{Gen}(\text{par})$, all messages $m \in \{0, 1\}^*$, we have $\text{Ver}(pk, m, \text{Sign}(sk, m)) = 1$.

Definition 2.8 (Multi-user Security). A signature scheme SIG is said to be (t, ε, N, Q_s) -MU-SUF-CMA secure (multi-user strongly unforgeable against chosen message attacks) if for all adversaries \mathcal{A} running in time at most t and making at most Q_s queries to the signing oracle,

$$\Pr \left[\begin{array}{l} \text{Ver}(pk_{i^*}, m^*, \sigma^*) = 1 \\ \wedge (i^*, m^*, \sigma^*) \notin \{(i_j, m_j, \sigma_j) \mid j \in [Q_s]\} \end{array} \middle| \begin{array}{l} \text{For } i = 1, \dots, N : (pk_i, sk_i) \stackrel{\boxtimes}{\leftarrow} \text{Gen}(\text{par}) \\ (i^*, m^*, \sigma^*) \stackrel{\boxtimes}{\leftarrow} \mathcal{A}^{\text{SIGN}(\cdot, \cdot)}(pk_1, \dots, pk_N) \end{array} \right] \leq \varepsilon,$$

where on the j -th query $(i_j, m_j) \in [N] \times \{0, 1\}^*$ ($j \in [Q_s]$) the signing oracle SIGN returns $\sigma_j \stackrel{\boxtimes}{\leftarrow} \text{Sign}(sk_{i_j}, m_j)$ to \mathcal{A} , i.e., a signature on message m_j under public-key pk_{i_j} .

We stress that an adversary in particular breaks multi-user security if he asks for a signature on message m under pk_1 and submits a valid forgery on the same message m under pk_2 .

The first condition in the probability statement of Definition 2.8 is called the correctness condition, the second condition is called the freshness condition. Definition 2.8 covers *strong* security in the sense that a new signature on a previously queried message is considered as a fresh forgery. For standard (non-strong) MU-UF-CMA security (multi-user unforgeability against chosen message attack) we modify the freshness condition in the experiment to $(i^*, m^*) \notin \{(i_j, m_j, \sigma_j) \mid j \in [Q_s]\}$, i.e., to break the scheme the adversary has to come up with a signature on a message-key pair which has not been queried to the signing oracle. We also define (t, ε, N) -MU-UF-KOA security (multi-user unforgeability against key only attack) as $(t, \varepsilon, N, 0)$ -MU-UF-CMA security, i.e. $Q_s = 0$, the adversary is not allowed to make any signing query

Definition 2.9 (Single-user Security). In the single-user setting, i.e. $N = 1$ users, (t, ε, Q_s) -SUF-CMA security (strong unforgeability against chosen message attacks) is defined as $(t, \varepsilon, 1, Q_s)$ -MU-SUF-CMA security. Similarly, standard (non-strong) (t, ε, Q_s) -UF-CMA security (unforgeability against chosen message attack) is defined as $(t, \varepsilon, 1, Q_s)$ -MU-UF-CMA security. Further, (t, ε) -UF-KOA security (unforgeability against key-only attack) is defined as $(t, \varepsilon, 1, 0)$ -MU-SUF-CMA security, i.e., $N = 1$ users and $Q_s = 0$ signing queries.

SECURITY IN THE RANDOM ORACLE MODEL. The security of identification and signature schemes containing a hash function can be analyzed in the random oracle model [BR93]. In this model hash values can only be accessed by an adversary through queries to an oracle H . On input x this oracle returns a uniformly random output $H(x)$ which is consistent with previous queries for input x . Using the random oracle model, the maximal number of queries to H becomes a parameter in the concrete security notions. For example, for $(t, \varepsilon, N, Q_s, Q_h)$ -MU-SUF-CMA security we consider all adversaries making at most Q_h queries to the random oracle. We make the convention that each query to the random oracle made during a signing query is counted as the adversary's random oracle query, meaning $Q_h \geq Q_s$.

2.4 Signatures from Identification Schemes

Let $\text{ID} := (\text{IGen}, \text{P}, \text{ChSet}, \text{V})$ be a canonical identification scheme. By the generalized Fiat-Shamir transformation [BP02], the signature scheme $\text{SIG}[\text{ID}] := (\text{Gen}, \text{Sign}, \text{Ver})$ from ID is defined as follows. par contains the system parameters of ID and a hash function $H : \{0, 1\}^* \rightarrow \text{ChSet}$.

<p>Gen(par): $(pk, sk) \stackrel{\boxtimes}{\leftarrow} \text{IGen}(\text{par})$ Return (pk, sk)</p>	<p>Sign(sk, m): $(R, St) \stackrel{\boxtimes}{\leftarrow} \text{P}_1(sk)$ $h = H(R, m)$ $s \stackrel{\boxtimes}{\leftarrow} \text{P}_2(sk, R, h, St)$ Return $\sigma = (R, s)$</p>	<p>Ver(pk, m, σ): Parse $\sigma = (R, s)$ $h = H(R, m)$ Return $\text{V}(pk, R, h, s)$</p>
---	---	---

In some variants of the Fiat-Shamir transform, the hash additionally inputs some public parameters, for example $h = H(pk, R, m)$.

We call ID *commitment-recoverable*, if $V(pk, R, h, s)$ first recomputes $R' = V'(pk, h, s)$ and then outputs 1 iff $R' = R$. For commitment-recoverable ID, we can define an alternative Fiat-Shamir transformation $SIG'[ID] := (\text{Gen}, \text{Sign}', \text{Ver}')$, where Gen is as in $SIG[ID]$. Algorithm $\text{Sign}'(sk, m)$ is defined as $\text{Sign}(sk, m)$ with the modified output $\sigma' = (h, s)$. Algorithm $\text{Ver}'(pk, m, \sigma')$ first parses $\sigma' = (h, s)$, then recomputes the commitment as $R' := V'(pk, h, s)$, and finally returns 1 iff $H(R', m) = h$.

Since $\sigma = (R, s)$ can be publicly transformed into $\sigma' = (h, s)$ and vice-versa, $SIG[ID]$ and $SIG'[ID]$ are equivalent in terms of security. On the one hand, the alternative Fiat-Shamir transform yields shorter signatures if $h \in \text{ChSet}$ has a smaller representation size than response s . On the other hand, signatures of the Fiat-Shamir transform maintain their algebraic structure which in some cases enables useful properties such as batch verification.

3 Security Implications

In this section we will prove the following two main results.

Theorem 3.1 (Main Theorem 1). *Suppose ID is SS, HVZK, RSR and has α bit min-entropy. If ID is (t, ε) -KR-KOA secure then $SIG[ID]$ is $(t', \varepsilon', Q_s, Q_h)$ -UF-CMA-secure and $(t'', \varepsilon'', N, Q_s, Q_h)$ -MU-UF-CMA-secure in the programmable random oracle model, where*

$$\begin{aligned} \frac{\varepsilon'}{t'} &\leq 6(Q_h + 1) \cdot \frac{\varepsilon}{t} + \frac{Q_s}{2^\alpha} + \frac{1}{|\text{ChSet}|}, \\ \frac{\varepsilon''}{t''} &\leq 24(Q_h + 1) \cdot \frac{\varepsilon}{t} + \frac{Q_s}{2^\alpha} + \frac{1}{|\text{ChSet}|}, \end{aligned}$$

and Q_s and Q_h are upper bounds on the number of signing and hash queries in the UF-CMA (or MU-UF-CMA) experiment, respectively.

The proof of Theorem 3.1 is obtained by combining Lemmas 3.5-3.10 below and using $Q_h \leq t' - 1$.

Theorem 3.2 (Main Theorem 2). *Suppose ID is HVZK, RSR and has α bit min-entropy. If $SIG[ID]$ is $(t, \varepsilon, Q_h + Q_s)$ -UF-KOA secure then $SIG[ID]$ is $(t', \varepsilon', N, Q_s, Q_h)$ -MU-UF-CMA secure in the programmable random oracle model, where*

$$\varepsilon' \leq 4\varepsilon + \frac{Q_h Q_s}{2^\alpha}, \quad t' \approx t$$

and Q_s, Q_h are upper bounds on the number of signing and hash queries in the MU-UF-CMA experiment, respectively.

The proof of Theorem 3.2 is obtained by combining Lemmas 3.9 and 3.10 below.

3.1 Multi-Instance Reset Lemma

We first state and prove a new reset lemma that we will later use in the proof of Theorem 3.1. It is presented in the style of Bellare and Neven's General Forking Lemma [BN06] and does not talk about signatures or identification protocols. It is a generalization to many parallel instances of the Reset Lemma [BP02], which is obtained by setting $N = 1$.

Lemma 3.3 (Multi-Instance Reset Lemma). *Fix an integer $N \geq 1$ and a non-empty set H . Let \mathcal{C} be a randomized algorithm that on input (I, h) returns a pair (b, σ) , where b is a bit and σ is called the side output. Let IG be a randomized algorithm that we call the input generator. The accepting probability of \mathcal{C} is defined as*

$$\text{acc} := \Pr[b = 1 \mid I \xleftarrow{\square} \text{IG}; h \xleftarrow{\square} H; (b, \sigma) \xleftarrow{\square} \mathcal{C}(I, h)]$$

The (multi-instance) reset algorithm $\mathcal{R}_{\mathcal{C}}$ associated to \mathcal{C} is the randomized algorithm that takes input I_1, \dots, I_N and proceeds as follows.

Algorithm $\mathcal{R}_C(I_1, \dots, I_N)$:

For $i \in [N]$:

 Pick random coins ρ_i

$h_i \xleftarrow{\boxplus} H$

$(b_i, \sigma_i) \xleftarrow{\boxplus} \mathcal{C}(I_i, h_i; \rho_i)$

If $b_1 = \dots = b_N = 0$ then return $(0, \epsilon, \epsilon)$

// Abort in Phase 1

Fix $i^* \in [N]$ such that $b_{i^*} = 1$

For $j \in [N]$:

$h'_j \xleftarrow{\boxplus} H$

$(b'_j, \sigma'_j) \xleftarrow{\boxplus} \mathcal{C}(I_{i^*}, h'_j; \rho_{i^*})$

If $\exists j^* \in [N] : (h_{i^*} \neq h'_{j^*} \text{ and } b'_{j^*} = 1)$ then return $(i^*, \sigma_{i^*}, \sigma'_{j^*})$

Else return $(0, \epsilon, \epsilon)$

// Abort in Phase 2

Let

$$\text{res} := \Pr[i^* \geq 1 \mid I_1, \dots, I_N \xleftarrow{\boxplus} \text{IG}; (i^*, \sigma, \sigma') \xleftarrow{\boxplus} \mathcal{R}_C(I_1, \dots, I_N)].$$

Then

$$\text{res} \geq \left(1 - \left(1 - \text{acc} + \frac{1}{|H|}\right)^N\right)^2.$$

Proof. For fixed instance I and coins ρ , we define the probabilities

$$\text{acc}(I, \rho) := \Pr_{h \xleftarrow{\boxplus} H} [b = 1 \mid (b, \sigma) \xleftarrow{\boxplus} \mathcal{A}(I, h; \rho)],$$

$$\text{res}(I, \rho) := \Pr_{h, h' \xleftarrow{\boxplus} H} [b = 1 \wedge b' = 1 \wedge h \neq h' \mid (b, \sigma) \xleftarrow{\boxplus} \mathcal{A}(I, h; \rho); (b', \sigma') \xleftarrow{\boxplus} \mathcal{A}(I, h'; \rho)].$$

As for fixed I, ρ , the two events $b = 1$ and $b' = 1$ are independent and we obtain

$$\text{res}(I, \rho) \geq \text{acc}(I, \rho) \cdot \left(\text{acc}(I, \rho) - \frac{1}{|H|}\right), \quad (1)$$

where the additive factor $\frac{1}{|H|}$ accounts for the fact that $\Pr[h' = h] = 1/|H|$. With the expectation taken over $I \xleftarrow{\boxplus} \text{IG}$ and random coins ρ , we bound

$$\begin{aligned} \mathbb{E}_{I, \rho} [\text{res}(I, \rho)] &\geq \mathbb{E}_{I, \rho} \left[\text{acc}(I, \rho) \cdot \left(\text{acc}(I, \rho) - \frac{1}{|H|}\right) \right] \\ &\geq \mathbb{E}_{I, \rho} [\text{acc}(I, \rho)] \cdot \left(\mathbb{E}_{I, \rho} [\text{acc}(I, \rho)] - \frac{1}{|H|}\right) \\ &= \text{acc} \left(\text{acc} - \frac{1}{|H|}\right). \end{aligned}$$

Above, we used (1), Jensen's inequality⁵ applied to the convex function $\varphi(X) := X \cdot (X - 1/|H|)$, and the fact that $\text{acc} = \mathbb{E}_{I, \rho} [\text{acc}(I, \rho)]$.

Next, consider the random variables b_{i^*} and b'_j ($j \in [N]$) as defined during in the execution of $\mathcal{R}_A(I_1, \dots, I_N)$. Using $\text{acc} = \Pr[b_{i^*} = 1]$ and $\Pr[b'_j = 1 \wedge b_{i^*} = 1] = \mathbb{E}_{I_{i^*}, \rho_{i^*}} [\text{res}(I_{i^*}, \rho_{i^*})]$, we obtain

$$\Pr[b'_j = 1 \mid b_{i^*} = 1] = \frac{\Pr[b'_j = 1 \wedge b_{i^*} = 1]}{\Pr[b_{i^*} = 1]} \geq \text{acc} - \frac{1}{|H|}.$$

Finally, we bound

$$\Pr[\text{no abort in phase 2} \mid \text{no abort in phase 1}] = 1 - \prod_{j=1}^N (1 - \Pr[b'_j = 1 \mid b_{i^*} = 1]) \geq 1 - \left(1 - \text{acc} + \frac{1}{|H|}\right)^N,$$

⁵Jensen's inequality states that if φ is a convex function and X is a random variable, then $\mathbb{E}[\varphi(X)] \geq \varphi(\mathbb{E}[X])$.

and

$$\Pr[\text{no abort in phase 1}] = 1 - \prod_{i=1}^N (1 - \Pr[b_i = 1]) = 1 - (1 - \text{acc})^N$$

to establish

$$\text{res} = \Pr[\text{no abort in phase 1} \wedge \text{no abort in phase 2}] \geq (1 - (1 - \text{acc} + \frac{1}{|H|})^N)^2.$$

This completes the proof. ■

3.2 Proof of the Main Theorems

Lemma 3.4 (XXX-KOA \rightarrow XXX-PA). *Let $\text{XXX} \in \{\text{KR}, \text{IMP}, \text{PIMP}\}$. If ID is $(t, \varepsilon, Q_{\text{CH}})$ -XXX-KOA secure and HVZK, then ID is $(\approx t, \varepsilon, Q_{\text{CH}}, Q_{\text{O}})$ -XXX-PA secure.*

Proof. Let \mathcal{A} be an adversary against the $(t, \varepsilon, Q_{\text{CH}}, Q_{\text{O}})$ -XXX-PA-security of ID. We now build an adversary \mathcal{B} against the $(t, \varepsilon, Q_{\text{CH}})$ -XXX-KOA security of ID, with (t, ε) as claimed.

CONSTRUCTION OF \mathcal{B} . Adversary \mathcal{B} inputs pk and runs \mathcal{A} on pk . Essentially, \mathcal{B} only has to simulate the TRAN oracle of the passive attack PA in the first phase. All queries to the CH oracle (for $\text{YYY} \in \{\text{IMP}, \text{PIMP}\}$) in the second phase are echoed by \mathcal{B} to its own CH oracle. Finally, \mathcal{B} outputs whatever \mathcal{A} outputs. A query to the TRAN oracle can be perfectly simulated by computing a simulated proof via $(R, h, s) \stackrel{\boxtimes}{\leftarrow} \text{Sim}(pk)$, and returning it. The running time of \mathcal{B} is that of \mathcal{A} plus roughly Q_{O} executions of Sim to simulate the TRAN oracle, which we ignore for simplicity. ■

Lemma 3.5 below proving that KR-KOA tightly implies IMP-KOA uses the Multi-Instance Reset Lemma and that takes advantage of ID's random self-reducibility (RSR).

