Optimal Security Proofs for Signatures from Identification Schemes

Eike Kiltz Daniel Masny Jiaxin Pan

Horst-Görtz Institute for IT Security and Faculty of Mathematics,
Ruhr-University Bochum, Germany
{eike.kiltz, daniel.masny, jiaxin.pan}@rub.de

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 rerandomizable and satisfies the weakest possible security notion (key-recoverability), then the implied signature scheme is unforgeability against chosen-message attacks in the multi-user setting in the random oracle model. The reduction loses a factor of roughly Q_h, the number of hash queries. Previous security reductions incorporated an additional multiplicative loss of N, the number of users in the system. As an important application of our framework, we obtain a concrete security treatment for Schnorr signatures.

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 via a meta-reduction.

Keywords: Signatures, Identification, Schnorr, tightness

1 Introduction

The Fiat-Shamir method [15] transforms canonical (i.e., three move) identification schemes into digital signature schemes. It yields very efficient signature schemes, the most popular among them being the Schnorr signature scheme [37].

Canonical Identification Schemes and the Fiat-Shamir Transform. A canonical identification scheme ID as formalized by Abdalla et al. [1] 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 (R, h, s), which we call the (identification) transcript. There is a large number of canonical identification schemes known (e.g. [15, 25, 11, 31, 37, 13, 22, 34, 33]). The Fiat-Shamir method transforms any such canonical identification scheme into a digital signature scheme $\mathsf{SIG}[\mathsf{ID}]$ 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 ($\mathsf{UF\text{-}CMA}$) [23] 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 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 ($\mathsf{MU\text{-}UF\text{-}CMA}$) the attacker obtains N independent public-keys and is said to break the scheme's security if he can produce (after obtaining Q 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 [20] 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 decrease 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 $N = 2^{30}$ 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 $\varepsilon = 2^{-80}$, then it can only be argued that the best algorithm breaking the multi-user security has success ratio $\varepsilon' = 2^{-80} \cdot 2^{30} = 2^{-50}$, which is not a safe security margin that defends against today's attackers.

Tightness. Generally, we call a security implication between two problems *tight*, if the success ratio $\mathsf{SR}(\mathcal{A})$ of any adversary attacking the first problem cannot decease by more than a small constant factor compared to the success ratio $\mathsf{SR}(\mathcal{B})$ of any adversary attacking the second problem. Here the success ratio $\mathsf{SR}(\mathcal{A})$ is defined as the quotient between the adversary's success probability and its running time [6]. 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. 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., [6]), our definition of tightness is sufficient for this purpose.

1.1 Our Contributions

This work contains a concrete and modular security analysis of signatures $\mathsf{SIG}[\mathsf{ID}]$ 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). Most known canonical identification schemes satisfy the above properties.

Security Notions. For identification schemes we consider XXX-YYY security, where $\mathsf{XXX} \in \{\mathsf{KR}, \mathsf{IMP}, \mathsf{PIMP}\}$ denotes the attacker's goal and $\mathsf{YYY} \in \{\mathsf{KOA}, \mathsf{PA}\}$ the attacker's capabilities. If the attacker's goal is key-recovery (KR), then 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. 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. By the above definitions we obtain $3 \times 2 = 6$ different security notions that that were all previously considered in the literature [36, 32, 1], except PIMP-YYY security.

Overview. We show via a chain of implications that KR-KOA-security (the weakest possible security notion for ID) implies multi-user unforgeability against chosen message attacks (MU-UF-CMA) of $\mathsf{SIG}[\mathsf{ID}]$. All our 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 the impossibility without this requirement. For example, our implication PIMP-$\mathsf{KOA} \to \mathsf{UF}$-$\mathsf{KOA}$ is in the random oracle model [7] and we show that the non-programmable random oracle model [17] is not sufficient to prove the same implication. Exactly one of our intermediate implications, namely IMP-$\mathsf{KOA} \to \mathsf{PIMP}$-$\mathsf{KOA}$ is non-tight, and we prove the impossibility of such a tight implication. Diagram in Figure 1 gives an overview of our results. We now discuss them 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* $\mathcal{A}$ *having success ratio* $\mathsf{SR}(\mathcal{A})$*, then* $\mathsf{SIG}[\mathsf{ID}]$ *is* UF-CMA*-secure in the random oracle model against any adversary* $\mathcal{B}$ *having success ratio* $\mathsf{SR}(\mathcal{B}) \approx \mathsf{SR}(\mathcal{A})/Q_h$*, where* Q_h *is the maximal number of* $\mathcal{B}$*'s random oracle queries.*

The proof of this theorem is done in four independent Lemmas 3.4, 3.5, 3.5, and 3.7 via intermediate security notions IMP-KOA, PIMP-KOA, and UF-KOA[1] security, see Figure 1.

[1] Unforgeability against key-only attack (UF-KOA-security) is like standard UF-CMA security, where there adversary is not allowed to ask any signing queries.

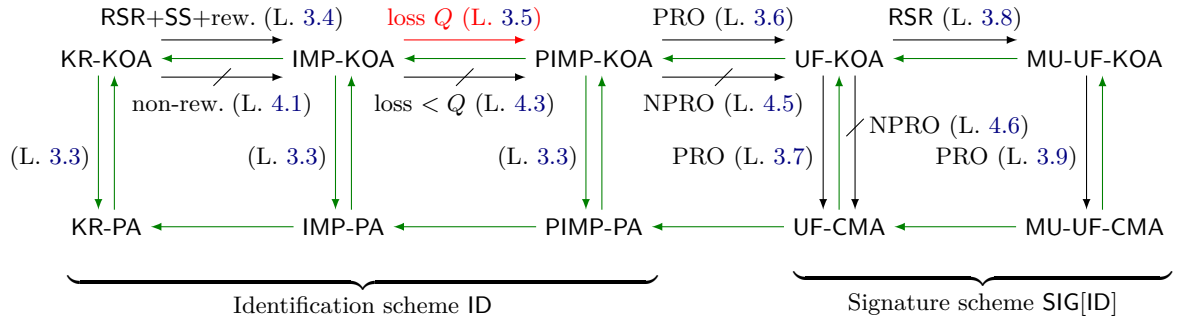

Figure 1: Overview of our notions and results for canonical identification schemes ID and their implied signature schemes SIG[ID]. $\mathsf{X} \xrightarrow{Z} \mathsf{Y}$ means that X-security implies Y-security under condition Z. Trivial implications are denoted with green arrows. All implications are 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. $\mathsf{X} \xrightarrow{Z\!\!\!/} \mathsf{Y}$ means that X-security does not imply Y-security unless they fulfill condition Z. 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).

We certainly do not claim any novelty of the above lemmas, nor a new proof technique. For example, the implication IMP-KOA $\to$ UF-CMA is already explicitly contained in [32] (and implicitly in the seminal paper by Pointcheval and Stern [36]). However, by our specific choice of the intermediate security notions, all four proofs are extremely simple and intuitive. In fact, none of our proofs requires the full power of the Forking Lemma [36]. Lemma 3.4 (KR-KOA $\to$ PIMP-KOA) is the only proof using rewinding and its analysis contains a simple application of Jensen's inequality. If ID is RSR, the implication is tight. We view identifying the intermediate security notions that allow for simple proofs as a conceptual contribution. In particular, 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 notion for ID that is tightly equivalent to MU-UF-CMA security of SIG[ID].

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 [38]; 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 [19] who only consider the implication IMP-PA $\to$ UF-CMA; Lemma 4.1 involves a new meta-reduction. All our impossibility results assume the reductions to be key-preserving [35] and are conditional in the sense that the existence of a reduction would imply that ID does not satisfy some other natural security property (e.g., Lemma 4.1 requires IMP-AA security, where AA stands for active attack).

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* $\mathcal{B}$ *having success ratio* $\mathsf{SR}(\mathcal{B})$*, then it is* MU-UF-CMA*-secure in the random oracle model against any adversary* $\mathcal{C}$ *having success ratio* $\mathsf{SR}(\mathcal{C}) \approx \mathsf{SR}(\mathcal{B})/4$*, independent of the number of users* N *in the multi-user scenario.*

This theorem improves the bound of previous generic reductions [20] by a factor of N. Following our modular approach, the theorem is proved in two steps via Lemmas 3.8 and 3.9. It makes use of the RSR property, meaning that from a given public key pk we can derive properly distributed $pk_1, \ldots, pk_N$ 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 3.8 uses the RSR property to prove that UF-KOA tightly implies MU-UF-KOA.

Lemma 3.9 is our main technical contribution and proves MU-UF-KOA $\to$ 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 UF-KOA $\rightarrow$ UF-CMA in the single user setting. In the single user setting, the reduction simulates the signing queries 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 keys are known is hidden from the adversary. 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 signature schemes in the discrete logarithm setting is the Schnorr signature scheme [37]. It is obtained via the Fiat-Shamir transform applied to the Schnorr identification scheme. 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. 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 example parameters for a security 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 Katz-Wang identification scheme [28] is a double-generator version of Schnorr. It is at least as security as Schnorr which means one cannot hope for a tight security reduction to the DLOG assumption. However, we can use a simple argument from [28] for a tight security proof of its PIMP-KOA security under the Decision Diffie-Hellman Assumption. By our framework, this implies a tight proof of (strong) MU-UF-CMA-security.

Other Signatures. Other canonical identification schemes with the required properties includes the ones by Guillou-Quisquater [24] and Okamoto [33]. Similar to Katz-Wang, for the Guillou-Quisquater scheme, we can use an argument from [2] for a tight proof of PIMP-KOA security under the Phi-hiding assumption. Alternatively, we can give a proof with loss Q under the Factoring assumption. Our framework also shows that this loss is unavoidable. For Okamoto's scheme, we can provide the same bounds as for Schnorr.

