The Fiat-Shamir Zoo: Relating the Security of Different Signature Variants

Matilda Backendal[1,2][0000−0002−8677−8301]✉, Mihir Bellare[1][0000−0002−8765−5573], Jessica Sorrell[1][0000−0001−9227−1032], and Jiahao Sun[1][0000−0002−1806−5000]

[1] Department of Computer Science and Engineering, University of California San Diego, USA. {mihir,jlsorrel}@eng.ucsd.edu, jis126@ucsd.edu.
[2] Faculty of Engineering, Lund University, Sweden
matilda.backendal.6668@student.lu.se.

Abstract. The Fiat-Shamir paradigm encompasses many different ways of turning a given identification scheme into a signature scheme. Security proofs pertain sometimes to one variant, sometimes to another. We systematically study three variants that we call the challenge (signature is challenge and response), commit (signature is commitment and response), and transcript (signature is challenge, commitment and response) variants. Our framework captures the variants via transforms that determine the signature scheme as a function of not only the identification scheme and hash function (to cover both standard and random oracle model hashing), but also what we call a signing algorithm, to cover both classical and with-abort signing. We relate the security of the signature schemes produced by these transforms, giving minimal conditions under which uf-security of one transfers to the other. To apply this comprehensively, we formalize linear identification schemes, show that many schemes in the literature are linear, and show that any linear scheme meets our conditions for the signature schemes given by the three transforms to have equivalent uf-security. Our results give a comprehensive picture of the Fiat-Shamir zoo and allow proofs of security in the literature to be transferred automatically from one variant to another.

1 Introduction

Ed25519 [13] is a fast signature scheme with widespread usage including in TLS 1.3, SSH, Signal, and Tor [22]. It is derived via the Fiat-Shamir paradigm [17] applied to the Schnorr identification scheme [28]. It is not alone; over the last three decades the Fiat-Shamir paradigm has been a popular way to obtain signature schemes, for reasons including the following: *Speed.* It yields some of our most efficient signature schemes. *Proofs.* The paradigm is backed by proofs of security [27, 1, 21]. *Extendability.* Classically used with number-theoretic schemes [17, 20, 28, 26], extensions of the paradigm now provide lattice-based schemes, some of which are proposed to NIST for post-quantum standards [23, 2, 16, 14].

However, referring, above, to "the" Fiat-Shamir paradigm is misleading, for the paradigm is not monolithic: It encompasses variant methods that, starting

from a given identification scheme, yield different signature schemes. This creates some confusion, with proofs in the literature pertaining sometimes to one variant, sometimes to another, yet being quoted without regard to which variant is being considered. Extensions such as signing with aborts [23, 2, 16, 14] bring further variants.

This paper aims to provide a systematic and comprehensive picture of the variants in a general setting, and give results relating their security under minimal assumptions. This allows us to leverage existing security proofs given for one variant [27, 1, 21], automatically transferring them to another, rather than prove security of different variants from scratch.

Background. An identification scheme ID is a 3-move interactive protocol. The prover, having public key pk and secret key sk, sends a commitment $\textsc{Ct}$, the verifier sends a random challenge $\textsc{Ch}$, the prover sends a response $\textsc{Rp}$, and the verifier computes a decision $d \leftarrow \mathsf{ID.V}(1^\lambda, pk, \textsc{Ct}, \textsc{Ch}, \textsc{Rp})$ to accept or reject, where 1^λ is the unary representation of the security parameter λ. In a signature scheme based on ID, the prover, now the signer, given message M, computes $\textsc{Ct}$ as before, sets $\textsc{Ch} \leftarrow \mathsf{F}(1^\lambda, pk, (\textsc{Ct}, M))$ to a hash of the commitment and message, computes $\textsc{Rp}$ and then returns a signature σ. We distinguish three variants with regard to what σ consists of. **(1)** In what we call the *transcript* variant [27], σ is $(\textsc{Ct}, \textsc{Ch}, \textsc{Rp})$. It is verified by checking that $\mathsf{ID.V}(1^\lambda, pk, \textsc{Ct}, \textsc{Ch}, \textsc{Rp}) = \mathsf{true}$ and $\textsc{Ch} = \mathsf{F}(1^\lambda, pk, (\textsc{Ct}, M))$. **(2)** In what we call the *commitment* variant [25, 1], σ is $(\textsc{Ct}, \textsc{Rp})$. It is verified by setting $\textsc{Ch} \leftarrow \mathsf{F}(1^\lambda, pk, (\textsc{Ct}, M))$ and checking that $\mathsf{ID.V}(1^\lambda, pk, \textsc{Ct}, \textsc{Ch}, \textsc{Rp}) = \mathsf{true}$. **(3)** In what we call the *challenge* variant [17, 28, 20, 26], σ is $(\textsc{Ch}, \textsc{Rp})$. This usually yields the shortest signatures but requires a *commitment reproducing algorithm* $\mathsf{ID.CR}$ that allows the verifier to reproduce $\textsc{Ct} \leftarrow \mathsf{ID.CR}(1^\lambda, pk, \textsc{Ch}, \textsc{Rp})$ and then check that $\textsc{Ch} = \mathsf{F}(1^\lambda, pk, (\textsc{Ct}, M))$.

The history of the various transforms is interesting. Fiat and Shamir (FS) [17], GQ [20], Schnorr [28] and Okamoto [26] all gave challenge-style signatures. However, the first security proofs, by Pointcheval and Stern (PS) [27], were for transcript-style signatures, which seem to originate with them. The proofs of Abdalla, An, Bellare and Namprempre (AABN) [1] were for commitment-style signatures, which seem to originate with Ohta and Okamoto (OO) [25]. The changes are (mostly) made silently: PS, OO, AABN (and subsequent literature) tend to refer to their results as establishing security of the FS, GQ, Schnorr and Okamoto schemes, but the proofs pertain to variants not only different from the original ones but in some cases also different from each other.

Questions. We would like a fuller picture, that given an identification scheme ID tells us, for each of the three variant signature schemes derived from ID, whether or not the variant is secure. The above-mentioned results do not directly yield this information. One approach to filling this gap would be to return to the techniques in prior proofs and directly try to prove security of each variant signature scheme. Given the complexity of the techniques, this would be tedious. Instead, we seek *relations between the variants*. This means that for each pair $\mathsf{DS}_\mathrm{x}, \mathsf{DS}_\mathrm{y}$ of variant signature schemes derived from a given identification scheme

ID, we want to determine an assumption or condition $\mathsf{A}_{\mathrm{x,y}}$ on ID under which the security of DS_{x} implies the security of DS_{y}. Then, if we know from prior work that DS_{x} is secure, and can establish that ID satisfies $\mathsf{A}_{\mathrm{x,y}}$, we can conclude that DS_{y} is secure too. This would leverage existing proofs in a modular way. We seek assumptions $\mathsf{A}_{\mathrm{x,y}}$ as weak as possible, both to maximize potential applicability and to understand, theoretically, what are the minimal requirements for a relation to hold.

The literature does contain claims about such relations [2, 21, 18], but (as we will discuss in more detail below) they are mostly informal, specific to particular schemes, or make assumptions we will show to be unnecessarily strong.

Our framework. We capture the variants via transforms that we call general to reflect a broader parameterization than in prior work. A *general Fiat-Shamir transform* $\mathbf{gFS}$ determines a signature scheme $\mathsf{DS} = \mathbf{gFS}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$ based on input parameters an identification scheme ID, a hash function F (allowed access to the random oracle H) and (most novel) a *signing algorithm* S (also allowed access to H). The signing algorithm takes $1^\lambda, pk, sk, M$ and returns either $\perp$ or an honest, accepting transcript $(\mathrm{Ct}, \mathrm{Ch}, \mathrm{Rp})$ satisfying $\mathrm{Ch} = \mathsf{F}^{\mathrm{H}}(1^\lambda, pk, (\mathrm{Ct}, M))$. But, beyond requiring this condition, we *do not prescribe how the signing algorithm operates*. To sign message M, run $T \leftarrow\!\!\$\ \mathsf{S}^{\mathrm{H}}(1^\lambda, pk, sk, M)$, and return $\perp$ as signature if $T = \perp$. Otherwise, parse T as $(\mathrm{Ct}, \mathrm{Ch}, \mathrm{Rp})$. Exactly what is returned as the signature σ, and how that signature is verified, depends on the transform. This is summarized for each of our three transforms $\mathbf{gFS}_{\mathrm{tr}}, \mathbf{gFS}_{\mathrm{ct}}, \mathbf{gFS}_{\mathrm{ch}}$ in Figure 1, reflecting the three variants discussed above. The schemes are shown in full in Figure 4. As we will see, the broad parameterization enhances applicability because our relations will hold for all choices of F, S.