Lemma 3.5 (KR-KOA $\xrightarrow{\text{rewinding}}$ IMP-KOA). *If ID is (t, ε) -KR-KOA secure, SS and RSR, then ID is (t', ε') -IMP-KOA secure, where for any $N \geq 1$,*

$$\varepsilon \geq (1 - (1 - \varepsilon' + \frac{1}{|\text{ChSet}|})^N)^2, \quad t \approx 2Nt'. \quad (2)$$

In particular, the two success ratios are related as

$$\frac{\varepsilon'}{t'} - \frac{1}{t'|\text{ChSet}|} \leq 6 \cdot \frac{\varepsilon}{t}. \quad (3)$$

We remark that without RSR, we can still obtain the weaker bounds $\varepsilon \geq \varepsilon'(\varepsilon' - \frac{1}{|\text{ChSet}|})$, $t \approx 2t'$.

Proof. We first show how to derive (3) from (2). If $\varepsilon' \leq 1/|\text{ChSet}|$, then (3) holds trivially. Assuming $\varepsilon' > 1/|\text{ChSet}|$, we set $N := (\varepsilon' - 1/|\text{ChSet}|)^{-1}$ to obtain $t \approx 2t'/(\varepsilon' - 1/|\text{ChSet}|)$ and $\varepsilon \geq (1 - \frac{1}{e})^2 \geq \frac{1}{3}$. Dividing ε by t yields (3).

To prove (2), let \mathcal{A} be an adversary against the (t', ε') -IMP-KOA-security of ID. We now build an adversary \mathcal{B} against the (t, ε) -KR-KOA security of ID, with (t, ε) as claimed in (2).

We use the Multi-Instance Reset Lemma (Lemma 3.3), where $H := \text{ChSet}$ and IG runs $(pk, sk) \stackrel{\boxtimes}{\leftarrow} \text{IGen}$ and returns pk as instance I . We first define adversary $\mathcal{C}(pk, h; \rho)$ that executes $\mathcal{A}(pk; \rho)$, answers \mathcal{A} 's single query R with h , and finally receives s from \mathcal{A} . If transcript (R, h, s) is valid with respect to pk (i.e., $V(pk, R, h, s) = 1$), \mathcal{C} returns $(b = 1, \sigma = (R, h, s))$; otherwise, it returns $(b = 0, \varepsilon)$. By construction, \mathcal{C} returns $b = 1$ iff \mathcal{A} is successful:

$$\text{acc} = \varepsilon'.$$

Adversary \mathcal{B} is defined as follows. For each $i \in [N]$, it uses the RSR property of ID to generate a fresh public key/trapdoor pair $(pk_i, \tau_i) \stackrel{\boxtimes}{\leftarrow} \text{Rerand}(pk)$. Next, it runs $(i^*, \sigma, \sigma') \stackrel{\boxtimes}{\leftarrow} \mathcal{R}_{\mathcal{C}}(pk_1, \dots, pk_N)$, with \mathcal{C} defined above. If $i^* \geq 1$, then both transcripts $\sigma = (R, h, s)$ and $\sigma' = (R, h', s')$ are valid with respect to pk_{i^*} and $h \neq h'$. \mathcal{B} uses the SS property of ID and computes $sk_{i^*} \leftarrow \text{Ext}(pk_{i^*}, R, h, s, h', s')$. Finally,

using the RSR property of ID, it returns $sk = \text{Derand}(pk_{i^*}, sk_{i^*}, \tau_{i^*})$ and terminates. By construction, \mathcal{B} is successful iff \mathcal{R}_C is. By Lemma 3.3 we can bound \mathcal{B} 's success probability as

$$\varepsilon = \text{res} \geq (1 - (1 - \varepsilon' + \frac{1}{|\text{ChSet}|})^N)^2.$$

The running time t of \mathcal{B} is that of \mathcal{R}_C , meaning $2Nt'$ plus the N times the time to run the `Rerand` and `Derand` algorithms of RSR plus the time to run the `Ext` algorithm of SS. We write $t \approx 2Nt'$ to indicate that this is the dominating running time of \mathcal{B} . ■

Lemma 3.6 (IMP-KOA $\xrightarrow{\text{loss } Q}$ PIMP-KOA). *If ID is (t, ε) -IMP-KOA secure, then ID is $(t', \varepsilon', Q_{\text{CH}})$ -PIMP-KOA secure, where*

$$\varepsilon' \leq Q_{\text{CH}} \cdot \varepsilon, \quad t' \approx t.$$

Proof. Let \mathcal{A} be an adversary against the $(t', \varepsilon', Q_{\text{CH}})$ -PIMP-KOA-security of ID. We now build an adversary \mathcal{B} against the (t, ε) -IMP-KOA security of ID, with (t, ε) as claimed.

CONSTRUCTION OF \mathcal{B} . First, \mathcal{B} obtains pk from its IMP-KOA experiment and forwards it to \mathcal{A} . Next, it picks $i^* \xleftarrow{\$} [Q_{\text{CH}}]$. On \mathcal{A} 's i -th query $\text{CH}_{\mathcal{A}}(R_i)$, it proceeds as follows. If $i \neq i^*$, then it returns $h_i \xleftarrow{\$} \text{ChSet}$. If $i = i^*$, then it defines $R := R_{i^*}$, makes a query $h \xleftarrow{\$} \text{CH}_{\mathcal{B}}(R)$ to its own challenge oracle, and returns $h_{i^*} := h$ to \mathcal{A} . Eventually, \mathcal{A} submits (i, s) and terminates. If $i \neq i^*$, then \mathcal{B} aborts. Otherwise, \mathcal{B} outputs s to its own experiment and terminates. Clearly, if $i = i^*$ then \mathcal{B} wins if \mathcal{A} wins. Since i^* is uniform in $[Q_{\text{CH}}]$ the probability that this happens is $1/Q_{\text{CH}}$. ■

Lemma 3.7 (PIMP-KOA $\xrightarrow{\text{PRO}}$ UF-KOA). *If ID is $(t, \varepsilon, Q_{\text{CH}})$ -PIMP-KOA secure, then $\text{SIG}[\text{ID}]$ is (t', ε', Q_h) -UF-KOA secure in the programmable random oracle model, where*

$$\varepsilon' = \varepsilon, \quad t' \approx t, \quad Q_h = Q_{\text{CH}} - 1.$$

Proof. Let \mathcal{A} be an adversary against the (t', ε', Q_h) -UF-KOA-security of $\text{SIG}[\text{ID}]$. We now build an adversary \mathcal{B} against the $(t, \varepsilon, Q_{\text{CH}})$ -PIMP-KOA security of ID, with $(t, \varepsilon, Q_{\text{CH}})$ as claimed.

CONSTRUCTION OF \mathcal{B} . First, \mathcal{B} obtains pk from its PIMP-KOA experiment which it forwards to \mathcal{A} . If \mathcal{A} makes a query (R_i, m_i) to the random oracle, \mathcal{B} returns $H(R_i, m_i)$ if it is already defined, otherwise \mathcal{B} makes a query $h_i \xleftarrow{\$} \text{CH}(R_i)$ and programs the random oracle $H(R_i, m_i) := h_i$. Eventually, \mathcal{A} submits a forgery $(m, \sigma = (R, s))$, and terminates. We assume that $(R, m) \in \{(R_i, m_i)\}$, i.e., $H(R, m)$ was queried by \mathcal{A} . If not, \mathcal{B} makes a dummy query to $H(R, m)$ which is simulated as described above. Hence, in total, there are $Q_{\text{CH}} := Q_h + 1$ queries to H . Let $i \in [Q_h + 1]$ be the unique index such that $(R_i, m_i) = (R, m)$. Adversary \mathcal{B} outputs (i, s_i) and terminates. Note that (R_i, h_i, s_i) is a valid transcript and hence breaks PIMP-KOA security iff \mathcal{A} 's forgery is valid, establishing $\varepsilon = \varepsilon'$. The running time of \mathcal{B} is roughly that of \mathcal{A} , hence $t' \approx t$. ■

The following lemma is a special case of Lemma 3.10 (with a slightly improved bound).

Lemma 3.8 (UF-KOA $\xrightarrow{\text{PRO}}$ UF-CMA). *Suppose ID is HVZK and has α bit min-entropy. If $\text{SIG}[\text{ID}]$ is (t, ε, Q_h) -UF-KOA secure, then $\text{SIG}[\text{ID}]$ is $(t', \varepsilon', Q_s, Q_h)$ -UF-CMA secure in the programmable random oracle model, where*

$$\varepsilon' \leq \varepsilon + \frac{Q_h Q_s}{2^\alpha}, \quad t' \approx t,$$

and Q_s, Q_h are upper bounds on the number of signing and hash queries in the UF-CMA experiment, respectively.

Lemma 3.9 (UF-KOA $\xrightarrow{\text{RSR}}$ MU-UF-KOA). *Suppose ID is RSR. If $\text{SIG}[\text{ID}]$ is (t, ε) -UF-KOA secure, then $\text{SIG}[\text{ID}]$ is (t', ε', N) -MU-UF-KOA secure, where*

$$\varepsilon' = \varepsilon, \quad t' \approx t.$$

Note that without the RSR property one can use the generic bounds from [GMS02] to obtain a non-tight bound with a loss of N .

Proof. Let \mathcal{A} be an algorithm that breaks (t', ε', N) -MU-UF-KOA security of SIG[ID]. We will describe an adversary \mathcal{B} invoking \mathcal{A} that breaks (t, ε) -UF-KOA security of SIG[ID] with (t, ε) as stated in the lemma. Adversary \mathcal{B} is executed in the UF-KOA experiment and obtains a public-key pk .

SIMULATION OF PUBLIC-KEYS INPUT TO \mathcal{A} . For each $i \in [N]$, \mathcal{B} generates $(pk_i, \tau_i) \stackrel{\boxplus}{\leftarrow} \text{Rerand}(pk)$ by using the RSR property of ID. Then \mathcal{B} runs \mathcal{A} on input (pk_1, \dots, pk_N) .

FORGERY. Eventually, \mathcal{A} will submit its forgery $(i^*, m^*, \sigma^* := (R^*, s^*))$ in the MU-UF-KOA experiment. \mathcal{B} computes $h^* = H(m^*, R^*)$ and runs $s \stackrel{\boxplus}{\leftarrow} \text{Tran}(pk, pk_{i^*}, \tau_{i^*}, (R^*, h^*, s^*))$. By the RSR property of ID, the random variables (pk, R^*, h^*, s) and $(pk_{i^*}, R^*, h^*, s^*)$ are identically distributed. If σ^* is a valid signature on message m^* under pk_{i^*} , then (R^*, s) is also a valid signature on m^* under pk . Thus, we have $\varepsilon = \varepsilon'$. The running time t of \mathcal{B} is t' plus the N times the time to run the Rerand and Tran algorithms of RSR. We again write $t \approx t'$. ■

Lemma 3.10 (MU-UF-KOA $\xrightarrow{\text{PRO}}$ MU-UF-CMA). *Suppose ID is HVZK and has α bit min-entropy. If SIG[ID] is (t, ε, N, Q_h) -MU-UF-KOA secure, then SIG[ID] is $(t', \varepsilon', N, Q_s, Q_h)$ -MU-UF-CMA secure in the programmable random oracle model, where*

$$\varepsilon' \leq 4\varepsilon + \frac{Q_h Q_s}{2^\alpha}, \quad t' \approx t,$$

and N is the number of users and Q_s and Q_h are upper bounds on the number of signing and hash queries in the MU-UF-CMA experiment, respectively.

Proof. Let \mathcal{A} be an algorithm that breaks $(t', \varepsilon', N, Q_s, Q_h)$ -MU-UF-CMA security of SIG[ID]. We will describe an adversary \mathcal{B} invoking \mathcal{A} that breaks (t, ε, N, Q_h) -MU-UF-KOA security of SIG[ID] with (t, ε) as stated in the lemma. Adversary \mathcal{B} is executed in the MU-UF-KOA experiment and obtains public-keys (pk_1, \dots, pk_N) , and has access to a random oracle H .

PREPARATION OF PUBLIC-KEYS. For each $i \in [N]$, adversary \mathcal{B} picks a secret bit $b_i \stackrel{\boxplus}{\leftarrow} \{0, 1\}$. If $b_i = 1$ then \mathcal{B} defines $pk'_i := pk_i$, else \mathcal{B} generates the key-pair $(pk'_i, sk'_i) \stackrel{\boxplus}{\leftarrow} \text{Gen}(\text{par})$ itself. We note that all simulated public-keys are correctly distributed.

Adversary \mathcal{B} runs \mathcal{A} on input (pk'_1, \dots, pk'_N) answering hash queries to random oracle H' and signing queries as follows.

SIMULATION OF HASH QUERIES. A hash query $H'(R, m)$ is answered by \mathcal{B} by querying its own hash oracle $H(R, m)$ and returning its answer.

SIMULATION OF SIGNING QUERIES. On \mathcal{A} 's j -th signature query (i_j, m_j) , \mathcal{B} returns a signature σ_j on message m_j under pk_{i_j} according to the following case distinction.

- **Case A:** $b_{i_j} = 0$. In that case sk'_{i_j} is known to \mathcal{B} and the signature is computed as $\sigma_j := (R_j, s_j) \stackrel{\boxplus}{\leftarrow} \text{Sign}(sk'_{i_j}, m_j)$. Note that this involves \mathcal{B} making a hash query and defining $H'(R_j, m_j) := H(R_j, m_j)$.
- **Case B:** $b_{i_j} = 1$. In that case sk'_{i_j} is unknown to \mathcal{B} and the signature is computed using the HVZK property of ID. Concretely, \mathcal{B} runs $(R_j, h_j, s_j) \stackrel{\boxplus}{\leftarrow} \text{Sim}(pk'_{i_j})$. If hash value $H'(R_j, m_j)$ was already defined (via one of \mathcal{A} 's hash/signing queries) and $H'(R_j, m_j) \neq h_j$, \mathcal{B} aborts. Otherwise, it defines the random oracle

$$H'(R_j, m_j) := h_j \tag{4}$$

and returns $\sigma_j := (R_j, s_j)$, which is a correctly distributed valid signatures on m_j under pk_{i_j} . Note that by (4), \mathcal{B} makes H and H' inconsistent, i.e., we have $H(R_j, m_j) \neq H'(R_j, m_j)$ with high probability. Also note that for each signing query, \mathcal{B} aborts with probability at most $Q_h/2^\alpha$ because R_j has min-entropy α . Since the number of signing queries is bounded by Q_s , \mathcal{B} aborts overall with probability at most $Q_h Q_s / 2^\alpha$.

FORGERY. Eventually, \mathcal{A} will submit its forgery $(i^*, m^*, \sigma^* := (R^*, s^*))$. We assume that it is a valid forgery in the MU-UF-CMA experiment, i.e., for $h^* = H'(R^*, m^*)$ we have $\mathbb{V}(pk'_{i^*}, R^*, h^*, s^*) = 1$. Furthermore, it satisfies the freshness condition, i.e.,

$$(i^*, m^*) \notin \{(i_j, m_j) : j \in [Q_s]\}. \tag{5}$$

After receiving \mathcal{A} 's forgery, \mathcal{B} computes a forgery for the MU-UF-KOA experiment according to the following case distinction.

- **Case 1:** There exists a $j \in [Q_s]$ such that $(m^*, R^*) = (m_j, R_j)$. (If there is more than one j , fix any of them.) In that case we have $h^* = h_j$ and furthermore $i^* \neq i_j$ by the freshness condition (5).
 - **Case 1a:** $(b_{i^*} = 1)$ and $(b_{i_j} = 0)$. Then the hash value $h^* = H'(R^*, m^*)$ was not programmed by \mathcal{B} in (4). That means $h^* = H'(R^*, m^*) = H(R^*, m^*)$ and \mathcal{B} returns $(i^*, m^*, (R^*, s^*))$ as a valid forgery to its MU-UF-KOA experiment.
 - **Case 1b:** $(b_{i^*} = b_{i_j})$ or $(b_{i^*} = 0 \wedge b_{i_j} = 1)$. Then \mathcal{B} aborts.

Note that in case 1 we always have $i^* \neq i_j$ and therefore \mathcal{B} does not abort with probability $1/4$ in which case it outputs a valid forgery.

- **Case 2:** For all $j \in [Q_s]$ we have: $(m^*, R^*) \neq (m_j, R_j)$.
 - **Case 2a:** $b_{i^*} = 1$. Then the hash value $h^* = H'(R^*, m^*)$ was not programmed by \mathcal{B} in (4). That means $h^* = H'(R^*, m^*) = H(R^*, m^*)$ and \mathcal{B} returns $(i^*, m^*, (R^*, s^*))$ as a valid forgery to its MU-UF-KOA experiment.
 - **Case 2b:** $b_{i^*} = 0$. Then \mathcal{B} aborts.

