

Lattice-based Signatures with Tight Adaptive Corruptions and More

Jiaxin Pan^{*1} 

Benedikt Wagner^{† 2} 

January 5, 2022

¹ Department of Mathematical Sciences,
NTNU - Norwegian University of Science and Technology, Trondheim, Norway
jiaxin.pan@ntnu.no

² CISPA Helmholtz Center for Information Security, Saarbrücken, Germany
benedikt.wagner@cispa.de

Abstract

We construct the *first* tightly secure signature schemes in the multi-user setting with adaptive corruptions from lattices. In stark contrast to the previous tight constructions whose security is solely based on number-theoretic assumptions, our schemes are based on the Learning with Errors (LWE) assumption which is supposed to be post-quantum secure. The security of our scheme is independent of the numbers of users and signing queries, and it is in the non-programmable random oracle model. Our LWE-based scheme is *compact*, namely, its signatures contain only a constant number of lattice vectors.

At the core of our construction are a new abstraction of the existing lossy identification (ID) schemes using dual-mode commitment schemes and a refinement of the framework by Diemert et al. (PKC 2021) which transforms a lossy ID scheme to a signature using sequential OR proofs. In combination, we obtain a tight generic construction of signatures from dual-mode commitments in the multi-user setting. Improving the work of Diemert et al., our new approach can be instantiated using not only the LWE assumption, but also an isogeny-based assumption. We stress that our LWE-based lossy ID scheme in the intermediate step uses a conceptually different idea than the previous lattice-based ones.

Of independent interest, we formally rule out the possibility that the aforementioned “ID-to-Signature” methodology can work tightly using parallel OR proofs. In addition to the results of Fischlin et al. (EUROCRYPT 2020), our impossibility result shows a qualitative difference between both forms of OR proofs in terms of tightness.

Keywords: Digital signatures, identification schemes, multi-user security, tightness, OR proofs, commitments, lattice, isogeny, impossibility result

1 Introduction

TIGHT SECURITY. Security of modern cryptographic constructions is established by security reductions. A reduction is an efficient algorithm \mathcal{R} that uses an efficient algorithm \mathcal{A} against the security of scheme X as a subroutine, and if \mathcal{A} can break the security of X , then \mathcal{R} can solve the computational problem Y . Thus, the hardness of Y implies the security of X . More precisely, we obtain $\epsilon_{\mathcal{A}}/t_{\mathcal{A}} \leq L \cdot \epsilon_{\mathcal{R}}/t_{\mathcal{R}}$, where \mathcal{A} runs in time $t_{\mathcal{A}}$ and has success probability $\epsilon_{\mathcal{A}}$, and \mathcal{R} runs in time $t_{\mathcal{R}}$ and has success probability $\epsilon_{\mathcal{R}}$. Here L is a polynomial in the security parameter λ , which we call the security loss. Asymptotically, any polynomial L is sufficient to show security. However, when we instantiate the scheme in a theoretically sound manner, the concrete L has impact on the setup of the system parameters. In particular, the smaller L is, the shorter the parameters will be. If L is a small constant, we call the reduction *tight* (e.g. [BBM00, BR09]). Many works (e.g. [CW13, BL16, GHKP18]) also consider a relaxed tightness notion, called “*almost tight*”, where L depends at most linearly on the security parameter λ . We do

*Supported by the Research Council of Norway under Project No. 324235.

†This work was done while the second author was a student at Karlsruhe Institute of Technology (Germany) and was doing an internship with the first author at NTNU (Norway).

not distinguish these two notions, but are precise about the security loss of our scheme in our security theorem and when we compare it with the related work.

SIGNATURES IN MULTI-USER SETTING. Digital Signatures play a central role in modern public-key cryptography. The standard security notion is unforgeability against chosen-message attacks [GMR88] (denoted by CMA security) which states that no efficient adversary can forge a signature on a new message after adaptively asking signatures for arbitrary messages. This is defined in a single-user setting where only one public key is involved. A seemingly more realistic notion is CMA security in the multi-user setting with adaptive corruptions (denoted by MU-CMA-Corr security). Here, adversary \mathcal{A} receives N public keys, can adaptively ask for signatures and additionally can corrupt some of the corresponding secret keys, and in the end it outputs a forgery for an uncorrupted user. This is also named MU-EUF-CMA^{corr} security in [BHJ⁺15, GJ18]. We note that there is a weaker notion of multi-user security considered in [KMP16, PR20] (MU-CMA security) where secret key corruptions are not allowed.

MU-CMA-Corr security is an interesting notion to consider. The most important reason is that MU-CMA-Corr security captures the actual security requirements of many applications that use digital signatures as a building block. A well-known example is authenticated key exchange (AKE) protocols which use signatures to authenticate protocol transcripts. Standard AKE security models (such as the Bellare-Rogaway [BR94] and Canetti-Krawczyk [CK01] models) are in multi-user settings and allow adversaries to corrupt signing keys of some honest users. In particular, the work of Bader et al. [BHJ⁺15] proposed the first tightly MU-CMA-Corr secure signature schemes and used it to construct the first tightly secure AKE protocol. The notion of MU-CMA-Corr has been used in many of its subsequent works [GJ18, LLGW20, JKRS21] for constructing more efficient AKE protocols, and the notion is also used to prove the tight security of real-world protocols [DJ20, DG21]. Tight security is of particular interest for these protocols, since they often have massive amount of users involved. Nevertheless, understanding and constructing efficient tightly MU-CMA-Corr secure signature schemes are fundamental research questions.

ON ACHIEVING TIGHT MU-CMA-Corr SECURITY. In general, CMA security can only *non-tightly* imply MU-CMA-Corr security by a guessing argument. The resulting reduction will lose a factor linear in the number of users, N . This is similar for the implication from MU-CMA to MU-CMA-Corr.

Many of the tightly secure signature schemes in the literature established their tightness in the weaker sense, namely, either tight CMA security (for instance, [BKPP15, GHKP18]) or tight MU-CMA security (for instance, [KMP16, PR20]). None of them will lead to a tightly MU-CMA-Corr secure scheme. Furthermore, Bader et al. [BJS16] even proved that tight MU-CMA-Corr is impossible to achieve if the signature satisfies certain properties. These properties are satisfied by most signature schemes, and thus constructing tightly MU-CMA-Corr secure signature schemes is very challenging.

To the best of our knowledge, signature schemes in [Bad14, BHJ⁺15, ABP19, GJ18, DGJL21, HJK⁺21] are the only exceptions with tight MU-CMA-Corr security. They all base their security on number-theoretic assumptions (such as the Diffie-Hellman assumption in pairing groups and φ -Hiding assumption), which leads to insecurity in the presence of a powerful quantum adversary. It is also worth mentioning that very recently Han et al. [HJK⁺21] identified a gap in the security proof of the compact and tightly MU-CMA-Corr-secure scheme in [BHJ⁺15] and closed this gap by following the blueprint of the pairing-based HIBE scheme in [LP19].

We highlight that the tight lattice-based signature schemes in [AFLT12, BKPP15, BL16] and isogeny-based scheme in [EKP20] are only in the single-user setting. It is not clear how to translate them tightly to the multi-user setting with adaptive corruptions. Hence, currently, there is no tightly MU-CMA-Corr secure signature scheme from post-quantum assumptions.

OUR GOAL AND ITS DIFFICULTIES. We aim at constructing *compact* lattice-based signature schemes with tight MU-CMA-Corr security. In this paper, “compact” means that the signature contains only a small constant number of lattice vectors and has size independent of the message length, which is in contrast to less efficient tree-based constructions.

As remarked above, there exist tight constructions of MU-CMA-Corr secure signature schemes. However, we argue why it is inherently difficult to extend them in realizing our goal:

- First, generic constructions in [Bad14, BHJ⁺15] and [ABP19, Section 9.2] require some extractability of the underlying proof system. Such a proof system is hard to construct in a compact and tightly secure manner using lattices. For instance, one can use the Unruh proof system [Unr15] that is

tightly secure and extractable, but its proof size is at least linear in the security parameter. This can only give us a scheme with linear-size signatures.

- Second, the tree-based construction from one-time signatures in [ABP19, Section 9.3] can give us a tight lattice-based construction, but it is not compact and has signature size linear in the message length.
- Third, in [DGJL21] a generic construction was proposed by transforming a lossy identification (ID) scheme [AFLT12] to a tightly MU-CMA-Corr secure signature scheme using the sequential OR proof technique [AOS02, FHJ20]. As pointed out by the authors, this transformation requires additional properties of the lossy ID scheme which are not obvious how to achieve using lattices.
- Last, the specific schemes in [GJ18, HJK⁺21] crucially rely on number-theoretic assumptions and the underlying algebraic structure. More precisely, [GJ18] requires the Decisional Diffie-Hellman (DDH) assumption and a proof system for the equality of discrete logarithms, and the compact scheme in [HJK⁺21] requires an algebraic MAC with affine structures.

1.1 Our Contributions

We construct the *first* compact lattice-based signature schemes with tight MU-CMA-Corr security in the random oracle model. Their security is based on the Learning with Errors (LWE) assumption, and their security loss is independent of the number of users and signing queries. Furthermore, our security proofs do not program a random oracle. We also give an instantiation of our approach in the isogeny setting to show its flexibility. Unfortunately, the resulting signature scheme in the isogeny setting is not compact.

We have three tight lattice-based schemes, and they are all constructed from our generic approach. One of them is almost tight, and the other two are fully tight. All three schemes have public key size and signature size independent of the message length. We note that our fully tight schemes (cf. Figs. 23 and 24) contain linearly (in λ) many lattice vectors in signatures, but independent of the message length. In Table 1 we compare the efficiency and concrete security of our schemes with some well-known efficient signature schemes in the random oracle model. Asymptotically, the signature size of our almost tight scheme is comparable to non-tight constructions, such as Lyubashevsky [Lyu12] and Ducas et al. [DDLL13], which require the rewinding technique. Due to the tightness of our scheme, it may have shorter signatures than these schemes. We stress that the main purpose of this work is taking the first theoretical step to study whether and how a tightly MU-CMA-Corr secure compact signature scheme from lattices is possible. We are optimistic that the efficiency of our schemes can be further improved.

Scheme	Assumption	Loss	$ \text{sk} $	$ \text{pk} $	$ \sigma $
GPV [GPV08]	SIS	N	T	M	mz
Lyu [Lyu12]	SIS	Q^N/Adv_A	mn	$n^2z + M$	$\omega(\log \lambda) + mz$
DDLL [DDLL13]	SIS	Q^N/Adv_A	M	$mn + M$	$m + n + mz$
AFLT [AFLT12]	RLWE	N	$2nz$	nz	$3nz$
KLS [KLS18]	MLWE	N	$2knz$	k^2nz	$3knz$
Ours (Fig. 8)	LWE	λ	$1 + T$	$4M$	$(4n + 2m)z$
Ours (Fig. 23)	LWE	1	$1 + T$	$2M$	$(2n^2 + 2nm)z$
Ours (Fig. 24)	LWE	1	$1 + T$	$2M$	$2n(M + T)$

Table 1: Overview of lattice-based signature schemes in the random oracle model. Here, Q denotes an upper bound on the number of signature and random oracle queries and λ is the security parameter. The security loss is up to constants and with respect to N -MU-CMA-Corr security. The modulus is denoted by $q = \text{poly}(n)$ and $M = n \cdot m \cdot \lceil \log q \rceil$ denotes the size of an $n \times m$ matrix, $m = \Theta(n \log q)$, T denotes the size of a trapdoor for such a matrix and z the size of an element in \mathbb{Z}_q .

Our schemes are constructed by a generic transformation that tightly turns a dual-mode commitment scheme into a MU-CMA-Corr secure signature. Our transformation contains two technical contributions, an abstraction of the existing lossy ID schemes and a refinement of the framework of Diemert et al. [DGJL21] which used the *sequential* OR proofs of Abe et al. [AOS02] and Fischlin et al. [FHJ20]. The




Figure 1: Overview of our construction. All implications are tight. New implications are marked with red, and new implications that implicitly exist in previous work are marked with blue. The assumption GADH is a generic assumption about group actions, capturing isogeny-based assumptions.

abstraction is a generic transformation from dual-mode commitment to lossy ID, and the existing lossy ID schemes [GQ90, CP93, AFLT12, ABP19] are concrete instantiations of our transformation. More importantly, this yields a new construction based on the LWE assumption using a conceptually new approach. Together with our refinement of the Diemert et al. framework, our tight lattice-based signature schemes are obtained.