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 [36] introduced the Forking Lemma as a tool to prove UF-CMA security of SIG[ID] from HVZK, SS and KR-KOA-security. Ohta and Okamoto [32] gave an alternative proof from IMP-KOA security and HVZK. Abdalla et al. [1] 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, [5] consider stronger security notions (e.g., IMP-AA and man-in-the middle security) for the Schnorr and GQ identification schemes. Abdalla et al. [2] show that lossy identification schemes tightly imply UF-CMA-secure signatures in the random oracle model.

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 [20] 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 [9] 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 [9] further expands on the insurmountable obstacle. Our Theorem 1.2 shows that there is such a tight security reduction for Schnorr signatures, reproving the GMLS theorem in the random oracle model.

Impossibility Results. In terms of impossibility results, Seurin [38], building on earlier work of [35, 21], 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 [18] even excludes a reduction from any non-interactive assumption.[2] Fukumitsu and Hasegawa [19], generalizing earlier work on Schnorr signatures [16, 35], 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 [9] 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 [10] 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 [12].) In [10] key-prefixing is advertized as "an inexpensive way to alleviate concerns that several public keys could be attacked simultaneously." Indeed, Bernstein [9] proves that single-user security of the original Schnorr signatures scheme tightly implies multi-user security of the key-prefixed variant of the scheme.

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 [26] 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 [26] and Struik [40] (among others) recommended against key prefixing for Schnorr. Shortly after, Bernstein [8] 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 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" [14]. As a consequence, key-prefixing is contained in the current draft for EdDSA [27]. In the light of our new results, we recommend to reconsider this decision.

2 Definitions

2.1 Preliminaries

For an integer p, define $[p] := \{1, \ldots, p\}$ and $\mathbb{Z}_p$ as the residual ring $\mathbb{Z}/p\mathbb{Z}$. If A is a set, then $a \xleftarrow{\$} A$ denotes picking a from A according to the uniform distribution. All our algorithms are probabilistic polynomial time unless states otherwise. If A is an algorithm, then $a \xleftarrow{\$} \mathsf{A}(b)$ 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 $a := (A)(b; \mathbf{t})$ meaning that the algorithm is executed on input b and randomness $\mathbf{t}$. Note that A's execution is now deterministic.

[2] The main result of the published paper [18] even excludes reduction from any *interactive* assumption (with special algebraic properties), but the proof turned out to be flawed.

2.2 Digital Signatures

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

Definition 2.1 (Digital Signature). *A digital signature scheme* SIG *is defined as a triple of probabilistic algorithms* $\mathsf{SIG} = (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Ver})$.

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

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

Definition 2.2 (Multi-user Security). *A signature scheme* SIG *is said to be* (t, ε, N, Q_s)-$\mathsf{MU\text{-}SUF\text{-}CMA}$ *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,*

$$\Pr\left[\begin{array}{l|l} \mathsf{Ver}(pk_{i^*}, m^*, \sigma^*) = 1 & \textit{For } i = 1, \ldots, N : (pk_i, sk_i) \xleftarrow{\$} \mathsf{Gen}(\mathsf{par}) \\ \wedge\ (i^*, m^*, \sigma^*) \notin \{(i_j, m_j, \sigma_j) \mid j \in [Q_s]\} & (i^*, m^*, \sigma^*) \xleftarrow{\$} \mathcal{A}^{\textsc{Sign}(\cdot,\cdot)}(pk_1, \ldots, pk_N) \end{array}\right] \leq \varepsilon,$$

where on the j*-th query* $(i_j, m_j) \in [N] \times \{0,1\}^*$ $(j \in [Q_s])$ *the signing oracle* $\textsc{Sign}$ *returns* $\sigma_j \xleftarrow{\$} \mathsf{Sign}(sk_{i_j}, m_j)$ *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 2.2 is called the correctness condition, the second condition is called the freshness condition. Definition 2.2 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) $\mathsf{MU\text{-}UF\text{-}CMA}$ security (multi-user unforgeablility against chosen message attack) we modify the freshness condition in the experiment to $(i^*, m^*) \notin \{(i_j, m_j,) \mid j \in [Q_s]\}$, 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)-$\mathsf{MU\text{-}UF\text{-}KOA}$ security (multi-user unforgeability against key only attack) as $(t, \varepsilon, N, 0)$-$\mathsf{MU\text{-}UF\text{-}CMA}$ security, i.e. $Q_s = 0$, the adversary is not allowed to make any signing query

Definition 2.3 (Single-user Security). *In the single-user setting, i.e.* $N = 1$ *users,* (t, ε, Q_s)-$\mathsf{SUF\text{-}CMA}$ *security (strong unforgeablility against chosen message attacks) is defined as* $(t, \varepsilon, 1, Q_s)$-$\mathsf{MU\text{-}SUF\text{-}CMA}$ *security. Similarly, standard (non-strong)* (t, ε, Q_s)-$\mathsf{UF\text{-}CMA}$ *security (unforgeablility against chosen message attack) is defined as* $(t, \varepsilon, 1, Q_s)$-$\mathsf{MU\text{-}UF\text{-}CMA}$ *security. Further,* (t, ε)-$\mathsf{UF\text{-}KOA}$ *security (unforgeablility against key-only attack) is defined as* $(t, \varepsilon, 1, 0)$-$\mathsf{MU\text{-}SUF\text{-}CMA}$ *security, i.e.,* $N = 1$ *users and* $Q_s = 0$ *signing queries.*

Security in the random oracle model. The security of signature scheme containing a hash function can be analyzed in the random oracle model [7]. 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 $(t, \varepsilon, N, Q_s, Q_h)$-$\mathsf{MU\text{-}SUF\text{-}CMA}$ 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 $Q_h \geq Q_s$.

2.3 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. Let par be common system parameters shared among all participants that we assume to be fixed.

Definition 2.4 (Canonical Identification Scheme). *A canonical identification scheme* ID *is defined as a tuple of* $\mathsf{ID} := (\mathsf{IGen}, \mathsf{P}, \mathsf{ChSet}, \mathsf{V})$*:*

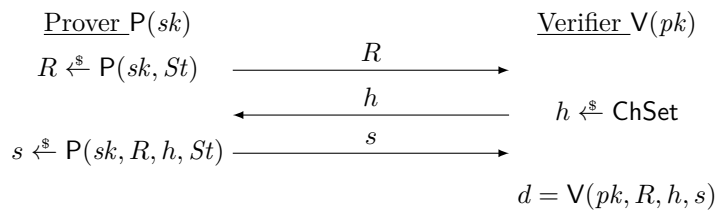

Figure 2: A canonical identification scheme and its transcript (R, h, s).

- *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* P *takes as input the secret key* sk*, a state* St *(which is initialized to the random coins* $\mathbf{t}$*), and the current conversation transcript and outputs the next message to be sent to the verifier.*
- *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 $(pk, sk) \in \mathsf{IGen}(\mathsf{par})$*, all* $R \in \mathsf{P}(sk, St)$*, all* $h \in \mathsf{ChSet}$ *and all* $s \in \mathsf{P}(sk, R, h, St)$*, we have* $\mathsf{V}(pk, R, h, s) = 1$*.*

An identification scheme ID is called *unique* if for all $(pk, sk) \in \mathsf{IGen}(\mathsf{par})$, $R \in \mathsf{P}(sk)$, $h \in \mathsf{ChSet}$, there exists at most one response s 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 *bis of min-entropy*, if for all $(pk, sk) \in \mathsf{IGen}(\mathsf{par})$, the commitment R 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 was honestly generated by the prover.

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

Definition 2.5 ((Parallel) Impersonation). *Let* $\mathsf{YYY} \in \{\mathsf{KOA}, \mathsf{PA}, \mathsf{AA}\}$*. A canonical identification* ID *is said to be* $(t, \varepsilon, Q_{\mathrm{CH}}, Q_{\mathrm{O}})$-$\mathsf{PIMP}$-$\mathsf{YYY}$ *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*,*

$$\Pr\left[\mathsf{V}(pk, R_{i^*}, h_{i^*}, s_{i^*}) = 1 \wedge\ i^* \in [Q_{\mathrm{CH}}]\} \;\middle|\; \begin{array}{l}(pk, sk) \xleftarrow{\$} \mathsf{IGen}(\mathsf{par}) \\ St \xleftarrow{\$} \mathcal{A}^{\mathrm{O}(\cdot)}(pk) \\ (i^*, s_{i^*}) \xleftarrow{\$} \mathcal{A}^{\mathrm{CH}(\cdot)}(pk)\end{array}\right] \leq \varepsilon,$$

where on the i*-th query* $\mathrm{CH}(R_i)$ *(*$i \in [Q_{\mathrm{CH}}]$*), the challenge oracle returns* $h_i \xleftarrow{\$} \mathsf{ChSet}$ *to* $\mathcal{A}$*.*[3] *Depending on* YYY*, oracle* O *is defined as follows.*

- *If* $\mathsf{YYY} = \mathsf{KOA}$ *(key-only attack), then* O *always returns* $\perp$*.*
- *If* $\mathsf{YYY} = \mathsf{AA}$ *(active attack), then* $\mathrm{O} := \textsc{Prover}$ *, where on the* j*-th empty query* $\textsc{Prover}(\epsilon)$ *(*$j \in Q_{\mathrm{O}}$*), the prover oracle returns* $R'_j \xleftarrow{\$} \mathsf{P}(sk, St')$ *to* $\mathcal{A}$*; on query* $\textsc{Prover}(j, h'_j)$*, the oracle returns* $s'_j \xleftarrow{\$} \mathsf{P}(sk, R'_j, h'_j, St')$*, if* h'_j *is already defined (and* $\perp$ *otherwise).*
- *If* $\mathsf{YYY} = \mathsf{AA}$ *(passive attack), then* $\mathrm{O} := \textsc{Prover}$*, where on the* j*-th empty query* $\textsc{Tran}(\epsilon)$ *(*$j \in Q_{\mathrm{O}}$*), the transcript oracle returns a transcript* (R'_j, h'_j, s'_j) *to* $\mathcal{A}$*, where* $R'_j \xleftarrow{\$} \mathsf{P}(sk, St')$*,* $h'_j \xleftarrow{\$} \mathsf{ChSet}$*;* $s'_j \xleftarrow{\$} \mathsf{P}(sk, R'_j, h'_j, St')$*.*

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

[3] On two queries $\mathrm{CH}(R_i)$ and $\mathrm{CH}(R_{i'})$ with the same input $R_i = R_{i'}$ the oracle returns two independent random challenges $h_i \xleftarrow{\$} \mathsf{ChSet}$ and $h_{i'} \xleftarrow{\$} \mathsf{ChSet}$.