Relations between security of signature schemes. The security attribute we consider for the signature schemes, hereafter called uf-security, is the usual unforgeability under chosen message attack [19] extended, due to growing recognition of its importance, to the multi-user setting [5, 21]. Now, given $\mathsf{ID}, \mathsf{F}, \mathsf{S}$, consider the three signature schemes $\mathsf{DS}_{\mathrm{x}} = \mathbf{gFS}_{\mathrm{x}}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$ for $\mathrm{x} \in \{\mathrm{tr}, \mathrm{ct}, \mathrm{ch}\}$. We seek relations between their uf-security, as discussed above. This means that for each (distinct) pair $\mathrm{x}, \mathrm{y} \in \{\mathrm{tr}, \mathrm{ct}, \mathrm{ch}\}$ we ask under what assumption $\mathsf{A}_{\mathrm{x,y}}$ the uf-security of DS_{x} implies the uf-security of DS_{y}.

Our results are summarized by the picture at the bottom of Figure 1. That $\mathsf{DS}_{\mathrm{tr}}$ and $\mathsf{DS}_{\mathrm{ct}}$ have equivalent uf-security is trivial. The interesting question is, does uf-security of one of $\mathsf{DS}_{\mathrm{ct}}, \mathsf{DS}_{\mathrm{ch}}$ imply uf-security of the other? The straight, barred arrows say that in general (that is, without any condition beyond basic completeness on the commitment reproducing algorithm) the answer is no. The curved, un-barred arrows say the answer is yes, under conditions on the commitment reproducing algorithm (formally, on the overlying identification scheme ID that includes this algorithm) that we give. Specifically, Theorem 2 says that if ID has a property we define and call soundness (SND) then, if $\mathsf{DS}_{\mathrm{ct}}$ is uf-secure, so is $\mathsf{DS}_{\mathrm{ch}}$. Theorem 4 says that if ID has a property we define and call consistency (CNS) then, if $\mathsf{DS}_{\mathrm{ch}}$ is uf-secure, so is $\mathsf{DS}_{\mathrm{ct}}$. SND-security asks that it be computationally hard to find a challenge and response such that the commitment

Signature Scheme	**Signature σ**	**To verify σ, check this:**
$\mathsf{DS}_{\mathrm{tr}} = \mathbf{gFS}_{\mathrm{tr}}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$	$(\mathrm{CT}, \mathrm{CH}, \mathrm{RP})$	$\mathsf{ID.V}(1^\lambda, pk, \mathrm{CT}, \mathrm{CH}, \mathrm{RP}) = \mathsf{true}$ $\mathrm{CH} = \mathsf{F}^{\mathrm{H}}(1^\lambda, pk, (\mathrm{CT}, M))$
$\mathsf{DS}_{\mathrm{ct}} = \mathbf{gFS}_{\mathrm{ct}}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$	$(\mathrm{CT}, \mathrm{RP})$	$\mathsf{ID.V}(1^\lambda, pk, \mathrm{CT}, \mathsf{F}^{\mathrm{H}}(1^\lambda, pk, (\mathrm{CT}, M)), \mathrm{RP}) = \mathsf{true}$
$\mathsf{DS}_{\mathrm{ch}} = \mathbf{gFS}_{\mathrm{ch}}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$	$(\mathrm{CH}, \mathrm{RP})$	$\mathrm{CH} = \mathsf{F}^{\mathrm{H}}(1^\lambda, pk, (\mathsf{ID.CR}(1^\lambda, pk, \mathrm{CH}, \mathrm{RP}), M))$

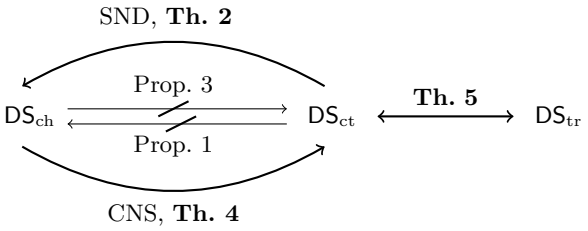

Fig. 1. Top: Signatures and verification in the signature schemes given by our transforms, where $\mathsf{ID.CR}$ is the commitment reproducing algorithm of ID. Signing of message M (not shown) is done by letting $(\mathrm{CT}, \mathrm{CH}, \mathrm{RP}) \leftarrow\!\!\$\ \mathsf{S}^{\mathrm{H}}(1^\lambda, pk, sk, M)$ and returning the shown σ. **Bottom:** Relations between uf-security of the signature schemes.

reproducing algorithm succeeds in returning a commitment but the resulting transcript is not accepting. CNS-security asks that it be computationally hard to create an accepting transcript in which the commitment is different from the one given by the commitment reproducing algorithm. The reductions underlying all our positive results are tight.

Breadth of applicability. The positive relations (un-barred arrows in Figure 1) hold for *all choices of hash function* F *and signing algorithm* S. This broadens applicability. With regard to hashing, it means we can transfer security in both the random oracle and the standard models: For $\mathrm{x}, \mathrm{y} \in \{\mathrm{tr}, \mathrm{ct}, \mathrm{ch}\}$, if DS_{x} provides uf-security with a random-oracle hash function then (assuming of course, as necessary, properties of ID as above) so does DS_{y}, but if DS_{x} provides uf-security with hash function SHA256, then so does DS_{y}. With regard to signing, this means that our framework captures both canonical and more modern variants of the Fiat-Shamir paradigm. For example, in the literature Fiat-Shamir with aborts [23, 2, 16, 14] is viewed as an extension of the canonical Fiat-Shamir paradigm. In our framework, the canonical and with-abort variants correspond simply to different choices of signing algorithm S (cf. Figure 4), so our results apply automatically to both.

We elaborate on the second point. We said above how the Fiat-Shamir paradigm prescribes signing a message M, which we now call the canonical way: generate CT as would the honest prover, set $\mathrm{CH} \leftarrow \mathsf{F}^{\mathrm{H}}(1^\lambda, pk, (\mathrm{CT}, M))$, generate RP as would the prover, then return σ computed from $\mathrm{CT}, \mathrm{CH}, \mathrm{RP}$ according to the variant (challenge, commit or transcript) of interest. This is captured for us by setting S to the canonical algorithm on the bottom left of Figure 4.

This works (yields a correct signature) if the identification scheme has perfect correctness. However, in the identification schemes from lattices [23, 2, 16, 14], the response can be $\perp$ with constant probability. So the process is modified to repeat picking $\text{Ct}, \text{Ch}, \text{Rp}$ as above until the conversation is accepting or some time bound is exceeded, which is called signing with aborts. (In this case, the signature schemes have imperfect correctness, returning $\perp$ with negligible probability.) The challenge, commit and transcript variants for the signature schemes exist here too, so the question of how their security relates arises again. We do not need to address this separately. It is captured for us, and addressed by the results noted above, simply by setting S to the algorithm on the bottom right of Figure 4. Choices of S beyond these two are possible as well, for potential further applications.

Perfect uniqueness. We have introduced the SND and CNS conditions on commitment reproducible identification schemes, showing that they suffice for transfer of uf-security between the signature variants. (SND allows the uf-security of DS_{ct} to imply that of DS_{ch}, and CNS the converse.) We also define a third condition called perfect uniqueness (P-UNIQ). It asks that a transcript $\text{Ct}, \text{Ch}, \text{Rp}$ be accepting if and only if the commitment reproducing algorithm $\mathsf{ID.CR}$ returns exactly Ct on inputs Ch, Rp. Figure 6 says that P-UNIQ implies both SND and CNS. Establishing P-UNIQ-ness of a commitment reproducible identification scheme ID is thus a simple path (and one we will often be able to use) to showing that all the signature variants derived from ID have equivalent uf-cma security. However, Figure 6 also says that P-UNIQ is a strictly stronger condition than SND or CNS. So for some commitment reproducible identification schemes, P-UNIQ may fail to be true, yet we might be able to directly establish SND and CNS to show equivalence of uf-security of the signature variants.