We stress that our approach is more general than Diemert et al.. To show this, we implement our approach with isogenies. For readability, we present our scheme using the (general) Group Action Diffie-Hellman assumption, which captures the post-quantum secure isogeny-based assumption used in [EKP20, Sto12], Decisional CSIDH. We detail our technical approach and show how it improves the existing literature in Section 1.2. We will mostly focus on the lattice-based construction for simplicity.

LIMITATION OF PARALLEL OR PROOFS. Complementing these positive results, we show the advantage of sequential OR proofs by formally proving the limitation of its natural counterpart, parallel OR proofs of Cramer et al. [CDS94], in constructing tightly secure signatures. More precisely, we prove that it is impossible to tightly turn an ID scheme into a MU-CMA-Corr secure signature using parallel OR proofs Cramer et al., if the underlying ID scheme satisfies some mild properties. We note that these properties are satisfied by many ID schemes, including the DDH-based lossy ID scheme [CP93]. We establish this impossibility result using meta-reduction techniques [Cor02, BJLS16, ABP19]. We note that our impossibility result does not apply to more generic but less efficient OR-proof-based tight construction in [BHJ⁺15], since they use the OR-proof ideas in a different manner.

Our result is very different to the previous impossibility results [Cor02, KK12, HJK12, BJLS16] about tight signatures, and it enriches our understanding on constructing tight signature schemes. More precisely, Bader et al. [BJLS16] show that, if a signature scheme has signatures that are either unique or rerandomizable over the whole signature space, it will not have a tight reduction. Here we note that the work of Bader et al. [BJLS16] summarized results [Cor02, KK12, HJK12]. Clearly, signature schemes from parallel OR proofs are neither unique nor rerandomizable. Thus, their approach cannot be directly applied here, while our work is the *first* tightness impossibility result applicable to non-unique and non-rerandomizable signatures.

1.2 Technical Details

We provide more details about our generic construction of tightly MU-CMA-Corr secure signatures. Our generic construction has two steps: It first transforms a dual-mode commitment scheme to a lossy ID scheme, and then from a lossy ID scheme to a MU-CMA-Corr secure signature scheme via sequential OR proofs. Both steps are tight. Fig. 1 gives an figurative overview of this framework.

OUR STARTING POINT: THE DIEMERT ET AL. (DGJL) APPROACH [DGJL21]. The DGJL approach transforms a lossy ID scheme into a tightly MU-CMA-Corr secure signature scheme using sequential OR proofs. A lossy ID scheme is a canonical three-move ID scheme (or, equivalently, a Σ protocol [Dam10]). Additionally, a lossy ID scheme has two sets of public keys, lossy keys and normal ones. It requires that under a lossy public key even an unbounded adversary cannot impersonate an honest user. For tight MU-CMA-Corr security, the DGJL approach required that, given multiple keys of a lossy ID scheme, it is hard to tell whether they all are lossy or normal. This is a property that can be tightly satisfied by the random self-reducibility of DDH- and φ -Hiding-based schemes in [CP93, GQ90], but not the lattice-based ones.

It is the main reason why their approach cannot be implemented from lattices. We call this property multi-key lossiness. Our main technical goal is to find a lattice-based lossy ID scheme with tight multi-key lossiness.

FROM DUAL-MODE COMMITMENT TO LOSSY ID. We take a closer look at the existing lattice-based lossy ID schemes, and they are based on the Ring-LWE [AFLT12] and Module-LWE [KLS18] assumptions. To tightly achieve multi-key lossiness, we need the random self-reducibility (RSR) of these structured LWE assumptions. Unfortunately, it is not known how to rerandomize these structured LWE instances. We suppose this is inherent, since if the RSR was possible then the hardness of Ring-LWE would not depend on the number of samples in the current worst-case to average-case reduction [LPR10]. However, for plain LWE assumption the number of samples does not influence security [Reg05, Pei09, BLP+13, GMPW20], i.e. we have RSR. Hence, we want to construct a lossy ID scheme based on the (plain) LWE assumption. A natural direction is to take the idea of these Ring-LWE and Module-LWE schemes and implement them directly using the plain LWE assumption. We suppose this cannot work, since in these schemes the ring structure is crucial for proving lossiness¹.

Instead, our approach uses a dual-mode commitment scheme which can be constructed from the plain LWE assumption. Roughly speaking, a dual-mode commitment scheme has two indistinguishable modes, hiding and binding. In the hiding mode, there exists a (private) trapdoor that can open a commitment to any message. In the binding mode, a commitment can be opened to only one message, which is a statistical property and similar to public-key encryption.

Our high-level idea can be described in a simple manner: The commitment key is the public key of the lossy ID scheme. The hiding commitment key is the normal public key of our lossy ID scheme, and the binding commitment key is the lossy key. In the protocol, a prover P holds the commitment trapdoor and its first move to the verifier V is a random commitment. After that, V returns a random message and asks P to open the previous commitment to the given message. If P sends back a valid opening for that in the third move, V will accept.

The correctness is implied by the hiding mode of the commitment scheme. In the binding mode (which is the lossy mode of the ID scheme), a commitment can only be opened to only one message, and thus even an unbounded adversary cannot successfully complete the interaction, since our message space is exponentially large.

We modify the Regev encryption scheme [Reg05] to construct this dual-mode commitment scheme. In particular, we are able to show that multiple hiding commitment keys are tightly indistinguishable from the binding ones, which implies tight multi-key lossiness of the resulting ID scheme. Interestingly, the resulting lossy ID scheme is the first lattice-based lossy ID scheme without using the rejection sampling technique [Lyu12].

Moreover, we show that many well-known lossy ID schemes [GQ90, CP93, AFLT12, ABP19, EKP20] are obtained from dual-mode commitment schemes. In particular, we give a new analysis of the isogeny-based scheme in [EKP20] to show that it is tightly multi-key lossy. It will give us the first tightly MU-CMA-Corr secure signature scheme from isogenies. We remark that this scheme is non-compact, since it requires parallel repetitions for soundness of the underlying ID scheme.

FROM LOSSY ID TO SIGNATURES. Equipped with our lattice-based lossy ID scheme, we can transform it to a tightly MU-CMA-Corr secure signature scheme using sequential OR proofs. We note that this cannot be done using parallel OR proofs by our impossibility result.

Our transformation follows the blueprint of the DGJL framework, but we adapt it to be suitable for our ID schemes. An important modification is our transformation requires universal honest-verifier zero-knowledge (uHVZK) property of the underlying lossy ID, instead of injective simulators as in [DGJL21]. This is more natural, as lossy ID schemes from dual-mode commitments do not necessarily have an injective simulator, but uHVZK. Our work shows that injective simulator is not necessary for tight MU-CMA-Corr security, but uHVZK is enough. Further, in contrast to [DGJL21], we allow the lossy keys to be correlated, which is necessary for the analysis of the isogeny-based scheme. Another (minor) adaptation is to tolerate correctness errors. This is a property which lattice-based constructions always have. Thus, our refinements make it possible to instantiate the DGJL framework based on a wider class of assumptions.

¹A trivial solution to argue lossiness with plain LWE is to have an ID scheme with single bit challenges, but that will result in a non-compact scheme with linear-size signatures, since for such an ID scheme we need to repeat $O(\lambda)$ times to get soundness (where λ is the security parameter).

Similar to the DGJL framework, our security proof does not program the random oracle. Different to them, our resulting signature scheme does not have strong MU-CMA-Corr security, but it can be tightly turned into a strongly secure scheme using one-time signatures [LM18] and the known transformation [SPW07].

OPEN PROBLEMS. We leave further improving the efficiency of our schemes as an open problem. Random oracles used in our proofs are classical, and it is an interesting direction to extend our approach in the quantum random oracle model, or even without random oracles. We also leave constructing tight and compact signatures from isogenies as an open problem.

2 Preliminaries

We denote the security parameter by $\lambda \in \mathbb{N}$. All algorithms will get 1^λ implicitly as input. A probabilistic algorithm \mathcal{A} is said to be PPT (probabilistic polynomial time) if its running time $\mathbf{T}(\mathcal{A})$ can be bounded by a polynomial in its input size. We make use of standard asymptotic notation for positive functions such as ω and O . A function $\nu : \mathbb{N} \rightarrow \mathbb{R}$ is negligible in its input λ if $\nu \in \lambda^{-\omega(1)}$. The term $\text{negl}(\lambda)$ always denotes a negligible function. If a function ν is at least $1 - \text{negl}(\lambda)$, we say that it is overwhelming. If \mathcal{D} is a distribution, we write $x \leftarrow \mathcal{D}$ to state that x is sampled from \mathcal{D} . If S is a finite set, the notation $x \xleftarrow{\$} S$ states that x is sampled uniformly random from S . The statistical distance of distributions $\mathcal{D}_1, \mathcal{D}_2$ on support \mathcal{X} is defined as $\frac{1}{2} \sum_{x \in \mathcal{X}} |\Pr[\mathcal{D}_1 = x] - \Pr[\mathcal{D}_2 = x]|$. If it is negligible in λ , we say the distributions are statistically close. The notation $y \leftarrow \mathbf{A}(x)$ means that the variable y is assigned to the output of algorithm \mathbf{A} on input x . Sometimes we make the randomness used by an algorithm explicit by writing $y = \mathbf{A}(x; r)$ if $r \in \{0, 1\}^*$ is \mathbf{A} 's randomness. If we want to state that y is a possible output of \mathbf{A} on input x , we write $y \in \mathbf{A}(x)$. In all code-based security games, numerical values are assumed to be implicitly initialized as 0, sets and lists as \emptyset . If \mathbf{G} is a game, we write $\mathbf{G}_{\Pi}^{\mathcal{A}}(1^\lambda) \Rightarrow b$ to state that the game \mathbf{G} outputs $b \in \{0, 1\}$ considering the adversary \mathcal{A} and the scheme Π . Whenever we deal with statistically negligible terms, we denote them by Greek letters, e.g. $\varepsilon_{\mathcal{A}}$. For computationally negligible terms we use notation like $\text{Adv}_{\mathcal{A}, \Pi}^{\mathbf{G}}(\lambda)$. Throughout the paper, we always denote the number of users or keys in a scheme by N . We implicitly assume that it is bounded by a polynomial in the security parameter.

Matrices and (column) vectors are written in bold letters. The Euclidean norm of a vector \mathbf{v} is denoted by $\|\mathbf{v}\|$, and the spectral norm of a matrix \mathbf{A} is denoted by $s_1(\mathbf{A})$. By $[n] := \{1, \dots, n\}$ we denote the set of the first n natural numbers.

We present the standard background on lattices in Supplementary Material Section A.

COMMITMENT SCHEMES. A dual-mode commitment scheme is a commitment scheme with two indistinguishable key generation modes, inducing statistically binding and hiding commitments, respectively. Additionally, the latter mode outputs a trapdoor that allows to open commitments to arbitrary messages.

Definition 2.1 (Dual-Mode Commitment Scheme). A dual-mode $(\varepsilon_b, \varepsilon_t, N)$ -commitment scheme is a tuple of PPT algorithms $\text{CMT} = (\text{Setup}, \text{TSetup}, \text{Gen}, \text{TGen}, \text{Com}, \text{TCom}, \text{Open}, \text{TCol})$ with the following syntax:

- $\text{Setup}(1^\lambda)$ outputs global system parameters par . We assume that par implicitly defines sets $\mathcal{K}, \mathcal{M}, \mathcal{C}, \mathcal{D}$ of keys, messages, commitments and decommitments, respectively. All algorithms related to CMT take at least implicitly par as input.
- $\text{Gen}(\text{par}, 1^N)$ outputs N commitment keys $\text{ck}_1, \dots, \text{ck}_N \in \mathcal{K}$.
- $\text{Com}(\text{ck}, m)$ outputs a commitment $c \in \mathcal{C}$ and a decommitment $\text{dc} \in \mathcal{D}$.
- $\text{Open}(\text{ck}, m, \text{dc}, c)$ is deterministic and outputs $b \in \{0, 1\}$.
- TSetup has the same output types as Setup and additionally implicitly defines a set \mathcal{T} of trapdoors.
- $\text{TGen}(\text{par})$ outputs a commitment key $\text{ck} \in \mathcal{K}$ and a trapdoor $\text{td} \in \mathcal{T}$.
- $\text{TCom}(\text{ck}, \text{td})$ outputs a commitment $c \in \mathcal{C}$ and a state St .
- $\text{TCol}(\text{ck}, \text{td}, St, m)$ outputs $\text{dc}' \in \mathcal{D}$.