Note that in case 2, \mathcal{B} does not abort with probability $1/2$ in which case it outputs a valid forgery.

Overall, \mathcal{B} returns a valid forgery of MU-UF-KOA experiment with probability

$$\varepsilon \geq \min \left\{ \frac{1}{4}, \frac{1}{2} \right\} \cdot \left(\varepsilon' - \frac{Q_h Q_s}{2^\alpha} \right) = \frac{1}{4} \left(\varepsilon' - \frac{Q_h Q_s}{2^\alpha} \right).$$

The running time of \mathcal{B} is that of \mathcal{A} plus the Q_s executions of Sim. We write $t' \approx t$. This completes the proof. ■

If s in ID is uniquely defined by (pk, R, h) (e.g., as in the Schnorr identification scheme), then one can show the above proof even implies MU-SUF-CMA security of SIG[ID]. The simulation of hash and signing queries is the same as in the above proof. Let (i^*, m^*, R^*, s^*) be \mathcal{A} 's forgery. The freshness condition of the MU-SUF-CMA experiment says that $(i^*, m^*, R^*, s^*) \notin \{(i_j, m_j, R_j, s_j) : j \in [Q_s]\}$. Together with the uniqueness of ID, this implies $(i^*, m^*, R^*) \notin \{(i_j, m_j, R_j) : j \in [Q_s]\}$. If $(i^*, m^*) \notin \{(i_j, m_j) : j \in [Q_s]\}$, then \mathcal{B} can break MU-UF-KOA security by the same case distinction as in the proof above. Otherwise, we have $R^* \notin \{R_j : j \in [Q_s]\}$, in which case we can argue as in case 2.

4 Impossibility Results

In this section, we show that Theorems 3.1 and 3.2 from the previous section are optimal in the sense that the security reduction requires: rewinding (Lemma 4.1), security loss of at least $O(Q)$ (Lemma 4.3) and programmability of random oracles (Lemmas 4.5 and 4.6).

Let X and Y be some hard cryptographic problems, defined through a (possibly) interactive experiment. A black-box reduction \mathcal{R} from X to Y is an algorithm that, given black-box access to an adversary \mathcal{A} breaking problem Y , breaks problem X . If X and Y are security notions for identification or signatures schemes, then a reduction \mathcal{R} is called key-preserving, if \mathcal{R} only makes calls to \mathcal{A} with the same pk that it obtained by its own problem X . All our reductions considered in this section are key-preserving.

Lemma 4.1 (KR-KOA $\xrightarrow{\text{non-rewind.}}$ IMP-KOA). *If there is a key-preserving reduction \mathcal{R} that $(t_{\mathcal{R}}, \varepsilon_{\mathcal{R}})$ -breaks KR-KOA security of ID with one-time black-box access to an adversary \mathcal{A} that $(t_{\mathcal{A}}, \varepsilon_{\mathcal{A}})$ -breaks IMP-KOA security of ID, then there exists an algorithm \mathcal{M} that $(t_{\mathcal{M}}, \varepsilon_{\mathcal{M}}, Q_0)$ -breaks IMP-AA security of ID, where*

$$\varepsilon_{\mathcal{M}} \geq \varepsilon_{\mathcal{R}} - \frac{1}{|\text{ChSet}|}, \quad t_{\mathcal{M}} \approx t_{\mathcal{R}}, \quad Q_0 = 1.$$

Proof. Assuming the existence of a key-preserving reduction \mathcal{R} as above, we construct a meta-reduction \mathcal{M} to break IMP-AA security of ID. \mathcal{M} gets the public key pk of the IMP-AA challenge as input and has oracle access to $O = \text{PROVER}$, black-box accesses to \mathcal{R} and simulates the adversary \mathcal{A} .

CONSTRUCTION OF $\mathcal{M}(pk)$. \mathcal{M} runs $\mathcal{R}(pk)$ and, upon receiving pk from \mathcal{R} , \mathcal{M} simulates the interaction with $\mathcal{A}(pk)$ by echoing \mathcal{M} 's queries to the IMP-AA experiment. That is, \mathcal{M} queries $R \stackrel{\boxplus}{\leftarrow} \text{PROVER}_1()$ provided by the IMP-AA experiment and returns R to \mathcal{R} . Upon receiving h from \mathcal{R} , \mathcal{M} queries $s \stackrel{\boxplus}{\leftarrow} \text{PROVER}_2(1, h)$ provided by the IMP-AA experiment and returns s to \mathcal{R} with probability $\varepsilon_{\mathcal{A}}$.




Figure 3: Meta-reduction \mathcal{M} uses \mathcal{R} to break the IMP-MIM security in n MIM rounds. n is the total amount of executions of \mathcal{A} performed by \mathcal{R} . For every MIM round $j \in [n]$, \mathcal{A} picks an $i_j^* \in [Q_{\text{CH}}]$ and forwards a valid response s_{j,i_j^*} . \mathcal{M} fails when \mathcal{R} fails, $(R, h, s) = (R_{i_j^*}, h_{i_j^*}, s_{i_j^*})$ for some $j \in [n]$. It also fails with some probability when \mathcal{A} gets rewinded on a different h_{j,i_j^*} after having requested s_{j,i_j^*} .

After receiving sk from \mathcal{R} , \mathcal{M} uses sk to impersonate a prover. First, \mathcal{M} computes $(R^*, St) \stackrel{\boxtimes}{\leftarrow} \text{P}_1(sk)$ and queries oracle $h^* \stackrel{\boxtimes}{\leftarrow} \text{CH}(R^*)$ provided by the IMP-AA experiment. Next, \mathcal{M} outputs $s^* \stackrel{\boxtimes}{\leftarrow} \text{P}_2(sk, R^*, h^*, St)$ and terminates.

By the correctness of ID, (R^*, h^*, s^*) is a valid transcript and $(R^*, h^*, s^*) \neq (R, h, s)$ with probability at least $1 - 1/|\text{ChSet}|$. We note that \mathcal{M} perfectly simulates an $(t_{\mathcal{A}}, \varepsilon_{\mathcal{A}})$ adversary against IMP-KOA security. Thus, we have $\varepsilon_{\mathcal{M}} \geq \varepsilon_{\mathcal{R}} - 1/|\text{ChSet}|$. ■

For our next impossibility result, we will require the following definition for identification schemes.

Definition 4.2 (Concurrent (Weak) Impersonation against Man-in-the-Middle Attacks). A canonical identification ID is said to be $(t, \varepsilon, Q_{\text{CH}}, Q_{\text{O}})$ -IMP-MIM secure (impersonation against man-in-the-middle attacks) if for all adversaries \mathcal{A} running in time at most t and adaptively making at most Q_{O} queries to the prover oracle PROVER and Q_{CH} queries to the challenge oracle CH ,

$$\Pr \left[\begin{array}{l} \vee(pk, R_{i^*}, h_{i^*}, s_{i^*}) = 1 \wedge (i^* \in [Q_{\text{CH}}]) \\ \wedge (R_{i^*}, h_{i^*}, s_{i^*}) \notin \{(R'_j, h'_j, s'_j) \mid j \in [Q_{\text{O}}]\} \end{array} \mid \begin{array}{l} (pk, sk) \stackrel{\boxtimes}{\leftarrow} \text{IGen}(\text{par}) \\ (i^*, s_{i^*}) \stackrel{\boxtimes}{\leftarrow} \mathcal{A}^{\text{PROVER}(\cdot), \text{CH}(\cdot)}(pk) \end{array} \right] \leq \varepsilon,$$

where oracles PROVER and CH are defined as in Definition 2.2. We define weak impersonation against man-in-the-middle attack (wIMP-MIM) by restricting $R_{i^*} \in \{R'_1, \dots, R'_{Q_{\text{O}}}\}$.

We remark that wIMP-MIM is a non-standard definition without any practical relevance, but it will only be used for showing negative results. The following generalizes a result by Seurin [Seu12] to canonical identification schemes.

Lemma 4.3 (IMP-KOA $\xrightarrow{\text{loss} < Q}$ PIMP-KOA). Suppose that ID has α bit min-entropy and there is a key-preserving reduction \mathcal{R} that $(t_{\mathcal{R}}, \varepsilon_{\mathcal{R}})$ -breaks IMP-KOA security of ID with n -time black-box access to an adversary \mathcal{A} that $(t_{\mathcal{A}}, \varepsilon_{\mathcal{A}}, Q_{\text{CH}})$ -breaks PIMP-KOA security of ID. Then there exists an algorithm \mathcal{M} that $(t_{\mathcal{M}}, \varepsilon_{\mathcal{M}}, 1, Q_{\text{O}} = nQ_{\text{CH}})$ -breaks IMP-MIM security of ID, where

$$\varepsilon_{\mathcal{M}} \geq \varepsilon_{\mathcal{R}} - \frac{n \ln((1 - \varepsilon_{\mathcal{A}})^{-1})}{Q_{\text{CH}}} - \frac{n}{|\text{ChSet}|} - \frac{n}{2^\alpha}, t_{\mathcal{M}} \approx t_{\mathcal{R}}.$$

Proof. Assuming the existence of a key-preserving reduction \mathcal{R} , we construct a meta-reduction \mathcal{M} to break IMP-MIM security of ID (see Figure 3). \mathcal{M} inputs public key pk of the IMP-MIM challenger, has

black-box accesses to \mathcal{R} and simulates the adversary \mathcal{A} while interacting within $Q_O = nQ_{\text{CH}}$ many MIM rounds.

W.l.o.g. we can assume that our adversary \mathcal{A} never accesses its random coins. Instead, it generates pseudorandomness directly using a PRF, where the key k of PRF is part of the description of \mathcal{A} . Adversary \mathcal{A} 's randomness is derived from its current view using the PRF. As we assume that \mathcal{R} has only black box access to \mathcal{A} , it can not access key k and hence it can not distinguish \mathcal{A} 's pseudorandom randomness from uniform randomness by observing the outputs of \mathcal{A} .

CONSTRUCTION OF $\mathcal{M}(pk)$. \mathcal{M} runs $\mathcal{R}(pk)$ who is interacting with a simulated $\mathcal{A}(pk)$. (Recall that \mathcal{R} is key-preserving, so it always executed \mathcal{A} on pk .) \mathcal{R} can execute \mathcal{A} at most n times and hence rewinds it at most $n - 1$ times to any desired state. In the simulation of \mathcal{A} described below we make the explicit convention that \mathcal{M} always keeps the simulation of \mathcal{A} consistent with previous executions. That is, as long as there exists a $j' < j$ such that for all $i' < i$, $h_{j',i'} = h_{j',i'}$, then \mathcal{M} will also use $R_{j,i} = R_{j',i}$ and $c_{j,i} = c_{j',i}$.

Upon receiving pk from \mathcal{R} , \mathcal{M} simulates the j -th execution or rewind ($j \in [n]$) of $\mathcal{A}(pk)$ as follows.

- First, \mathcal{M} sets a flag $b_j := 0$. The flag b_j will be switched to 1 once \mathcal{M} has obtained one valid transcript from the PROVER oracle.
- To simulate the i -th query to the challenge oracle ($i \in [Q_{\text{CH}}]$), \mathcal{M} starts an interaction with a new prover: \mathcal{M} calls $R_{j,i} \stackrel{\boxtimes}{\leftarrow} \text{PROVER}_1()$ and forwards it to \mathcal{R} , which will reply with an arbitrary $h_{j,i} \in \text{ChSet}$. If $b_j = 1$, \mathcal{M} sets $c_{j,i} := 0$. Otherwise, \mathcal{M} flips a biased coin $c_{j,i}$ with $\Pr[c_{j,i} = 1] = \mu$, where μ will be defined later.

Case 1: $c_{j,i} = 1$. If there is an index $j' < j$ with $R_{j',i} = R_{j,i}$, $h_{j',i} \neq h_{j,i}$, and $c_{j',i} = 1$, then \mathcal{M} aborts its attempt to break IMP-MIM security of ID. Otherwise, it defines $i_j^* := i$ and requests $s_{j,i_j^*} \stackrel{\boxtimes}{\leftarrow} \text{PROVER}_2(I_{j,i_j^*}, h_{j,i_j^*})$, where $I_{j,i_j^*} = (j - 1) \cdot Q_{\text{CH}} + i_j^*$ refers to the I_{j,i_j^*} 's query to the PROVER₁ oracle from which \mathcal{M} obtained challenge h_{j,i_j^*} . Note that \mathcal{M} now obtained one transcript $(R_{j,i_j^*}, h_{j,i_j^*}, s_{j,i_j^*})$ from the PROVER oracle and therefore sets $b_j := 1$.

Case 2: $c_{j,i} = 0$. \mathcal{M} does nothing.

- After Q_{CH} simulated challenge queries, \mathcal{M} sets $(i_j^*, s_{j,i_j^*}) := (\perp, \perp)$ if i_j^* is undefined. Finally, \mathcal{M} returns (i_j^*, s_{j,i_j^*}) to \mathcal{R} .

This completes the simulation of the j -th execution of \mathcal{A} .

At some point \mathcal{R} makes a query $\text{CH}(R)$, which \mathcal{M} forwards to its own CH, receiving h . Finally, \mathcal{R} outputs s and terminates. \mathcal{M} also outputs s and terminates. This completes the description of \mathcal{M} .

ANALYSIS OF \mathcal{M} . We define Bad_1 as the event that the transcript (R, h, s) output by \mathcal{R} does not satisfy the freshness condition $(R, h, s) \notin \{(R_{j,i}, h_{j,i}, s_{j,i}) \mid (j, i) \in [n] \times [Q_{\text{CH}}]\}$ of the IMP-MIM security experiment. Note that $s_{j,i} \neq \perp$ only if $i = i_j^*$ and therefore we only need to consider the case $i = i_j^*$.

$$\begin{aligned} \Pr[\text{Bad}_1] &= \Pr[\exists j \in [n] : (R, h, s) = (R_{j,i_j^*}, h_{j,i_j^*}, s_{j,i_j^*})] \\ &\leq \Pr[\exists j \in [n] : (R, h) = (R_{j,i_j^*}, h_{j,i_j^*})]. \end{aligned}$$

We let $(j_0, i_0) \in [n] \times [Q_{\text{CH}}]$ be the unique pair of indices such that \mathcal{R} makes its single query $\text{CH}(R)$ after receiving R_{j_0, i_0} but before receiving R_{j_0, i_0+1} .

$$\Pr[\exists j \in [n] : (R, h) = (R_{j,i_j^*}, h_{j,i_j^*})] \leq \Pr[\exists (j, i_j^*) \neq (j_0, i_0) : (R, h) = (R_{j,i_j^*}, h_{j,i_j^*})] \quad (6)$$

$$+ \Pr[(j, i_j^*) = (j_0, i_0)]. \quad (7)$$

We bound the probabilities (6) and (7) individually. To bound (7), only a single query is considered. Therefore

$$\Pr[(j, i_j^*) = (j_0, i_0)] = \Pr[c_{j_0, i_0} = 1] \leq \mu.$$

To bound (6), we define a natural order on the set $[n] \times [Q_{\text{CH}}]$ via $(j, i) < (j_0, i_0)$ iff $R_{j,i}$ was received before R_{j_0, i_0} , i.e., $(j - 1)Q_{\text{CH}} + i < (j_0 - 1)Q_{\text{CH}} + i_0$. Note that \mathcal{R} chooses h_{j,i_j^*} for $(j, i_j^*) < (j_0, i_0)$ before seeing $h \stackrel{\boxtimes}{\leftarrow} \text{ChSet}$. Furthermore, R is fixed for $(j, i_j^*) > (j_0, i_0)$ while $R_{j,i_j^*} \stackrel{\boxtimes}{\leftarrow} \text{PROVER}_1()$ has at least α

bits of min-entropy. Therefore by splitting the probabilities and using a union bound

$$\begin{aligned}
& \Pr[\exists(j, i_j^*) \neq (j_0, i_0) : (R, h) = (R_{j, i_j^*}, h_{j, i_j^*})] \\
& \leq \Pr[\exists(j, i_j^*) < (j_0, i_0) : (R, h) = (R_{j, i_j^*}, h_{j, i_j^*})] + \Pr[\exists(j, i_j^*) > (j_0, i_0) : (R, h) = (R_{j, i_j^*}, h_{j, i_j^*})] \\
& \leq \Pr[\exists(j, i_j^*) < (j_0, i_0) : h = h_{j, i_j^*}] + \Pr[\exists(j, i_j^*) > (j_0, i_0) : R = R_{j, i_j^*}] \\
& \leq \frac{j_0}{|\text{ChSet}|} + \frac{n - j_0 + 1}{2^\alpha} \leq \frac{n}{|\text{ChSet}|} + \frac{n}{2^\alpha}.
\end{aligned}$$