Linear identification schemes. We'd like to know whether identification schemes in the literature meet our conditions (P-UNIQ, or SND, CNS as necessary). However, there are many schemes, and new ones keep appearing, and testing them individually is tedious. Instead, we formalize *linear* identification schemes and show that any linear identification scheme is (unconditionally) P-UNIQ. Our results thus say that the three variant signature schemes emanating from any linear identification scheme have equivalent uf-security.

We then show that classical identification schemes like FS [17], Sch [28], GQ [20] and Ok [26] are linear. We also show that the Ly lattice based identification scheme of [23] is linear. Since proofs of uf-security exist for at least one signature variant for all these identification schemes, we can conclude that all three variants are uf-secure.

Lyubashevsky [24] directly gives a lattice-based signature scheme that he does not derive via the FS paradigm. (Indeed the paper presents no identification scheme.) We show how to capture it in our framework as $\mathbf{gFS}_{\text{ch}}[\mathsf{ID}, \mathsf{F}, \mathsf{SA}_{\mathsf{ID},\mathsf{F},t}]$ where SA_t is the abort-based signing algorithm on the bottom right of Figure 4 and ID is an identification scheme that we define and show is linear. This means we can define the other variant signature schemes and transfer the proofs of [24] to them.

As the above indicates, the concept of linear identification schemes serves also to unify the literature, showing that what look like different schemes are in fact instances of one underlying scheme. We see this as something that was understood but not, until now, formalized.

Due to lack of space, the material on linear identification schemes is entirely omitted from this proceedings version and can be found in our full version [4].

WHICH VARIANT SHOULD ONE USE? Our work is about relating the security of the different signature variants. The question of which variant to prefer in usage is orthogonal, and the answer differs from case to case. We discuss the choices briefly. The challenge variant $\mathbf{gFS}_{\mathrm{ch}}$ usually yields the shortest signatures (examples where this is true are FS [17], GQ [20], Sch in the group of integers modulo a prime [28] and Ly [23]) but requires that ID be commitment reproducible (meaning, there exists a commitment reproducing algorithm ID.CR) which is not always true. When ID is not commitment reproducible, one can use $\mathbf{gFS}_{\mathrm{ct}}$. Here, in some cases (like Sch over elliptic curve groups) the signature size stays as small as with $\mathbf{gFS}_{\mathrm{ch}}$, but in other cases, it might grow. The transcript variant $\mathbf{gFS}_{\mathrm{tr}}$ is also an option for usage when commitment reproducibility is lacking, but there seems no practical reason for this, because signatures are always shorter with $\mathbf{gFS}_{\mathrm{ct}}$. We consider $\mathbf{gFS}_{\mathrm{tr}}$ in this paper because it was the variant for which the seminal work of Pointcheval and Stern [27] gave proofs.

Of course performance (including signature size) is just one criterion with regard to a choice for usage. Another is security proofs. The general results in the literature give proofs for $\mathbf{gFS}_{\mathrm{ct}}$ [1] and $\mathbf{gFS}_{\mathrm{tr}}$ [27], not $\mathbf{gFS}_{\mathrm{ch}}$. Our framework and results can be used to transfer them to the (usually more efficient) $\mathbf{gFS}_{\mathrm{ch}}$.

RELATED WORK. Kiltz, Masny, and Pan [21] briefly note that $\mathsf{DS}_{\mathrm{ch}}, \mathsf{DS}_{\mathrm{ct}}$ are equivalent in terms of uf-security assuming the verification algorithm has a certain property. This seems to be equivalent to the identification scheme being P-UNIQ. Figure 6 shows that the SND and CNS properties that allow us to establish the same equivalence are implied by, and strictly weaker than, P-UNIQ, making our results stronger. Also their results are for the canonical signing algorithm, while ours are for an arbitrary one. Abdalla, Fouque, Lyubashevsky, and Tibouchi [2] give results for commitment-style signatures with aborts, saying that these transfer to the challenge style for their schemes because "the commitment is uniquely determined by the challenge and response." The phrase in quotes is not too precise but the intent is likely P-UNIQ. Galbraith, Petit, and Silva [18] show that, for their particular scheme, under weak conditions on commitment reproducibility, security of the commit version implies security of a version that is like the challenge one except that signature verification additionally checks that the verifier accepts the transcript. This is added verification cost compared to the classical Fiat-Shamir style challenge variant, which is the version we consider and which does not have such a check.

We view our work as unifying, systematizing and formalizing long-standing understanding, scattered observations and folklore. Nothing in this paper is very novel or technically difficult. Our hope is that it fills some gaps and can be a point of reference for variants of Fiat-Shamir signatures.

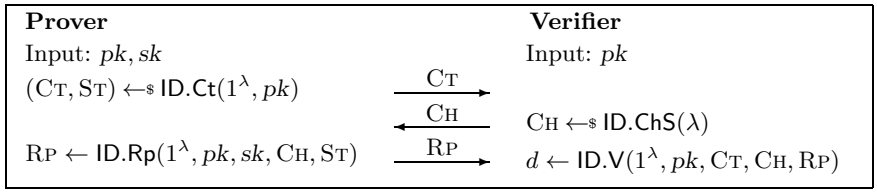

Fig. 2. Operation of an identification scheme ID.

EXTENSIONS. Beyond basic (uf-cma) signature schemes, identification schemes have been used to build identity-based signatures [8], blind signatures [15, 27] leakage-resilient signatures [3, 6], double authentication preventing signatures [10] and beyond. Returning to the basic setting, variants of the Fiat-Shamir paradigm offering better concrete security have been considered [9]. In all these places and settings, the commit, challenge and transcript variants arise. One can ask how their security relates, and extend our framework and results to answer this question.

2 Basic definitions

NOTATION. We let ε denote the empty string. If Z is a string then $|Z|$ denotes its length. If X is a finite set, we let $x \leftarrow\!\!\$\ X$ denote picking an element of X uniformly at random and assigning it to x, and we let $|X|$ denote the size of X. We use $\perp$ (bot) as a special symbol to denote rejection, and it is assumed to not be in $\{0,1\}^*$. Both inputs and outputs to algorithms can be $\perp$. We adopt the convention that if any input to an algorithm is $\perp$, then its output is $\perp$ as well. By $\lambda \in \mathbb{N}$ we denote the security parameter, and by 1^λ its unary representation. Recall that a function $\nu\colon \mathbb{N} \to \mathbb{R}$ is negligible if for every positive polynomial $p\colon \mathbb{N} \to \mathbb{R}$ there is a $\lambda_p \in \mathbb{N}$ such that $\nu(\lambda) \leq 1/p(\lambda)$ for all $\lambda \geq \lambda_p$.

Algorithms may be randomized unless otherwise indicated. Running time is worst case. "PT" stands for "polynomial time," whether for a randomized algorithm or a deterministic one. If A is an algorithm, we let $y \leftarrow A^{\text{O}_1,\cdots}(x_1, \ldots; \omega)$ denote running A on inputs $x_1, \ldots$ and coins ω, with oracle access to $\text{O}_1, \ldots$, and assigning the output to y. By $y \leftarrow\!\!\$\ A^{\text{O}_1,\cdots}(x_1, \ldots)$ we denote picking ω at random and letting $y \leftarrow A^{\text{O}_1,\cdots}(x_1, \ldots; \omega)$. We let $[A^{\text{O}_1,\cdots}(x_1, \ldots)]$ denote the set of all possible outputs of A when run on inputs $x_1, \ldots$ and with oracle access to $\text{O}_1, \ldots$. An adversary is an algorithm.

We use the code-based game-playing framework of [12]. (See Figure 5 for an example.) By $\Pr[\mathbf{G}]$ we denote the probability that the execution of game $\mathbf{G}$ results in the game returning true. We adopt the convention that the running time of an adversary executing with some game refers to the worst case execution time of the game with the adversary, meaning the time taken for oracles to compute replies to queries is included. The random oracle (RO) model [11] is captured by inclusion in the game of a procedure H that implements a variable output length RO. See for example Figure 3.