Definition 2.6 (Key-recovery). *Let* $\mathsf{YYY} \in \{\mathsf{KOA}, \mathsf{PA}, \mathsf{AA}\}$. *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,

$$\Pr\left[(sk^*, pk) \in \mathsf{IGen}(\mathsf{par}) \;\middle|\; \begin{array}{l} (pk, sk) \xleftarrow{\$} \mathsf{IGen}(\mathsf{par}) \\ sk^* \xleftarrow{\$} \mathcal{A}^{\mathrm{O}(\cdot)}(pk) \end{array} \right] \leq \varepsilon,$$

where O *is defined as in Definition 2.5. The winning condition* $(sk^*, pk) \in \mathsf{IGen}(\mathsf{par})$ *means that the tuple* (sk^*, pk) *is in the support of* $\mathsf{IGen}(\mathsf{par})$*, i.e., that* sk^* *is a valid secret-key with respect to* pk.

Definition 2.7 (Special Soundness). *A canonical identification* ID *is said to be* SS *(special sound) if there there exists an extractor algorithm* Ext *such that, for all* $(pk, sk) \in \mathsf{IGen}(\mathsf{par})$*, given any two accepting transcripts* (R, h, s) *and* (R, h', s') *(where* $h \neq h'$*), we have* $\Pr[(sk^*, pk) \in \mathsf{IGen}(\mathsf{par}) \mid sk^* \xleftarrow{\$} \mathsf{Ext}(pk, R, h, s, h', s')] = 1$.

Definition 2.8 (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* $(pk, sk) \in \mathsf{IGen}(\mathsf{par})$*:*

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

Definition 2.9 (Honest-verifier Zero-knowledge). *A canonical identification* ID *is said to be* HVZK *(honest-verifier zero-knowledge) if there exists an algorithm* Sim *that, given public key* pk *and* $h \xleftarrow{\$} \mathsf{ChSet}$*, outputs* (R, s) *such that* (R, h, s) *is a real transcript with respect to* pk.

2.4 Signatures from Identification Schemes

Let $\mathsf{ID} := (\mathsf{IGen}, \mathsf{P}, \mathsf{ChSet}, \mathsf{V})$ be a canonical identification scheme. By the generalized Fiat-Shamir transformation [5], the signature scheme $\mathsf{SIG}[\mathsf{ID}] := (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Ver})$ from ID is defined as follows. par contains the system parameters of ID and a hash function $H : \{0,1\}^* \to \mathsf{ChSet}$.

$\underline{\mathsf{Gen}(\mathsf{par})}$:	$\underline{\mathsf{Sign}(sk, m)}$:	$\underline{\mathsf{Ver}(sk, m, \sigma)}$:
$(pk, sk) \xleftarrow{\$} \mathsf{IGen}(\mathsf{par})$	$R \xleftarrow{\$} \mathsf{P}(sk, St)$	Parse $\sigma = (R, s)$
Return (pk, sk)	$h = H(R, m)$	$h = H(R, m)$
	$s \xleftarrow{\$} \mathsf{P}(sk, R, h, St)$	Return $\mathsf{V}(pk, R, h, s)$
	Return $\sigma = (R, s)$	

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 reconstructive, if $\mathsf{V}(pk, R, h, s)$ first recomputes the commitment via $R' = \mathsf{V}'(pk, h, s)$ and then outputs 1 iff $R' = R$. For reconstructive ID, we can define an alternative Fiat-Shamir transformation $\mathsf{SIG}'[\mathsf{ID}] := (\mathsf{Gen}, \mathsf{Sign}', \mathsf{Ver}')$. Algorithm $\mathsf{Sign}'(sk, m)$ is defined as $\mathsf{Sign}(sk, m)$, but outputs $\sigma' = (h, s)$. Algorithm $\mathsf{Ver}'(pk, m, \sigma')$ first parses $\sigma' = (h, s)$, then recomputes $R' := \mathsf{V}'(pk, h, s)$, and finally returns 1 iff $H(R', m) = h$. Since $\sigma = (R, s)$ can publicly transformed into $\sigma' = (h, s)$ and vice-cersa, $\mathsf{SIG}[\mathsf{ID}]$ and $\mathsf{SIG}[\mathsf{ID}']$ are equivalent in terms of security. On the one hand, the alternative Fiat-Shamir transform yields shorter signatures if $h \in \mathsf{ChSet}$ has a smaller representation than response s. On the other hand, signatures of Fiat-Shamir transform maintain their algebraic structure which in some cases of ID enables 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*

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

The proof of Theorem 3.1 is obtained by combining Lemmas 3.4-3.9 and using $Q_h \leq t'-1$.

Theorem 3.2 (Main Theorem 2). *Suppose* SIG[ID] *is* HVZK, RSR *and has* α *bit min-entropy. If* SIG[ID] *is* (t,ε,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*

$$\varepsilon' \leq 4\varepsilon+\frac{Q_hQ_s}{2^\alpha}, \qquad t' \approx t$$

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 3.8 and 3.9.

3.1 Proof of the Main Theorems

Lemma 3.3 (XXX-KOA $\rightarrow$ XXX-PA). *Let* $\mathsf{XXX} \in \{\mathsf{KR},\mathsf{IMP},\mathsf{PIMP}\}$. *If* ID *is* $(t,\varepsilon,Q_{\mathrm{CH}})$-XXX-KOA *secure and* HVZK*, then* ID *is* $(\approx t,\varepsilon,Q_{\mathrm{CH}},Q_{\mathrm{O}})$-XXX-KOA *secure.*

Proof. Let $\mathcal{A}$ be an adversary against the $(t,\varepsilon,Q_{\mathrm{CH}},Q_{\mathrm{O}})$-XXX-KOA-security of ID. We now build an adversary $\mathcal{B}$ against the $(t,\varepsilon,Q_{\mathrm{O}})$-XXX-KR 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 Prover oracle of the passive attack PA in the first phase, all queries to the Ch oracle (for $\mathsf{YYY} \in \{\mathsf{IMP},\mathsf{PIMP}\}$) 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 Prover oracle can be perfectly simulated by picking $h \xleftarrow{\$} \mathsf{ChSet}$, computing $(R,s) \xleftarrow{\$} \mathsf{Sim}(pk,h)$ a simulated proof, and returning a transcript (R,h,s). The running time of $\mathcal{B}$ is that of $\mathcal{A}$ plus roughly Q_{O} executions of Sim to simulate the Prover oracle, which we ignore for simplicity. ■

Lemma 3.4 below can be viewed as a generalization of Bellare and Palacio's Reset Lemma [5] that takes advantage of random self-reducibility. The IMP-KOA security experiment is quite simple is relatively simple and the adversary makes exactly one query to the challenge oracle Ch. Therefore the reduction in the proof of the lemma does not have to guess which of the Q_{CH} many challenges the adversary is using to break security. This is the reason why its proof is considerably simpler than the corresponding previous proofs analyzing the security of identification/signature schemes using rewinding, for example the Forking Lemma [36, 4] or the proofs in [38, 32, 30].

Lemma 3.4 (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* $N>0$*,*

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

In particular, the two success ratios are related as

$$\frac{\varepsilon'}{t'} \leq 6\cdot\frac{\varepsilon}{t}+\frac{1}{t'|\mathsf{ChSet}|}. \tag{2}$$

We remark that without the RSR property we can still prove the theorem for $N=1$, i.e., $\varepsilon \geq \varepsilon'(\varepsilon'-\frac{1}{|\mathsf{ChSet}|})$, $t \approx 2t'$.

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

To prove (1), 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 (1).

Construction of $\mathcal{B}$. In phase 1, for each $i \in [N]$, $\mathcal{B}$ does the following. it picks random tape $\mathbf{t}_i$, runs $(pk_i, \mathbf{a}_i) \xleftarrow{\$} \mathsf{Rerand}(pk)$ and executes $\mathcal{A}(pk_i; \mathbf{t}_i)$. On query R_i, $\mathcal{B}$ answers with $h_i \xleftarrow{\$} \mathsf{ChSet}$ to obtain s_i from $\mathcal{A}$. If any of $\mathcal{A}$'s executions produces a valid transcript, i.e., if there exists an index $i^* \in [N]$ such that transcript $(R_{i^*}, h_{i^*}, s_{i^*})$ is a valid transcript with respect to pk_{i^*}, then $\mathcal{B}$ continues its execution. Otherwise, it aborts.