We say that CMT is ρ -complete if for all $\text{par} \in \text{TSetup}(1^\lambda)$, $(\text{ck}, \text{td}) \in \text{TGen}(\text{par})$, $m \in \mathcal{M}$ we have that

$$\Pr[\text{Open}(\text{ck}, m, \text{dc}, c) = 1 \mid (c, \text{dc}) \leftarrow \text{Com}(\text{ck}, m)] \geq \rho.$$

Finally, the following security properties should hold:

- **Key Indistinguishability:** The following advantage is negligible for all PPT algorithms \mathcal{A} :

$$\begin{aligned} & \text{Adv}_{\mathcal{A}, \text{CMT}}^{N\text{-keydist}}(\lambda) := \\ & \left| \Pr \left[\mathcal{A}(\text{par}, \text{ck}_1, \dots, \text{ck}_N) = 1 \mid \begin{array}{l} \text{par} \leftarrow \text{Setup}(1^\lambda), \\ (\text{ck}_1, \dots, \text{ck}_N) \leftarrow \text{Gen}(\text{par}, 1^N) \end{array} \right] \right. \\ & \left. - \Pr \left[\mathcal{A}(\text{par}, \text{ck}_1, \dots, \text{ck}_N) = 1 \mid \begin{array}{l} \text{par} \leftarrow \text{TSetup}(1^\lambda), \\ (\text{ck}_i, \text{td}_i) \leftarrow \text{TGen}(\text{par}), i \in [N] \end{array} \right] \right|. \end{aligned}$$

- ε_t -**Trapdoor Property:** For all $\text{par} \in \text{TSetup}(1^\lambda)$, $(\text{ck}, \text{td}) \in \text{TGen}(\text{par})$, $m \in \mathcal{M}$ the following distributions have statistical distance at most ε_t :

$$\{(c, m, \text{dc}) \mid (c, \text{dc}) \leftarrow \text{Com}(\text{ck}, m)\}$$

and

$$\{(c, m, \text{dc}) \mid (c, St) \leftarrow \text{TCom}(\text{ck}, \text{td}), \text{dc} \leftarrow \text{TCol}(\text{ck}, \text{td}, St, m)\}.$$

- (ε_b, N) -**Statistically Binding:** The following probability is at most ε_b :

$$\Pr \left[\exists i \in [N], c \in \mathcal{C}, m \neq m' \in \mathcal{M} : \begin{array}{l} \exists \text{dc} \in \mathcal{D} : \text{Open}(\text{ck}_i, m, \text{dc}, c) = 1 \\ \wedge \exists \text{dc}' \in \mathcal{D} : \text{Open}(\text{ck}_i, m', \text{dc}', c) = 1 \end{array} \right],$$

where the probability is taken over

$$\text{par} \leftarrow \text{Setup}(1^\lambda), (\text{ck}_1, \dots, \text{ck}_N) \leftarrow \text{Gen}(\text{par}, 1^N).$$

SIGNATURE SCHEMES. We define the standard notion of signature schemes and their security.

Definition 2.2 (Digital Signature Scheme). A signature scheme is a tuple of PPT algorithms $\text{SIG} = (\text{Setup}, \text{Gen}, \text{Sig}, \text{Ver})$, where

- $\text{Setup}(1^\lambda)$ outputs global system parameters par . We assume that par implicitly defines sets $\mathcal{K}_p, \mathcal{K}_s, \mathcal{M}, \mathcal{S}$ of public keys, secret keys, messages and signatures, respectively. All algorithms related to SIG take at least implicitly par as input.
- $\text{Gen}(\text{par})$ outputs public and secret key $(\text{pk}, \text{sk}) \in \mathcal{K}_p \times \mathcal{K}_s$.
- $\text{Sig}(\text{sk}, m)$ returns a signature $\sigma \in \mathcal{S}$.
- $\text{Ver}(\text{pk}, m, \sigma)$ is deterministic and returns $b \in \{0, 1\}$.

We say that SIG is ρ -complete, if for all $\text{par} \in \text{Setup}(1^\lambda)$, all $(\text{pk}, \text{sk}) \in \text{Gen}(\text{par})$, all $m \in \mathcal{M}$ we have

$$\Pr[\text{Ver}(\text{pk}, m, \sigma) = 1 \mid \sigma \leftarrow \text{Sig}(\text{sk}, m)] \geq \rho.$$

Definition 2.3 (Multi-User Security). Consider a signature scheme $\text{SIG} = (\text{Setup}, \text{Gen}, \text{Sig}, \text{Ver})$, let $N \in \mathbb{N}$ be a natural number and consider the game $N\text{-MU-CMA-Corr}$ given in Fig. 2. We say that SIG is $N\text{-MU-CMA-Corr}$ secure, if for every PPT adversary \mathcal{A} the following advantage is negligible in λ :

$$\text{Adv}_{\mathcal{A}, \text{SIG}}^{N\text{-MU-CMA-Corr}}(\lambda) := \Pr \left[N\text{-MU-CMA-Corr}_{\text{SIG}}^{\mathcal{A}}(\lambda) \Rightarrow 1 \right].$$

In addition, the notion $N\text{-MU-CMA}$ is defined similarly, but \mathcal{A} does not get access to the oracle KEY .

IDENTIFICATION SCHEMES. Here, we introduce identification schemes and their properties, where we extend the notions of [AFLT12, KLS18] to the multi-user setting.

<p>Game $N\text{-MU-CMA-Corr}_{\text{SIG}}^{\mathcal{A}}(\lambda)$</p> <pre style="margin: 0;"> 01 par \leftarrow Setup(1^λ) 02 for $i \in [N]$: $(\text{pk}_i, \text{sk}_i) \leftarrow$ Gen(par) 03 O := (SIG, KEY) 04 $(i^*, m^*, \sigma^*) \leftarrow \mathcal{A}^{\text{O}}(\text{par}, (\text{pk}_i)_{i=1}^N)$ 05 if $i^* \in \mathcal{L}_{id}$: return 0 06 if $\exists \sigma : (i^*, m^*, \sigma) \in \mathcal{L}_m$: return 0 07 return Ver($\text{pk}_{i^*}, m^*, \sigma^*$) </pre>	<p>Oracle KEY(i)</p> <pre style="margin: 0;"> 08 $\mathcal{L}_{id} := \mathcal{L}_{id} \cup \{i\}$ 09 return sk_i </pre> <p>Oracle SIG(i, m)</p> <pre style="margin: 0;"> 10 $\sigma \leftarrow$ Sig(sk_i, m) 11 $\mathcal{L}_m := \mathcal{L}_m \cup \{(i, m, \sigma)\}$ 12 return σ </pre>
---	---

Figure 2: The games **MU-CMA**, **MU-CMA-Corr** for a signature scheme **SIG** and an adversary \mathcal{A} . The shaded statement is only executed in game **MU-CMA-Corr**.

Definition 2.4 (Canonical Identification Scheme). A canonical identification scheme **ID** is defined as a tuple of PPT algorithms $\text{ID} := (\text{ISetup}, \text{IGen}, \text{P} := (\text{P}_1, \text{P}_2), \text{V})$, with the following properties:

- $\text{ISetup}(1^\lambda)$ outputs global system parameters **par**. We assume that **par** implicitly defines a set **ChSet**, the set of challenges and sets **CmtSet**, **RspSet**. All algorithms related to **ID** take at least implicitly **par** as input.
- $\text{IGen}(\text{par})$ returns public and secret key (pk, sk) .
- $\text{P} := (\text{P}_1, \text{P}_2)$ is split into two algorithms. $\text{P}_1(\text{sk})$ returns a commitment $\text{cmt} \in \text{CmtSet}$ and a state St ; $\text{P}_2(\text{sk}, \text{ch}, St)$ returns a response $\text{rs} \in \text{RspSet}$.
- $\text{V}(\text{pk}, \text{cmt}, \text{ch}, \text{rs})$ is deterministic and outputs $b \in \{0, 1\}$.

Given **ID** as above, we define transcript generation as follows:

<p>Alg Tran($\text{pk}, \text{sk}, \text{ch}$)</p> <pre style="margin: 0;"> 01 $(\text{cmt}, St) \leftarrow \text{P}_1(\text{sk}), \text{rs} \leftarrow \text{P}_2(\text{sk}, \text{ch}, St)$ 02 if $\text{rs} = \perp$: $(\text{cmt}, \text{ch}) = (\perp, \perp)$ 03 return $(\text{cmt}, \text{ch}, \text{rs})$ </pre>
--

We say that **ID** is ρ -complete, if for all $\text{par} \in \text{ISetup}(1^\lambda)$, all $(\text{pk}, \text{sk}) \in \text{IGen}(\text{par})$ we have

$$\Pr \left[\text{V}(\text{pk}, \text{cmt}, \text{ch}, \text{rs}) = 1 \mid \begin{array}{l} \text{ch} \xleftarrow{\$} \text{ChSet} \\ (\text{cmt}, \text{ch}, \text{rs}) \leftarrow \text{Tran}(\text{par}, \text{pk}, \text{sk}, \text{ch}) \end{array} \right] \geq \rho.$$

From now on, without loss of generality, we assume that **V** accepts an honestly generated transcript if and only if $\text{P}_2(\text{sk}, \text{ch}, St) \neq \perp$. This can be assumed as the algorithm P_2 can call **V** to check the transcript itself before returning **rs**.

For the following definitions, we let $\text{ID} = (\text{ISetup}, \text{IGen}, \text{P} = (\text{P}_1, \text{P}_2), \text{V})$ be a canonical identification scheme.

Definition 2.5 (Special Honest Verifier Zero-Knowledge). We say that **ID** is ε_{zk} -special honest verifier zero-knowledge (HVZK) if there is a PPT algorithm **Sim** such that for all $\text{par} \in \text{ISetup}(1^\lambda)$, all $(\text{pk}, \text{sk}) \in \text{IGen}(\text{par})$ the following distributions have statistical distance at most ε_{zk} :

$$\{(\text{cmt}, \text{ch}, \text{rs}) \leftarrow \text{Tran}(\text{pk}, \text{sk}, \text{ch}) \mid \text{ch} \xleftarrow{\$} \text{ChSet}\}$$

and

$$\{(\text{cmt}, \text{ch}, \text{rs}) \mid \text{ch} \xleftarrow{\$} \text{ChSet}, (\text{cmt}, \text{rs}) \leftarrow \text{Sim}(\text{pk}, \text{ch})\}.$$

We also introduce a slightly stronger version of HVZK, called universal special honest verifier zero-knowledge (uHVZK), where the distributions should be the same for every challenge. Clearly, uHVZK implies HVZK.

Definition 2.6 (Universal Special Honest Verifier Zero-Knowledge). We say that ID is ε_{zk} -universal special honest verifier zero-knowledge (uHVZK) if there is a PPT algorithm Sim such that for all $\text{par} \in \text{ISetup}(1^\lambda)$, all $(\text{pk}, \text{sk}) \in \text{IGen}(\text{par})$ and all $\text{ch} \in \text{ChSet}$ the following distributions have statistical distance at most ε_{zk} :

$$\{(\text{cmt}, \text{ch}, \text{rs}) \leftarrow \text{Tran}(\text{pk}, \text{sk}, \text{ch})\} \text{ and } \{(\text{cmt}, \text{ch}, \text{rs}) \mid (\text{cmt}, \text{rs}) \leftarrow \text{Sim}(\text{pk}, \text{ch})\}.$$

Definition 2.7 (Multi-Key Lossiness). Let N be a natural number. We say that ID is (ε_{mkl}, N) -multi-key lossy, if there exists a PPT algorithm LIGen which takes the number of users 1^N as input and returns system parameters par and public keys $\text{pk}_1, \dots, \text{pk}_N$ such that the following holds:

- For every PPT algorithm \mathcal{D} , the following advantage is negligible in λ :

$$\begin{aligned} \text{Adv}_{\mathcal{D}, \text{ID}}^{N\text{-keydist}}(\lambda) := & \\ & \left| \Pr \left[\mathcal{D}(\text{par}, \text{pk}_1, \dots, \text{pk}_N) = 1 \mid \begin{array}{l} \text{par} \leftarrow \text{ISetup}(1^\lambda) \\ (\text{pk}_i, \text{sk}_i) \leftarrow \text{IGen}(\text{par}), i \in [N] \end{array} \right] \right. \\ & \left. - \Pr \left[\mathcal{D}(\text{par}, \text{pk}_1, \dots, \text{pk}_N) = 1 \mid (\text{par}, \text{pk}_1, \dots, \text{pk}_N) \leftarrow \text{LIGen}(1^N) \right] \right|. \end{aligned}$$

- The following inequality holds:

$$\mathbb{E} \left[\max_{i \in [N]} \max_{\text{cmt}} \Pr [\exists \text{rs} \in \text{RspSet} : \text{V}(\text{pk}_i, \text{cmt}, \text{ch}, \text{rs}) = 1] \right] \leq \varepsilon_{mkl},$$

where we take the expectation, maximum and probability over

$$(\text{par}, \text{pk}_1, \dots, \text{pk}_N) \leftarrow \text{LIGen}(1^N), \text{cmt} \in \text{CmtSet}, \text{ch} \xleftarrow{\$} \text{ChSet},$$

respectively. That is, if the keys are generated in this lossy way, for every unbounded adversary the advantage of successfully completing the protocol with respect to any user is bounded by ε_{mkl} .