IDENTIFICATION SCHEMES. An identification scheme ID (called a canonical identification scheme in [1]) specifies several algorithms and associated quantities, as follows. In an initialization phase, via $(pk, sk) \leftarrow\!\!\$\ \mathsf{ID.Kg}(1^\lambda)$, the prover runs the key-generation algorithm $\mathsf{ID.Kg}$ on input the unary representation 1^λ of the security parameter λ to obtain a public key pk and a private key sk, both of which she stores. It is assumed that the verifier is in possession of pk. (In practice this is likely done via certificates, but that is not in the scope of the identification scheme.) Identification then operates as depicted in Figure 2. Via $(\textsc{Ct}, \textsc{St}) \leftarrow\!\!\$\ \mathsf{ID.Ct}(1^\lambda, pk)$, the prover generates a *commitment* $\textsc{Ct}$ and corresponding private state $\textsc{St}$. The verifier sends a challenge $\textsc{Ch} \leftarrow\!\!\$\ \mathsf{ID.ChS}(\lambda)$ drawn at random from the *challenge space* $\mathsf{ID.ChS}(\lambda) = \{0,1\}^{\mathsf{ID.ChL}(\lambda)}$ where $\mathsf{ID.ChL}\colon \mathbb{N} \to \mathbb{N}$ is the *challenge length* function associated to ID. The prover's *response* $\textsc{Rp} \leftarrow \mathsf{ID.Rp}(1^\lambda, pk, sk, \textsc{Ch}, \textsc{St})$ is computed via a deterministic algorithm $\mathsf{ID.Rp}$. The verifier's *decision* $d \leftarrow \mathsf{ID.V}(1^\lambda, pk, \textsc{Ct}, \textsc{Ch}, \textsc{Rp})$, which is either $\mathsf{true}, \mathsf{false}$ or $\perp$, is also computed deterministically. Algorithms $\mathsf{ID.Kg}$, $\mathsf{ID.Ct}$, $\mathsf{ID.Rp}$, $\mathsf{ID.V}$ are required to be PT.

The *honest-transcript generating function* $\mathbf{HTR}_{\mathsf{ID},\lambda}$ associated to ID and $\lambda \in \mathbb{N}$ takes input $(pk, sk) \in [\mathsf{ID.Kg}(1^\lambda)]$, and returns a transcript of a conversation between the honest prover and the verifier, as follows:

$\underline{\mathbf{HTR}_{\mathsf{ID},\lambda}(pk, sk)}$
$(\textsc{Ct}, \textsc{St}) \leftarrow\!\!\$\ \mathsf{ID.Ct}(1^\lambda, pk)$; $\textsc{Ch} \leftarrow\!\!\$\ \mathsf{ID.ChS}(\lambda)$; $\textsc{Rp} \leftarrow \mathsf{ID.Rp}(1^\lambda, pk, sk, \textsc{Ch}, \textsc{St})$
Return $(\textsc{Ct}, \textsc{Ch}, \textsc{Rp})$

For $\lambda \in \mathbb{N}$ and $(pk, sk) \in [\mathsf{ID.Kg}(1^\lambda)]$, we define the *set of accepting transcripts*

$$\mathbf{ACC}_{\mathsf{ID},\lambda}(pk) = \{\ (\textsc{Ct}, \textsc{Ch}, \textsc{Rp})\ :\ \mathsf{ID.V}(1^\lambda, pk, \textsc{Ct}, \textsc{Ch}, \textsc{Rp}) = \mathsf{true}\ \}\ .$$

Correctness, for most schemes, is simple, saying that honest transcripts are always accepting: formally, for all $\lambda \in \mathbb{N}$ and all $(pk, sk) \in [\mathsf{ID.Kg}(1^\lambda)]$ we have $[\mathbf{HTR}_{\mathsf{ID},\lambda}(pk, sk)] \subseteq \mathbf{ACC}_{\mathsf{ID},\lambda}(pk)$. We call this perfect correctness. However we will need to also consider a relaxation where there is a correctness error, and this has to be carefully formulated. We say that ID has correctness error $\nu\colon \mathbb{N} \to \mathbb{R}$ if for all $\lambda \in \mathbb{N}$ and all $(pk, sk) \in [\mathsf{ID.Kg}(1^\lambda)]$ we have $\Pr[(\textsc{Ct}, \textsc{Ch}, \textsc{Rp}) \notin \mathbf{ACC}_{\mathsf{ID},\lambda}(pk)] \leq \nu(\lambda)$, where the probability is over $(\textsc{Ct}, \textsc{Ch}, \textsc{Rp}) \leftarrow\!\!\$\ \mathbf{HTR}_{\mathsf{ID},\lambda}(pk, sk)$. This captures the requirement that the verifier accepts with probability at least $1 - \nu(\lambda)$ in an interaction with the honest prover. Some commonly occurring choices for ν are a constant, like $\nu(\cdot) = 1/2$, or a negligible function, and in the latter case we say that ID has negligible correctness error.

SIGNATURE SCHEMES. A (digital) signature scheme DS specifies several algorithms and associated quantities, as follows. In an initialization phase, via $(pk, sk) \leftarrow\!\!\$\ \mathsf{DS.Kg}(1^\lambda)$, the signer runs the PT key-generation algorithm $\mathsf{DS.Kg}$ on input 1^λ to obtain a public key pk and a private key sk, both of which she stores. It is assumed that the verifier is in possession of pk. (As with identification, how this happens is not in the scope of the signature scheme.) Via $\sigma \leftarrow\!\!\$\ \mathsf{DS.Sign}^{\mathrm{H}}(1^\lambda, pk, sk, M)$, the signer generates a signature σ of a message

Game $\mathbf{G}^{\mathrm{uf}}_{\mathsf{DS},\mathcal{A}}(\lambda)$	$\underline{\mathrm{Sign}(i, M)}$
$n \leftarrow 0$; $S \leftarrow \emptyset$	$\sigma \leftarrow\!\!\$\ \mathsf{DS.Sign}^{\mathrm{H}}(1^\lambda, pk_i, sk_i, M)$
$(M, \sigma, i) \leftarrow\!\!\$\ \mathcal{A}^{\mathrm{New},\mathrm{Sign},\mathrm{H}}(1^\lambda)$	$S \leftarrow S \cup \{(i, M)\}$
$d \leftarrow \mathsf{DS.V}^{\mathrm{H}}(1^\lambda, pk_i, M, \sigma)$	Return σ
Return $(d = \mathsf{true}) \wedge ((i, M) \notin S)$	
	$\underline{\mathrm{New}()}$
$\underline{\mathrm{H}(W, \ell)}$	$n \leftarrow n + 1$
If $\mathrm{HT}[W, \ell] = \perp$ then $\mathrm{HT}[W, \ell] \leftarrow\!\!\$\ \{0,1\}^\ell$	$(pk_n, sk_n) \leftarrow\!\!\$\ \mathsf{DS.Kg}(1^\lambda)$
Return $\mathrm{HT}[W, \ell]$	Return pk_n

Fig. 3. Game for UF-CMA security of digital signature schemes in the multi-user setting.

$M \in \{0,1\}^*$. Via $d \leftarrow \mathsf{DS.V}^{\mathrm{H}}(1^\lambda, pk, M, \sigma)$, a verifier can deterministically obtain a decision regarding whether σ is a valid signature of M under pk. The signing and verifying algorithms have oracle access to the random oracle H and are required to be PT. We say that DS has correctness error ν: $\mathbb{N} \to \mathbb{R}$ if, for all $\lambda \in \mathbb{N}$, all $(pk, sk) \in [\mathsf{DS.Kg}(1^\lambda)]$ and all $M \in \{0,1\}^*$ we have $\Pr[\mathsf{DS.V}^{\mathrm{H}}(1^\lambda, pk, M, \mathsf{DS.Sign}^{\mathrm{H}}(1^\lambda, pk, sk, M)) \neq \mathsf{true}] \leq \nu(\lambda)$, where the probability is over the random choices of H and the coins of $\mathsf{DS.Sign}$. We say correctness is perfect if $\nu(\cdot) = 0$.