Overall, this yields

$$\Pr[\text{Bad}_1] \leq \frac{n}{|\text{ChSet}|} + \mu + \frac{n}{2^\alpha}.$$

Next, we define Bad_2 as the event that \mathcal{M} aborts. By a union bound we get

$$\begin{aligned}
\Pr[\text{Bad}_2] &= \Pr[\exists j \in [n], j' < j, i \in [Q_{\text{CH}}] : R_{j', i} = R_{j, i} \wedge h_{j', i} \neq h_{j, i} \wedge c_{j, i} = c_{j', i} = 1] \\
&= \Pr[\exists j \in [n], j' < j : R_{j', i_j^*} = R_{j, i_j^*} \wedge h_{j', i_j^*} \neq h_{j, i_j^*} \wedge c_{j', i_j^*} = 1] \\
&\leq \Pr[\exists j \in [n], j' < j : c_{j', i_j^*} = 1] \leq (n - 1)\mu.
\end{aligned}$$

CHOICE OF μ . We now choose μ such that on one side \mathcal{A} forges with probability $\varepsilon_{\mathcal{A}}$ and on the other side the probability that Bad_1 or Bad_2 happen is bounded. We set

$$\mu = 1 - (1 - \varepsilon_{\mathcal{A}})^{1/Q_{\text{CH}}}$$

for a desired success probability $0 < \varepsilon_{\mathcal{A}} < 1$ of \mathcal{A} and Q_{CH} queries. Note that for an execution $j \in [n]$ that unless for all $i \in [Q_{\text{CH}}]$ we have $c_{j, i} = 0$, \mathcal{A} will always send a valid transcript and break the PIMP-KOA security. Let $\bar{\mu} := (1 - \mu)$. For any execution $j \in [n]$, \mathcal{A} has success probability

$$\Pr[\exists i \in [Q_{\text{CH}}] : c_{j, i} = 1] = \sum_{k=1}^{Q_{\text{CH}}} \mu(1 - \mu)^{k-1} = \sum_{k=1}^{Q_{\text{CH}}} (\bar{\mu}^{k-1} - \bar{\mu}^k) = 1 - (1 - \mu)^{Q_{\text{CH}}} = \varepsilon_{\mathcal{A}}.$$

Finally, we can bound the success probability of \mathcal{M}

$$\Pr[\text{Bad}_1 \wedge \text{Bad}_2] \leq n \cdot \mu + \frac{n}{|\text{ChSet}|} + \frac{n}{2^\alpha} \leq \frac{n \ln((1 - \varepsilon_{\mathcal{A}})^{-1})}{Q_{\text{CH}}} + \frac{n}{|\text{ChSet}|} + \frac{n}{2^\alpha},$$

where the bound $\mu \leq \ln((1 - \varepsilon_{\mathcal{A}})^{-1})/Q_{\text{CH}}$ was proved in [Seu12, Lemma 1]. Therefore we have

$$\varepsilon_{\mathcal{M}} \geq \varepsilon_{\mathcal{R}} - \frac{n \ln((1 - \varepsilon_{\mathcal{A}})^{-1})}{Q_{\text{CH}}} - \frac{n}{|\text{ChSet}|} - \frac{n}{2^\alpha}, \quad t_{\mathcal{M}} \approx t_{\mathcal{R}} \approx nt_{\mathcal{A}}$$

which concludes the proof of the lemma. \blacksquare

For a precise analysis of the function $\ln((1 - \varepsilon_{\mathcal{A}})^{-1})$, we refer to [Seu12]. For our purpose, it is sufficient that for a concrete choice of $\varepsilon_{\mathcal{A}}$, there is a constant c such that $c \cdot \varepsilon_{\mathcal{A}} = \ln((1 - \varepsilon_{\mathcal{A}})^{-1})$. Hence Lemma 4.3 gives roughly $\varepsilon_{\mathcal{M}} \geq \varepsilon_{\mathcal{R}} - c \cdot n/Q_{\text{CH}} \cdot \varepsilon_{\mathcal{A}}$ for a suitable choice of $\varepsilon_{\mathcal{A}}$. Therefore $\varepsilon_{\mathcal{R}}$ can be at most $c \cdot n/Q_{\text{CH}} \cdot \varepsilon_{\mathcal{A}}$. Otherwise \mathcal{M} would break IMP-MIM security of ID with $\varepsilon_{\mathcal{M}} > 0$.

It is easy to see that the meta-reduction of the proof of Lemma 4.3 just forwards all $R_{j, i}$ received during the Man-in-the-Middle attack and R sent by \mathcal{R} . So if \mathcal{R} is furthermore randomness-preserving, i.e., it chooses $R \in \{R_{1,1}, \dots, R_{n, Q_{\text{CH}}}\}$, then \mathcal{M} attacks wIMP-MIM-security of ID. This observation (formalized in the following corollary) is important since the Schnorr identification scheme is wIMP-MIM but not IMP-MIM-secure (cf. Section 5.1).

Corollary 4.4. *If ID has α bit min-entropy and there exists a key- and randomness-preserving reduction \mathcal{R} that $(t_{\mathcal{R}}, \varepsilon_{\mathcal{R}})$ -breaks IMP-KOA security of ID with n -time black-box access to an adversary \mathcal{A} that $(t_{\mathcal{A}}, \varepsilon_{\mathcal{A}}, Q_{\text{CH}})$ -breaks PIMP-KOA security of ID, then there exists an algorithm \mathcal{M} that $(t_{\mathcal{M}}, \varepsilon_{\mathcal{M}}, 1, Q_{\text{O}} = nQ_{\text{CH}})$ -breaks wIMP-MIM security of ID, where*

$$\varepsilon_{\mathcal{M}} \geq \varepsilon_{\mathcal{R}} - \frac{n \ln((1 - \varepsilon_{\mathcal{A}})^{-1})}{Q_{\text{CH}}} - \frac{n}{|\text{ChSet}|} - \frac{n}{2^\alpha}, \quad t_{\mathcal{M}} \approx t_{\mathcal{R}}.$$

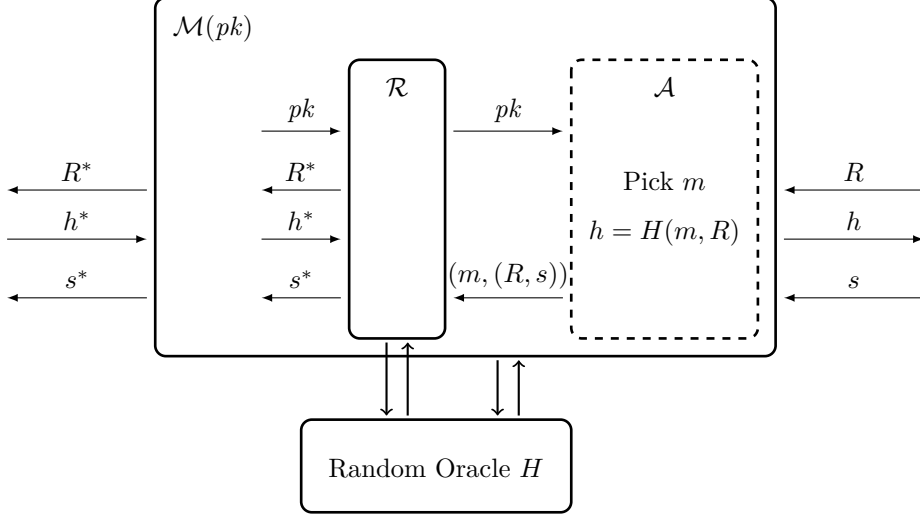


Figure 4: Meta-reduction \mathcal{M} runs \mathcal{R} to break IMP-AA security in the non-programmable random oracle model, where both \mathcal{M} and \mathcal{R} have oracle access to the same external random oracle H . \mathcal{M} simulates an adversary \mathcal{A} that breaks UF-KOA security of SIG[ID] (which is in the dashed box) and answers the oracle queries of \mathcal{R} .

Lemma 4.5 (IMP-KOA $\xrightarrow{\text{NPRO}}$ UF-KOA). *If there exists a key-preserving reduction \mathcal{R} in the non-programmable random oracle (NPRO) model that $(t_{\mathcal{R}}, \varepsilon_{\mathcal{R}})$ -breaks IMP-KOA security of ID with n -time black-box access to an adversary \mathcal{A} that $(t_{\mathcal{A}}, \varepsilon_{\mathcal{A}}, Q_h)$ -breaks UF-KOA security of SIG[ID], then there exists an algorithm \mathcal{M} that $(t_{\mathcal{M}}, \varepsilon_{\mathcal{M}}, 1)$ -breaks IMP-AA-security of ID, where*

$$\varepsilon_{\mathcal{M}} \geq \varepsilon_{\mathcal{R}} - \frac{1}{|\text{ChSet}|}, t_{\mathcal{M}} \approx t_{\mathcal{R}}.$$

Proof. Assuming the existence of a key-preserving reduction \mathcal{R} as above, we construct a meta-reduction \mathcal{M} to break IMP-AA security of ID. Figure 4 gives a pictorial overview of it \mathcal{M} . We can w.l.o.g. assume that $n = 1$ since \mathcal{M} only interacts with \mathcal{A} by sending the fixed pk . \mathcal{M} obtains the public key pk from the IMP-AA experiment and has oracle access to PROVER, black-box accesses to \mathcal{R} and simulates the adversary \mathcal{A} . Additionally, both \mathcal{M} and \mathcal{R} get access to the same external random oracle H , in the NPRO model.

CONSTRUCTION OF $\mathcal{M}(pk)$. \mathcal{M} runs $\mathcal{R}(pk)$ and, upon receiving pk from \mathcal{R} , \mathcal{M} simulates $\mathcal{A}(pk)$ as follows. First, \mathcal{M} queries $R \stackrel{\boxtimes}{\leftarrow} \text{PROVER}_1()$ to the IMP-AA experiment. Next, \mathcal{M} picks an arbitrary message m , queries $h = H(m, R)$ to the random oracle, and $s \stackrel{\boxtimes}{\leftarrow} \text{PROVER}_2(1, h)$ to the IMP-AA experiment. With probability $\varepsilon_{\mathcal{A}}$, \mathcal{M} returns $(m, (R, s))$ as a forgery to \mathcal{R} .

Upon receiving a challenge query $\text{CH}_{\mathcal{R}}(R^*)$ query from \mathcal{R} , \mathcal{M} answers with $h^* \stackrel{\boxtimes}{\leftarrow} \text{CH}_{\mathcal{M}}(R^*)$, where $\text{CH}_{\mathcal{M}}$ is provided by the IMP-AA experiment. Finally, \mathcal{R} outputs s^* to break IMP-KOA security and terminates. \mathcal{M} also outputs s^* to its IMP-AA experiment and terminates. We note that $(R^*, h^*, s^*) = q(R, h, s)$ with probability $1/|\text{ChSet}|$, since h^* is a random challenge chosen by the IMP-AA experiment and h is the response of a random oracle query. Thus, if s^* breaks IMP-KOA security, then s^* breaks IMP-AA security. Moreover, \mathcal{M} perfectly simulates an adversary that $(t_{\mathcal{A}}, \varepsilon_{\mathcal{A}}, Q_h)$ -breaks UF-KOA security. This establishes $\varepsilon_{\mathcal{M}} \geq \varepsilon_{\mathcal{R}} - 1/|\text{ChSet}|$. ■

By Lemma 3.6, Lemma 4.5 implies that there is no reduction from PIMP-KOA to UF-KOA in the non-programmable random oracle model.

The following simple lemma actually holds for any signature scheme SIG.

Lemma 4.6 (UF-KOA $\xrightarrow{\text{NPRO}}$ UF-CMA). *Suppose that there is a key-preserving reduction \mathcal{R} in the non-programmable random oracle (NPRO) model that $(t_{\mathcal{R}}, \varepsilon_{\mathcal{R}}, Q_h)$ -breaks UF-KOA security of SIG with n -time black-box access to an adversary \mathcal{A} that $(\varepsilon_{\mathcal{A}}, t_{\mathcal{A}}, Q_s, Q_h)$ -breaks UF-CMA security of SIG. Then there exists an algorithm \mathcal{M} that $(\varepsilon_{\mathcal{M}}, t_{\mathcal{M}})$ -breaks UF-KOA security of SIG, where*

$$\varepsilon_{\mathcal{M}} \geq \varepsilon_{\mathcal{R}}, t_{\mathcal{M}} \approx t_{\mathcal{R}}.$$

Proof. Assuming the existence of a key-preserving reduction \mathcal{R} as above, we construct a meta-reduction \mathcal{M} to break UF-KOA security of SIG[ID]. Similar to the proof of the last lemma, we can assume w.l.o.g. $n = 1$. \mathcal{M} gets the public key pk from the UF-KOA experiment and simulates an adversary \mathcal{A} . Additionally, both \mathcal{M} and \mathcal{R} get access to the same external random oracle H , in the NPRO model.

CONSTRUCTION OF $\mathcal{M}(pk)$. \mathcal{M} runs $\mathcal{R}(pk)$ and, upon receiving pk from \mathcal{R} , \mathcal{M} make a signing query on $m \xleftarrow{\$} \mathcal{M}$ to \mathcal{R} . Upon receiving the signature $\sigma = (R, s)$, \mathcal{M} terminates and returns (m, σ) as a UF-KOA forgery. As both \mathcal{M} and \mathcal{R} have access to the same random oracle, (m, σ) is also a valid forgery in UF-KOA experiment. (This argument does not hold if the random oracle is programmable.) Thus, we have $\varepsilon_{\mathcal{M}} \geq \varepsilon_{\mathcal{R}}$. ■

Remark 4.7. All the reductions considered in this section are key-preserving which is the main restriction of our results. If pk and R are elements from some multiplicative group \mathbb{G} of prime order p , then we can extend our previous techniques to exclude the larger class of algebraic reductions. A reduction is algebraic, if for all group elements h output by the reduction, their respective representation is known. That is, if at some point of its execution the reduction holds group elements $g_1, \dots, g_n \in \mathbb{G}$ and outputs a new group element h , then it also knows its representation meaning it also outputs $(\alpha_1, \dots, \alpha_n) \in \mathbb{Z}_p^n$ satisfying $h = \prod g_i^{\alpha_i}$. Note that key-preserving and randomness-preserving reductions are a special case of algebraic reductions.

5 Instantiations

In this section we consider three important identification schemes, namely the ones by Schnorr [Sch91], by Katz-Wang [KW03, CP93] and by Guillou-Quisquater [GQ90]. We use our framework to derive tight security bounds and concrete parameters for the corresponding Schnorr/Katz-Wang/Guillou-Quisquater signature schemes.

5.1 Schnorr Identification/Signature Scheme

5.1.1 Schnorr's Identification Scheme

The well-known Schnorr's identification scheme is one of the most important instantiations of our framework. For completeness we show that Schnorr's identification has large min-entropy, special soundness (SS), honest-verifier zero-knowledge (HVZK), random-self reducibility (RSR) and key-recovery security (KR-KOA) based on the discrete logarithm problem (DLOG). Moreover, based on the one-more discrete logarithm problem (OMDL), Schnorr's identification is actively secure (IMP-AA) and weakly secure against man-in-the-middle attack (wIMP-MIM).

Let $\text{par} := (p, g, \mathbb{G})$ be a set of system parameters, where $\mathbb{G} = \langle g \rangle$ is a cyclic group of prime order p with a hard discrete logarithm problem. Examples of groups \mathbb{G} include appropriate subgroups of certain elliptic curve groups, or subgroups of \mathbb{Z}_q^* . The Schnorr identification scheme $\text{ID}_{\mathcal{S}} := (\text{IGen}, \text{P}, \text{ChSet}, \text{V})$ is defined as follows.

<p><u>IGen(par):</u> $sk := x \xleftarrow{\\$} \mathbb{Z}_p$ $pk := X = g^x$ $\text{ChSet} := \{0, 1\}^n$ Return (pk, sk)</p> <p><u>V(pk, R, h, s):</u> If $R = g^s \cdot X^{-h}$ then return 1 Else return 0</p>	<p><u>P₁(sk):</u> $r \xleftarrow{\\$} \mathbb{Z}_p; R = g^r$ $St := r$ Return (R, St)</p> <p><u>P₂(sk, R, h, St):</u> Parse $St = r$ Return $s = x \cdot h + r \bmod p$</p>
--	---

We recall the DLOG assumption.