Note that N -multi-key lossiness for $N = 1$ is just lossiness as defined in [AFLT12].

Remark 2.8 (Correlation of Lossy Keys). Note that in our definition of multi-key lossiness, we define one algorithm that outputs N lossy keys, whereas the definition in [DGJL21] is with regards to N keys that are generated via N independent invocations of the lossy key generator. We claim that our definition is more general, as it also captures the possibility that the N lossy keys are somehow correlated. As long as the expectation in our definition is bounded, this correlation is not a problem. In fact, in some cases it is only possible to tightly achieve key indistinguishability if the lossy keys are correlated, see our instantiation from group actions in Supplementary Material Section E.

3 Tight Signatures from Sequential OR Proofs, revisited

In this section we will generically construct a signature scheme with tight security in presence of adaptive corruptions. First, we show that sequential OR proofs can be used to construct signatures with this strong form of security from lossy identification schemes. Then, we introduce a new generic construction of lossy identification schemes from dual-mode commitments.

3.1 Generic Construction of Signatures in the Multi-User Setting

Let $\text{ID} := (\text{ISetup}, \text{IGen}, \text{P} := (\text{P}_1, \text{P}_2), \text{V})$ be a canonical identification scheme with challenge set ChSet . We use $\ell \in \mathbb{N}$ to model multiple attempts to compute a signature for schemes with non-perfect completeness. Assuming that ID is uHVZK, we construct a signature scheme $\text{SIG}_s[\text{ID}, \text{H}, \ell]$ with random oracle $\text{H} : \{0, 1\}^* \rightarrow \text{ChSet}$ and message space $\{0, 1\}^*$ using the sequential OR proof technique as defined in Fig. 3.

Intuitively, in the sequential OR proof signature, the challenge of one instance is computed as the hash of the commitment of the other instance. To break the circularity, the HVZK simulator is used on the instance for which the signer does not know a secret key. Note that the construction is a combination of the constructions in [AFLT12, FHJ20], in a sense that we combine the sequential OR proof from [FHJ20] with the lossy identification framework and the repetition as in [AFLT12].

<u>Alg Gen(par)</u>	<u>Alg Sig(sk, m)</u>
01 $(pk_0, sk_0) \leftarrow \text{IGen}(\text{par})$	11 $ctr := 0$
02 $(pk_1, sk_1) \leftarrow \text{IGen}(\text{par})$	12 while $ctr \leq \ell \wedge (rs_0 = \perp \vee rs_1 = \perp)$:
03 $b \xleftarrow{\$} \{0, 1\}, sk := (b, sk_b)$	13 $ctr := ctr + 1$
04 $pk := (pk_0, pk_1)$	14 $(cmt_b, St_b) \leftarrow P_1(sk_b)$
05 return (pk, sk)	15 $ch_{1-b} \leftarrow H(b, pk, cmt_b, m)$
<u>Alg Ver(pk, m, σ)</u>	16 $(cmt_{1-b}, rs_{1-b}) \leftarrow \text{Sim}(pk_{1-b}, ch_{1-b})$
06 $ch_1 \leftarrow H(0, pk, cmt_0, m)$	17 $ch_b \leftarrow H(1-b, pk, cmt_{1-b}, m)$
07 $ch_0 \leftarrow H(1, pk, cmt_1, m)$	18 $rs_b \leftarrow P_2(sk_b, ch_b, St_b)$
08 $v_0 \leftarrow V(pk_0, cmt_0, ch_0, rs_0)$	19 if $rs_0 = \perp \vee rs_1 = \perp$: return \perp
09 $v_1 \leftarrow V(pk_1, cmt_1, ch_1, rs_1)$	20 return $\sigma := (cmt_0, cmt_1, rs_0, rs_1)$
10 return $(v_0 \wedge v_1)$	

Figure 3: The signature scheme $\text{SIG}_s[\text{ID}, H, \ell] = (\text{Setup}, \text{Gen}, \text{Sig}, \text{Ver})$ for a canonical identification scheme $\text{ID} := (\text{ISetup}, \text{IGen}, P := (P_1, P_2), V)$ with HVZK simulator Sim , where $\text{Setup} := \text{ISetup}$.

Lemma 3.1 (Completeness). *Let $\text{ID} := (\text{ISetup}, \text{IGen}, P, V)$ be a ε_{zk} -HVZK canonical identification scheme with HVZK simulator Sim . Let $\ell \in \mathbb{N}$. If ID is ρ -complete, then $\text{SIG}_s[\text{ID}, H, \ell]$ is ρ' -complete, where $\rho' \geq 1 - (1 - \rho^2 - \rho\varepsilon_{zk})^\ell$.*

Proof. Consider an honest signature $\sigma \leftarrow \text{Sig}(sk, m)$. Note that by the definition of completeness of identification schemes, $\text{Ver}(pk, m, \sigma)$ will accept if and only if there is a $i \in [\ell]$, such that in the i -th iteration of the loop in the signing algorithm, $rs_0 \neq \perp$ and $rs_1 \neq \perp$. Call these event E_i^0, E_i^1 . Set $E_i := E_i^0 \cap E_i^1$. Now consider an arbitrary iteration and the calls $(cmt_{1-b}, rs_{1-b}) \leftarrow \text{Sim}(pk_{1-b}, ch_{1-b})$ and $rs_b \leftarrow P_2(sk_b, ch_b, St_b)$. In the random oracle model, these are independent. By definition of HVZK, the probability that $rs_{1-b} \neq \perp$ is ε_{zk} -close to the probability that $rs_b \neq \perp$. Then, as all iterations are independent, we can bound the probability that after ℓ iterations, no valid signature was found by

$$\prod_{i=1}^{\ell} \Pr[\neg E_i] = \prod_{i=1}^{\ell} (1 - \Pr[E_i^0 \cap E_i^1]) \leq (1 - \rho \cdot (\rho + \varepsilon_{zk}))^\ell = (1 - \rho^2 - \rho\varepsilon_{zk})^\ell.$$

□

Theorem 3.2 (Security). *Let ID be a canonical identification scheme. If ID is ε_{zk} -uHVZK and (ε_{mkl}, N) -multi-key lossy for negligible $\varepsilon_{zk}, \varepsilon_{mkl}$, then $\text{SIG}_s[\text{ID}, H, \ell]$ is N -MU-CMA-Corr secure, with a tight reduction. More precisely, for any adversary \mathcal{A} making at most Q_S signing queries, Q_C secret key queries and Q_H hash queries (including the indirect ones induced by signing queries), there exists an adversary \mathcal{D} such that $\mathbf{T}(\mathcal{D}) \approx \mathbf{T}(\mathcal{A})$ and*

$$\text{Adv}_{\mathcal{A}, \text{SIG}_s[\text{ID}, H, \ell]}^{N\text{-MU-CMA-Corr}}(\lambda) \leq 2 \cdot \text{Adv}_{\mathcal{D}, \text{ID}}^{N\text{-keydist}}(\lambda) + 2 \cdot (Q_H + 2)^2 \cdot \varepsilon_{mkl} + 3 \cdot \ell \cdot Q_S \cdot \varepsilon_{zk}.$$

Proof. Let $\text{SIG} := \text{SIG}_s[\text{ID}, H, \ell]$. We prove the statement via a sequence of games \mathbf{G}_i . In every game i , we denote the advantage of an adversary \mathcal{A} as $\text{Adv}_i(\mathcal{A}) := \Pr[\mathbf{G}_i^{\mathcal{A}} \Rightarrow 1]$. Games $\mathbf{G}_1 - \mathbf{G}_4$ are formally presented in Fig. 4. Let \mathcal{A} be an adversary in the game $\text{MU-CMA-Corr}_{\text{SIG}}^{\mathcal{A}}(\lambda)$ making Q_H, Q_S, Q_C hash, signing and corruption queries, respectively. Game \mathbf{G}_0 is the original game $\text{MU-CMA-Corr}_{\text{SIG}}^{\mathcal{A}}(\lambda)$.

Game \mathbf{G}_1 : In game \mathbf{G}_1 , the adversary only wins with a forgery $(cmt_0^*, cmt_1^*, rs_0^*, rs_1^*)$ for user i^* and message m^* if the original winning conditions are satisfied and the adversary queries both $(0, pk_{i^*}, cmt_0^*, m^*)$ and $(1, pk_{i^*}, cmt_1^*, m^*)$ to the random oracle. We note that for every adversary \mathcal{A} winning \mathbf{G}_0 , we can build an adversary \mathcal{B} winning \mathbf{G}_1 with the same advantage, making $Q_H + 2$ hash queries. Namely, \mathcal{B} just simulates \mathcal{A} and before forwarding \mathcal{A} 's forgery to its own game, \mathcal{B} makes the two required queries. Hence it follows that

$$\text{Adv}_{\mathcal{A}, \text{SIG}}^{N\text{-MU-CMA-Corr}}(\lambda) = \text{Adv}_0(\mathcal{A}) = \text{Adv}_1(\mathcal{B}).$$

Game \mathbf{G}_2 : Game \mathbf{G}_2 is as \mathbf{G}_1 , but to answer signing queries for user i the game uses both secret keys $\text{sk}_{i,0}, \text{sk}_{i,1}$ instead of using the uHVZK simulator Sim . Note that we can transition from \mathbf{G}_1 to \mathbf{G}_2 using $Q_S \cdot \ell$ intermediate games $\mathbf{G}_{1,j,j'}$, where this change is applied for the first j signing queries in the first j' iterations of the loop. Each transition only changes the advantage by a negligible amount, which follows from the statistical uHVZK property of ID. In particular, even if a random oracle value $\text{ch}_{1-b} = \text{H}(b, \text{pk}, \text{cmt}_b, \text{m})$ for which we want to apply uHVZK is defined yet, the universal HVZK notion implies that the distribution of signatures is close to the one in the previous game. We have

$$|\text{Adv}_1(\mathcal{B}) - \text{Adv}_2(\mathcal{B})| \leq \ell \cdot Q_S \cdot \varepsilon_{\text{zk}}.$$

Note that in \mathbf{G}_2 a successful adversary \mathcal{B} does not learn anything about b_{i^*} anymore, where i^* is the user of \mathcal{B} 's final forgery. This is because \mathcal{B} is not allowed to query the key of identity i^* and all signature queries for identity i^* are independent of b_{i^*} , as the b_{i^*} and the $1 - b_{i^*}$ part of the signature are both computed using the corresponding secret key.