Our security metric for signatures, called uf-security, is the usual unforgeability under chosen-message attack [19], but in the multi-user setting, due to increasing recognition of the importance of the latter [5, 21]. Consider game $\mathbf{G}^{\mathrm{uf}}_{\mathsf{DS},\mathcal{A}}(\lambda)$ in Figure 3 associated to signature scheme DS and adversary $\mathcal{A}$. By calling the New oracle, the adversary can initialize a new user (signer), obtaining her public key. The number of users n, being the number of queries to New, is thus under the adversary's control. Via the Sign oracle, the adversary can mount its chosen-message attack, obtaining a signature on a message of its choice under a user of its choice. The adversary eventually outputs a pointer $i \in \{1, \ldots, n\}$ to a user, a message M, and a claimed signature of M under pk_i, winning if the signature is valid and non-trivial. We let $\mathbf{Adv}^{\mathrm{uf}}_{\mathsf{DS},\mathcal{A}}(\lambda) = \Pr[\mathbf{G}^{\mathrm{uf}}_{\mathsf{DS},\mathcal{A}}(\lambda)]$ be the probability that the game returns true. We say that DS is uf-secure if the function $\mathbf{Adv}^{\mathrm{uf}}_{\mathsf{DS},\mathcal{A}}(\cdot)$ is negligible for all PT adversaries $\mathcal{A}$.

3 Transforms and signature relations

The FS transforms are usually viewed as turning an identification scheme into a signature scheme in the random oracle model. Our general transforms take not only an identification scheme, but a hash function F, so that both standard model and random oracle model hash functions are covered. More novel, they take a description S of a signing process, to cover the fact that FS has been used in settings with and without abort. We begin with commitment reproducibility, needed for some of the transforms, then discuss the other parameters, and then specify the transforms. We then define the SND and CNS security notions for

commitment reproducible identification schemes that allow us to relate the security of the schemes emanating from the different general transforms. Finally we study relations between security notions for commitment reproducible identification schemes.

COMMITMENT REPRODUCIBILITY. A *commitment reproducing algorithm* for identification scheme ID is a deterministic, PT algorithm $\mathsf{ID.CR}$ that returns an output in $\{0,1\}^* \cup \{\perp\}$. We require the following completeness condition: for all $\lambda \in \mathbb{N}$, all $(pk, sk) \in [\mathsf{ID.Kg}(1^\lambda)]$ and all $(\mathrm{CT}, \mathrm{CH}, \mathrm{RP}) \in [\mathbf{HTR}_{\mathsf{ID},\lambda}(pk, sk)] \cap \mathbf{ACC}_{\mathsf{ID},\lambda}(pk)$ we have $\mathrm{CT} = \mathsf{ID.CR}(1^\lambda, pk, \mathrm{CH}, \mathrm{RP})$. Completeness says that the commitment in an accepting transcript of an interaction between the honest prover and the verifier is uniquely determined by the challenge and response, and moreover can be computed from them in PT by the commitment reproducing algorithm. An identification scheme ID is *commitment reproducible* if it specifies (in addition to the quantities it already specifies as per Section 2) a commitment reproducing algorithm $\mathsf{ID.CR}$ that satisfies the completeness condition.

Commitment reproducibility is enough to define the $\mathbf{gFS}_{\mathrm{ch}}$ transform, but further attributes $(\mathrm{SND}, \mathrm{CNS})$ will be necessary to establish relations between uf-security of the signature schemes.

HASHING. The **gFS** transforms use a hash function. Most of our results hold regardless of the choice of the hash function, in particular both when it is a standard-model function and when it is a random oracle. To capture this formally, we define a hash function as a deterministic algorithm F that may have access to a random oracle H. It is *compatible* with identification scheme ID if $\mathsf{F}^{\mathrm{H}}(1^\lambda, pk, x) \in \mathsf{ID.ChS}(\lambda)$ for all $\lambda \in \mathbb{N}$, all $(pk, sk) \in [\mathsf{ID.Kg}(1^\lambda)]$, all x and all H. In our usage, $x = (\mathrm{CT}, M)$ will consist of a commitment and message. By setting $\mathsf{F}^{\mathrm{H}}(1^\lambda, pk, x) = \mathrm{H}((1^\lambda, pk, x), \ell(\lambda))$ for some $\ell\colon \mathbb{N} \to \mathbb{N}$ we can cover the case where the hash function is a random oracle, but we can also, for example, set $\mathsf{F}^{\mathrm{H}}(1^\lambda, pk, x) = \mathrm{SHA256}((1^\lambda, pk, x))$ to cover schemes where the hash function has been instantiated via SHA256.

SIGNING. Let ID be an identification scheme, and F a hash function compatible with it. A *signing algorithm* compatible with ID and F is a PT algorithm S that operates as $T \leftarrow\!\!\$\ \mathsf{S}^{\mathrm{H}}(1^\lambda, pk, sk, M)$. We require that if $T \neq \perp$ then it parses as $(\mathrm{CT}, \mathrm{CH}, \mathrm{RP}) \leftarrow T$ satisfying $\mathrm{CH} = \mathsf{F}^{\mathrm{H}}(1^\lambda, pk, (\mathrm{CT}, M))$ and $(\mathrm{CT}, \mathrm{CH}, \mathrm{RP}) \in [\mathbf{HTR}_{\mathsf{ID},\lambda}(pk, sk)] \cap \mathbf{ACC}_{\mathsf{ID},\lambda}(pk)$. That is, a non-$\perp$ signature is an honest, accepting transcript in which the challenge is the hash of the commitment and message. We say that S has signing error $\nu\colon \mathbb{N} \to \mathbb{R}$ if $\Pr[\mathsf{S}^{\mathrm{H}}(1^\lambda, pk, sk, M) = \perp] \leq \nu(\lambda)$ for all $\lambda \in \mathbb{N}$, all $(pk, sk) \in [\mathsf{ID.Kg}(1^\lambda)]$ and all $M \in \{0,1\}^*$, where the probability is over the coins of S and H.

On the bottom left of Figure 4 is the canonical signing algorithm $\mathsf{SC}_{\mathsf{ID},\mathsf{F}}$. This is the classical choice, representing the usual, prescribed way to generate FS signatures. When ID has perfect correctness, $\mathsf{SC}_{\mathsf{ID},\mathsf{F}}$ has zero signing error. On the right is a signing with aborts algorithm $\mathsf{SA}_{\mathsf{ID},\mathsf{F},t}$ as per [23], where $t\colon \mathbb{N} \to \mathbb{N}$ is a polynomial. This may be used when ID has imperfect correctness. It tries to generate an honest, accepting transcript, returning $\perp$ if it fails after $t(\cdot)$ attempts.

$\underline{\mathsf{DS}_{\rm tr}.\mathsf{Sign}^{\rm H}(1^\lambda, pk, sk, M)}$	$\underline{\mathsf{DS}_{\rm tr}.\mathsf{V}^{\rm H}(1^\lambda, pk, M, \sigma)}$
$T \leftarrow\!\!\$\ \mathsf{S}^{\rm H}(1^\lambda, pk, sk, M)$ If $(T = \bot)$ then return $\bot$ $(\textsc{Ct}, \textsc{Ch}, \textsc{Rp}) \leftarrow T$ $\sigma \leftarrow (\textsc{Ct}, \textsc{Ch}, \textsc{Rp})$; Return σ	$(\textsc{Ct}, \textsc{Ch}, \textsc{Rp}) \leftarrow \sigma$ $d_0 \leftarrow \mathsf{ID.V}(1^\lambda, pk, \textsc{Ct}, \textsc{Ch}, \textsc{Rp})$ $d_1 \leftarrow (\textsc{Ch} = \mathsf{F}^{\rm H}(1^\lambda, pk, (\textsc{Ct}, M)))$ Return $(d_0 \wedge d_1)$

$\underline{\mathsf{DS}_{\rm ct}.\mathsf{Sign}^{\rm H}(1^\lambda, pk, sk, M)}$	$\underline{\mathsf{DS}_{\rm ct}.\mathsf{V}^{\rm H}(1^\lambda, pk, M, \sigma)}$
$T \leftarrow\!\!\$\ \mathsf{S}^{\rm H}(1^\lambda, pk, sk, M)$ If $(T = \bot)$ then return $\bot$ $(\textsc{Ct}, \textsc{Ch}, \textsc{Rp}) \leftarrow T$ $\sigma \leftarrow (\textsc{Ct}, \textsc{Rp})$; Return σ	$(\textsc{Ct}, \textsc{Rp}) \leftarrow \sigma$ $\textsc{Ch} \leftarrow \mathsf{F}^{\rm H}(1^\lambda, pk, (\textsc{Ct}, M))$ Return $\mathsf{ID.V}(1^\lambda, pk, \textsc{Ct}, \textsc{Ch}, \textsc{Rp})$