Definition 5.1 (Discrete Logarithm Assumption). *The discrete logarithm problem DLOG is (t, ε) -hard in $\text{par} = (p, g, \mathbb{G})$ if for all adversaries \mathcal{A} running in time at most t ,*

$$\Pr [g^x = X \mid X \stackrel{\boxtimes}{\leftarrow} \mathbb{G}; x \stackrel{\boxtimes}{\leftarrow} \mathcal{A}(X)] \leq \varepsilon.$$

Lemma 5.2. *ID_S is a canonical identification with $\alpha = \log p$ bit min-entropy and it is unique, has special soundness (SS), honest-verifier zero-knowledge (HVZK) and is random-self reducible (RSR). Moreover, if DLOG is (t, ε) -hard in $\text{par} = (p, g, \mathbb{G})$ then ID_S is (t, ε) -KR-KOA secure.*

Proof. The correctness of ID_S is straightforward to verify. We note that R in $(R, St) \stackrel{\boxtimes}{\leftarrow} \text{P}_1(sk)$ is uniformly random over \mathbb{G} . Hence, ID_S has $\log |\mathbb{G}| = \log p$ bit min-entropy. We show the other properties as follows.

UNIQUENESS. For all $(X, x) \in \text{IGen}(\text{par})$, $(R := g^r, St := r) \in \text{P}_1(sk)$ and $h \in \{0, 1\}^n$, the value $s \in \mathbb{Z}_p$ satisfying $g^s = X^h R \Leftrightarrow s = xh + r$ is uniquely defined.

SPECIAL SOUNDNESS (SS). Given two accepting transcripts (R, h, s) and (R, h', s') with $h \neq h'$, we define an extractor algorithm $\text{Ext}(X, R, h, s, h', s') := x^* := (s - s')/(h - h')$ such that, for all $(X := g^x, x) \in \text{IGen}(\text{par})$, we have $\Pr[g^{x^*} = X] = 1$, since we have $R = g^s X^{-h} = g^{s'} X^{-h'}$ and then $X = g^{(s-s')/(h-h')}$.

HONEST-VERIFIER ZERO-KNOWLEDGE (HVZK). Given public key X , we let $\text{Sim}(X)$ first sample $h \stackrel{\boxtimes}{\leftarrow} \{0, 1\}^n$ and $s \stackrel{\boxtimes}{\leftarrow} \mathbb{Z}_p$ and then output $(R := g^s X^{-h}, h, s)$. Clearly, (R, h, s) is a real transcript, since s is uniformly random over \mathbb{Z}_p and R is the unique value satisfying $R = g^s X^{-h}$.

RANDOM-SELF REDUCIBILITY (RSR). Algorithm Rerand and two deterministic algorithm Derand and Tran are defined as follows:

- $\text{Rerand}(X)$ chooses $\tau' \stackrel{\boxtimes}{\leftarrow} \mathbb{Z}_p$ and outputs $(X' := X \cdot g^{\tau'}, \tau')$. We have that, for all $(X, x) \in \text{IGen}(\text{par})$, X' is uniform and has the same distribution as X'' , where $(X'', x'') \stackrel{\boxtimes}{\leftarrow} \text{IGen}(\text{par})$.
- $\text{Derand}(X, X', x', \tau')$ outputs $x^* = x' - \tau'$. We have, for all $(X', \tau') \stackrel{\boxtimes}{\leftarrow} \text{Rerand}(X := g^x)$ and $(X', x') \in \text{IGen}(\text{par})$, $X' = g^{x'}$ and $x' = x + \tau'$ and thus $x^* = x$.
- $\text{Tran}(X, X', \tau', (R', h', s'))$ outputs $s = s' - \tau' \cdot h'$. We have, for all $(X', \tau') \in \text{Rerand}(X := g^x)$, if (R', h', s') is valid with respect to $X' := g^{x+\tau'}$ then $s = s' - \tau' \cdot h' = (x + \tau')h' + r - \tau' \cdot h' = xh' + r$ and (R', h', s) is valid with respect to X .

KEY-RECOVERY AGAINST KEY-ONLY ATTACK (KR-KOA). KR-KOA-security for ID is exactly the DLOG assumption. ■

We recall the OMDL assumption.

Definition 5.3 (One-more Discrete Logarithm Assumption [BNPS03]). *We says that OMDL is (t, ε, Q) -hard in $\text{par} = (p, g, \mathbb{G})$ if for all adversaries \mathcal{A} running in time at most t and adaptively making at most Q queries to the discrete logarithm oracle DL,*

$$\Pr \left[\text{For } i \in [Q + 1] : X_i = g^{x_i} \mid \begin{array}{l} X_1, \dots, X_{Q+1} \stackrel{\boxtimes}{\leftarrow} \mathbb{G} \\ (x_1, \dots, x_{Q+1}) \stackrel{\boxtimes}{\leftarrow} \mathcal{A}^{\text{DL}(\cdot)}(X_1, \dots, X_{Q+1}) \end{array} \right] \leq \varepsilon,$$

where on input arbitrary group element Y the discrete logarithm oracle DL returns $y \in \mathbb{Z}_p$ such that $g^y = Y$.

Lemma 5.4 (Theorem 5.1 in [BP02]). *If the OMDL problem is (t, ε, Q) -hard then ID_S is (t', ε', Q_0) -IMP-AA secure, where $\varepsilon' \leq \sqrt{\varepsilon} + 1/p$, $t \approx 2t'$, and $Q_0 = Q$.*

We now show that the Schnorr identification scheme is weakly IMP-MIM secure based on one-more discrete logarithm assumption.

Lemma 5.5. *If OMDL problem is (t, ε, Q) -hard then ID_S is $(t', \varepsilon', Q_{\text{Ch}}, Q_0)$ -wIMP-MIM secure, where*

$$\varepsilon' = \varepsilon, \quad t' \approx t, \quad Q_0 = Q.$$

Proof. Let \mathcal{A} be an algorithm that breaks $(t', \varepsilon', Q_{\text{Ch}}, Q_0)$ -wIMP-MIM security of ID_S . We will describe an adversary \mathcal{B} invoking \mathcal{A} that (t, ε, Q) -breaks OMDL with (t, ε, Q) as stated in the theorem. Adversary \mathcal{B} obtains X_1, \dots, X_{Q+1} , and has access to a discrete logarithm oracle DL. \mathcal{B} runs \mathcal{A} on input $pk := X := X_{Q+1}$ and answers the adaptive PROVER and CH queries as follows:

- On the j -th $\text{PROVER}_1()$ query ($j \in [Q_0]$), \mathcal{B} returns $R'_j := X_j$.
- On the j -th $\text{PROVER}_2(j, h'_j)$ query, \mathcal{B} queries and returns $s'_j = \text{DL}(X^{h'_j} \cdot R'_j)$.
- On the i -th $\text{CH}(R_i)$ query, \mathcal{B} chooses a random $h_i \xleftarrow{\boxtimes} \text{ChSet}$ and returns h_i .

Note that weak IMP-MIM security requires that for all i, j , we have $R_i = R'_j$ for some R'_j previously returned by the $\text{PROVER}_1()$ oracle.

Eventually, \mathcal{A} returns (i^*, s_{i^*}) and terminates. We can assume that \mathcal{A} has made the queries $\text{PROVER}_2(j, h'_j)$ for all $j \in [Q_0]$. If not \mathcal{B} makes the dummy query $\text{PROVER}_2(j, h'_j)$ for an arbitrary $h'_j \neq h_{i^*}$ to obtain a valid transcript (R'_j, h'_j, s'_j) for all $j \in [Q_0]$. So in total, \mathcal{B} made exactly Q_0 calls to the DL oracle.

\mathcal{A} wins if $(R_{i^*}, h_{i^*}, s_{i^*})$ is a valid transcript, $(R_{i^*}, h_{i^*}, s_{i^*}) \notin \{(R'_j, h'_j, s'_j) \mid j \in [Q_0]\}$, and $R_{i^*} = R'_{j^*}$, for some index j^* . (If there exists more than one index j^* , we fix an arbitrary one.) From the above observations we conclude that \mathcal{B} knows two valid transcripts, $(R_{i^*}, h_{i^*}, s_{i^*})$ and $(R'_{j^*}, h'_{j^*}, s'_{j^*})$ satisfying $(h_{i^*}, s_{i^*}) \neq (h'_{j^*}, s'_{j^*})$. From the two valid transcripts, \mathcal{B} can reconstruct $sk = x_{Q+1}$ using the special soundness of the Schnorr identification scheme. Furthermore, since $(R'_j, h'_j, s'_j) = (X_j, h'_j, s'_j)$ is a valid transcript and x_{Q+1} is known, \mathcal{B} can compute $x_j = s'_j - x_{Q+1}h'_j$ for all $j \in [Q]$. Finally, \mathcal{B} returns (x_1, \dots, x_{Q+1}) , breaks OMDL problem with $\varepsilon = \varepsilon'$ and $t \approx t'$. ■

We now define the Q -interactive discrete-logarithm problem which precisely models PIMP-KOA-security for ID_S , where $Q = Q_0$ is the number of parallel impersonation rounds.

Definition 5.6 (Q -IDLOG). *The interactive discrete-logarithm assumption Q -IDLOG is said to be (t, ε) -hard in $\text{par} = (p, g, \mathbb{G})$ if for all adversaries \mathcal{A} running in time at most t and making at most Q queries to the challenge oracle CH ,*

$$\Pr \left[s \in \{xh_i + r_i \mid i \in [Q]\} \mid \begin{array}{l} x \xleftarrow{\boxtimes} \mathbb{Z}_p; X = g^x \\ s \xleftarrow{\boxtimes} \mathcal{A}^{\text{CH}(\cdot)}(X) \end{array} \right] \leq \varepsilon,$$

where on the i -th query $\text{CH}(g^{r_i})$ ($i \in [Q]$), the challenge oracle returns $h_i \xleftarrow{\boxtimes} \mathbb{Z}_p$ to \mathcal{A} .

In Appendix A we prove that in the generic group model, the Q -IDLOG problem in groups of prime-order p is at least $(2t^2/p, t)$ -hard. Note that the bound is independent of Q .

5.1.2 Schnorr's Signature scheme

Let $H : \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a hash function with $n < \log_2(p)$. As ID_S is commitment-recoverable we can use the alternative Fiat-Shamir transformation to obtain the Schnorr signature scheme $\text{Schnorr} := (\text{Gen}, \text{Sign}, \text{Ver})$.

<p>Gen(par): $sk := x \xleftarrow{\boxtimes} \mathbb{Z}_p$ $pk := X = g^x$ Return (pk, sk)</p>	<p>Sign(sk, m): $r \xleftarrow{\boxtimes} \mathbb{Z}_p; R = g^r$ $h = H(R, m)$ $s = x \cdot h + r \pmod p$ $\sigma = (h, s) \in \{0, 1\}^n \times \mathbb{Z}_p$ Return σ</p>	<p>Ver(pk, m, σ): Parse $\sigma = (h, s) \in \{0, 1\}^n \times \mathbb{Z}_p$ $R = g^s X^{-h}$ If $h = H(R, m)$ then return 1 Else return 0.</p>
---	---	---

The DLOG problem is tightly equivalent to the 1-IDLOG problem by Lemma 3.5. Assuming the OMDL problem is hard, Schnorr is wIMP-MIM-secure and by Corollary 4.4 there cannot exist a tight implication $1\text{-IDLOG} \rightarrow Q\text{-IDLOG}$ meaning the bound from Lemma 3.6 is optimal. By Lemmas 3.7 and 3.8, the Q -IDLOG problem is tightly equivalent to SUF-CMA-security of Schnorr in the programmable ROM. The latter is only tightly equivalent to MU-SUF-CMA-security in the programmable ROM (via Lemmas 3.9 and 3.10). Lemma B.1 in the Appendix improves this by proving that SUF-CMA security is tightly equivalent to MU-SUF-CMA-security in the standard model. Figure 5 summarizes the modular security implications for Schnorr.

We derive the following concrete security implications.




Figure 5: Security relations for the Schnorr signature scheme. All implications except the red one are tight.

Lemma 5.7. *If DLOG is (t, ε) -hard in $\text{par} = (p, g, \mathbb{G})$ then Schnorr is $(t', \varepsilon', Q_s, Q_h)$ -SUF-CMA secure and $(t'', \varepsilon'', N, Q_s, Q_h)$ -MU-SUF-CMA secure in the programmable random oracle model, where*

$$\begin{aligned}
\frac{\varepsilon'}{t'} &\leq 6(Q_h + 1) \cdot \frac{\varepsilon}{t} + \frac{Q_s}{p} + \frac{1}{2^n}, \\
\frac{\varepsilon''}{t''} &\leq 12(Q_h + 1) \cdot \frac{\varepsilon}{t} + \frac{Q_s}{p} + \frac{1}{2^n},
\end{aligned}$$

Lemma 5.8. *If Q_h -IDLOG is (t, ε) -hard in par then Schnorr is $(t', \varepsilon', N, Q_s, Q_h)$ -MU-SUF-CMA secure in the programmable random oracle model, where*

$$\varepsilon' \leq 2\varepsilon + \frac{Q_h Q_s}{p}, \quad t' \approx t.$$

We leave it an open problem to come up with a more natural hard problem over par that tightly implies Q -IDLOG (and hence MU-SUF-CMA-security of Schnorr). Note that according to [FJS14], the hard problem has to have at least one round of interaction.

The interpretation for the multi-user security of Schnorr over elliptic-curve groups is as follows. It is well-known that a group of order p providing k -bits security against the DLOG problem requires $\log p \geq 2k$. If one requires provable security guarantees for Schnorr under DLOG, then one has to increase the group size by $\approx \log(Q_h)$ bits. Reasonable upper bounds for $\log Q_h$ are between 40 and 80. However, the generic lower bound of Theorem A.1 indicates that the only way to attack Schnorr in the sense of UF-KOA (and hence to attack Q -IDLOG) is to break the DLOG problem. In that case using groups with $\log p \approx 2k$ already gives security guarantees for Schnorr.

5.2 Chaum-Pedersen Identification/Katz-Wang Signature Scheme

5.2.1 Chaum-Pedersen Identification Scheme

Let $\text{par} := (p, g_1, g_2, \mathbb{G})$ be a set of system parameters, where $\mathbb{G} = \langle g_1 \rangle = \langle g_2 \rangle$ is a cyclic group of prime order p . The Chaum-Pedersen identification scheme $\text{ID}_{\text{CP}} := (\text{IGen}, \text{P}, \text{ChSet}, \text{V})$ is defined as follows.

<p><u>IGen(par):</u> $sk := x \stackrel{\boxtimes}{\leftarrow} \mathbb{Z}_p$ $pk := (X_1, X_2) = (g_1^x, g_2^x)$ $\text{ChSet} := \{0, 1\}^n$ Return (pk, sk)</p> <p><u>V(pk, R = (R₁, R₂), h, s):</u> If $R_1 = g^s \cdot X_1^{-h}$ and $R_2 = g^s \cdot X_2^{-h}$ then return 1 Else return 0</p>	<p><u>P₁(sk):</u> $r \stackrel{\boxtimes}{\leftarrow} \mathbb{Z}_p$; $R = (R_1, R_2) = (g_1^r, g_2^r)$ $St := r$ Return (R, St)</p> <p><u>P₂(sk, R, h, St):</u> Parse $St = r$ Return $s = x \cdot h + r \bmod p$</p>
---	--

We recall the DDH assumption.

Definition 5.9 (Decision Diffie-Hellman Assumption). *The Decision Diffie-Hellman problem DDH is (t, ε) -hard in $\text{par} = (p, g_1, g_2, \mathbb{G})$ if for all adversaries \mathcal{A} running in time at most t ,*

$$|\Pr [1 \stackrel{\boxtimes}{\leftarrow} \mathcal{A}(g_1^x, g_2^x) \mid x \stackrel{\boxtimes}{\leftarrow} \mathbb{Z}_p] - \Pr [1 \stackrel{\boxtimes}{\leftarrow} \mathcal{A}(g_1^{x_1}, g_2^{x_2}) \mid x_1 \stackrel{\boxtimes}{\leftarrow} \mathbb{Z}_p; x_2 \stackrel{\boxtimes}{\leftarrow} \mathbb{Z}_p \setminus \{x_1\}]| \leq \varepsilon.$$

Clearly, all security results of Schnorr carry over to the Chaum-Pedersen identification scheme, i.e., ID_{CP} is at least as secure as ID_{S} . That also means that we cannot hope for tight PIMP-KOA security from the DLOG assumption. Instead, for the Chaum-Pedersen identification scheme, we give a direct tight proof of PIMP-KOA security under the DDH assumption which we extracted from [KW03].