In phase 2, $\mathcal{B}$ fixes i^* and, for each $j \in [N]$, it does the following. It executes $\mathcal{A}(pk_{i^*}; \mathbf{t}_{i^*})$. Adversary $\mathcal{A}$ will always query R_{i^*}, which $\mathcal{B}$ answers with $h'_j \xleftarrow{\$} \mathsf{ChSet} \setminus \{h_{i^*}\}$ to obtain s'_j from $\mathcal{A}$. If any of $\mathcal{A}$'s executions produces a valid transcript, i.e., if there exists an index $j^* \in [N]$ such that transcript $(R_{i^*}, h'_{j^*}, s'_{j^*})$ is a valid transcript with respect to pk_{i^*}, then $\mathcal{B}$ continues its execution. Otherwise, it aborts.

Finally, $\mathcal{B}$ uses the SS property of ID and computes $sk_{i^*} \leftarrow \mathsf{Ext}(pk_{i^*}, R_{i^*}, h_{i^*}, s_{i^*}, h'_{j^*}, s'_{j^*})$. By the RSR property of ID, it returns $sk \leftarrow \mathsf{Derand}(pk_{i^*}, sk_{i^*}, \mathbf{a}_{i^*})$ and terminates.

Success Probability of $\mathcal{B}$. For each $i \in [N]$, pk_i is a properly distributed public-key and

$$\Pr[\mathsf{V}(pk_i, R_i, h_i, s_i) = 1] = \varepsilon'.$$

Therefore,

$$\Pr[\text{no abort in phase 1}] = 1 - (1 - \varepsilon')^N. \tag{3}$$

Next, for each $i, j \in [N]$ and fixed pk_i and $\mathbf{t}_i$, we define

$$\begin{aligned} q_{i,j} = q_{i,j}(pk_i, \mathbf{t}_i) &:= \Pr_{h_i}[\mathsf{V}(pk_i, R_i, h_i, s_i) = 1], \\ p_{i,j} = p_{i,j}(pk_i, \mathbf{t}_i) &:= \Pr_{h_i, h'_j}[\mathsf{V}(pk_i, R_i, h_i, s_i) = 1 \wedge \mathsf{V}(pk_i, R_i, h'_j, s'_j) = 1]. \end{aligned}$$

Note that the value R_i deterministically depends on pk_i and $\mathbf{t}_i$, the value s_i deterministically depends on $pk_i, \mathbf{t}_i$, and h_i, and the value s'_j deterministically depends on $pk_i, \mathbf{t}_i$, and h'_j. Therefore, the probability-space of $p_{i,j}$ is the collection of random variables (h_i, h'_j), where $h_i \xleftarrow{\$} \mathsf{ChSet}$ and $h'_j \xleftarrow{\$} \mathsf{ChSet} \setminus \{h_i\}$. Since $\Pr[h_i = h'_j] = \frac{1}{|\mathsf{ChSet}|}$, we have

$$p_{i,j} \geq q_{i,j} \cdot \left(q_{i,j} - \frac{1}{|\mathsf{ChSet}|} \right).$$

We bound the expectation of $p_{i,j}$ as

$$\begin{aligned} \mathsf{E}_{\mathbf{t}_i, pk_i}[p_{i,j}] &\geq \mathsf{E}_{\mathbf{t}_i, pk_i}\left[q_{i,j} \cdot \left(q_{i,j} - \frac{1}{|\mathsf{ChSet}|} \right) \right] \\ &\geq \mathsf{E}_{\mathbf{t}_i, pk_i}[q_{i,j}] \cdot \left(\mathsf{E}_{\mathbf{t}_i, pk_i}[q_{i,j}] - \frac{1}{|\mathsf{ChSet}|} \right) \\ &= \varepsilon' \left(\varepsilon' - \frac{1}{|\mathsf{ChSet}|} \right). \end{aligned}$$

In the last inequation we used Jensen's inequality[4] applied to the convex function $\varphi(q_{i,j}) := q_{i,j}(q_{i,j} - 1/|\mathsf{ChSet}|)$ for the constant $|\mathsf{ChSet}|$. Finally, we bound the probability $\varepsilon_{i,j}$ that transcript (R_i, h'_j, s'_j) is valid with respect to pk_i conditioned on the event that transcript (R_i, h_i, s_i) is valid with respect to pk_i.

$$\begin{aligned} \varepsilon_{i,j} &= \Pr[\mathsf{V}(pk_i, R_i, h'_j, s'_j) = 1 \mid \mathsf{V}(pk_i, R_i, h_i, s_i) = 1] \\ &= \frac{\Pr[\mathsf{V}(pk_i, R_i, h'_j, s'_j) = 1 \wedge \mathsf{V}(pk_{i^*}, R_i, h_i, s_i) = 1]}{\Pr[\mathsf{V}(pk_i, R_i, h_i, s_i) = 1]} \\ &= \frac{\mathsf{E}_{\mathbf{t}_i, pk_i}[p_{i,j}]}{\varepsilon'} \geq \varepsilon' - \frac{1}{|\mathsf{ChSet}|} \end{aligned}$$

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

Using $\varepsilon_{i^*,j} \geq \varepsilon' - 1/|\mathsf{ChSet}|$, we obtain

$$\Pr[\text{no abort in phase 2} \mid \text{no abort in phase 1}] = 1 - (1 - \varepsilon' + \frac{1}{|\mathsf{ChSet}|})^N. \tag{4}$$

As $\mathcal{B}$ is successful if it does not abort, by (3), (4) we obtain

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

The running time t of $\mathcal{B}$ is $2Nt'$ 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 3.5 (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*

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

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 $\text{CH}(R_i)$, it proceeds as follows. If $i \neq i^*$, then it return $h_i \xleftarrow{\$} \mathsf{ChSet}$. If $i = i^*$, then it defines $R := R_{i^*}$, makes a query $h \xleftarrow{\$} \text{CH}(R)$ to its own oracle, and returns $h_{i^*} := h$ to $\mathcal{A}$. Finally, $\mathcal{A}$ submits (i,s). If $i \neq i^*$, then $\mathcal{B}$ aborts. Otherwise, it outputs s to its 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 3.6 (PIMP-KOA $\xrightarrow{\text{PRO}}$ UF-KOA)**.** *If* ID *is* $(t,\varepsilon,Q_{\text{CH}})$-PIMP-KOA *secure, then* SIG[ID] *is* (t',ε',Q_h)-UF-KOA *secure in the programmable random oracle model, where*

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

Proof. Let $\mathcal{A}$ be an adversary against the (t',ε',Q_h)-UF-KOA-security of 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 and which it forwards to $\mathcal{A}$. If $\mathcal{A}$ makes a query (R_i,m_i) to the random oracle, $\mathcal{B}$ makes a query $h_i \xleftarrow{\$} \text{CH}(R_i)$ and programs the random oracle $H(R_i,m_i) := h_i$. Finally, $\mathcal{A}$ submits a forgery (m,σ), where $\sigma = (R,s)$. We assume that $(R,m) \in \{(R_i,m_i)\}$, i.e., $H(R,m)$ was queries 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 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 3.9 (with a slightly improved bound).

Lemma 3.7 (UF-KOA $\xrightarrow{\text{PRO}}$ UF-CMA)**.** *Suppose* ID *is* HVZK *and has* α *bit min-entropy. If* SIG[ID] *is* (t,ε,Q_h)-UF-KOA *secure, then* SIG[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 3.8 (UF-KOA $\xrightarrow{\text{RSR}}$ MU-UF-KOA)**.** *Suppose* ID *is* RSR*. If* SIG[ID] *is* (t,ε)-UF-KOA *secure, then* SIG[ID] *is* (t',ε',N)-MU-UF-KOA *secure, where*

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

Again, without the RSR property one can use the generic bounds from [20] 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 theorem. 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 $(pk_i, \mathbf{a}_i) \xleftarrow{\$} \mathsf{Rerand}(pk)$ by using the RSR property of ID. Then $\mathcal{B}$ runs $\mathcal{A}$ on input $(pk_1, \ldots, pk_N)$.

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 $s \xleftarrow{\$} \mathsf{Tran}(pk, pk_{i^*}, \mathbf{a}_{i^*}, (R^*, h^*, s^*))$. 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 3.9 (MU-UF-KOA $\xrightarrow{\textbf{PRO}}$ 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 theorem. Adversary $\mathcal{B}$ is executed in the MU-UF-KOA experiment and obtains public-keys $(pk_1, \ldots, pk_N)$, and has access to a random oracle H.

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

Adversary $\mathcal{B}$ runs $\mathcal{A}$ on input $(pk'_1, \ldots, pk'_N)$ 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 $\sigma_j := (R_j, s_j) \xleftarrow{\$} \mathsf{Sign}(sk'_{i_j}, m_j)$. Note that this involves $\mathcal{B}$ making a query $H'(R_j, m_j)$.
- 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, $\mathcal{B}$ samples $h_j \xleftarrow{\$} \mathsf{ChSet}$ and runs $(R_j, s_j) \xleftarrow{\$} \mathsf{Sim}(pk'_{i_j}, h_j)$. 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

$$H'(R_j, m_j) := h_j \tag{5}$$

  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 (5), $\mathcal{B}$ 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 $h^* = H'(m^*, R^*)$ we have $\mathsf{V}(pk'_{i^*}, R^*, h^*, s^*) = 1$. Furthermore, it satisfies the freshness condition, i.e.,

$$(i^*, m^*) \notin \{(i_j, m_j) : j \in [Q_s]\}. \tag{6}$$

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 and $h^* = h_j$ and furthermore $i^* \neq i_j$ by the freshness condition (6).
  - 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 (5). 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 (5). 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 are key-preserving.