Game \mathbf{G}_3 : In \mathbf{G}_3 we add the additional winning condition that \mathcal{B} 's hash query $(1 - b_{i^*}, \text{pk}_{i^*}, \text{cmt}_{1-b_{i^*}}^*, \text{m}^*)$ occurs before any query $(b_{i^*}, \text{pk}_{i^*}, \text{cmt}_{b_{i^*}}^*, \text{m}^*)$. As \mathcal{B} has no information about b_{i^*} and is required to query both, we have

$$\text{Adv}_2(\mathcal{B}) = 2 \cdot \text{Adv}_3(\mathcal{B}).$$

Game \mathbf{G}_4 : From \mathbf{G}_3 to \mathbf{G}_4 we reverse the changes we made from \mathbf{G}_1 to \mathbf{G}_2 . Namely, signing queries will be answered as in the original signing algorithm, by using Sim for the $1 - b_i$ instance. As before, uHVZK of ID implies that

$$|\text{Adv}_3(\mathcal{B}) - \text{Adv}_4(\mathcal{B})| \leq \ell \cdot Q_S \cdot \varepsilon_{\text{zk}}.$$

Game \mathbf{G}_5 : In game \mathbf{G}_5 , we change the way the parameters par and keys $\text{pk}_{i,1-b_i}$ are generated to

$$(\text{par}, \text{pk}_{1,1-b_1}, \dots, \text{pk}_{N,1-b_N}) \leftarrow \text{LIGen}(1^N).$$

Note that this is possible as we do not need the secret keys $\text{sk}_{i,1-b_i}$ here. Via a direct reduction \mathcal{D} that obtains $(\text{par}, \text{pk}_{1,1-b_1}, \dots, \text{pk}_{N,1-b_N})$ as input we can show

$$|\text{Adv}_4(\mathcal{B}) - \text{Adv}_5(\mathcal{B})| \leq \text{Adv}_{\mathcal{D}, \text{ID}}^{N\text{-keydist}}(\lambda).$$

That is, after getting the input, \mathcal{D} constitutes the rest of the simulation as in \mathbf{G}_4 and \mathbf{G}_5 and finally outputs 1 iff \mathcal{B} wins. In combination, we get

$$\left| \text{Adv}_{\mathcal{A}, \text{SIG}}^{N\text{-MU-CMA-Corr}}(\lambda) - 2 \cdot \text{Adv}_5(\mathcal{B}) \right| \leq 3 \cdot \ell \cdot Q_S \cdot \varepsilon_{\text{zk}} + 2 \cdot \text{Adv}_{\mathcal{D}, \text{ID}}^{N\text{-keydist}}(\lambda) \quad (1)$$

Finally, we will bound $\text{Adv}_5(\mathcal{B})$ via a statistical argument using multi-key lossiness. Note that due to the changes we made in \mathbf{G}_3 , \mathcal{B} is now restricted to submitting $\text{cmt}_{1-b_{i^*}}^*$ before learning the random challenge $\text{ch}_{1-b_{i^*}}^* \leftarrow \text{H}(b_{i^*}, \text{pk}_{i^*}, \text{cmt}_{b_{i^*}}^*, \text{m}^*)$ for some user i^* . To finally win, \mathcal{B} has to find $\text{rs}_{1-b_{i^*}}^*$ such that $\text{V}(\text{pk}_{i^*,1-b_{i^*}}, \text{cmt}_{b_{i^*}}^*, \text{ch}_{1-b_{i^*}}^*, \text{rs}_{1-b_{i^*}}^*) = 1$. As all the $\text{pk}_{i,1-b_i}$ are generated via $\text{LIGen}(1^N)$, multi-key lossiness tells us that \mathcal{B} 's advantage can be bounded: \mathcal{B} can try at most $Q_H + 2$ pairs for i^* , $\text{cmt}_{1-b_{i^*}}^*$ and will receive at most $Q_H + 2$ different challenges from the random oracle. For any fixed parameters, keys, $i \in [N]$, $\text{cmt} \in \text{CmtSet}$, denote the probability that \mathcal{B} can forge with these parameters and keys, for user $i^* = i$ and using $\text{cmt}_{1-b_{i^*}}^* = \text{cmt}$ by $p_{i,\text{cmt}}$. A union bound yields:

$$\begin{aligned} p_{i,\text{cmt}} &\leq \Pr_{\text{ch}_j} [\exists j \in [Q_H + 2] : \exists \text{rs} \in \text{RspSet} : \text{V}(\text{pk}_{i,1-b_i}, \text{cmt}, \text{ch}_j, \text{rs}) = 1] \\ &\leq (Q_H + 2) \cdot \Pr_{\text{ch}} [\exists \text{rs} \in \text{RspSet} : \text{V}(\text{pk}_{i,1-b_i}, \text{cmt}, \text{ch}, \text{rs}) = 1], \end{aligned}$$

where we take the first probability over $\text{ch}_j \xleftarrow{\$} \text{ChSet}$ for all $j \in [Q_H + 2]$ and the second over $\text{ch} \xleftarrow{\$} \text{ChSet}$. Further, for fixed keys and parameters, we can upper bound the advantage of the adversary by replacing each try $i^*, \text{cmt}_{1-b_{i^*}}^*$ by the optimal try. Averaging over the keys and parameters yields

$$\begin{aligned} \text{Adv}_5(\mathcal{B}) &\leq \mathbb{E} \left[(Q_H + 2) \cdot \max_{i \in [N]} \max_{\text{cmt}} p_{i, \text{cmt}} \right] \leq (Q_H + 2) \cdot \mathbb{E} \left[\max_{i \in [N]} \max_{\text{cmt}} p_{i, \text{cmt}} \right] \\ &\leq (Q_H + 2)^2 \cdot \mathbb{E} \left[\max_{i \in [N]} \max_{\text{cmt}} \Pr_{\text{ch}} [\exists \text{rs} \in \text{RspSet} : \text{V}(\text{pk}_{i, 1-b_i}, \text{cmt}, \text{ch}, \text{rs}) = 1] \right] \\ &\leq (Q_H + 2)^2 \cdot \varepsilon_{\text{mkl}}, \end{aligned}$$

where the inner maximum is over $\text{cmt} \in \text{CmtSet}$ and we take the probability over $\text{ch} \xleftarrow{\$} \text{ChSet}$. We can now conclude by using Eq. (1) with

$$\text{Adv}_{\mathcal{A}, \text{SIG}}^{N\text{-MU-CMA-Corr}}(\lambda) \leq 2 \cdot (Q_H + 2)^2 \cdot \varepsilon_{\text{mkl}} + 3 \cdot \ell \cdot Q_S \cdot \varepsilon_{\text{zk}} + 2\text{Adv}_{\mathcal{D}, \text{ID}}^{N\text{-keydist}}(\lambda).$$

Game $\mathbf{G}_1, \mathbf{G}_2, \mathbf{G}_3, \mathbf{G}_4$	Oracle $\text{SIG}(i, m)$
01 $\text{par} \leftarrow \text{Setup}(1^\lambda)$	24 $\text{ctr} := 0$
02 for $i \in [N]$:	25 while $\text{ctr} \leq \ell \wedge (\text{rs}_0 = \perp \vee \text{rs}_1 = \perp)$:
03 $(\text{pk}_{i,0}, \text{sk}_{i,0}) \leftarrow \text{IGen}(\text{par})$	26 $\text{ctr} := \text{ctr} + 1$
04 $(\text{pk}_{i,1}, \text{sk}_{i,1}) \leftarrow \text{IGen}(\text{par})$	27 $(\text{cmt}_{b_i}, \text{St}_{b_i}) \leftarrow \text{P}_1(\text{sk}_{i,b_i})$
05 $b_i \xleftarrow{\$} \{0, 1\}$	28 $\text{ch}_{1-b_i} \leftarrow \text{H}(b_i, \text{pk}_i, \text{cmt}_{b_i}, m)$
06 $\text{pk}_i := (\text{pk}_{i,0}, \text{pk}_{i,1})$	29 $(\text{cmt}_{1-b_i}, \text{rs}_{1-b_i})$
07 $\text{sk}_i := (b_i, \text{sk}_{i,b_i})$	$\leftarrow \text{Sim}(\text{pk}_{i,1-b_i}, \text{ch}_{1-b_i})$
08 $\text{O} := (\text{H}, \text{SIG}, \text{KEY})$	30 $(\text{cmt}_{1-b_i}, \text{St}_{1-b_i}) \leftarrow \text{P}_1(\text{sk}_{i,1-b_i})$
09 $(i^*, m^*, \sigma^*) \leftarrow \mathcal{B}^{\text{O}}(\text{par}, (\text{pk}_i)_{i=1}^N)$	31 $\text{ch}_{b_i} \leftarrow \text{H}(1-b_i, \text{pk}_i, \text{cmt}_{1-b_i}, m)$
10 if $i^* \in \mathcal{L}_{\text{id}}$: return 0	32 $\text{rs}_{b_i} \leftarrow \text{P}_2(\text{sk}_{i,b_i}, \text{ch}_{b_i}, \text{St}_{b_i})$
11 if $\exists \sigma : (i^*, m^*, \sigma) \in \mathcal{L}_m$:	33 $\text{rs}_{1-b_i} \leftarrow \text{P}_2(\text{sk}_{i,1-b_i}, \text{ch}_{1-b_i}, \text{St}_{1-b_i})$
12 return 0	34 if $\text{rs}_0 = \perp \vee \text{rs}_1 = \perp$: return \perp
13 let $\sigma^* = (\text{cmt}_0^*, \text{cmt}_1^*, \text{rs}_0^*, \text{rs}_1^*)$	35 $\sigma := (\text{cmt}_0, \text{cmt}_1, \text{rs}_0, \text{rs}_1)$
14 if $\mathcal{L}_h[0, \text{pk}_{i^*}, \text{cmt}_0^*, m^*] = \perp$:	36 $\mathcal{L}_m := \mathcal{L}_m \cup \{(i, m, \sigma)\}$
15 return 0	37 return σ
16 if $\mathcal{L}_h[1, \text{pk}_{i^*}, \text{cmt}_1^*, m^*] = \perp$:	Oracle $\text{H}(b, \text{pk}, \text{cmt}, m)$
17 return 0	38 if $\mathcal{L}_h[b, \text{pk}, \text{cmt}, m] = \perp$:
18 $(t_0, h_0) := \mathcal{L}_h[0, \text{pk}_{i^*}, \text{cmt}_0^*, m^*]$	39 $h \xleftarrow{\$} \text{ChSet}$
19 $(t_1, h_1) := \mathcal{L}_h[1, \text{pk}_{i^*}, \text{cmt}_1^*, m^*]$	40 $t := t + 1$
20 if $(t_{1-b_{i^*}} > t_{b_{i^*}})$: return 0	41 $\mathcal{L}_h[b, \text{pk}, \text{cmt}, m] := (t, h)$
21 return $\text{Ver}(\text{pk}_{i^*}, m^*, \sigma^*)$	42 return h
Oracle $\text{KEY}(i)$	43 $(t', h) := \mathcal{L}_h[b, \text{pk}, \text{cmt}, m]$
22 $\mathcal{L}_{\text{id}} := \mathcal{L}_{\text{id}} \cup \{i\}$	44 return h
23 return sk_i	

Figure 4: Games $\mathbf{G}_1, \mathbf{G}_2, \mathbf{G}_3, \mathbf{G}_4$ for adversary \mathcal{B} in the proof of MU-CMA-Corr security of $\text{SIG}_s[\text{ID}, \text{H}, \ell]$. Shaded statements are executed in \mathbf{G}_3 and \mathbf{G}_4 , statements in boxes are executed in \mathbf{G}_2 and \mathbf{G}_3 . □

Similar to the above result, we can show that the Fiat-Shamir transformation applied to a multi-key lossy identification scheme leads to a tightly secure signature scheme in the multi-user setting without corruptions. We postpone this result to Supplementary Material Section B.

3.2 Generic Construction of Lossy Identification Schemes

In this section we show a relation between (multi-key) lossy identification schemes and dual-mode commitments. Note that it is well-known how to use canonical identification schemes to build standard

Game G_5	
<pre> 01 for $i \in [N]$: $b_i \leftarrow^s \{0, 1\}$ 02 $(\text{par}, (\text{pk}_{i,1-b_i})_{i \in [N]}) \leftarrow \text{LGen}(1^N)$ 03 for $i \in [N]$: 04 $(\text{pk}_{i,b_i}, \text{sk}_{i,b_i}) \leftarrow \text{IGen}(\text{par})$ 05 $\text{pk}_i := (\text{pk}_{i,0}, \text{pk}_{i,1})$ 06 $\text{sk}_i := (b_i, \text{sk}_{i,b_i})$ 07 $O := (\text{H}, \text{SIG}, \text{KEY})$ 08 $(i^*, m^*, \sigma^*) \leftarrow \mathcal{B}^O(\text{par}, (\text{pk}_i)_{i=1}^N)$ 09 if $i^* \in \mathcal{L}_{id}$: return 0 </pre>	<pre> 10 if $\exists \sigma : (i^*, m^*, \sigma) \in \mathcal{L}_m$: return 0 11 let $\sigma^* = (\text{cmt}_0^*, \text{cmt}_1^*, \text{rs}_0^*, \text{rs}_1^*)$ 12 if $\mathcal{L}_h[0, \text{pk}_{i^*}, \text{cmt}_0^*, m^*] = \perp$: 13 return 0 14 if $\mathcal{L}_h[1, \text{pk}_{i^*}, \text{cmt}_1^*, m^*] = \perp$: 15 return 0 16 $(t_0, h_0) := \mathcal{L}_h[0, \text{pk}_{i^*}, \text{cmt}_0^*, m^*]$ 17 $(t_1, h_1) := \mathcal{L}_h[1, \text{pk}_{i^*}, \text{cmt}_1^*, m^*]$ 18 if $(t_{1-b_{i^*}} > t_{b_{i^*}})$: return 0 19 return $\text{Ver}(\text{pk}_{i^*}, m^*, \sigma^*)$ </pre>

Figure 5: Game G_5 for adversary \mathcal{B} in the proof of MU-CMA-Corr security of $\text{SIG}_s[\text{ID}, \text{H}, \ell]$. The oracles $\text{H}, \text{SIG}, \text{KEY}$ are as in G_4 , see Fig. 4. The highlighted line shows what changed from G_4 to G_5 .

commitment schemes [HL10]. This section shows that this can be used to understand lossy identification in a novel way. In combination with the result from the previous section, we obtain an N -MU-CMA-Corr secure signature scheme from a dual-mode commitment in a tight way. Let $\text{CMT} = (\text{Setup}, \text{TSetup}, \text{Gen}, \text{TGen}, \text{Com}, \text{TCom}, \text{Open}, \text{TCol})$ be a dual-mode commitment with message space \mathcal{M} . We construct a canonical identification scheme $\text{ID}[\text{CMT}]$ in Fig. 6.