Lemma 5.10. ID_{CP} is a canonical identification scheme with $\alpha = \log p$ bit min-entropy and it is unique, has special soundness (SS), honest-verifier zero-knowledge (HVZK) and is random-self reducible (RSR). Moreover, if DDH is (t, ε) -hard in $\text{par} = (p, g_1, g_2, \mathbb{G})$ then ID_{CP} is $(t', \varepsilon', Q_{\text{CH}})$ -PIMP-KOA secure, where $t \approx t'$ and $\varepsilon \geq \varepsilon' - Q_{\text{CH}}/2^n$.

Proof. The proof of SS, HVZK, uniqueness, and RSR is the same as in ID_{S} .

To prove PIMP-KOA-security under DDH, let \mathcal{A} be an adversary that $(t', \varepsilon', Q_{\text{CH}})$ -breaks PIMP-KOA security. We build an adversary \mathcal{B} against the (t, ε) -hardness of DDH as follows. Adversary \mathcal{B} inputs (X_1, X_2) and defines $pk := (X_1, X_2)$. On the i -th challenge query $\text{Ch}(R_{i,1}, R_{i,2})$, it returns $h_i \xleftarrow{\boxtimes} \mathbb{Z}_p$. Eventually, \mathcal{A} returns $i^* \in [Q_{\text{CH}}]$ and s_{i^*} and terminates. Finally, \mathcal{B} outputs $d := \mathcal{V}(pk, R_{i^*}, h_{i^*}, s_{i^*})$.

ANALYSIS OF \mathcal{B} . If $(X_1, X_2) = (g_1^x, g_2^x)$, then \mathcal{B} perfectly simulates the PIMP-KOA game and hence $\Pr[d = 1 \mid (X_1, X_2) = (g_1^x, g_2^x)] = \varepsilon'$. If $(X_1, X_2) = (g_1^{x_1}, g_2^{x_2})$ with $x_1 \neq x_2$, then we claim that even a computationally unbounded \mathcal{A} can only win with probability $Q_{\text{CH}}/2^n$, i.e., $\Pr[d = 1 \mid (X_1, X_2) = (g_1^{x_1}, g_2^{x_2})] \leq Q_{\text{CH}}/2^n$.

It remains to prove the claim. For each index $i \in [Q_{\text{CH}}]$, \mathcal{A} first commits to $R_{i,1} = g_1^{r_{i,1}}$ and $R_{i,2} = g_2^{r_{i,2}}$ (for arbitrary $r_{i,1}, r_{i,2} \in \mathbb{Z}_p$) and can only win if there exists an $s_i \in \mathbb{Z}_p$ such that

$$\begin{aligned} r_{i,1} + h_i x_1 &= s_i = r_{i,2} + h_i x_2 \\ \Leftrightarrow h_i &= \frac{r_{i,2} - r_{i,1}}{x_1 - x_2} \end{aligned}$$

where $h_i \xleftarrow{\boxtimes} \{0, 1\}^n$ is chosen independently of $r_{i,1}, r_{i,2}$. This happens with probability at most $1/2^n$, so by the union bound we obtain the bound $Q_{\text{CH}}/2^n$, as claimed. ■

5.2.2 Katz-Wang Signature scheme

Let $H : \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a hash function with $n < \log_2(p)$. As ID_{CP} is commitment-recoverable we can use the alternative Fiat-Shamir transformation to obtain a signature scheme which is known as the Katz-Wang signature scheme $\text{KW} := (\text{Gen}, \text{Sign}, \text{Ver})$.

<p><u>Gen(par):</u> $sk := x \xleftarrow{\boxtimes} \mathbb{Z}_p$ $pk := (X_1, X_2) = (g_1^x, g_2^x)$ Return (pk, sk)</p>	<p><u>Sign(sk, m):</u> $r \xleftarrow{\boxtimes} \mathbb{Z}_p; R = (R_1, R_2) = (g_1^r, g_2^r)$ $h = H(R, m)$ $s = x \cdot h + r \pmod p$ $\sigma = (h, s) \in \{0, 1\}^n \times \mathbb{Z}_p$ Return σ</p>	<p><u>Ver(pk, m, σ):</u> Parse $\sigma = (h, s) \in \{0, 1\}^n \times \mathbb{Z}_p$ $R = g^s X^{-h}$ If $h = H(R, m)$ then return 1 Else return 0.</p>
---	--	---

By our results we obtain the following concrete security statements, where the first bound matches [KW03, Theorem 1].

Lemma 5.11. If DDH is (t, ε) -hard in $\text{par} = (p, g_1, g_2, \mathbb{G})$ then KW is $(t', \varepsilon', Q_s, Q_h)$ -SUF-CMA secure and $(t'', \varepsilon'', N, Q_s, Q_h)$ -MU-SUF-CMA secure in the programmable random oracle model, where

$$\begin{aligned} \frac{\varepsilon'}{t'} &\leq \frac{\varepsilon}{t} + \frac{Q_s}{p} + \frac{1}{2^n}, \\ \frac{\varepsilon''}{t''} &\leq 4 \cdot \frac{\varepsilon}{t} + \frac{Q_s}{p} + \frac{1}{2^n}. \end{aligned}$$

This bound shows that groups with $\log p \approx 2k$ and $n \approx k$ provides k bits security for KW.

5.3 Guillou-Quisquater Identification/Signature Scheme

5.3.1 Background on RSA and Notations

For $n \in \mathbb{N}$ we denote by $\mathbb{P}_{n/2}$ the set of all $n/2$ -bit primes and $\text{RSA}_n := \{(N = pq, p, q) \mid p, q \in \mathbb{P}_{n/2}, p \neq q\}$. Let $\phi(N) := (p-1)(q-1)$ be Euler's totient function for $(N, p, q) \in \text{RSA}_n$. Let \mathcal{R} be a relation on p and q . By $\text{RSA}_n[\mathcal{R}]$ we denote the subset of RSA_n for that the relation \mathcal{R} holds on p and q .

Let $n \in \mathbb{N}$ and $0 < c < \frac{1}{4}$ be a constant. We define the following two distributions.

$$\begin{aligned} \mathcal{I}_{n,c} &:= \{(N, e) \mid e \stackrel{\boxtimes}{\leftarrow} \mathbb{P}_{cn}; (N, p, q) \stackrel{\boxtimes}{\leftarrow} \text{RSA}_n[\text{gcd}(e, \phi(N) = 1)]\} \\ \mathcal{L}_{n,c} &:= \{(N, e) \mid e \stackrel{\boxtimes}{\leftarrow} \mathbb{P}_{cn}; (N, p, q) \stackrel{\boxtimes}{\leftarrow} \text{RSA}_n[p = 1 \bmod e, p \neq 1 \bmod e^2, q \neq 1 \bmod e]\} \end{aligned}$$

Using the above notation, we recall the Phi-hiding assumption [CMS99, KOS10, KK12].

Definition 5.12 (Phi-hiding Assumption). *The Phi-hiding problem $\phi\text{-H}_{n,c}$ is (t, ε) -hard if for all adversaries \mathcal{A} running in time at most t ,*

$$|\Pr[1 \stackrel{\boxtimes}{\leftarrow} \mathcal{A}(N, e) \mid (N, e) \stackrel{\boxtimes}{\leftarrow} \mathcal{I}_{n,c}] - \Pr[1 \stackrel{\boxtimes}{\leftarrow} \mathcal{A}(N, e) \mid (N, e) \stackrel{\boxtimes}{\leftarrow} \mathcal{L}_{n,c}]| \leq \varepsilon.$$

5.3.2 Guillou-Quisquater Identification Scheme

Let $\text{par} = (N, e) \stackrel{\boxtimes}{\leftarrow} \mathcal{I}_{n,c}$ be system parameters. The Guillou-Quisquater identification scheme $\text{ID}_{\text{GQ}} := (\text{IGen}, \text{P}, \text{ChSet}, \text{V})$ is defined as follows, where $\mathbb{Z}_N^* := \{y \in \mathbb{Z}_N \mid \text{gcd}(y, N) = 1\}$.

<p>IGen(par): $sk := x \stackrel{\boxtimes}{\leftarrow} \mathbb{Z}_N^*$ $pk := X := x^e \bmod N$ $\text{ChSet} := \mathbb{Z}_e$ Return (pk, sk)</p> <p>V(pk, R, h, s): If $R = s^e \cdot X^{-h} \bmod N$ and $((R, s) \in \mathbb{Z}_N^* \times \mathbb{Z}_N^*)$ then return 1 Else return 0</p>	<p>P₁(sk): $r \stackrel{\boxtimes}{\leftarrow} \mathbb{Z}_N^*; R = r^e \bmod N$ $St := r$ Return (R, St)</p> <p>P₂(sk, R, h, St): Parse $St = r$ Return $s = x^h \cdot r \bmod N$</p>
---	---

It is easy to prove IMP-KOA security of ID_{GQ} under the standard RSA assumption. Using our framework this implies MU-UF-CMA security of the implies GQ signature scheme, with an unavoidable security loss of Q_h . Under the $\phi\text{-H}_{n,c}$ assumption we can, however, give a direct tight proof of PIMP-KOA security, which is similar to [ABP13].

Lemma 5.13. *ID_{GQ} is a canonical identification scheme with $\alpha = \log(\phi(N))$ bit min-entropy and it is unique, has special sound (SS), honest-verifier zero-knowledge (HVZK) and random-self reducible (RSR). Moreover, if $\phi\text{-H}_{n,c}$ is (t, ε) -hard then ID_{GQ} is $(t', \varepsilon', Q_{\text{CH}})$ -PIMP-KOA secure, where $t \approx t'$ and $\varepsilon \geq \varepsilon' - (Q_{\text{CH}} + 1)/e \geq \varepsilon' - (Q_{\text{CH}} + 1)/2^{cn}$.*

Proof. The correctness of ID_{GQ} is straightforward to verify. We note that $R \stackrel{\boxtimes}{\leftarrow} \text{P}_1(sk)$ is uniformly random over \mathbb{Z}_N^* . Hence, ID_{GQ} has $\log |\mathbb{Z}_N^*| = \log(\phi(N))$ bit min-entropy. We show the other properties as follows.

UNIQUENESS. For all $(X, x) \in \text{IGen}(\text{par})$, $R := r^e \in \text{P}_1(sk)$ and $h \in \mathbb{Z}_e$, the value $s \in \mathbb{Z}_N^*$ satisfying $s^e = X^h R \bmod N \Leftrightarrow s = x^h \cdot r \bmod N$ is uniquely defined, since $\text{gcd}(e, \phi(N)) = 1$.

SPECIAL SOUNDNESS (SS). Given two accepting transcripts (R, h, s) and (R, h', s') with $h \neq h'$ (wlog. let $h > h'$), we have $s^e X^h = R = s'^e X^{h'} \bmod N$ and $(s/s')^e = X^{h-h'} \bmod N$. Since $h, h' \in \mathbb{Z}_e$, $\text{gcd}(e, h-h') = 1$. Applying the extended Euclidean algorithm we can compute $A, B \in \mathbb{Z}_N^*$ such that

$$Ae + B(h-h') = \text{gcd}(e, h-h') = 1.$$

Then we define an extractor algorithm $\text{Ext}(X, R, h, s, h', s') := x^* := X^A (s/s')^B$ such that, for all $(X := x^e \bmod N, x) \in \text{IGen}(\text{par})$, we have $\Pr[(x^*)^e = X \bmod N] = 1$, since we have $(x^*)^e = (X^A (s/s')^B)^e = X^{Ae} (s/s')^{Be} = X^{Ae} X^{B(h-h')} = X$.

HONEST-VERIFIER ZERO-KNOWLEDGE (HVZK). Given a public key $pk = X$, we let $\text{Sim}(pk)$ first sample $h \xleftarrow{\boxtimes} \mathbb{Z}_e$ and $s \xleftarrow{\boxtimes} \mathbb{Z}_N^*$ and then output $(R := s^e X^{-h} \bmod N, h, s)$. Clearly, (R, h, s) is a real transcript, since (h, s) is uniformly random over $\mathbb{Z}_e \times \mathbb{Z}_N^*$ and R is the unique value satisfying $R = s^e X^{-h} \bmod N$.

RANDOM-SELF REDUCIBILITY (RSR). Algorithm Rerand and two deterministic algorithm Derand and Tran are defined as follows:

- $\text{Rerand}(X_1)$ chooses $\tau_2 \xleftarrow{\boxtimes} \mathbb{Z}_N^*$, computes $X_2 := X_1 \cdot \tau_2^e \bmod N$ and returns (X_2, τ_2) . We have that, for all $(X_1, x_1) \in \text{IGen}(\text{par})$, X_2 is uniform and has the same distribution as X_3 , where $(X_2, \tau_2) \xleftarrow{\boxtimes} \text{Rerand}(X_1)$ and $(X_3, x_3) \xleftarrow{\boxtimes} \text{IGen}(\text{par})$.
- $\text{Derand}(X_1, X_2, x_2, \tau_2)$ outputs $x^* = x_2/\tau_2 \bmod N$. We have, for all $(X_2, \tau_2) \xleftarrow{\boxtimes} \text{Rerand}(X_1 := x_1^e \bmod N)$ with $(X_2, x_2) \in \text{IGen}(\text{par})$, $X_2 = x_2^e \bmod N$ and $x_2 = x_1 \cdot \tau_2 \bmod N$ and thus $x^* = x_1 \bmod N$.
- $\text{Tran}(X_1, X_2, \tau_2, (R_2 := r_2^e \bmod N, h_2, s_2))$ outputs $s_1 = s_2/\tau_2^{h_2} \bmod N$. We have, for all $(X_2, \tau_2) \in \text{Rerand}(X_1 := x_1^e \bmod N)$, if (R_2, h_2, s_2) is valid with respect to $X_2 := (x_1 \cdot \tau_2)^e \bmod N$ then $s_1 = s_2/\tau_2^{h_2} = (x_1 \tau_2)^{h_2} \cdot r_2/\tau_2^{h_2} = x_1^{h_2} \cdot r_2 \bmod N$ and (R_2, h_2, s_1) is valid with respect to X_1 .

PIMP-KOA SECURITY. Let \mathcal{A} be an adversary that $(t', \varepsilon', Q_{\text{CH}})$ -breaks PIMP-KOA-security. We build an adversary \mathcal{B} against the (t, ε) -hardness of ϕ - $\text{H}_{n,c}$ as follows. Adversary \mathcal{B} inputs (N, e) , chooses $X \xleftarrow{\boxtimes} \mathbb{Z}_N^*$ and defines $pk := X$. On the i -th challenge query $\text{CH}(R_i)$, it returns $h_i \xleftarrow{\boxtimes} \mathbb{Z}_e$. Eventually, \mathcal{A} returns $i^* \in [Q_{\text{CH}}]$ and s_{i^*} and terminates. Finally, \mathcal{B} outputs $d := \mathbb{V}(pk, R_{i^*}, h_{i^*}, s_{i^*})$.

ANALYSIS OF \mathcal{B} . If $(N, e) \xleftarrow{\boxtimes} \mathcal{I}_{n,c}$, then \mathcal{B} perfectly simulates the PIMP-KOA game and hence $\Pr[d = 1 \mid (N, e) \xleftarrow{\boxtimes} \mathcal{I}_{n,c}] = \varepsilon'$. If $(N, e) \xleftarrow{\boxtimes} \mathcal{L}_{n,c}$ with $\gcd(N, e) \neq 1$, then we claim that even a computationally unbounded \mathcal{A} can only win with probability $(Q_{\text{CH}} + 1)/e$, i.e., $\Pr[d = 1 \mid (N, e) \xleftarrow{\boxtimes} \mathcal{L}_{n,c}] \leq (Q_{\text{CH}} + 1)/e$.