Lemma 4.1 (KR-KOA $\xrightarrow{\text{non-rewind.}}\!\!\!\!\!\!\not\qquad$ IMP-KOA)**.** *If there is a public-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 $(t_\mathcal{M}, \varepsilon_\mathcal{M}, Q_\mathrm{O})$-breaks* IMP-AA *security of* ID*, where*

$$\varepsilon_\mathcal{M} \geq \varepsilon_\mathcal{R} - \frac{1}{|\mathsf{ChSet}|},\ t_\mathcal{M} \approx t_\mathcal{R}, Q_\mathrm{O} = 1.$$

Proof. Assuming the existence of a public-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 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 $\mathcal{A}(pk)$ as follows. First, $\mathcal{M}$ queries $R \xleftarrow{\$}$ Prover() from the IMP-AA challenger $\mathcal{C}_{\mathsf{IMP\text{-}AA}}$ and returns R to $\mathcal{R}$. Upon

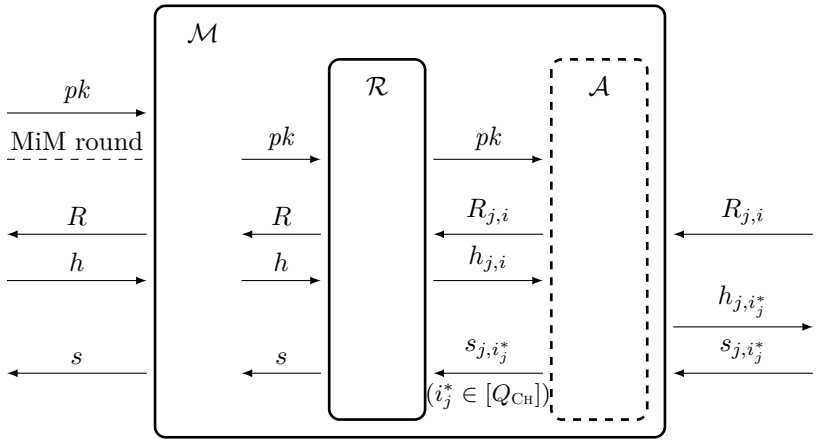

Figure 3: Meta-reduction $\mathcal{M}$ uses $\mathcal{R}$ to break the IMP-MIM security in n MIM rounds. n is the total amount of executions of $\mathcal{A}$ performed by $\mathcal{R}$. For every MIM round $j \in [n]$, $\mathcal{A}$ picks an $i_j^* \in [Q_{\text{CH}}]$ and forwards a valid response s_{j,i_j^*}. $\mathcal{M}$ 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^*}.

receiving h from $\mathcal{R}$, $\mathcal{M}$ queries $s \xleftarrow{\$} \text{PROVER}(1, h)$ from the IMP-AA challenger $\mathcal{C}_{\mathsf{IMP\text{-}AA}}$ and returns s to $\mathcal{R}$ with probability $\varepsilon_{\mathcal{A}}$.

After receiving sk from $\mathcal{R}$, $\mathcal{M}$ follows the protocol honestly by using sk and breaks IMP-AA security of ID: $\mathcal{M}$ computes $(R^*, St) \xleftarrow{\$} \mathsf{P}(sk, St)$ and returns R^* to the IMP-AA security challenger $\mathcal{C}_{\mathsf{IMP\text{-}AA}}$; upon receiving h^* from $\mathcal{C}_{\mathsf{IMP\text{-}AA}}$, $\mathcal{M}$ computes $(s^*, St) \xleftarrow{\$} \mathsf{P}(sk, R^*, h^*, St)$.

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

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*,*

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

where oracles PROVER *and* CH *are defined as in Definition 2.5. We define weak impersonation against man-in-the-middle attack (*wIMP-MIM*) by restricting* $R_{i^*} \in \{R'_1, \ldots, R'_{Q_{\text{O}}}\}$.

The following generalizes a result by Seurin [38] to canonical identification schemes.

Lemma 4.3 (IMP-KOA $\xrightarrow{\textbf{loss} <\textbf{Q}}$ PIMP-KOA)**.** *Suppose that* ID *has* α *bit min-entropy and there is a public-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*

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

We note that the Schnorr identification scheme is wIMP-MIM but not IMP-MIM-secure (cf. Section 5.1).

Proof. Assuming the existence of a public-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 $Q_{\text{O}} = nQ_{\text{Ch}}$ 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 public-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 of 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 $R_{j,i} \xleftarrow{\$}$ Prover() and forwards it to $\mathcal{R}$, which will reply with an arbitrary $h_{j,i} \in \mathsf{ChSet}$. 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 $s_{j,i_j^*} \xleftarrow{\$}$ Prover$((j-1)\cdot Q_{\text{Ch}} + i_{j^*}, h_{j,i_j^*})$. 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 $(i_j^*, s_{j,i_j^*}) := (\bot, \bot)$ 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 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 $s_{j,i} \neq \bot$ only if $i = i_j^*$ and therefore to consider the case when $i = i_j^*$.

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

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

$$\begin{aligned} \Pr[\exists j \in [n] : (R, h) = (R_{j,i_j^*}, h_{j,i_j^*})] \leq\ & \Pr[\exists (j, i_j^*) \neq (j_0, i_0) : (R, h) = (R_{j,i_j^*}, h_{j,i_j^*})] && (7) \\ & + \Pr[(j, i_j^*) = (j_0, i_0)]. && (8) \end{aligned}$$

We bound the probabilities (7) and (8) individually. To bound (8), 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 (7), 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 $h \xleftarrow{\$} \mathsf{ChSet}$. Furthermore, R is fixed for $(j, i_j^*) > (j_0, i_0)$ while $R_{j,i_j^*} \xleftarrow{\$}$ Prover(). Therefore by

using a union bound

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

Overall, this yields

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

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

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

Chioce 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

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

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 $i \in [Q_{\mathrm{CH}}]$ we have $c_{j,i} = 0$, $\mathcal{A}$ will always send a valid transcript and break the PIMP-KOA security. Let $\overline{\mu} := (1-\mu)$. For any execution $j \in [n]$, $\mathcal{A}$ has success probability

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

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

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

where the bound $\mu \leq \ln((1-\varepsilon_{\mathcal{A}})^{-1})/Q_{\mathrm{CH}}$ was proved in [38, Lemma 1]. Therefore we have

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

which concludes the proof of the lemma. ■

For a precise analysis of the function $\ln\left((1-\varepsilon_{\mathcal{A}})^{-1}\right)$, we refer to [38]. For our purpose, it is sufficient that for a concrete choice of $\varepsilon_{\mathcal{A}}$, there is a constant c such that $c \cdot \varepsilon_{\mathcal{A}} = \ln\left((1-\varepsilon_{\mathcal{A}})^{-1}\right)$. Hence Lemma 4.3 gives roughly $\varepsilon_{\mathcal{M}} \geq \varepsilon_{\mathcal{R}} - c \cdot n/Q_{\mathrm{CH}} \cdot \varepsilon_{\mathcal{A}}$ for a suitable choice of $\varepsilon_{\mathcal{A}}$. Therefore $\varepsilon_{\mathcal{R}}$ can be at most $c \cdot n/Q_{\mathrm{CH}} \cdot \varepsilon_{\mathcal{A}}$. 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 send by $\mathcal{R}$. So if $\mathcal{R}$ is furthermore randomness-preserving, i.e., it chooses $R \in \{R_{1,1}, \ldots, R_{n,Q_{\mathrm{CH}}}\}$, then $\mathcal{M}$ attacks wIMP-MIM-security of ID. This observation is formalized in the following corollary.

Corollary 4.4. *If* ID *has* α *bit min-entropy and there exists a public 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* $(t_{\mathcal{A}}, \varepsilon_{\mathcal{A}}, Q_{\mathrm{CH}})$*-breaks* PIMP-KOA *security of* ID*, then there exists an algorithm* $\mathcal{M}$ *that* $(t_{\mathcal{M}}, \varepsilon_{\mathcal{M}}, 1, Q_{\mathrm{O}} = nQ_{\mathrm{CH}})$*-breaks* wIMP-MIM *security of* ID*, where*

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

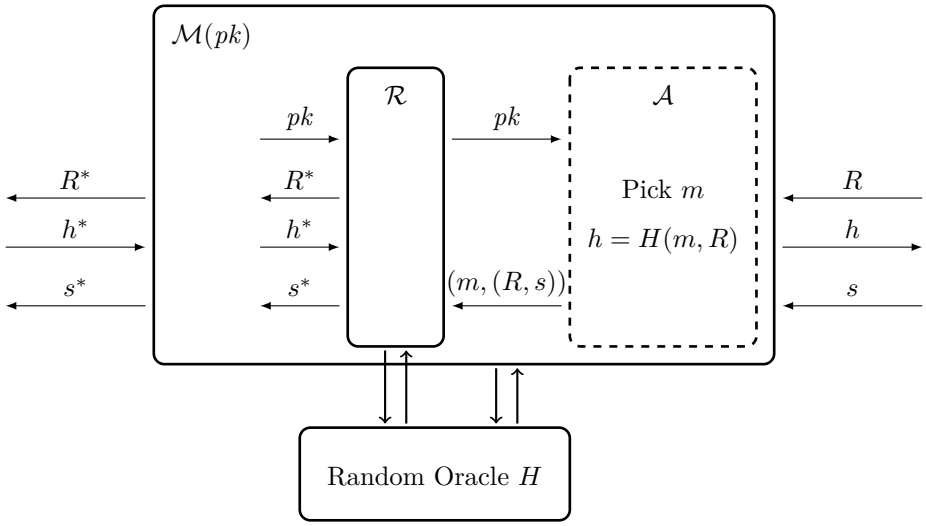

Figure 4: Meta-reduction $\mathcal{M}$ runs $\mathcal{R}$ to break IMP-AA security in the non-programmable random oracle model, where both $\mathcal{M}$ and $\mathcal{R}$ have oracle access to the same external random oracle H. $\mathcal{M}$ simulates an adversary $\mathcal{A}$ that $(t_{\mathcal{A}}, \varepsilon_{\mathcal{A}}, Q_h)$-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 $\xrightarrow{\text{NPRO}\not}$ UF-KOA)**.** *If there exists a public 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 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*

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

Proof. Assuming the existence of a public 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 $\mathcal{M}$. $\mathcal{M}$ obtains the public key pk of the IMP-AA challenge 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 $R \xleftarrow{\$}$ Prover() and returns R to $\mathcal{R}$. Next, $\mathcal{M}$ picks an arbitrary message m, queries $h = H(m, R)$ and $s \xleftarrow{\$}$ Prover$(1, h)$. With probability $\varepsilon_{\mathcal{A}}$ $\mathcal{M}$ returns $(m, (R, s))$ as a forgery to $\mathcal{R}$.