The intuition is that the prover sends a random commitment and is challenged with a random element from the message space. Then the prover needs to open the commitment for the challenge message. If the prover knows the trapdoor of the dual-mode commitment, this is no problem. On the other hand, if the commitment key is in binding mode, opening the commitment for the challenge message is infeasible.

<pre> Alg LGen(1^N) 01 par $\leftarrow \text{Setup}(1^\lambda)$ 02 $(\text{ck}_1, \dots, \text{ck}_N) \leftarrow \text{Gen}(\text{par}, 1^N)$ 03 for $i \in [N]$: $\text{pk}_i := \text{ck}_i$ 04 return (par, $\text{pk}_1, \dots, \text{pk}_N$) </pre>	<pre> Alg P₂(sk, ch, St) 07 dc $\leftarrow \text{TCol}(\text{pk}, \text{sk}, \text{St}, \text{ch})$ 08 return dc </pre>
<pre> Alg P₁(sk = td) 05 $(c, St) \leftarrow \text{TCom}(\text{pk}, \text{sk})$ 06 return (cmt := c, St) </pre>	<pre> Alg V(pk, cmt, ch, rs) 09 c := cmt, m := ch, dc := rs 10 return Open(pk, m, dc, c) </pre>
	<pre> Alg Sim(pk, ch) 11 (c, dc) $\leftarrow \text{Com}(\text{pk}, \text{ch})$ 12 return (cmt := c, rs := dc) </pre>

Figure 6: The identification scheme $\text{ID}[\text{CMT}] = (\text{ISetup} := \text{TSetup}, \text{IGen} := \text{TGen}, \text{P}, \text{V})$ with challenge set $\text{ChSet} := \mathcal{M}$ and related algorithms Sim, LGen for a given dual-mode commitment $\text{CMT} = (\text{Setup}, \text{TSetup}, \text{Gen}, \text{TGen}, \text{Com}, \text{TCom}, \text{Open}, \text{TCol})$ with message space \mathcal{M} .

Lemma 3.3 (uHVZK and Completeness). *If CMT is a ρ -complete dual-mode $(\varepsilon_{\text{bind}}, \varepsilon_{\text{trap}}, N)$ -commitment scheme, then $\text{ID}[\text{CMT}]$ is ε_{zk} -uHVZK and ρ' -complete, where $\varepsilon_{\text{zk}} \leq \varepsilon_{\text{trap}}$ and $\rho' \geq \rho - \varepsilon_{\text{trap}}$.*

Proof. By definition of a dual-mode commitment scheme, the following distributions have statistical distance at most $\varepsilon_{\text{trap}}$ for any $m \in \mathcal{M}$:

$$\{(c, m, \text{dc}) \mid (c, \text{dc}) \leftarrow \text{Com}(\text{ck}, m)\}$$

and

$$\{(c, m, \text{dc}) \mid (c, St) \leftarrow \text{TCom}(\text{ck}, \text{td}), \text{dc} \leftarrow \text{TCol}(\text{ck}, \text{td}, St, m)\},$$

and the former is exactly the distribution output by Sim on input $\text{ch} = m$, and the latter is exactly the distribution of a real transcript using m as the challenge. The completeness of CMT now implies that V accepts a simulated transcript output by Sim with probability at least ρ . Thus, a real transcript will be accepted with probability at least $\rho - \varepsilon_{\text{trap}}$, which finishes the proof. \square

Lemma 3.4 (Multi-Key Lossiness). *If CMT is a dual-mode $(\varepsilon_{\text{bind}}, \varepsilon_{\text{trap}}, N)$ -commitment scheme, then ID[CMT] is $(\varepsilon_{\text{mkl}}, N)$ -multi-key lossy, where*

$$\varepsilon_{\text{mkl}} \leq \varepsilon_{\text{bind}} + 1/|\mathcal{M}|.$$

In particular, for every PPT algorithm \mathcal{A} there exists a PPT algorithm \mathcal{B} , such that $\mathbf{T}(\mathcal{B}) \approx \mathbf{T}(\mathcal{A})$ and

$$\text{Adv}_{\mathcal{A}, \text{ID}[\text{CMT}]}^{N\text{-keydist}}(\lambda) \leq \text{Adv}_{\mathcal{B}, \text{CMT}}^{N\text{-keydist}}(\lambda).$$

Proof. As $(\text{ISetup}, \text{IGen}) = (\text{TSetup}, \text{TGen})$ and LIGen combines the outputs of Setup and Gen, distinguishing lossy and honest keys of ID[CMT] is exactly equivalent to distinguishing commitment keys generated via Setup, Gen and TSetup, TGen. Thus, the reduction \mathcal{B} is trivial. It remains to show the statement about ε_{mkl} . To this end, let $(\text{par}, \text{pk}_1, \dots, \text{pk}_N) \leftarrow \text{LIGen}(1^N)$, which is the same as writing

$$\text{par} \leftarrow \text{Setup}(1^\lambda), (\text{ck}_1, \dots, \text{ck}_N) \leftarrow \text{Gen}(\text{par}, 1^N).$$

Define the event \mathbf{E} of finding a collision for some $i \in [N]$ as

$$\begin{aligned} \mathbf{E} := & (\exists i \in [N], c \in \mathcal{C}, m, m' \in \mathcal{M}, dc, dc' \in \mathcal{D} : \\ & m \neq m' \wedge \text{Open}(\text{ck}_i, m, dc, c) = 1 \wedge \text{Open}(\text{ck}_i, m', dc', c) = 1). \end{aligned}$$

By definition of the (multi-key) binding property, we know that $\Pr[\mathbf{E}] \leq \varepsilon_{\text{bind}}$. We can rewrite this event \mathbf{E} in terms of ID[CMT]:

$$\begin{aligned} \exists i \in [N], \text{cmt} \in \text{CmtSet}, \text{ch}, \text{ch}' \in \text{ChSet}, \text{rs}, \text{rs}' \in \text{RspSet} : \\ \text{ch} \neq \text{ch}' \wedge \mathbf{V}(\text{pk}_i, \text{cmt}, \text{ch}, \text{rs}) = \mathbf{V}(\text{pk}_i, \text{cmt}, \text{ch}', \text{rs}') = 1. \end{aligned}$$

Define the random variable W as

$$W := \max_{i \in [N]} \max_{\text{cmt} \in \text{CmtSet}} \Pr_{\text{ch} \leftarrow \mathcal{S}\text{-ChSet}} [\exists \text{rs} \in \text{RspSet} : \mathbf{V}(\text{pk}_i, \text{cmt}, \text{ch}, \text{rs}) = 1].$$

Then, note that $\neg \mathbf{E}$ implies that for any $i \in [N]$ and $\text{cmt} \in \text{CmtSet}$ there is at most one challenge such that there is a valid response for it (with respect to pk_i). Hence

$$\mathbb{E}[W \mid \neg \mathbf{E}] \leq 1/|\mathcal{M}|.$$

To finish our proof, we need to bound the expectation of W :

$$\begin{aligned} \mathbb{E}[W] &= \mathbb{E}[W \mid \mathbf{E}] \Pr[\mathbf{E}] + \mathbb{E}[W \mid \neg \mathbf{E}] \Pr[\neg \mathbf{E}] \leq 1 \cdot \varepsilon_{\text{bind}} + \mathbb{E}[W \mid \neg \mathbf{E}] \cdot 1 \\ &\leq \varepsilon_{\text{bind}} + 1/|\mathcal{M}|. \end{aligned}$$

□

4 Instantiations

In the previous sections we showed how to tightly transform any (multi-key) dual-mode commitment scheme into a signature scheme with security in presence of corruptions. We will now construct such dual-mode commitment schemes based on a variety of assumptions, including LWE and isogenies.

4.1 Instantiation based on LWE

Our scheme CMT_{LWE} based on the LWE assumption is presented in Fig. 7. It is inspired by the classical lattice cryptosystem by Regev [Reg05] and its extension to multiple bits from [PVW08]. It makes use of parameters $n, m \in \mathbb{N}$ and $q \in \mathbb{P}$ and a parameter $k \in \mathbb{N}, k \in \Theta(\lambda)$, as well as Gaussian widths $s_0, s > 0$. For the trapdoor algorithms (see Lem. A.3) to work, we need to ensure that

$$\begin{aligned} m &\geq 3(n+k) \lceil \log q \rceil \\ s &\geq C_1 \cdot \sqrt{s_0^2 C_0^2 (\sqrt{m-w} + \sqrt{w})^2 + 1} \cdot \omega(\sqrt{\log(n+k)}), \end{aligned}$$

where $w = (n+k) \lceil \log q \rceil$. Additionally, we need a parameter $0 < \alpha < 1$ with $\alpha < 1/(4sm)$ and $\alpha q \geq 2\sqrt{n}$, which is used for setting up statistically binding keys.

<p>Alg TGen(par)</p> <p>01 $(\mathbf{A}, \mathbf{T}_\mathbf{A}) \leftarrow \text{GenTrap}(1^{n+k}, 1^m, s_0, q)$</p> <p>02 $\text{ck} := \mathbf{A} \in \mathbb{Z}_q^{(n+k) \times m}, \text{td} := \mathbf{T}_\mathbf{A}$</p> <p>03 return (ck, td)</p> <p>Alg Gen(par, 1^N)</p> <p>04 $\bar{\mathbf{S}} \xleftarrow{s} \mathbb{Z}_q^{n \times k}$</p> <p>05 for $i \in [N]$:</p> <p>06 $\bar{\mathbf{A}}_i \xleftarrow{s} \mathbb{Z}_q^{n \times m}, \bar{\mathbf{E}}_i \leftarrow D_{\mathbb{Z}, \alpha q}^{m \times k}$</p> <p>07 $\text{ck}_i := \mathbf{A}_i := \begin{bmatrix} \bar{\mathbf{A}}_i \\ \bar{\mathbf{S}}^t \bar{\mathbf{A}}_i + \bar{\mathbf{E}}_i^t \end{bmatrix}$</p> <p>08 return (ck₁, ..., ck_N)</p> <p>Alg Com(ck, m)</p> <p>09 $\mathbf{z} \leftarrow D_{\mathbb{Z}^m, s}, \mathbf{u} := \mathbf{Az} + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m} \end{bmatrix}$</p> <p>10 if $\ \mathbf{z}\ > s \cdot \sqrt{m}$: return \perp</p> <p>11 return (c := u, dc := z)</p>	<p>Alg TCom(ck, td)</p> <p>12 $\mathbf{u} \xleftarrow{s} \mathbb{Z}_q^{n+k}$</p> <p>13 return (u, St := u)</p> <p>Alg TCol(ck, td, St, m)</p> <p>14 $\mathbf{y} := \mathbf{u} - \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m} \end{bmatrix}$</p> <p>15 $\mathbf{z} \leftarrow \text{SampleD}(\mathbf{A}, \mathbf{T}_\mathbf{A}, \mathbf{y}, s)$</p> <p>16 if $\ \mathbf{z}\ > s \cdot \sqrt{m}$: return \perp</p> <p>17 return z</p> <p>Alg Open(ck, m, z, u)</p> <p>18 if $\ \mathbf{z}\ > s \cdot \sqrt{m}$: return 0</p> <p>19 if $\mathbf{Az} + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m} \end{bmatrix} \neq \mathbf{u}$: return 0</p> <p>20 return 1</p>
--	---

Figure 7: The dual-mode commitment $\text{CMT}_{\text{LWE}} = (\text{Setup}, \text{TSetup}, \text{Gen}, \text{TGen}, \text{Com}, \text{TCom}, \text{Open}, \text{TCol})$ with message space $\mathcal{M} = \{0, 1\}^k$, where $\text{Setup} = \text{TSetup}$ sets parameters par as in the text.