It remains to prove the claim. Let $\mathcal{R}_e := \{X \mid \exists x \in \mathbb{Z}_N^* : X = x^e \bmod N\}$ be the set of all e -th residues in \mathbb{Z}_N^* . For $X, R \in \mathbb{Z}_N^*$ we first analyze

$$p(X, R) := \Pr_{h \xleftarrow{\boxtimes} \mathbb{Z}_e} [\exists s \in \mathbb{Z}_N^* : s^e = X^h \cdot R \bmod N].$$

- Case 1: $X \in \mathcal{R}_e$. Then $p(X, R) \leq 1$ by choosing $s := (X^{1/e})^h \cdot R^{1/e}$ if $R \in \mathcal{R}_e$.
- Case 2: $X \notin \mathcal{R}_e$. Then $p(X, R) \leq 1/e$ because

$$\begin{aligned} p(X, R)(p(X, R) - 1/e) &= \Pr_{h \neq \hat{h}} [\exists s, \hat{s} : s^e = X^h \cdot R \bmod N \wedge \hat{s}^e = X^{\hat{h}} \cdot R \bmod N] \\ &= \Pr_{h \neq \hat{h}} [\exists s, \hat{s} : (s/\hat{s})^e = X^{h-\hat{h}} \bmod N] = 0. \end{aligned}$$

The last equality holds, since $\gcd(e, h - \hat{h}) = 1$ which implies that $(s/\hat{s})^{e/(h-\hat{h})}$ is an e -th residue and cannot equal to $X \notin \mathcal{R}_e$.

Using the bounds on $p(X, R)$ we obtain

$$\begin{aligned} \Pr[\mathcal{A} \text{ wins}] &= \Pr[\mathcal{A} \text{ wins} \mid X \in \mathcal{R}_e] \Pr[X \in \mathcal{R}_e] + \Pr[\mathcal{A} \text{ wins} \mid X \notin \mathcal{R}_e] \cdot \Pr[X \notin \mathcal{R}_e] \\ &\leq \frac{1}{e} + (1 - \frac{1}{e}) \Pr_{h_1, \dots, h_{Q_{\text{CH}}}} [\bigvee \exists s_i : s_i^e = X^{h_i} \cdot R_i \bmod N \mid X \notin \mathcal{R}_e] \\ &\leq \frac{1}{e} + (1 - \frac{1}{e}) \frac{Q_{\text{CH}}}{e} \\ &\leq \frac{Q_{\text{CH}} + 1}{e} \end{aligned}$$

This completes the proof of the claim. ■

5.3.3 Guillou-Quisquater Signature Scheme

Let $H : \{0, 1\}^* \rightarrow \mathbb{Z}_e$ be a hash function. As ID_{GQ} is commitment-recoverable we can use the alternative Fiat-Shamir transformation to obtain the Guillou-Quisquater signature scheme $\text{GQ} := (\text{Gen}, \text{Sign}, \text{Ver})$.

Gen(par): $sk := x \stackrel{\boxtimes}{\leftarrow} \mathbb{Z}_N^*$ $X := x^e \bmod N$ $pk := X$ Return (pk, sk)	Sign(sk, m): $r \stackrel{\boxtimes}{\leftarrow} \mathbb{Z}_N^*$; $R = r^e \bmod N$ $h = H(R, m)$ $s = x^h \cdot r \bmod N$ $\sigma = (h, s) \in \mathbb{Z}_e \times \mathbb{Z}_N^*$ Return σ	Ver(pk, m, σ): Parse $\sigma = (h, s) \in \mathbb{Z}_e \times \mathbb{Z}_N^*$ $R = s^e X^{-h} \bmod N$ If $h = H(R, m)$ and $R \in \mathbb{Z}_N^*$ then return 1 Else return 0.
---	--	--

By our results we obtain the following concrete security statements.

Lemma 5.14. *If ϕ -H $_{n,c}$ is (t, ε) -hard then GQ is $(t', \varepsilon', Q_s, Q_h)$ -SUF-CMA secure and $(t'', \varepsilon'', N, Q_s, Q_h)$ -MU-SUF-CMA secure in the programmable random oracle model, where*

$$\frac{\varepsilon'}{t'} \leq \frac{\varepsilon}{t} + \frac{Q_s}{2^{n-2}} + \frac{3}{2^{cn}},$$

$$\frac{\varepsilon''}{t''} \leq 4 \cdot \frac{\varepsilon}{t} + \frac{Q_s}{2^{n-2}} + \frac{3}{2^{cn}}.$$

Acknowledgments

Eike Kiltz was supported in part by ERC Project ERCC (FP7/615074). Daniel Masny was supported by the DFG Research Training Group GRK 1817/1. Jiaxin Pan was supported by the DFG Research Training Group GRK 1817/1 and in part by the DFG grant HO 4534/4-1.

References

- [AABN02] Michel Abdalla, Jee Hea An, Mihir Bellare, and Chanathip Namprempre. From identification to signatures via the Fiat-Shamir transform: Minimizing assumptions for security and forward-security. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 418–433. Springer, Heidelberg, April / May 2002. (Cited on page 1, 2, 5.)
- [ABP13] Michel Abdalla, Fabrice Ben Hamouda, and David Pointcheval. Tighter reductions for forward-secure signature schemes. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 292–311. Springer, Heidelberg, February / March 2013. (Cited on page 4, 24.)
- [AFLT12] 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. (Cited on page 5.)
- [BDL⁺11] Daniel J. Bernstein, Niels Duif, Tanja Lange, Peter Schwabe, and Bo-Yin Yang. High-speed high-security signatures. In Bart Preneel and Tsuyoshi Takagi, editors, *CHES 2011*, volume 6917 of *LNCS*, pages 124–142. Springer, Heidelberg, September / October 2011. (Cited on page 5.)
- [Ber15a] Daniel Bernstein. [Cfrg] key as message prefix => multi-key security. <https://mailarchive.ietf.org/arch/msg/cfrg/44gJyZlZ7-myJqWkChhpEF1KE9M>, 2015. (Cited on page 5.)
- [Ber15b] Daniel J. Bernstein. Multi-user Schnorr security, revisited. Cryptology ePrint Archive, Report 2015/996, 2015. <http://eprint.iacr.org/>. (Cited on page 5.)
- [Bet88] Thomas Beth. Efficient zero-knowledge identification scheme for smart cards. In C. G. Günther, editor, *EUROCRYPT'88*, volume 330 of *LNCS*, pages 77–84. Springer, Heidelberg, May 1988. (Cited on page 1.)
- [BGLS03] Dan Boneh, Craig Gentry, Ben Lynn, and Hovav Shacham. Aggregate and verifiably encrypted signatures from bilinear maps. In Eli Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 416–432. Springer, Heidelberg, May 2003. (Cited on page 5.)

- [BM91] Ernest F. Brickell and Kevin S. McCurley. An interactive identification scheme based on discrete logarithms and factoring. In Ivan Damgård, editor, *EUROCRYPT'90*, volume 473 of *LNCS*, pages 63–71. Springer, Heidelberg, May 1991. (Cited on page 1.)
- [BN06] Mihir Bellare and Gregory Neven. Multi-signatures in the plain public-key model and a general forking lemma. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 06*, pages 390–399. ACM Press, October / November 2006. (Cited on page 3, 9.)
- [BNPS03] Mihir Bellare, Chanathip Namprempre, David Pointcheval, and Michael Semanko. The one-more-RSA-inversion problems and the security of Chaum’s blind signature scheme. *Journal of Cryptology*, 16(3):185–215, June 2003. (Cited on page 20.)
- [BP02] Mihir Bellare and Adriana Palacio. GQ and Schnorr identification schemes: Proofs of security against impersonation under active and concurrent attacks. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 162–177. Springer, Heidelberg, August 2002. (Cited on page 3, 5, 8, 9, 20.)
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In V. Ashby, editor, *ACM CCS 93*, pages 62–73. ACM Press, November 1993. (Cited on page 2, 8.)
- [BR96] Mihir Bellare and Phillip Rogaway. The exact security of digital signatures: How to sign with RSA and Rabin. In Ueli M. Maurer, editor, *EUROCRYPT'96*, volume 1070 of *LNCS*, pages 399–416. Springer, Heidelberg, May 1996. (Cited on page 2.)
- [BR09] 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. (Cited on page 2.)
- [Bro15] Dan Brown. [Cfrg] key as message prefix => multi-key security. <http://www.ietf.org/mail-archive/web/cfrg/current/msg07336.html>, 2015. (Cited on page 6.)
- [CGH98] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited (preliminary version). In *30th ACM STOC*, pages 209–218. ACM Press, May 1998. (Cited on page 2.)
- [CKMS16] Sanjit Chatterjee, Neal Koblitz, Alfred Menezes, and Palash Sarkar. Another look at tightness II: Practical issues in cryptography. Cryptology ePrint Archive, Report 2016/360, 2016. <http://eprint.iacr.org/>. (Cited on page 2.)
- [CMS99] Christian Cachin, Silvio Micali, and Markus Stadler. Computationally private information retrieval with polylogarithmic communication. In Jacques Stern, editor, *EUROCRYPT'99*, volume 1592 of *LNCS*, pages 402–414. Springer, Heidelberg, May 1999. (Cited on page 24.)
- [CP93] David Chaum and Torben P. Pedersen. Wallet databases with observers. In Ernest F. Brickell, editor, *CRYPTO'92*, volume 740 of *LNCS*, pages 89–105. Springer, Heidelberg, August 1993. (Cited on page 4, 19.)
- [FF13] 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. (Cited on page 5.)
- [FH15] Masayuki Fukumitsu and Shingo Hasegawa. Black-box separations on Fiat-Shamir-type signatures in the non-programmable random oracle model. In Javier Lopez and Chris J. Mitchell, editors, *ISC 2015*, volume 9290 of *LNCS*, pages 3–20. Springer, Heidelberg, September 2015. (Cited on page 3, 5.)

- [FJS14] 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. (Cited on page 5, 22.)
- [FLR⁺10] Marc Fischlin, Anja Lehmann, Thomas Ristenpart, Thomas Shrimpton, Martijn Stam, and Stefano Tessaro. Random oracles with(out) programmability. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 303–320. Springer, Heidelberg, December 2010. (Cited on page 2.)
- [FS87] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO’86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987. (Cited on page 1.)
- [Gal04] David Galindo. The exact security of pairing based encryption and signature schemes. Based on a talk at Workshop on Provable Security, INRIA, Paris, 2004. <http://www.dgalindo.es/galindoEcrypt.pdf>, 2004. (Cited on page 2.)
- [GBL08] 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. (Cited on page 5.)
- [Gir91] Marc Girault. An identity-based identification scheme based on discrete logarithms modulo a composite number (rump session). In Ivan Damgård, editor, *EUROCRYPT’90*, volume 473 of *LNCS*, pages 481–486. Springer, Heidelberg, May 1991. (Cited on page 1.)
- [GJKW07] Eu-Jin Goh, Stanislaw Jarecki, Jonathan Katz, and Nan Wang. Efficient signature schemes with tight reductions to the Diffie-Hellman problems. *Journal of Cryptology*, 20(4):493–514, October 2007. (Cited on page 1, 4.)
- [GMR88] 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. (Cited on page 1.)
- [GMS02] Steven D. Galbraith, John Malone-Lee, and Nigel P. Smart. Public key signatures in the multi-user setting. *Inf. Process. Lett.*, 83(5):263–266, 2002. (Cited on page 1, 3, 5, 12.)
- [GQ90] Louis C. Guillou and Jean-Jacques Quisquater. A “paradoxical” indentity-based signature scheme resulting from zero-knowledge. In Shafi Goldwasser, editor, *CRYPTO’88*, volume 403 of *LNCS*, pages 216–231. Springer, Heidelberg, August 1990. (Cited on page 1, 4, 19.)
- [Ham15] Mike Hamburg. Re: [Cfrg] EC signature: next steps. <https://mailarchive.ietf.org/arch/msg/cfrg/af170b60rLyNZUHBMOPWxcDrVRI>, 2015. (Cited on page 5.)
- [JL] S. Josefsson and I. Liusvaara. Edwards-curve digital signature algorithm (EdDSA), October 7, 2015. <https://tools.ietf.org/html/draft-irtf-cfrg-eddsa-00>. (Cited on page 6.)
- [KK12] 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. (Cited on page 24.)
- [KOS10] Eike Kiltz, Adam O’Neill, and Adam Smith. Instantiability of RSA-OAEP under chosen-plaintext attack. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 295–313. Springer, Heidelberg, August 2010. (Cited on page 24.)
- [KW03] 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. (Cited on page 1, 4, 19, 23.)
- [Mau05] Ueli M. Maurer. Abstract models of computation in cryptography (invited paper). In Nigel P. Smart, editor, *10th IMA International Conference on Cryptography and Coding*, volume 3796 of *LNCS*, pages 1–12. Springer, Heidelberg, December 2005. (Cited on page 29.)

- [MS90] Silvio Micali and Adi Shamir. An improvement of the Fiat-Shamir identification and signature scheme. In Shafi Goldwasser, editor, *CRYPTO'88*, volume 403 of *LNCS*, pages 244–247. Springer, Heidelberg, August 1990. (Cited on page 1.)
- [Oka93] Tatsuaki Okamoto. Provably secure and practical identification schemes and corresponding signature schemes. In Ernest F. Brickell, editor, *CRYPTO'92*, volume 740 of *LNCS*, pages 31–53. Springer, Heidelberg, August 1993. (Cited on page 1.)
- [OO98] Kazuo Ohta and Tatsuaki Okamoto. On concrete security treatment of signatures derived from identification. In Hugo Krawczyk, editor, *CRYPTO'98*, volume 1462 of *LNCS*, pages 354–369. Springer, Heidelberg, August 1998. (Cited on page 2, 3, 5.)
- [OS91] H. Ong and Claus-Peter Schnorr. Fast signature generation with a Fiat-Shamir-like scheme. In Ivan Damgård, editor, *EUROCRYPT'90*, volume 473 of *LNCS*, pages 432–440. Springer, Heidelberg, May 1991. (Cited on page 1.)
- [PS00] David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, 2000. (Cited on page 2, 3, 4.)
- [PV05] 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. (Cited on page 3, 5.)
- [Sch91] Claus-Peter Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991. (Cited on page 1, 4, 19.)
- [Seu12] 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. (Cited on page 3, 5, 15, 17.)
- [Sho97] Victor Shoup. Lower bounds for discrete logarithms and related problems. In Walter Fumy, editor, *EUROCRYPT'97*, volume 1233 of *LNCS*, pages 256–266. Springer, Heidelberg, May 1997. (Cited on page 29.)
- [Str15] Rene Struik. Re: [Cfrg] EC signature: next steps. <https://mailarchive.ietf.org/arch/msg/cfrg/TOWH1DSzB-PfDGK8qEXtF3iC6Vc>, 2015. (Cited on page 5.)

A Hardness of Q -IDLOG in the Generic Group Model

In the generic group model for the discrete logarithm setting [Sho97, Mau05], group operations in group \mathbb{G} can only be carried out via an oracle $\mathcal{O}_{\mathbb{G}}$. Since (\mathbb{G}, \cdot) of order p is isomorphic to $(\mathbb{Z}_p, +)$, elements from \mathbb{G} are internally identified with elements from \mathbb{Z}_p . The oracle maintains a list that initially contains the elements $(1, C_1 = 1)$ (the generator), and $(x, C_x = 2)$ for $x \stackrel{\boxplus}{\leftarrow} \mathbb{Z}_p$, and a counter i that counts the number of entries in the list and is initialized to 2. During the execution of the experiment, the list contains entries of the form (a, C_a) , where $a \in \mathbb{Z}_p$ and $C_a \in \mathbb{N}$ is a counter. On input of two counters $C_a, C_b \in [c] \times [c]$, the oracle looks up the internal values (a, C_a) and (b, C_b) , and computes $z = a + b \bmod p$. If there already exists a tuple (z, C_z) in the list, then counter C_z is output. Otherwise, the counter i is increased by 1, the tuple $(z, C_z := i)$ is stored in the list, and the counter C_z is output.