Upon receiving a Ch(R^*) query from $\mathcal{R}$, $\mathcal{M}$ breaks IMP-AA security as follows: $\mathcal{M}$ forwards R^* to the challenger $\mathcal{C}_{\mathsf{IMP\text{-}AA}}$; after receiving h^* from $\mathcal{C}_{\mathsf{IMP\text{-}AA}}$, $\mathcal{M}$ returns h^* as the answer of the Ch(R^*) query; when $\mathcal{R}$ responds with s^* to break IMP-KOA security, $\mathcal{M}$ forwards s^* to $\mathcal{C}_{\mathsf{IMP\text{-}AA}}$. We note that $h^* = h$ with probability $1/|\mathsf{ChSet}|$, since h^* is a random challenge from $\mathcal{C}_{\mathsf{IMP\text{-}AA}}$ and h is a respond of a random oracle query. Thus, if s^* breaks IMP-KOA security, then s^* breaks IMP-AA security, since $(R^*, h^*, s^*) \neq (R, h, s)$. Moreover, $\mathcal{M}$ perfectly simulates an adversary that $(t_{\mathcal{A}}, \varepsilon_{\mathcal{A}})$-breaks UF-KOA security. Thus, we have $\varepsilon_{\mathcal{M}} \geq \varepsilon_{\mathcal{R}} - 1/|\mathsf{ChSet}|$. ■

By Lemma 3.5, Lemma 4.5 implies that there is no reduction from PIMP-KOA to UF-KOA in the non-programmable random oracle model.

Lemma 4.6 (UF-KOA $\xrightarrow{\text{NPRO}\not}$ UF-CMA)**.** *Suppose that there is a public-key preserving reduction $\mathcal{R}$ in the non-programmable random oracle (NPRO) model that $(t_{\mathcal{R}}, \varepsilon_{\mathcal{R}}, Q_s, Q_h)$-breaks* UF-KOA *security of* SIG[ID] *with black-box access to an adversary $\mathcal{A}$ that $(\varepsilon_{\mathcal{A}}, t_{\mathcal{A}}, Q_h)$-breaks* UF-CMA *security of* SIG[ID]*. Then there exists an algorithm $\mathcal{M}$ that $(\varepsilon_{\mathcal{M}}, t_{\mathcal{M}})$-breaks* UF-KOA *security of* SIG[ID]*, where*

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

Proof. Assuming the existence of a public key preserving reduction $\mathcal{R}$ as above, we construct a meta-reduction $\mathcal{M}$ to break UF-KOA security of SIG[ID]. $\mathcal{M}$ gets the public key pk from UF-KOA challenger $\mathcal{C}_{\mathsf{UF\text{-}KOA}}$ 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 $\mathcal{R}$, $\mathcal{M}$ 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 $\mathcal{R}$ access to the same random oracle, (m, σ) is a valid forgery respond to $\mathcal{C}_{\mathsf{UF\text{-}KOA}}$. 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 downside of our results. In case of the Schnorr identification/signature scheme we can extend our techniques to exclude the larger class of algebraic reductions. A reduction is algebraic over some multiplicative group $\mathbb{G}$ or prime-order p, if for all group elements h output by the reduction, their respective representation is known. That is, if the reduction holds group elements $g_1, \ldots, g_n \in \mathbb{G}$ and outputs a new group element h, then it also provides $(\alpha_1, \ldots, \alpha_n) \in \mathbb{Z}_p^n$ satisfying $h = \prod g_i^{\alpha_i}$.

5 Example Instantiations

In this section we consider two examples of important identification schemes, namely the ones by Schnorr [37] and by Katz-Wang [28]. We use our framework to derive tight security bounds and concrete parameters for the corresponding Schnorr/Katz-Wang 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 examples 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) [5] and weakly secure against man-in-the-middle attack (wIMP-MIM) (Lemma 5.5).

Let $\mathsf{par} := (p, g, \mathbb{G})$ 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 $\mathsf{ID}_\mathsf{S} := (\mathsf{IGen}, \mathsf{P}, \mathsf{ChSet}, \mathsf{V})$ is defined as follows.

$\underline{\mathsf{IGen}(\mathsf{par})}$:	$\underline{\mathsf{P}(sk, St)}$:
$sk := x \xleftarrow{\$} \mathbb{Z}_p$	$r \xleftarrow{\$} \mathbb{Z}_p$; $R = g^r$
$pk := X = g^x$	$St := St \cup \{(R, r)\}$
$\mathsf{ChSet} := \{0,1\}^n$; $St := \emptyset$	Return (R, St)
Return (pk, sk)	
	$\underline{\mathsf{P}(sk, R, h, St)}$:
$\underline{\mathsf{V}(pk, R, h, s)}$:	If $(R, \cdot) \notin St$ then return $\perp$
If $R = g^s \cdot X^{-h}$ then return 1	Let $(R, r) \in St$
Else return 0.	Return $s = x \cdot h + r \bmod p$

We recall the DLOG and OMDL assumptions.

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

$$\Pr\left[\ g^x = X \ \middle|\ X \xleftarrow{\$} \mathbb{G}; x \xleftarrow{\$} \mathcal{A}(X)\ \right] \leq \varepsilon.$$

Lemma 5.2. *Let* $\mathsf{ID}_\mathsf{S} := (\mathsf{IGen}, \mathsf{P}, \mathsf{ChSet}, \mathsf{V})$ *be the Schnorr identification scheme as defined above.* 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* $\mathsf{par} = (p, g, \mathbb{G})$ *then* ID_S *is* (t, ε)-KR-KOA *secure.*

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

Uniqueness. For all $(X, x) \in \mathsf{IGen}(\mathsf{par})$, $R := g^r \in \mathsf{P}(sk, St)$ and $h \in \{0,1\}^n$, the value $s \in Z_p$ 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 $\mathsf{Ext}(X, R, h, s, h', s') := x^* := (s - s')/(h - h')$ such that, for all $(X := g^x, x) \in \mathsf{IGen}(\mathsf{par})$, 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 and $h \in \{0,1\}^n$, we let $\mathsf{Sim}(X, h)$ first sample $s \xleftarrow{\$} \mathbb{Z}_p$ and then output $(R := g^s X^{-h}, s)$. 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:

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

Definition 5.3 (One-more Discrete Logarithm Assumption [3]). *We says that* OMDL *is* (t, ε, Q)*-hard in* $\mathsf{par} = (p, g, \mathbb{G})$ *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*,*

$$\Pr\left[\textit{For } i \in [Q+1] : X_i = g^{x_i} \;\middle|\; \begin{array}{l} X_1, \ldots, X_{Q+1} \xleftarrow{\$} \mathbb{G} \\ (x_1, \ldots, x_{Q+1}) \xleftarrow{\$} \mathcal{A}^{\mathrm{Dl}(\cdot)}(X_1, \ldots, X_{Q+1}) \end{array} \right] \leq \varepsilon,$$

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 [5]). *If the* OMDL *problem is* (t, ε, Q)*-hard then* ID_S *is* $(t', \varepsilon', Q_\mathrm{O})$-IMP-AA *secure, where* $\varepsilon' \leq \sqrt{\varepsilon} + 1/p$*,* $t \approx 2t'$*, and* $Q_\mathrm{O} = Q$*.*

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* $(t', \varepsilon', Q_\mathrm{Ch}, Q_\mathrm{O})$-wIMP-MIM *secure, where*

$$\varepsilon' = \varepsilon, \quad t' \approx t, \quad Q_\mathrm{O} = Q.$$

Proof. Let $\mathcal{A}$ be an algorithm that breaks $(t', \varepsilon', Q_\mathrm{Ch}, Q_\mathrm{O})$-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 $X_1, \ldots, X_{Q+1}$, 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 Prover() query ($j \in [Q_\mathrm{O}]$) $\mathcal{B}$ returns $R'_j := X_j$.
- On the j-th Prover(i, h'_j) query, $\mathcal{B}$ queries and returns $s'_j = \mathrm{Dl}(X^{h'_j} \cdot R'_j)$.
- On the i-th Ch(R_i) query, $\mathcal{B}$ chooses a random $h_i \xleftarrow{\$} \mathsf{ChSet}$ and returns h_i. Note that $R_i = R'_j$ for some R'_j previously returned by the Prover() oracle.

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

$\mathcal{A}$ wins if $(R_{i^*}, h_{i^*}, s_{i^*})$ is a valid transcript, $(R_{i^*}, h_{i^*}, s_{i^*}) \notin \{(R'_j, h'_j, s'_j) \mid j \in [Q_O]\}$, 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 $(R'_{j^*} = R_{i^*}, h'_{j^*}, s'_{j^*})$ 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 $(x_1, \ldots, x_{Q+1})$, breaks OMDL problem with $\varepsilon = \varepsilon'$ and $t \approx t'$. ■

We now define the interactive discrete-logarithm problem which models $\mathsf{PIMP\text{-}KOA}$-security for $\mathsf{ID_S}$.