$\underline{\mathsf{DS}_{\rm ch}.\mathsf{Sign}^{\rm H}(1^\lambda, pk, sk, M)}$	$\underline{\mathsf{DS}_{\rm ch}.\mathsf{V}^{\rm H}(1^\lambda, pk, M, \sigma)}$
$T \leftarrow\!\!\$\ \mathsf{S}^{\rm H}(1^\lambda, pk, sk, M)$ If $(T = \bot)$ then return $\bot$ $(\textsc{Ct}, \textsc{Ch}, \textsc{Rp}) \leftarrow T$ $\sigma \leftarrow (\textsc{Ch}, \textsc{Rp})$; Return σ	$(\textsc{Ch}, \textsc{Rp}) \leftarrow \sigma$ $\textsc{Ct} \leftarrow \mathsf{ID.CR}(1^\lambda, pk, \textsc{Ch}, \textsc{Rp})$ If $(\textsc{Ct} = \bot)$ then return false Return $(\textsc{Ch} = \mathsf{F}^{\rm H}(1^\lambda, pk, (\textsc{Ct}, M)))$

$\underline{\text{Algorithm } \mathsf{SC}^{\rm H}_{\mathsf{ID},\mathsf{F}}(1^\lambda, pk, sk, M)}$	$\underline{\text{Algorithm } \mathsf{SA}^{\rm H}_{\mathsf{ID},\mathsf{F},t}(1^\lambda, pk, sk, M)}$
$(\textsc{Ct}, \textsc{St}) \leftarrow\!\!\$\ \mathsf{ID.Ct}(1^\lambda, pk)$ $\textsc{Ch} \leftarrow \mathsf{F}^{\rm H}(1^\lambda, pk, (\textsc{Ct}, M))$ $\textsc{Rp} \leftarrow \mathsf{ID.Rp}(1^\lambda, pk, sk, \textsc{Ch}, \textsc{St})$ Return $(\textsc{Ct}, \textsc{Ch}, \textsc{Rp})$	$d \leftarrow \mathsf{false}$; $i \leftarrow 0$ While $(d = \mathsf{false}$ and $i < t(\lambda))$ do: $\quad i \leftarrow i + 1$ $\quad (\textsc{Ct}, \textsc{St}) \leftarrow\!\!\$\ \mathsf{ID.Ct}(1^\lambda, pk)$ $\quad \textsc{Ch} \leftarrow \mathsf{F}^{\rm H}(1^\lambda, pk, (\textsc{Ct}, M))$ $\quad \textsc{Rp} \leftarrow \mathsf{ID.Rp}(1^\lambda, pk, sk, \textsc{Ch}, \textsc{St})$ $\quad d \leftarrow \mathsf{ID.V}(1^\lambda, pk, \textsc{Ct}, \textsc{Ch}, \textsc{Rp})$ If $(d = \mathsf{true})$ then return $(\textsc{Ct}, \textsc{Ch}, \textsc{Rp})$ Else return $\bot$

Fig. 4. Top three panels show signing and verifying algorithms of the signature schemes $\mathsf{DS}_{\rm tr}$, $\mathsf{DS}_{\rm ct}$ and $\mathsf{DS}_{\rm ch}$ obtained by applying the $\mathbf{gFS}_{\rm tr}, \mathbf{gFS}_{\rm ct}$ and $\mathbf{gFS}_{\rm ch}$ transforms, respectively, to identification scheme ID, hash function F and signing algorithm S. Bottom panel shows examples of signing algorithms.

If ID has correctness error a (non-zero) constant $\nu(\cdot) = \varepsilon < 1$, then setting $t(\lambda)$, to, say, $\lceil \log^2(\lambda) \cdot \log(1/\varepsilon) \rceil$ will result in $\mathsf{SA}_{\mathsf{ID},\mathsf{F},t}$ having negligible signing error in the case that F is a random oracle. For other choices of F, the correctness error of $\mathsf{SA}_{\mathsf{ID},\mathsf{F},t}$ would have to be evaluated directly (this seems to be somewhat glossed over in prior work) but for practical choices of F we expect it to still be about ν by the random oracle paradigm [11]. Our transforms will not pin down a particular way of generating signatures, but rather allow that to be specified by a signing algorithm S that they take as input. This allows our results to cover many different types of signing.

THE **gFS** TRANSFORMS. Let ID be an identification scheme, F a hash function compatible with it, and S a signing algorithm compatible with both. The $\mathbf{gFS}_{\rm tr}$ transform associates to $\mathsf{ID}, \mathsf{F}, \mathsf{S}$ the signature scheme $\mathsf{DS}_{\rm tr} = \mathbf{gFS}_{\rm tr}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$

whose algorithms are specified in the first panel in Figure 4. The $\mathbf{gFS}_{\mathrm{ct}}$ transform associates to $\mathsf{ID}, \mathsf{F}, \mathsf{S}$ the signature scheme $\mathsf{DS}_{\mathrm{ct}} = \mathbf{gFS}_{\mathrm{ct}}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$ whose algorithms are specified in the second panel in Figure 4. Assuming additionally that ID is commitment reproducible, and letting $\mathsf{ID.CR}$ be its commitment reproducing algorithm, the $\mathbf{gFS}_{\mathrm{ch}}$ transform associates to $\mathsf{ID}, \mathsf{F}, \mathsf{S}$ the signature scheme $\mathsf{DS}_{\mathrm{ch}} = \mathbf{gFS}_{\mathrm{ch}}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$ whose algorithms are specified in the third panel of Figure 4. Although this is not explicitly indicated in the code, note that in all cases, as per our general conventions, the signature verification algorithm returns $\perp$ if its input signature σ is $\perp$. The correctness error of a signature scheme $\mathsf{DS} = \mathbf{gFS}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$ given by one of our transforms is just the signing error of the signing algorithm S. So, for example, if ID has perfect correctness and $\mathsf{S} = \mathsf{SC}_{\mathsf{ID},\mathsf{F}}$, then DS has perfect correctness.

Attributes of the commitment reproducing algorithm. Security of the different variants of the FS transform will rely on different properties of commitment reproducible identification schemes that we now introduce. In the following let ID be a commitment reproducible identification scheme.

The strongest attribute is what we call *Perfect Uniqueness* (P-UNIQ). It asks that for all $\lambda \in \mathbb{N}$, all $(pk, sk) \in [\mathsf{ID.Kg}(1^\lambda)]$, all $\mathrm{Ch} \in \mathsf{ID.ChS}(\lambda)$ and all Ct, Rp that are not $\perp$ we have: $\mathsf{ID.V}(1^\lambda, pk, \mathrm{Ct}, \mathrm{Ch}, \mathrm{Rp}) = \mathsf{true}$ if and only if $\mathrm{Ct} = \mathsf{ID.CR}(1^\lambda, pk, \mathrm{Ch}, \mathrm{Rp})$. Figure 6 says the SND, CNS attributes we define next are implied by P-UNIQ, but strictly weaker than it.

We now introduce *soundness*. To understand it, we start with *Perfect Soundness* (P-SND). This asks that for all $\lambda \in \mathbb{N}$, all $(pk, sk) \in [\mathsf{ID.Kg}(1^\lambda)]$, all $\mathrm{Ch} \in \mathsf{ID.ChS}(\lambda)$ and all Rp we have: If $\mathrm{Ct} \leftarrow \mathsf{ID.CR}(1^\lambda, pk, \mathrm{Ch}, \mathrm{Rp})$ is not $\perp$ then $\mathsf{ID.V}(1^\lambda, pk, \mathrm{Ct}, \mathrm{Ch}, \mathrm{Rp}) = \mathsf{true}$. SND-security is a computational relaxation of this, asking that it be computationally hard to create a challenge and response where commitment reproducibility succeeds but the transcript is rejecting. This is formalized in game $\mathbf{G}^{\mathrm{snd}}_{\mathsf{ID},\mathcal{A}}(\lambda)$ in Figure 5. Via oracle New, the adversary can initialize a user (we are in the multi-user setting) and obtain not only its public key but also its secret key. It outputs a challenge $\mathrm{Ch} \in \mathsf{ID.ChS}(\lambda)$ and response Rp, as well as a pointer to some user $i \in \{1, \ldots, n\}$. It wins if the commitment reproducing algorithm, given $pk_i, \mathrm{Ch}, \mathrm{Rp}$, returns a non-$\perp$ value but the corresponding transcript is rejected by the verifier. Let $\mathbf{Adv}^{\mathrm{snd}}_{\mathsf{ID},\mathcal{A}}(\lambda) = \Pr[\mathbf{G}^{\mathrm{snd}}_{\mathsf{ID},\mathcal{A}}(\lambda)]$. We say that ID is SND-secure if the function $\mathbf{Adv}^{\mathrm{snd}}_{\mathsf{ID},\mathcal{A}}(\cdot)$ is negligible for every PT adversary $\mathcal{A}$.