Theorem A.1. *Let \mathbb{G} be a group of prime order p . Then, every adversary \mathcal{A} against Q -IDLOG in the generic group model making at most $Q_{\mathbb{G}}$ queries to the group oracle $\mathcal{O}_{\mathbb{G}}$, has success probability at most*

$$\varepsilon \leq \frac{(Q_{\mathbb{G}} + 2)^2}{2p} + \frac{2Q}{p}.$$

Proof. Let \mathcal{A} be an adversary against Q -IDLOG in the generic group model. In the proof we will simulate the list with entries of the form $(z(\mathbf{x}), C_{z(\mathbf{x})})$, where z is a polynomial of degree one in some variable \mathbf{x} . As we will see, our simulation will sometimes fail. Initially, the counter is set to $i = 2$ and the list contains the elements $(1, C_1 = 1)$ and $(\mathbf{x}, C_{\mathbf{x}} = 2)$, where \mathbf{x} is a variable. After \mathcal{A} has finished its execution, \mathbf{x}

will be assigned a value $x \xleftarrow{\boxtimes} \mathbb{Z}_p$. \mathcal{A} is invoked on input $C_1 = 1$ and $C_x = 2$. During its execution, \mathcal{A} can query oracle \mathcal{O}_G on $(C_{a(\mathbf{x})}, C_{b(\mathbf{x})}) \in [i] \times [i]$. \mathcal{O}_G first computes the polynomial $z(\mathbf{x}) = a(\mathbf{x}) + b(\mathbf{x})$. If $(z(\mathbf{x}), C_{z(\mathbf{x})})$ is not in the list, \mathcal{O}_G increments counter i and adds $(z(\mathbf{x}), C_{z(\mathbf{x})} := i)$ to the list. Finally, \mathcal{O}_G outputs $C_{z(\mathbf{x})}$. In total, \mathcal{A} makes Q_G queries to this oracle and we denote by $(z_i(\mathbf{x}), i)$ the i -entry in the list ($i \in [Q_G + 2]$).

Furthermore, \mathcal{A} can make queries to $\text{CH}(j)$, for some counter $j \in [c]$, which is answered with $h_j \xleftarrow{\boxtimes} \mathbb{Z}_p$. For $j \in [Q]$, we denote by $(r_j(\mathbf{x}) = a_j\mathbf{x} + b_j, C_{r_j(\mathbf{x})})$ the polynomial associated to the j -th query to the CH oracle. Eventually, \mathcal{A} outputs $s \in \mathbb{Z}_p$ and terminates. Next, $x \xleftarrow{\boxtimes} \mathbb{Z}_p$ is chosen and \mathcal{A} wins if there is a $j \in [Q]$ such that $s = (h_j + a_j)x + b_j$.

We remark that we simulate the \mathcal{O}_G perfectly, if none of the distinct polynomials $z_i(\mathbf{x})$ collide when evaluated on input x . We define Bad as the event that this is the case, i.e. there exist an $i \neq \ell \in [Q_G]$ such that the polynomials $z_i(\mathbf{x}), z_\ell(\mathbf{x})$ are distinct but $z_i(x) = z_\ell(x)$. By a union bound we first bound

$$\begin{aligned} \Pr[\text{Bad}] &= \Pr_x[(\exists i, \ell \in [Q_G] \times [Q_G] : z_i(\mathbf{x}) \neq z_\ell(\mathbf{x}) \wedge z_i(x) = z_\ell(x))] \\ &\leq \binom{Q_G + 2}{2} \frac{1}{p} \leq \frac{(Q_G + 2)^2}{2p}. \end{aligned}$$

The success probability ε of \mathcal{A} can be bounded as

$$\begin{aligned} \varepsilon &\leq \Pr[\text{Bad} \vee \exists j \in [Q] : s = (h_j + a_j)x + b_j] \\ &\leq \Pr[\text{Bad}] + \Pr[\exists j \in [Q] : s = (h_j + a_j)x + b_j] \\ &\leq \frac{(Q_G + 2)^2}{2p} + \Pr_x[\exists j \in [Q] : s = (h_j + a_j)x + b_j \mid h_j \neq -a_j] + \Pr_{h_1, \dots, h_Q}[\exists j \in [Q] : h_j = -a_j] \\ &\leq \frac{(Q_G + 2)^2}{2p} + \frac{2Q}{p} \end{aligned}$$

This completes the proof. \blacksquare

B From Single-user to multi-user security for Schnorr

The following result provides an alternative way to prove (a slightly tighter version of) Theorem 3.2 for the special case of Schnorr signatures. It proves that SUF-CMA security tightly implies MU-SUF-CMA security in the standard model. Note that we require strong security but using Lemma 3.8, UF-KOA security tightly implies SUF-CMA (and hence MU-SUF-CMA) security of Schnorr in the random oracle model.

Lemma B.1 (SUF-CMA \Rightarrow MU-SUF-CMA). *If Schnorr is (t, ε, Q_s) -SUF-CMA secure then, for any $N \geq 1$, Schnorr is $(t', \varepsilon', N, Q_s)$ -MU-SUF-CMA secure, where*

$$\varepsilon' \leq 2\varepsilon + \frac{Q_s^2}{p}, \quad t' \approx t,$$

Q_s is an upper bounds on the number of signing queries and N is the number of users.

Proof. Let \mathcal{A} be an adversary that breaks $(t', \varepsilon', N, Q_s)$ -MU-SUF-CMA security of Schnorr. We construct an adversary \mathcal{B} that breaks (t, ε, Q_s) -SUF-CMA security of Schnorr. Adversary \mathcal{B} is executed in the SUF-CMA experiment. It obtains a public-key $pk = X = g^x$ and has access to a signing oracle SIGN.

SIMULATION OF PUBLIC-KEYS. First, for each $i \in [N]$, adversary \mathcal{B} picks $a_i \xleftarrow{\boxtimes} \mathbb{Z}_p$, secret bits $b_i \xleftarrow{\boxtimes} \{0, 1\}$, and computes

$$pk_i = X_i := X^{b_i} \cdot g^{a_i}. \quad (8)$$

That is, if $b_i = 0$, then $sk_i = a_i$ is known to \mathcal{B} ; if $b_i = 1$ then $sk_i = x + a_i$ is unknown to \mathcal{B} . Note that the public-keys are correctly distributed. Next, \mathcal{B} runs \mathcal{A} on input (pk_1, \dots, pk_N) answering signing queries as follows.

SIMULATION OF SIGNATURE QUERIES. On \mathcal{A} 's j -th signing query $(i_j, m_j) \in [N] \times \{0, 1\}^*$, \mathcal{B} is supposed to return a signature σ_j on message m_j under pk_{i_j} . Those are computed by adversary \mathcal{B} according to the following case distinction.

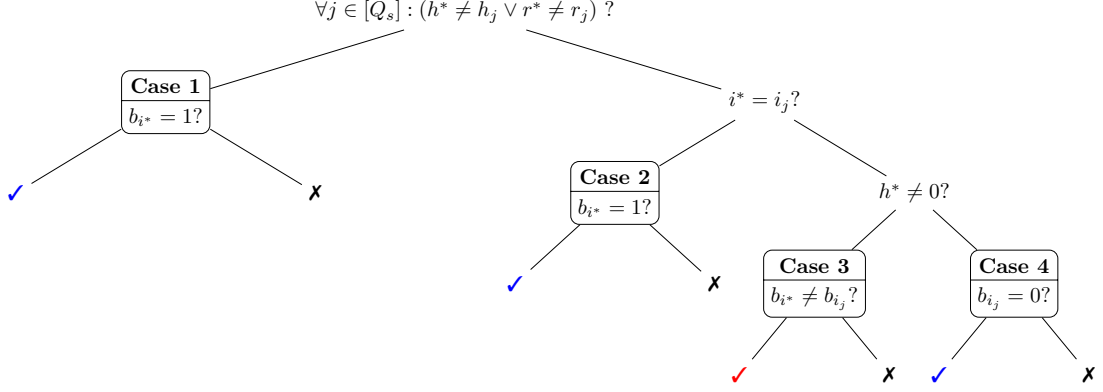


Figure 6: Overview of the case distinction in the proof of Lemma B.1. Each node contains a condition. If the condition is satisfied then we continue to the left child, otherwise to the right child. A leaf denotes either a good case (getting a valid SUF-CMA forgery, marked with “✓”, or extracting the secret-key, marked with “✓”) or a bad case, marked with “✗” (in which we abort).

- Case A: $b_{i_j} = 0$. In that case $sk_{i_j} = a_{i_j}$ is known to \mathcal{B} and the signature is computed as $\sigma_j := (h_j, s_j) \stackrel{\boxplus}{\leftarrow} \text{Sign}(sk_{i_j}, m_j)$.
- Case B: $b_{i_j} = 1$. In that case $sk_{i_j} = x + a_{i_j}$ is unknown to \mathcal{B} and the signature is computed using \mathcal{B} 's signing oracle by first querying $(h_j, \hat{s}_j) \stackrel{\boxplus}{\leftarrow} \text{SIGN}(m_j)$. Then $\sigma_j = (h_j, s_j := \hat{s}_j + a_{i_j} h_j)$ is a valid signature on message m_j under pk_{i_j} . Indeed, $\text{Ver}(pk_{i_j}, m_j) = 1$ because $H(g^{s_j} X_{i_j}^{-h_j}, m_j) = H(g^{\hat{s}_j} X^{-h_j}, m_j) = h_j$.

Adversary \mathcal{B} returns $\sigma_j = (h_j, s_j)$ which in both cases is a correctly distributed valid signature. For future reference we also define $R_j := g^{s_j} X_{i_j}^{-h_j}$ and by (8)

$$r_j := \log_g(R_j) = s_j - (b_{i_j} x + a_{i_j}) h_j. \quad (9)$$

We assume that

$$\forall k \neq j \in [Q_s] : r_k \neq r_j. \quad (10)$$

Since s_j and hence r_j are uniform elements from \mathbb{Z}_p , condition (10) is not satisfied with probability at most Q_s^2/p . Note that the simulation of the public-keys and the signing queries do not leak any information about the secret bits b_i .

FORGERY. Eventually, \mathcal{A} will submit a forgery $(i^*, m^*, \sigma^* := (h^*, s^*))$ and terminate. For the remainder of this proof we assume σ^* is a correct signature on m^* under pk_{i^*} , i.e., for $R^* := g^{s^*} X_{i^*}^{-h^*}$ it holds that $H(R^*, m^*) = h^*$. Using (8) the correctness condition can be equivalently expressed as

$$r^* := \log_g(R^*) = s^* - (b_{i^*} x + a_{i^*}) h^*. \quad (11)$$

Furthermore we assume that σ^* is a valid fresh forgery in the MU-SUF-CMA experiment:

$$(i^*, m^*, h^*, s^*) \notin \{(i_j, m_j, h_j, s_j) \mid j \in [Q_s]\}. \quad (12)$$

After receiving \mathcal{A} 's forgery, \mathcal{B} is supposed to compute its own valid forgery under $pk = X$. To this end, \mathcal{B} defines the set of all indices j such that it queried m_j to its signing oracle $\mathcal{J} := \{j \in [Q_s] \mid b_{i_j} = 1\}$ and makes the following case distinction. A pictorial overview of all cases is given in Figure 6.

- Case 1: For all $j \in [Q_s]$ we have: $h^* \neq h_j$ or $r^* \neq r_j$,
 - Case 1a: $b_{i^*} = 1$. Then for $\hat{s}^* := s^* - a_{i^*} h^*$ we have

$$H(g^{\hat{s}^*} X^{-h^*}, m^*) = H(g^{s^*} X_{i^*}^{-h^*}, m^*) = h^*$$

and hence

$$\hat{\sigma}^* := (h^*, \hat{s}^*)$$

is a correct signature on message m^* under $pk = X$. It remains to show that $\hat{\sigma}^*$ is a fresh strong forgery in the SUF-CMA experiment.

On the one hand, if $h^* \notin \{h_1, \dots, h_{Q_s}\}$, we directly obtain $\hat{\sigma}^* = (h^*, \hat{s}^*) \notin \{(h_j, \hat{s}_j) \mid j \in \mathcal{J}\}$ (the set of all signatures obtained from the SUF-CMA signing oracle) which means that $(m^*, \hat{\sigma}^*)$ satisfies the freshness condition of the SUF-CMA experiment. On the other hand, if the set \mathcal{J}^* of indices $j \in [Q_s]$ such that $h_j = h^*$ is non-empty, then we will use the condition $r^* \neq r_j$ to show that the corresponding \hat{s}_j values are all distinct from \hat{s}^* . Indeed, for all $k \in \mathcal{J}^* \cap \mathcal{J}$ we have $\hat{s}_k = r_k + xh^* \neq r^* + xh^*$ and therefore $\hat{s}^* = r^* + xh^* \notin \{\hat{s}_k \mid k \in \mathcal{J}^* \cap \mathcal{J}\}$. For all $k \in \mathcal{J} \setminus \mathcal{J}^*$ we have $h^* \neq h_k$ and therefore $h^* \notin \{h_k \mid k \in \mathcal{J} \setminus \mathcal{J}^*\}$. Consequently, $\hat{\sigma}^* = (h^*, \hat{s}^*) \notin \{(h_k, \hat{s}_k) \mid k \in \mathcal{J}\}$ and $(m^*, \hat{\sigma}^*)$ satisfies the freshness condition of the SUF-CMA experiment.

– Case 1b: $b_{i^*} = 0$. Then \mathcal{B} aborts.

Note that in case 1, \mathcal{B} aborts with probability exactly 1/2. If it does not abort, it outputs a valid strong forgery.

- Case 2: There exists a $j \in [Q_s]$ such that $h^* = h_j$ and $r^* = r_j$ and $i^* = i_j$.

Note that if j exists it is uniquely defined by (10).

– Case 2a: $b_{i^*} = 1$. As in case 1a,

$$\hat{\sigma}^* := (h^*, \hat{s}^* := s^* - a_{i^*} h^*)$$

is a correct signature on message m^* under $pk = X$. By $r^* = r_j$ and $h^* = h_j$ we obtain $(h^*, s^*) = (h_j, s_j)$. Since we also have $i^* = i_j$, \mathcal{A} 's freshness condition (12) implies $m^* \neq m_j$ meaning that $\hat{\sigma}^*$ is a valid fresh forgery in the SUF-CMA experiment.

– Case 2b: $b_{i^*} = 0$. Then \mathcal{B} aborts.

Note that in case 2, \mathcal{B} aborts with probability exactly 1/2. If it does not abort, it outputs a valid strong forgery.

- Case 3: There exists a $j \in [Q_s]$ such that $h^* = h_j \neq 0$ and $r^* = r_j$ and $i^* \neq i_j$.

Note that if j exists it is uniquely defined by (10).

– Case 3a: $b_{i_j} \neq b_{i^*}$. By (9) and (11) we obtain two equations in the intermediates (r^*, x)

$$\begin{aligned} r^* &= s^* - (b_{i^*} x + a_{i^*}) h^* \\ r^* &= s_j - (b_{i_j} x + a_{i_j}) h^*, \end{aligned}$$

from which \mathcal{B} can extract the single-user scheme's secret-key $x = \log_g(X)$ as

$$x := ((s^* - s_j)(h^*)^{-1} + a_{i_j} - a_{i^*}) \cdot (b_{i^*} - b_{i_j})^{-1}.$$

Using $sk = x$, \mathcal{B} computes a valid forgery on any fresh message.

– Case 3b: $b_{i_j} = b_{i^*}$. Then \mathcal{B} aborts.

Note that in case 3, since $b_{i^*} \neq b_{i_j}$, \mathcal{B} aborts with probability exactly 1/2. If it does not abort, it outputs a valid strong forgery.

- Case 4: There exists a $j \in [Q_s]$ such that $h_j = h^* = 0$ and $r^* = r_j$ and $i^* \neq i_j$.⁶

Again, if j exists it is uniquely defined by (10).

⁶By assuming the hash function to be zero-resistant we may as well discard this case.

– Case 4a: $b_{i_j} = 0$. Then

$$\hat{\sigma}^* := (0, s^*)$$

is a correct signature on m^* under $pk = X$. For all $k \neq j$ with $h_k = h^* = 0$ we have by (10) $r^* \neq r_k$ and therefore $s^* = r^* \neq r_k = \hat{s}_k$. This means that $\hat{\sigma}^* = (0, s^*) = (0, r^*) \notin \{(h_k, \hat{s}_k) \mid k \in \mathcal{J}\}$ (the set of all signatures obtained from the SUF-CMA signing oracle). Therefore $(m^*, \hat{\sigma}^*)$ satisfies the freshness condition of the SUF-CMA experiment.

– Case 4b: $b_{i_j} = 1$. Then \mathcal{B} aborts.

Note that in case 4, \mathcal{B} aborts with probability exactly $1/2$. If it does not abort, it outputs a valid strong forgery.

Overall, \mathcal{B} returns a fresh strong forgery $(m^*, \hat{\sigma}^*)$ under $pk = X$ with probability $\varepsilon = \frac{1}{2}(\varepsilon' - \frac{Q_s^2}{p})$. Adversary \mathcal{B} makes at most Q_s signing queries (in expectation only $Q_s/2$). Its running time is that of \mathcal{A} plus some additional small computation for each signing query and each user (which we neglect), hence $t' \approx t$. ■

We remark that due to forgery cases 1 and 4 our reduction requires strong SUF-CMA security and does not work with standard UF-CMA security.