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

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

where on the i-th query $\textsc{Ch}(g^{r_i})$ $(i \in [Q])$, the challenge oracle returns $h_i \xleftarrow{\$} \mathbb{Z}_p$ 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 $H : \{0,1\}^* \to \{0,1\}^n$ be a hash function with $n < \log_2(p)$. As $\mathsf{ID_S}$ is reconstructible we can use the alternative Fiat-Shamir transformation to obtain the Schnorr signature scheme $\mathsf{Schnorr} := (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Ver})$.

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

By Theorem 3.1 and the results from this section, we obtain concrete bounds for Schnorr's single-user and multi-user security.

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

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

The DLOG problem is tightly equivalent to the 1-IDLOG problem by Lemma 3.4. Assuming the OMDL problem is hard, $\mathsf{Schnorr}$ is $\mathsf{wIMP\text{-}MIM}$-secure and by Corollary 4.4 there cannot exist a tight implication 1-$\mathsf{IDLOG} \to Q$-IDLOG. Furthermore, by Theorem 3.2, the Q-IDLOG problem is tightly equivalent to $\mathsf{MU\text{-}SUF\text{-}CMA}$-security of $\mathsf{Schnorr}$.

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

$$\varepsilon' \leq 4\varepsilon + \frac{Q_h Q_s}{p}, \qquad t' \approx t.$$

We leave it an open problem to come up with a more natural hard problem over par that tightly implies Q-IDLOG (and hence $\mathsf{MU\text{-}SUF\text{-}CMA}$-security of $\mathsf{Schnorr}$). Note that according to [18], the hard problem has to have at least one round of interaction.

5.1.3 Concrete parameters

In this section we derive parameters for Schnorr providing k-bit security in the multi-user setting. Following [6], for k-bit security one requires $(\varepsilon', t', N, Q_s, Q_h)$-MU-SUF-CMA security with $\frac{\varepsilon'}{t'} \leq 2^{-k}$.

The following lemma assumes that a generic algorithm (for example, the Pollard-rho algorithm) is the best possible algorithm to break discrete logarithms in group $\mathbb{G}$. This is generally believed to be true for prime-order subgroups of elliptic curves.

Lemma 5.9. *Let* Schnorr *be instantiated with* $\mathsf{par} = (p, g, \mathbb{G}, H)$*, where p is a prime and $H : \{0,1\}^* \to \{0,1\}^n$. If a generic algorithm is the best possible algorithm to break discrete logarithms in group $\mathbb{G}$, then* Schnorr *provides k-bits security in the multi-user setting if*

$$\log p \geq 2k + \log(Q_h) + c'_{\mathsf{dl}}, \quad n \geq k+1,$$

*where c'_{dl} is a constant that only depends on the generic algorithm. Furthermore, if a generic algorithm is the best possible algorithm to break the Q-*IDLOG *problem in* par*, then* Schnorr *provides k-bits security in the multi-user setting if*

$$\log p \geq 2k + c''_{\mathsf{dl}},$$

where c''_{dl} is a constant that only depends on the generic algorithm.

Proof. Assuming a generic algorithm is the best possible algorithm to compute discrete logarithms, means that DLOG in group $\mathbb{G}$ of prime-order p is $(\varepsilon = c_{\mathsf{dl}} \cdot t^2/p, t)$-hard, for any time bound t, where c_{dl} is a fixed constant that only depends on the specific choice of the generic algorithm.

We assume that the adversary makes $Q_h > 3$ hash queries. Define the constant $c'_{\mathsf{dl}} := 6 + \log(c_{\mathsf{dl}})$. Plugging in the parameters from Lemma 5.7 and using $Q_s \leq t \leq 2^k$ we obtain

$$\begin{aligned} \frac{\varepsilon'}{t'} &\leq 24(Q_h+1)\frac{\varepsilon}{t} + \frac{Q_s}{p} + \frac{1}{2^n} \\ &\leq 32 Q_h c_{\mathsf{dl}} \frac{t}{p} + \frac{1}{2^n} \\ &\leq \frac{t}{2^{2k+1}} + \frac{1}{2^{k+1}} \leq 2^{-k} \end{aligned}$$

which proves the first part of the statement.

A similar computation can be done to prove the second part using Theorem A.1 saying that the best generic algorithm against Q-IDLOG has a success ratio of at most $\frac{2t^2}{p}$ ■

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 provable security guarantees for Schnorr.

5.2 Katz-Wang Identification/Signature Scheme

5.2.1 Katz-Wang Identification Scheme

Let $\mathsf{par} := (p, g_1, g_2, \mathbb{G})$ 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 Katz-Wang identification scheme $\mathsf{ID}_{\mathsf{KW}} := (\mathsf{IGen}, \mathsf{P}, \mathsf{ChSet}, \mathsf{V})$ is defined as follows.

$\underline{\mathsf{IGen}(\mathsf{par})}$:	$\underline{\mathsf{P}(sk, St)}$:
$sk := x \xleftarrow{\$} \mathbb{Z}_p$	$r \xleftarrow{\$} \mathbb{Z}_p$; $R = (R_1, R_2) = (g_1^r, g_2^r)$
$pk := (X_1, X_2) = (g_1^x, g_2^x)$	$St := St \cup \{(R, r)\}$
$\mathsf{ChSet} := \{0,1\}^n$; $St := \emptyset$	Return (R, St)
Return (pk, sk)	
	$\underline{\mathsf{P}(sk, R, h, St)}$:
$\underline{\mathsf{V}(pk, R = (R_1, R_2), h, s)}$:	If $(R, \cdot) \notin St$ then return $\perp$
If $R_1 = g^s \cdot X_1^{-h}$ and $R_2 = g^s \cdot X_2^{-h}$ then return 1	Let $(R, r) \in St$
Else return 0.	Return $s = x \cdot h + r \bmod p$

We recall the DDH assumption.

Definition 5.10 (Decision Diffie-Hellman Assumption). *The Decision Diffie-Hellman problem* DDH *is* (t,ε)*-hard in* $\mathsf{par} = (p, g_1, g_2, \mathbb{G})$ *if for all adversaries* $\mathcal{A}$ *running in time at most* t*,*

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

Clearly, all security results of Schnorr carry over to the Katz-Wang identification scheme, i.e., $\mathsf{ID}_{\mathsf{KW}}$ is at least as secure as ID. That also means that we cannot hope for tight PIMP-KOA security from the DLOG assumption. Instead, for the Katz-Wang identification scheme, we give a direct tight proof of PIMP-KOA security under the DDH assumption.

Lemma 5.11. $\mathsf{ID}_{\mathsf{KW}}$ *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* $\mathsf{par} = (p, g_1, g_2, \mathbb{G})$ *then* $\mathsf{ID}_{\mathsf{KW}}$ *is* $(t', \varepsilon', Q_{\mathrm{CH}})$-PIMP-KOA *secure, where* $t \approx t'$ *and* $\varepsilon \geq \varepsilon' - Q_{\mathrm{CH}}/2^n$*.*

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

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

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 $Q_{\mathrm{CH}}/2^n$, i.e., $\Pr[d = 1 \mid (X_1, X_2) = (g_1^{x_1}, g_2^{x_2})] \leq Q_{\mathrm{CH}}/2^n$.

It remains to prove the claim. For each index $i \in [Q_{\mathrm{CH}}]$, $\mathcal{A}$ first commits to $R_{i,1} = g_1^{r_1}$ and $R_{i,2} = g_2^{r_2}$ (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

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

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

5.2.2 Katz-Wang Signature scheme

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

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

By our results we obtain the following concrete security statements, where the first bound matches [28, Theorem 1].

Lemma 5.12. *If* DDH *is* (t,ε)*-hard in* $\mathsf{par} = (p, g_1, g_2, \mathbb{G})$ *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*

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

With a similar computation as in the case of Schnorr, one can compute concrete parameters for k-bits security assuming that a generic algorithm is the best method to attack the DDH assumption in par. If $\log p \geq 2k + c'''_{\mathsf{dl}}$ and $n \geq k + 1$, then KW is MU-SUF-CMA-secure, where c'''_{dl} is a constant that only depends on the generic algorithm.

References

[1] M. Abdalla, J. H. An, M. Bellare, and C. Namprempre. From identification to signatures via the Fiat-Shamir transform: Minimizing assumptions for security and forward-security. In L. R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 418–433. Springer, Heidelberg, Apr. / May 2002. (Cited on page 1, 2, 4.)

[2] M. Abdalla, P.-A. Fouque, V. Lyubashevsky, and M. Tibouchi. Tightly-secure signatures from lossy identification schemes. In D. Pointcheval and T. Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 572–590. Springer, Heidelberg, Apr. 2012. (Cited on page 4.)

[3] M. Bellare, C. Namprempre, D. Pointcheval, and M. 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 19.)

[4] M. Bellare and G. Neven. Multi-signatures in the plain public-key model and a general forking lemma. In A. Juels, R. N. Wright, and S. Vimercati, editors, *ACM CCS 06*, pages 390–399. ACM Press, Oct. / Nov. 2006. (Cited on page 9.)

[5] M. Bellare and A. Palacio. GQ and Schnorr identification schemes: Proofs of security against impersonation under active and concurrent attacks. In M. Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 162–177. Springer, Heidelberg, Aug. 2002. (Cited on page 4, 8, 9, 18, 19.)