We turn to *consistency*. Again, to understand it we start with *Perfect Consistency* (P-CNS). This asks that for all $\lambda \in \mathbb{N}$, all $(pk, sk) \in [\mathsf{ID.Kg}(1^\lambda)]$, all $\mathrm{Ch} \in \mathsf{ID.ChS}(\lambda)$ and all Ct, Rp we have: If $\mathrm{Ct} \neq \mathsf{ID.CR}(1^\lambda, pk, \mathrm{Ch}, \mathrm{Rp})$ then $\mathsf{ID.V}(1^\lambda, pk, \mathrm{Ct}, \mathrm{Ch}, \mathrm{Rp}) \neq \mathsf{true}$. CNS-security is a computational relaxation of this, asking that it be computationally hard to create an accepting transcript in which the commitment is different from the one given by the commitment reproducing algorithm. This is formalized using game $\mathbf{G}^{\mathrm{cns}}_{\mathsf{ID},\mathcal{A}}(\lambda)$ in Figure 5. Via oracle New, the adversary can initialize a user and obtain both its keys. It outputs $\mathrm{Ct}, \mathrm{Ch}, \mathrm{Rp}$ with $\mathrm{Ch} \in \mathsf{ID.ChS}(\lambda)$ and a pointer to some user $i \in \{1, \ldots, n\}$. It wins if the transcript is accepting but the commitment repro-

Game $\mathbf{G}^{\mathrm{snd}}_{\mathsf{ID},\mathcal{A}}(\lambda)$	Game $\mathbf{G}^{\mathrm{cns}}_{\mathsf{ID},\mathcal{A}}(\lambda)$
$n \leftarrow 0$	$n \leftarrow 0$
$(\mathrm{CH}, \mathrm{RP}, i) \leftarrow\!\!\$\ \mathcal{A}^{\mathrm{New}}(1^\lambda)$	$(\mathrm{CT}_1, \mathrm{CH}, \mathrm{RP}, i) \leftarrow\!\!\$\ \mathcal{A}^{\mathrm{New}}(1^\lambda)$
$\mathrm{CT} \leftarrow \mathsf{ID.CR}(1^\lambda, pk_i, \mathrm{CH}, \mathrm{RP})$	$\mathrm{CT}_0 \leftarrow \mathsf{ID.CR}(1^\lambda, pk_i, \mathrm{CH}, \mathrm{RP})$
$d \leftarrow \mathsf{ID.V}(1^\lambda, pk_i, \mathrm{CT}, \mathrm{CH}, \mathrm{RP})$	$d_1 \leftarrow \mathsf{ID.V}(1^\lambda, pk_i, \mathrm{CT}_1, \mathrm{CH}, \mathrm{RP})$
Return $(d = \mathsf{false}) \wedge (\mathrm{CT} \neq \bot)$	Return $(d_1 = \mathsf{true}) \wedge (\mathrm{CT}_0 \neq \mathrm{CT}_1)$
New()	New()
$n \leftarrow n + 1$; $(pk_n, sk_n) \leftarrow\!\!\$\ \mathsf{ID.Kg}(1^\lambda)$	$n \leftarrow n + 1$; $(pk_n, sk_n) \leftarrow\!\!\$\ \mathsf{ID.Kg}(1^\lambda)$
Return (pk_n, sk_n)	Return (pk_n, sk_n)

Fig. 5. Games defining soundness (SND-security) and consistency (CNS-security) of a commitment reproducible identification scheme ID.

ducing algorithm returns a commitment different from the one in the transcript. Let $\mathbf{Adv}^{\mathrm{cns}}_{\mathsf{ID},\mathcal{A}}(\lambda) = \Pr[\mathbf{G}^{\mathrm{cns}}_{\mathsf{ID},\mathcal{A}}(\lambda)]$. We say that ID is CNS-secure if the function $\mathbf{Adv}^{\mathrm{cns}}_{\mathsf{ID},\mathcal{A}}(\cdot)$ is negligible for every PT adversary $\mathcal{A}$.

For convenience of our reductions, the definitions of soundness and consistency are in the multi-user setting. A standard hybrid argument shows that single user security (captured as security relative to adversaries making only one call to New) implies multi-user security. This reduction is not tight, the advantage degrading linearly in the number of queries to New. When we say that the results in our paper are underlain by tight reductions we mean that the reductions in Theorems 2 and 4 are tight to the assumptions made in these theorems, which are the multi-user versions of SND and CNS, respectively.

Signature scheme relations. We give the formal result statements underlying the picture at the bottom of Figure 1. The proofs are in [4]. We start with whether uf-security of $\mathsf{DS}_{\mathrm{ct}}$ implies that of $\mathsf{DS}_{\mathrm{ch}}$. The following Proposition says that in general (meaning, with no conditions on the commitment reproducing algorithm other than completeness) the answer is "no." Theorem 2 will show that SND-security of ID suffices to make the answer "yes." For simplicity the Proposition sets the signing algorithm to the canonical one, but the Theorem holds for *all* signing algorithms.

Proposition 1. *Let* ID^* *be a commitment reproducible identification scheme and* F *a hash function compatible with* ID^*. *Assume signature scheme* $\mathsf{DS}^*_{\mathrm{ct}} = \mathbf{gFS}_{\mathrm{ct}}[\mathsf{ID}^*, \mathsf{F}, \mathsf{SC}_{\mathsf{ID}^*,\mathsf{F}}]$ *is uf-secure. Then there is a commitment reproducible identification scheme* ID *such that* F *is compatible with* ID *and (1)* $\mathsf{DS}_{\mathrm{ct}} = \mathbf{gFS}_{\mathrm{ct}}[\mathsf{ID}, \mathsf{F}, \mathsf{SC}_{\mathsf{ID},\mathsf{F}}]$ *is uf-secure but (2)* $\mathsf{DS}_{\mathrm{ch}} = \mathbf{gFS}_{\mathrm{ch}}[\mathsf{ID}, \mathsf{F}, \mathsf{SC}_{\mathsf{ID},\mathsf{F}}]$ *is not uf secure.*

If ID has the stronger property of being SND-secure, then uf-security of $\mathsf{DS}_{\mathrm{ct}}$ does transfer to $\mathsf{DS}_{\mathrm{ch}}$. Note that ID as constructed in the proof of Proposition 1 is *not* SND-secure, so there is no contradiction. Hence the Proposition can also be viewed as showing that the SND-security assumption is necessary for the following Theorem. For conciseness, the theorem statement is asymptotic, but it is underlain by a tight reduction explicitly stated and proved in [4].

Theorem 2. *Let* ID *be a commitment reproducible identification scheme,* F *a hash function compatible with* ID *and* S *a signing algorithm compatible with* ID, F*. Let* $\mathsf{DS}_{\rm ct} = \mathbf{gFS}_{\rm ct}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$ *and* $\mathsf{DS}_{\rm ch} = \mathbf{gFS}_{\rm ch}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$*. Assume* ID *is SND-secure and* $\mathsf{DS}_{\rm ct}$ *is uf-secure. Then* $\mathsf{DS}_{\rm ch}$ *is uf-secure.*

This result holds regardless of F, S, meaning no (extra) conditions are put on these, which means we cover both canonical and with-abort signing via the choices of S shown in Figure 4.

We turn to the converse, asking whether uf-security of $\mathsf{DS}_{\rm ch}$ implies that of $\mathsf{DS}_{\rm ct}$. Analogously to the above, Proposition 3 says that in general the answer is "no," and Theorem 4 says that it becomes "yes" assuming ID is CNS-secure.