Lemma 4.1 (Completeness, Trapdoor Property). *The scheme CMT_{LWE} is ρ -complete and satisfies the ε_t -trapdoor property with $\rho \geq 1 - \text{negl}(\lambda)$ and $\varepsilon_t \leq \text{negl}(\lambda)$.*

Proof. Let $(\text{ck} = \mathbf{A}, \text{td} = \mathbf{T}_\mathbf{A}) \leftarrow \text{TGen}(\text{par})$. First, we show that commitments and decommitments generated using the trapdoor are accepted with overwhelming probability, then we show the trapdoor property. In combination, this also implies completeness.

First, let $(\mathbf{u}, St) \leftarrow \text{TCom}(\text{ck}, \text{td}), \mathbf{m} \in \{0, 1\}^k, \text{TCol}(\text{ck}, \text{td}, St, \mathbf{m})$. The properties of GenTrap ensure that \mathbf{A} is statistically close to uniform. By the definition of algorithm TCol and algorithm SampleD we have that \mathbf{z} is distributed statistically close to $D_{\Lambda_{\bar{\mathbf{y}}}^\perp(\mathbf{A}), s}$, where $\mathbf{y} = \mathbf{u} - [\mathbf{0}^t \lfloor q/2 \rfloor \cdot \mathbf{m}^t]^t$. It follows by definition of $\Lambda_{\bar{\mathbf{y}}}^\perp(\mathbf{A})$ that we have

$$\mathbf{Az} = \mathbf{y} = \mathbf{u} - \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m} \end{bmatrix} \implies \mathbf{Az} + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m} \end{bmatrix} = \mathbf{u},$$

and with overwhelming probability (by Lem. A.2) $\|\mathbf{z}\| \leq s \cdot \sqrt{m}$ (implying that the transcript is not \perp), which makes Open accept.

For the second part, note that the aborting condition $\|\mathbf{z}\| > s \cdot \sqrt{m}$ is given in Com and in the execution of TCom, TCol , hence we only have to show that for every \mathbf{m} the distributions

$$\mathcal{D}_1 := \left\{ (\mathbf{u}, \mathbf{m}, \mathbf{z}) \mid \mathbf{u} \xleftarrow{s} \mathbb{Z}_q^{n+k}, \mathbf{z} \leftarrow \text{SampleD}(\mathbf{A}, \mathbf{T}_\mathbf{A}, \mathbf{u} - \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m} \end{bmatrix}, s) \right\}$$

and

$$\mathcal{D}_2 := \left\{ (\mathbf{u}, \mathbf{m}, \mathbf{z}) \mid \mathbf{u} := \mathbf{Az} + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m} \end{bmatrix}, \mathbf{z} \leftarrow D_{\mathbb{Z}^m, s} \right\}$$

are statistically close. Notice that in both distributions, \mathbf{u} is uniquely determined by \mathbf{m} and $\mathbf{y} := \mathbf{u} - [\mathbf{0}^t \lfloor q/2 \rfloor \cdot \mathbf{m}^t]^t$ and \mathbf{y} by \mathbf{m} and \mathbf{u} , which means we can instead bound the statistical distance between

$$\mathcal{D}'_1 := \{ (\mathbf{y}, \mathbf{z}) \mid \mathbf{y} \xleftarrow{s} \mathbb{Z}_q^{n+k}, \mathbf{z} \leftarrow \text{SampleD}(\mathbf{A}, \mathbf{T}_\mathbf{A}, \mathbf{y}, s) \}$$

and

$$\mathcal{D}'_2 := \{ (\mathbf{y}, \mathbf{z}) \mid \mathbf{y} := \mathbf{Az}, \mathbf{z} \leftarrow D_{\mathbb{Z}^m, s} \}.$$

Standard lattice trapdoor techniques (see Lem. A.1 and A.3) imply that these are statistically close, which finishes the proof. \square

Lemma 4.2 (Key Indistinguishability). *Let $N = \text{poly}(\lambda)$ be a natural number. Then CMT_{LWE} satisfies key indistinguishability, under the $\text{LWE}_{n,q,D_{\mathbb{Z},\alpha q}}$ assumption, where for every PPT algorithm \mathcal{A} there exists a PPT algorithm \mathcal{B} , such that $\mathbf{T}(\mathcal{B}) \approx \mathbf{T}(\mathcal{A})$ and*

$$\text{Adv}_{\mathcal{A}, \text{CMT}_{\text{LWE}}}^{N\text{-keydist}}(\lambda) \leq k \cdot \text{Adv}_{\mathcal{B}}^{\text{LWE}_{n,q,D_{\mathbb{Z},\alpha q}}}(\lambda) + \text{negl}(\lambda).$$

Proof. By the properties of lattice trapdoor generation (Lem. A.3), we know that the distribution of N commitment keys independently generated via TGen is statistically close to N uniformly random matrices $\mathbf{A}_i \in \mathbb{Z}_q^{(n+k) \times m}$, $i \in [N]$. Now write each matrix as

$$\mathbf{A}_i = \begin{bmatrix} \bar{\mathbf{A}}_i \\ \mathbf{B}_i^t \end{bmatrix}, i \in [N].$$

Combining these matrices as the columns of a block matrix of dimension $(n+k) \times Nm$, we see that the distribution of N such commitment keys is statistically close to the distribution

$$\left\{ \begin{bmatrix} \bar{\mathbf{A}} \\ \mathbf{B}^t \end{bmatrix} \middle| \bar{\mathbf{A}} \leftarrow_{\mathbb{Z}_q}^{s} \mathbb{Z}_q^{n \times Nm}, \mathbf{B} \leftarrow_{\mathbb{Z}_q}^{s} \mathbb{Z}_q^{Nm \times k} \right\},$$

which, by applying LWE k -many times, is computationally indistinguishable from

$$\left\{ \begin{bmatrix} \bar{\mathbf{A}} \\ \bar{\mathbf{S}}^t \bar{\mathbf{A}} + \bar{\mathbf{E}}^t \end{bmatrix} \middle| \bar{\mathbf{A}} \leftarrow_{\mathbb{Z}_q}^{s} \mathbb{Z}_q^{n \times Nm}, \bar{\mathbf{S}} \leftarrow_{\mathbb{Z}_q}^{s} \mathbb{Z}_q^{n \times k}, \bar{\mathbf{E}} \leftarrow D_{\mathbb{Z},\alpha q}^{Nm \times k} \right\}.$$

But this is exactly the distribution output by $\text{Gen}(\text{par}, 1^N)$, which can be seen by considering chunks \mathbf{A}_i of m columns per user. \square

Lemma 4.3 (Binding Property). *For any $N = \text{poly}(\lambda)$ the scheme CMT_{LWE} is (ε_b, N) -statistically binding, with $\varepsilon_b \leq \text{negl}(\lambda)$.*

Proof. Consider the random experiment

$$\text{par} \leftarrow \text{Setup}(1^\lambda), (\text{ck}_1, \dots, \text{ck}_N) \leftarrow \text{Gen}(\text{par}, 1^N).$$

Fix some user $i \in [N]$ and some commitment \mathbf{u} . We show that with high probability, there is at most one challenge \mathbf{m} for which there is a decommitment \mathbf{z} that makes Open accept: Consider the matrix $\bar{\mathbf{S}} \in \mathbb{Z}_q^{n \times k}$ used in Gen and set $\mathbf{S} := [-\bar{\mathbf{S}}^t \mid \mathbf{I}_k] \in \mathbb{Z}_q^{k \times (n+k)}$. Then we have $\mathbf{S}\mathbf{A}_i = \bar{\mathbf{E}}_i^t$. Now consider accepting pairs $(\mathbf{u}, \mathbf{m}, \mathbf{z}), (\mathbf{u}, \mathbf{m}', \mathbf{z}')$ of commitment, message and decommitment and denote $\mathbf{A} := \mathbf{A}_i, \mathbf{E} := \bar{\mathbf{E}}_i$ for simplicity. Let \mathbf{e}_j denote the j -th column of \mathbf{E} for $j \in [k]$. By definition of Open , we have $\|\mathbf{z}\|, \|\mathbf{z}'\| \leq s\sqrt{m}$ and

$$\mathbf{A}\mathbf{z} + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m} \end{bmatrix} = \mathbf{u} = \mathbf{A}\mathbf{z}' + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m}' \end{bmatrix}.$$

Multiplying with \mathbf{S} from the left this implies

$$\mathbf{E}^t \mathbf{z} + \lfloor q/2 \rfloor \cdot \mathbf{m} = \mathbf{E}^t \mathbf{z}' + \lfloor q/2 \rfloor \cdot \mathbf{m}' \implies \lfloor q/2 \rfloor \cdot \mathbf{m} - \lfloor q/2 \rfloor \cdot \mathbf{m}' = \mathbf{E}^t (\mathbf{z}' - \mathbf{z}).$$

Looking at the absolute value of each coordinate $j \in [k]$ of this equality individually we see that

$$\{ \lfloor q/2 \rfloor, 0 \} \ni \lfloor q/2 \rfloor \cdot m_j - \lfloor q/2 \rfloor \cdot m'_j = |\mathbf{e}_j^t (\mathbf{z}' - \mathbf{z})| \leq 2s\sqrt{m} \|\mathbf{e}_j\| \leq 2s\alpha qm,$$

where the last inequality holds with overwhelming probability, as $\mathbf{e}_j \leftarrow D_{\mathbb{Z},\alpha q}^m$. By our assumption $\alpha < 1/(4sm)$, this term is less than $q/2$, hence it is 0. This means that $m_j = m'_j$. In summary, we have that with overwhelming probability there is only one message \mathbf{m} for which there exists a decommitment \mathbf{z} that makes Open accept. This holds for any i and any \mathbf{u} and the claim follows. \square

<p>Alg Gen(par)</p> 01 $(\mathbf{A}_0, \mathbf{T}_0) \leftarrow \text{GenTrap}(1^{2n}, 1^m, s_0, q)$ 02 $(\mathbf{A}_1, \mathbf{T}_1) \leftarrow \text{GenTrap}(1^{2n}, 1^m, s_0, q)$ 03 $b \xleftarrow{s} \{0, 1\}, \text{sk} := (b, \mathbf{T}_b)$ 04 $\text{pk} := (\mathbf{A}_0, \mathbf{A}_1)$ 05 return (pk, sk) <hr/> <p>Alg Ver(pk, m, $\sigma = (\mathbf{u}_0, \mathbf{u}_1, \mathbf{z}_0, \mathbf{z}_1)$)</p> 06 $\mathbf{m}_1 \leftarrow \text{H}(0, \text{pk}, \mathbf{u}_0, \mathbf{m})$ 07 $\mathbf{m}_0 \leftarrow \text{H}(1, \text{pk}, \mathbf{u}_1, \mathbf{m})$ 08 if $\ \mathbf{z}_0\ > s \cdot \sqrt{m}$: return 0 09 if $\ \mathbf{z}_1\ > s \cdot \sqrt{m}$: return 0 10 if $\mathbf{A}_0 \mathbf{z}_0 + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m}_0 \end{bmatrix} \neq \mathbf{u}_0$: return 0 11 if $\mathbf{A}_1 \mathbf{z}_1 + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m}_1 \end{bmatrix} \neq \mathbf{u}_1$: return 0 12 return 1	<p>Alg Sig(sk, m)</p> 13 let $\sigma = (b, \mathbf{T}_b)$ 14 $\mathbf{u}_b \xleftarrow{s} \mathbb{Z}_q^{2n}$ 15 $\mathbf{m}_{1-b} \leftarrow \text{H}(b, \text{pk}, \mathbf{u}_b, \mathbf{m})$ 16 $\mathbf{z}_{1-b} \leftarrow D_{\mathbb{Z}^m, s}$ 17 $\mathbf{u}_{1-b} := \mathbf{A}_{1-b} \mathbf{z}_{1-b} + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m}_{1-b} \end{bmatrix}$ 18 $\mathbf{m}_b \leftarrow \text{H}(1-b, \text{pk}, \mathbf{u}_{1-b}, \mathbf{m})$ 19 $\mathbf{y}_b := \mathbf{u}_b - \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m}_b \end{bmatrix}$ 20 $\mathbf{z}_b \leftarrow \text{SampleD}(\mathbf{A}_b, \mathbf{T}_b, \mathbf{y}_b, s)$ 21 if $\ \mathbf{z}_0\ > s \cdot \sqrt{m}$: return \perp 22 if $\ \mathbf{z}_1\ > s \cdot \sqrt{m}$: return \perp 23 return $\sigma := (\mathbf{u}_0, \mathbf{u}_1, \mathbf{z}_0, \mathbf{z}_1)$
--	---

Figure 8: The signature scheme $\text{SIG}_s[\text{ID}[\text{CMT}_{\text{LWE}}], \text{H}, 1] = (\text{Setup}, \text{Gen}, \text{Sig}, \text{Ver})$, where **Setup** sets parameters as in Section 4.1.