[6] M. Bellare and T. Ristenpart. Simulation without the artificial abort: Simplified proof and improved concrete security for Waters' IBE scheme. In A. Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 407–424. Springer, Heidelberg, Apr. 2009. (Cited on page 2, 21.)

[7] M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In V. Ashby, editor, *ACM CCS 93*, pages 62–73. ACM Press, Nov. 1993. (Cited on page 2, 6.)

[8] D. Bernstein. [Cfrg] key as message prefix => multi-key security. https://mailarchive.ietf.org/arch/msg/cfrg/44gJyZlZ7-myJqWkChhpEF1KE9M, 2015. (Cited on page 5.)

[9] D. J. Bernstein. Multi-user Schnorr security, revisited. Cryptology ePrint Archive, Report 2015/996, 2015. http://eprint.iacr.org/. (Cited on page 5.)

[10] D. J. Bernstein, N. Duif, T. Lange, P. Schwabe, and B.-Y. Yang. High-speed high-security signatures. In B. Preneel and T. Takagi, editors, *CHES 2011*, volume 6917 of *LNCS*, pages 124–142. Springer, Heidelberg, Sept. / Oct. 2011. (Cited on page 5.)

[11] T. 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.)

[12] D. Boneh, C. Gentry, B. Lynn, and H. Shacham. Aggregate and verifiably encrypted signatures from bilinear maps. In E. Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 416–432. Springer, Heidelberg, May 2003. (Cited on page 5.)

[13] E. F. Brickell and K. S. McCurley. An interactive identification scheme based on discrete logarithms and factoring. In I. Damgård, editor, *EUROCRYPT'90*, volume 473 of *LNCS*, pages 63–71. Springer, Heidelberg, May 1991. (Cited on page 1.)

[14] D. Brown. [Cfrg] key as message prefix => multi-key security. http://www.ietf.org/mail-archive/web/cfrg/current/msg07336.html, 2015. (Cited on page 5.)

[15] A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. In A. M. Odlyzko, editor, *CRYPTO'86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, Aug. 1987. (Cited on page 1.)

[16] M. Fischlin and N. Fleischhacker. Limitations of the meta-reduction technique: The case of Schnorr signatures. In T. Johansson and P. Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 444–460. Springer, Heidelberg, May 2013. (Cited on page 5.)

[17] M. Fischlin, A. Lehmann, T. Ristenpart, T. Shrimpton, M. Stam, and S. Tessaro. Random oracles with(out) programmability. In M. Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 303–320. Springer, Heidelberg, Dec. 2010. (Cited on page 2.)

[18] N. Fleischhacker, T. Jager, and D. Schröder. On tight security proofs for Schnorr signatures. In P. Sarkar and T. Iwata, editors, *ASIACRYPT 2014, Part I*, volume 8873 of *LNCS*, pages 512–531. Springer, Heidelberg, Dec. 2014. (Cited on page 5, 20.)

[19] M. Fukumitsu and S. Hasegawa. Black-box separations on Fiat-Shamir-type signatures in the non-programmable random oracle model. In J. Lopez and C. J. Mitchell, editors, *ISC 2015*, volume 9290 of *LNCS*, pages 3–20. Springer, Heidelberg, Sept. 2015. (Cited on page 3, 5.)

[20] S. D. Galbraith, J. Malone-Lee, and N. P. Smart. Public key signatures in the multi-user setting. *Inf. Process. Lett.*, 83(5):263–266, 2002. (Cited on page 1, 3, 4, 12.)

[21] S. Garg, R. Bhaskar, and S. V. Lokam. Improved bounds on security reductions for discrete log based signatures. In D. Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 93–107. Springer, Heidelberg, Aug. 2008. (Cited on page 5.)

[22] M. Girault. An identity-based identification scheme based on discrete logarithms modulo a composite number (rump session). In I. Damgård, editor, *EUROCRYPT'90*, volume 473 of *LNCS*, pages 481–486. Springer, Heidelberg, May 1991. (Cited on page 1.)

[23] S. Goldwasser, S. Micali, and R. L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, Apr. 1988. (Cited on page 1.)

[24] L. C. Guillou and J.-J. Quisquater. A practical zero-knowledge protocol fitted to security microprocessor minimizing both trasmission and memory. In C. G. Günther, editor, *EUROCRYPT'88*, volume 330 of *LNCS*, pages 123–128. Springer, Heidelberg, May 1988. (Cited on page 4.)

[25] L. C. Guillou and J.-J. Quisquater. A "paradoxical" indentity-based signature scheme resulting from zero-knowledge. In S. Goldwasser, editor, *CRYPTO'88*, volume 403 of *LNCS*, pages 216–231. Springer, Heidelberg, Aug. 1990. (Cited on page 1.)

[26] M. Hamburg. Re: [Cfrg] EC signature: next steps. https://mailarchive.ietf.org/arch/msg/cfrg/af170b6OrLyNZUHBMOPWxcDrVRI, 2015. (Cited on page 5.)

[27] 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 5.)

[28] J. Katz and N. Wang. Efficiency improvements for signature schemes with tight security reductions. In S. Jajodia, V. Atluri, and T. Jaeger, editors, *ACM CCS 03*, pages 155–164. ACM Press, Oct. 2003. (Cited on page 4, 18, 22.)

[29] U. M. Maurer. Abstract models of computation in cryptography (invited paper). In N. P. Smart, editor, *10th IMA International Conference on Cryptography and Coding*, volume 3796 of *LNCS*, pages 1–12. Springer, Heidelberg, Dec. 2005. (Cited on page 25.)

[30] S. Micali and L. Reyzin. Improving the exact security of digital signature schemes. *Journal of Cryptology*, 15(1):1–18, 2002. (Cited on page 9.)

[31] S. Micali and A. Shamir. An improvement of the Fiat-Shamir identification and signature scheme. In S. Goldwasser, editor, *CRYPTO'88*, volume 403 of *LNCS*, pages 244–247. Springer, Heidelberg, Aug. 1990. (Cited on page 1.)

[32] K. Ohta and T. Okamoto. On concrete security treatment of signatures derived from identification. In H. Krawczyk, editor, *CRYPTO'98*, volume 1462 of *LNCS*, pages 354–369. Springer, Heidelberg, Aug. 1998. (Cited on page 2, 3, 4, 9.)

[33] T. Okamoto. Provably secure and practical identification schemes and corresponding signature schemes. In E. F. Brickell, editor, *CRYPTO'92*, volume 740 of *LNCS*, pages 31–53. Springer, Heidelberg, Aug. 1993. (Cited on page 1, 4.)

[34] H. Ong and C.-P. Schnorr. Fast signature generation with a Fiat-Shamir-like scheme. In I. Damgård, editor, *EUROCRYPT'90*, volume 473 of *LNCS*, pages 432–440. Springer, Heidelberg, May 1991. (Cited on page 1.)

[35] P. Paillier and D. Vergnaud. Discrete-log-based signatures may not be equivalent to discrete log. In B. K. Roy, editor, *ASIACRYPT 2005*, volume 3788 of *LNCS*, pages 1–20. Springer, Heidelberg, Dec. 2005. (Cited on page 3, 5.)

[36] D. Pointcheval and J. Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, 2000. (Cited on page 2, 3, 4, 9.)

[37] C.-P. Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991. (Cited on page 1, 4, 18.)

[38] Y. Seurin. On the exact security of Schnorr-type signatures in the random oracle model. In D. Pointcheval and T. Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 554–571. Springer, Heidelberg, Apr. 2012. (Cited on page 3, 5, 9, 14, 16.)

[39] V. Shoup. Lower bounds for discrete logarithms and related problems. In W. Fumy, editor, *EUROCRYPT'97*, volume 1233 of *LNCS*, pages 256–266. Springer, Heidelberg, May 1997. (Cited on page 25.)

[40] R. 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 [39, 29], 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 $x \xleftarrow{\$} \mathbb{Z}_p$, 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 $z = a + b$. 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, in the generic group model, Q-IDLOG is (t, ε)-hard where*

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

and $Q_{\mathbb{G}}$ is the amount of queries to $\mathcal{O}_{\mathbb{G}}$.

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 $x \xleftarrow{\$} \mathbb{Z}_p$. $\mathcal{A}$ is invoked on input $C_1 = 1$ and $C_{\mathbf{x}} = 2$. During its execution, $\mathcal{A}$ can query oracle $\mathcal{O}_{\mathbb{G}}$ on $(C_{a(\mathbf{x})}, C_{b(\mathbf{x})}) \in [i] \times [i]$. $\mathcal{O}_{\mathbb{G}}$ 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, $\mathcal{O}_{\mathbb{G}}$ increments counter i and adds $(z(\mathbf{x}), C_{z(\mathbf{x})} := i)$ to the list. Finally,

$\mathcal{O}_{\mathbb{G}}$ outputs $C_{z(\mathbf{x})}$. In total, $\mathcal{A}$ makes $Q_{\mathbb{G}}$ queries to this oracle and we denote by $(z_i(\mathbf{x}), i)$ the i-entry in the list $(i \in [Q_{\mathbb{G}} + 2])$.

Furthermore, $\mathcal{A}$ can make queries to $\textsc{Ch}(j)$, for some counter $j \in [c]$, which is answered with $h_j \xleftarrow{\$} \mathbb{Z}_p$. 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 $\textsc{Ch}$ oracle. Eventually, $\mathcal{A}$ outputs $s \in \mathbb{Z}_p$ and terminates. Next, $x \xleftarrow{\$} \mathbb{Z}_p$ 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 $\mathcal{O}_{\mathbb{G}}$ 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 $i \neq \ell \in [Q_{\mathbb{G}}]$ 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

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

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

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

This completes the proof. ■