Proposition 3. *Let* ID^* *be a commitment reproducible identification scheme and* F *a hash function compatible with* ID^**. Assume signature scheme* $\mathsf{DS}^*_{\rm ch} = \mathbf{gFS}_{\rm ch}[\mathsf{ID}^*, \mathsf{F}, \mathsf{SC}_{\mathsf{ID}^*,\mathsf{F}}]$ *is uf-secure. Then there is a commitment reproducible identification scheme* ID *such that* F *is compatible with* ID *and (1)* $\mathsf{DS}_{\rm ch} = \mathbf{gFS}_{\rm ch}[\mathsf{ID}, \mathsf{F}, \mathsf{SC}_{\mathsf{ID},\mathsf{F}}]$ *is uf-secure but (2)* $\mathsf{DS}_{\rm ct} = \mathbf{gFS}_{\rm ct}[\mathsf{ID}, \mathsf{F}, \mathsf{SC}_{\mathsf{ID},\mathsf{F}}]$ *is not uf secure.*

Theorem 4. *Let* ID *be a commitment reproducible identification scheme,* F *a hash function compatible with* ID *and* S *a signing algorithm compatible with* ID, F*. Let* $\mathsf{DS}_{\rm ct} = \mathbf{gFS}_{\rm ct}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$ *and* $\mathsf{DS}_{\rm ch} = \mathbf{gFS}_{\rm ch}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$*. Assume* ID *is CNS-secure and* $\mathsf{DS}_{\rm ch}$ *is uf-secure. Then* $\mathsf{DS}_{\rm ct}$ *is uf-secure.*

Recall that the interest of $\mathbf{gFS}_{\rm tr}$ is that the first proofs were for this variant [27]. However the following says it is equivalent in uf-security to $\mathbf{gFS}_{\rm ct}$.

Theorem 5. *Let* ID *be an identification scheme,* F *a hash function compatible with* ID *and* S *a signing algorithm compatible with* ID, F*. Let* $\mathsf{DS}_{\rm ct} = \mathbf{gFS}_{\rm ct}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$ *and* $\mathsf{DS}_{\rm tr} = \mathbf{gFS}_{\rm tr}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$*. Then* $\mathsf{DS}_{\rm ct}$ *is uf-secure if and only if* $\mathsf{DS}_{\rm tr}$ *is uf-secure.*

Identification relations. We have defined several attributes of commitment reproducing identification schemes: P-UNIQ, P-SND, SND, P-CNS, CNS. Figure 6 determines the relations between the five notions, in the style introduced by [7]. An arrow XX $\rightarrow$ YY is an implication: *every* commitment reproducible identification scheme that has property XX also has property YY. A barred arrow XX $\not\rightarrow$ YY is a separation: *there exists* a commitment reproducible identification scheme having property XX but not having property YY. Proofs of the relations in Figure 6 are in [4].

The picture shows a minimal set of implications and separations but determines the relation between *any* two nodes. For example, does P-CNS imply P-SND? No, because if it did we would get a path from P-CNS to SND, contradicting that shown separation.

What emanates from the relations? Recall we have seen that if $\mathsf{DS}_{\rm ct} = \mathbf{gFS}_{\rm ct}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$ and $\mathsf{DS}_{\rm ch} = \mathbf{gFS}_{\rm ch}[\mathsf{ID}, \mathsf{F}, \mathsf{S}]$ then SND suffices for uf-security of $\mathsf{DS}_{\rm ct}$ to imply that of $\mathsf{DS}_{\rm ch}$, and CNS suffices for the converse. Figure 6 says that P-UNIQ would also suffice for (both) these conclusions, but that SND, CNS are

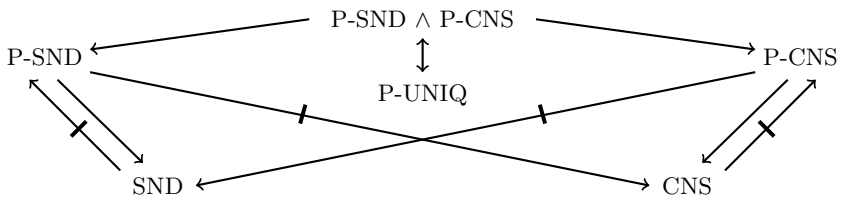

Fig. 6. Relations between security notions for commitment reproducible identification scheme. Arrows denote implications and barred arrows denote separations.

strictly weaker assumptions. It also says that SND, CNS are distinct; neither implies the other. In fact even P-SND does not imply CNS, and P-CNS does not imply SND. So the conditions required for uf-security to transfer across $\mathsf{DS}_{\mathrm{ch}}$ and $\mathsf{DS}_{\mathrm{ct}}$ are not symmetric.

Acknowledgments

The first and fourth authors were supported in part by Scott Klemmer and the CSE Undergraduate Summer Research Internship program at the Department of Computer Science and Engineering, University of California San Diego. The second author was supported in part by NSF grants CNS-1717640 and CNS-1526801, a gift from Microsoft corporation and ERC Project ERCC (FP7/615074). The third author was supported in part by NSF grant CNS-1528068. The second author thanks Tom Ristenpart for asking about the security of the different variants of Fiat-Shamir signatures. We thank the NordSec 2018 reviewers for their comments.

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. *EUROCRYPT 2002*.
2. M. Abdalla, P.-A. Fouque, V. Lyubashevsky, and M. Tibouchi. Tightly-secure signatures from lossy identification schemes. *EUROCRYPT 2012*.
3. J. Alwen, Y. Dodis, and D. Wichs. Leakage-resilient public-key cryptography in the bounded-retrieval model. *CRYPTO 2009*.
4. M. Backendal, M. Bellare, J. Sorrell, and J. Sun. The Fiat-Shamir zoo: Relating the security of different signature variants. Cryptology ePrint Archive, Report 2018/775.
5. M. Bellare, A. Boldyreva, and S. Micali. Public-key encryption in a multi-user setting: Security proofs and improvements. *EUROCRYPT 2000*.
6. M. Bellare and W. Dai. Defending against key exfiltration: Efficiency improvements for big-key cryptography via large-alphabet subkey prediction. *ACM CCS 17*.
7. M. Bellare, A. Desai, D. Pointcheval, and P. Rogaway. Relations among notions of security for public-key encryption schemes. *CRYPTO'98*.

8. M. Bellare, C. Namprempre, and G. Neven. Security proofs for identity-based identification and signature schemes. *Journal of Cryptology*, 22(1):1–61, Jan. 2009.
9. M. Bellare, B. Poettering, and D. Stebila. From identification to signatures, tightly: A framework and generic transforms. *ASIACRYPT 2016, Part II*.
10. M. Bellare, B. Poettering, and D. Stebila. Deterring certificate subversion: Efficient double-authentication-preventing signatures. *PKC 2017, Part II*.
11. M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. *ACM CCS 93*.
12. M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. *EUROCRYPT 2006*.
13. D. J. Bernstein, N. Duif, T. Lange, P. Schwabe, and B.-Y. Yang. High-speed high-security signatures. *CHES 2011*.
14. N. Bindel, S. Akleylek, E. Alkim, P. S. L. M. Barreto, J. Buchmann, E. Eaton, G. Gutoski, J. Kramer, P. Longa, H. Polat, J. E. Ricardini, and G. Zanon. qTESLA. Technical report, National Institute of Standards and Technology, 2017.
15. S. Brands. Untraceable off-line cash in wallets with observers. *CRYPTO'93*.
16. L. Ducas, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehle. CRYSTALS – dilithium: Digital signatures from module lattices. Cryptology ePrint Archive, Report 2017/633.
17. A. Fiat and A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. *CRYPTO'86*.
18. S. D. Galbraith, C. Petit, and J. Silva. Identification protocols and signature schemes based on supersingular isogeny problems. *ASIACRYPT 2017, Part I*.
19. 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.
20. L. C. Guillou and J.-J. Quisquater. A practical zero-knowledge protocol fitted to security microprocessor minimizing both trasmission and memory. *EUROCRYPT'88*.
21. E. Kiltz, D. Masny, and J. Pan. Optimal security proofs for signatures from identification schemes. *CRYPTO 2016, Part II*.
22. LANIX. Things that use Ed25519, Aug. 2018. `https://ianix.com/pub/curve25519-deployment.html`.
23. V. Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. *ASIACRYPT 2009*.
24. V. Lyubashevsky. Lattice signatures without trapdoors. *EUROCRYPT 2012*.
25. K. Ohta and T. Okamoto. On concrete security treatment of signatures derived from identification. *CRYPTO'98*.
26. T. Okamoto. Provably secure and practical identification schemes and corresponding signature schemes. *CRYPTO'92*.
27. D. Pointcheval and J. Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, 2000.
28. C.-P. Schnorr. Efficient signature generation by smart cards. *Journal of Cryptology*, 4(3):161–174, 1991.