To satisfy all the requirements of the previous analysis, we can set

$$\begin{aligned}
k &:= n, & m &:= 6n \lceil \log q \rceil, & \alpha &:= \frac{1}{5C^*} m^{-3/2} \cdot \omega(\sqrt{\log n})^{-2}, \\
4n^3 &\leq q \leq n^4, & s_0 &= \omega(\sqrt{\log n}), & s &:= C^* \cdot \sqrt{m} \cdot \omega(\sqrt{\log n})^2,
\end{aligned}$$

where $C^* := \sqrt{8} \cdot C_0 \cdot C_1$ is chosen such that s satisfies the requirement. Then especially the hardness of LWE is supported by worst-case to average case reductions, i.e. $\alpha q \geq 2\sqrt{n}$. Also, Bertrand's postulate implies that there is such a prime number q between $4n^3$ and $8n^3$, which is upper bounded by n^4 for all reasonable n .

Remark 4.4 (On Complete Tightness). Let us sketch two variants of turning the above ideas into a completely tight scheme. The first variant is to start with the single bit version of the above scheme, i.e. use $k = 1$. Unfortunately, with such a constant message space, the statement of Lem. 3.4 becomes useless and lossiness is not guaranteed anymore. The solution is to repeat $\Theta(n)$ many instances with the same key in parallel and to accept only if all of the instances accept. Then uHVZK can be seen for each instance independently and our message space is large enough to apply Lem. 3.4. The second variant is to use commitments resulting from [GSW13, GVW15] instead of the Regev-based construction we used here. In this variant a commitment for $\mathbf{x} \in \{0, 1\}^k$ with decommitment \mathbf{R} is $\mathbf{C} := \mathbf{A}\mathbf{R} + \mathbf{x}^t \otimes \mathbf{G}$. It can be proven that this is also a dual-mode commitment scheme, using the same ideas we used here. We postpone a formal description of these variants to Supplementary Material Section D.

We will now instantiate our generic construction in Section 3 with the dual-mode commitment scheme CMT_{LWE} . As it has negligible completeness error, $\ell = 1$ repetition of the sequential OR proof is sufficient. The final tightly N -MU-CMA-Corr secure signature scheme is presented in Fig. 8. Note that signatures contain a linear number of elements from \mathbb{Z}_q . For completeness, the signature schemes based on the completely tight dual-mode commitments mentioned above are formally presented in Supplementary Material Section D.

4.2 Instantiation based on Isogenies

We show how to instantiate our approach in the isogeny setting. In [EKP20] a lossy identification scheme is based on an isogeny assumption is presented. Our new analysis shows that this can be obtained from a dual-mode commitment scheme. More importantly, we are able to show tight multi-user security. Here, we use the subtle fact that our definition allows lossy keys to be correlated. Applying our approach leads to the first tightly MU-CMA-Corr secure signature scheme based on isogenies. We postpone our analysis to Supplementary Material Section E.

4.3 Instantiation based on Other Assumptions

In Supplementary Material Section C, we show that the previously known lossy ID schemes [GQ90, CP93, AFLT12, KMP16] are concrete instantiations of our transformation in Section 3.2. Applying our approach, we obtain tightly secure signature schemes based on the classical assumptions MDDH and φ -Hiding.

5 Impossibility Result for Parallel OR Proofs

In this section, we consider a canonical identification scheme $ID = (I\text{Setup}, I\text{Gen}, P := (P_1, P_2), V)$ with challenge set ChSet and a random oracle $H : \{0, 1\}^* \rightarrow \text{ChSet}$. Recall that sequential OR proofs can be used to construct MU-CMA-Corr secure signatures in a tight way (see the previous sections). Here, we show that a similar tight result for parallel OR proofs $\text{SIG}_p[ID, H]$ defined in Fig. 9 is unlikely. For

Alg Gen(par)	Alg Sig(sk, m)
01 $(pk_0, sk_0) \leftarrow I\text{Gen}(\text{par})$	10 $(cmt_b, St_b) \leftarrow P_1(\text{par}, sk_b)$
02 $(pk_1, sk_1) \leftarrow I\text{Gen}(\text{par})$	11 $ch_{1-b} \xleftarrow{\$} \text{ChSet}$
03 $b \xleftarrow{\$} \{0, 1\}, sk := (b, sk_b)$	12 $(cmt_{1-b}, rs_{1-b}) \leftarrow \text{Sim}(pk_{1-b}, ch_{1-b})$
04 return $(pk := (pk_0, pk_1), sk)$	13 $ch \leftarrow H(pk, cmt_0, cmt_1, m)$
Alg Ver(pk, m, σ)	14 $ch_b := ch \oplus ch_{1-b}$
05 $ch \leftarrow H(pk, cmt_0, cmt_1, m)$	15 $rs_b \leftarrow P_2(sk_b, ch_b, St_b)$
06 if $ch_0 \oplus ch_1 \neq ch$: return 0	16 if $rs_0 = \perp \vee rs_1 = \perp$: return \perp
07 $v_0 \leftarrow V(pk_0, cmt_0, ch_0, rs_0)$	17 $\sigma := (cmt_0, cmt_1, ch_0, ch_1, rs_0, rs_1)$
08 $v_1 \leftarrow V(pk_1, cmt_1, ch_1, rs_1)$	18 return σ
09 return $(v_0 \wedge v_1)$	

Figure 9: The signature scheme $\text{SIG}_p[ID, H] = (\text{Setup}, \text{Gen}, \text{Sig}, \text{Ver})$ for a canonical identification scheme $ID := (I\text{Setup}, I\text{Gen}, P := (P_1, P_2), V)$ with HVZK simulator Sim , where $\text{Setup} := I\text{Setup}$.

simplicity, we assume perfect completeness and hence only $\ell = 1$ repetition of the signing procedure. We will consider reductions without rewinding that use the adversary as a black box. First, we fix an intermediate security notion and the assumptions about the underlying identification scheme. After that we state and prove our impossibility result.

SECURITY NOTION AND ASSUMPTIONS. We will now define a security notion for digital signature scheme, which is weaker than N -MU-CMA-Corr security. Here, the adversary can only corrupt statically and can not ask for signatures. To be more precise, for a given signature scheme, the security game picks N (distinct) public keys pk_i and corresponding secret keys sk_i and sends all public keys to the adversary. Then the adversary can pick an index $j \in [N]$ and gets all sk_i , except sk_j from the game. Finally, the adversary has to return a valid forgery (m^*, σ^*) for pk_j . Note that there is a straightforward tight reduction, showing that if SIG is N -MU-CMA-Corr secure, then it is also N -MU-CMA-S secure. Thus, to prove that there is no tight proof of N -MU-CMA-Corr security of a signature scheme SIG , it is sufficient to show the same for N -MU-CMA-S security.

Game N-MU-CMA-S$_{\text{SIG}}^A(\lambda)$
01 $\text{par} \leftarrow \text{Setup}(1^\lambda)$
02 for $i \in [N]$: $(pk_i, sk_i) \leftarrow \text{Gen}(\text{par})$ // Assume pk_i 's pairwise distinct
03 $(j, St_A) \leftarrow \mathcal{A}_1(\text{par}, (pk_i)_{i \in [N]})$
04 if $j \notin [N]$: return 0
05 $(m^*, \sigma^*) \leftarrow \mathcal{A}_2(St_A, (sk_i)_{i \in [N] \setminus \{j\}})$
06 return $\text{Ver}(pk_j, m^*, \sigma^*)$

Figure 10: Game MU-CMA-S for a signature scheme $\text{SIG} = (\text{Setup}, \text{Gen}, \text{Sig}, \text{Ver})$, used in the proof of the impossibility result in Section 5. We assume that the keys pk_1, \dots, pk_N are pairwise distinct.

Definition 5.1 (Static Multi-User Security). Let $\text{SIG} = (\text{Setup}, \text{Gen}, \text{Sig}, \text{Ver})$ be a signature scheme and $N \in \mathbb{N}$ be a natural number. Consider the game MU-CMA-S given in Fig. 10. We say that SIG is N - MU-CMA-S secure, if for every PPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ the following advantage is negligible in λ :

$$\text{Adv}_{\mathcal{A}, \text{SIG}}^{N\text{-MU-CMA-S}}(\lambda) := \Pr \left[N\text{-MU-CMA-S}_{\text{SIG}}^{\mathcal{A}}(\lambda) \Rightarrow 1 \right].$$

Next, we define some properties the underlying identification scheme ID should have, in order to apply our impossibility result. These are similar to the ones defined in [BJLS16]. However, in our case they need to hold for the underlying identification scheme and not for the resulting signature scheme as it would be required for applying the result of [BJLS16] directly. For the rest of the section, we denote the set of secret keys for a given public key pk with respect to some parameters par , which should be clear from the context, of an identification scheme by $\mathcal{SK}(\text{pk})$. More formally $\mathcal{SK}(\text{pk}) := \{\text{sk} \mid (\text{pk}, \text{sk}) \in \text{IGen}(\text{par})\}$.

Definition 5.2 (Verifiability). Let $\text{ID} = (\text{ISetup}, \text{IGen}, \text{P}, \text{V})$ be a canonical identification scheme. We say that ID is parameter-verifiable if there is a deterministic polynomial time algorithm VerP such that for all par :

$$\text{VerP}(\text{par}) = 1 \iff \text{par} \in \text{ISetup}(1^\lambda).$$

Further, we say that ID is key-verifiable if there is a deterministic polynomial time algorithm VerK such that for all $\text{par} \in \text{ISetup}(1^\lambda)$ and pk, sk :

$$\text{VerK}(\text{par}, \text{pk}, \text{sk}) = 1 \iff (\text{pk}, \text{sk}) \in \text{IGen}(\text{par}).$$

Definition 5.3 (Key-Rerandomization). Let $\text{ID} = (\text{ISetup}, \text{IGen}, \text{P}, \text{V})$ be a canonical identification scheme. We say that ID is key-rerandomizable if there is a PPT algorithm RerandK such that for all $\text{par} \in \text{ISetup}(1^\lambda)$ and all $(\text{pk}, \text{sk}) \in \text{IGen}(\text{par})$ the key $\text{sk}' \leftarrow \text{RerandK}(\text{par}, \text{pk}, \text{sk})$ is distributed uniformly over $\mathcal{SK}(\text{pk})$.

We note that these properties are quite natural and are satisfied for example by the Chaum-Pedersen (CP) lossy identification scheme [CP93], which is easy to see.

Example 5.4 The parameters of the CP scheme are the description of a cyclic group \mathbb{G} of prime order p and two generators $g_1, g_2 \in \mathbb{G}$. To check the validity of these parameters, one simply has to check that g_1 and g_2 are not the identity element and that p is prime. Hence, CP is parameter-verifiable. The secret key is a single exponent $x \in \mathbb{Z}_p$, sampled uniformly at random, and the public key is $(X, Y) := (g_1^x, g_2^x)$. Given x, g_1, g_2, X, Y it is trivial to check if this relation is satisfied, showing key-verifiability. Moreover, such an x is unique for given X, Y, g_1, g_2 , which implies that CP is also key-rerandomizable.

REDUCTION SYNTAX. Before defining reductions, we need to define the underlying problem, where we follow the notation in [AGO11, BJLS16].

Definition 5.5 (Non-Interactive Problem). A non-interactive computational problem is a triple of algorithms $\text{NIP} = (\text{T}, \text{V}, \text{U})$, where

- $\text{T}(1^\lambda)$ takes the security parameter as input and outputs an instance c and a witness w .
- $\text{U}(c)$ takes an instance c as input and outputs a candidate solution s .
- $\text{V}(c, w, s)$ takes an instance c , a witness w and a candidate solution s as input and outputs a bit $b \in \{0, 1\}$.

For any algorithm \mathcal{A} taking z bits of randomness, we define the advantage

$$\text{Adv}_{\mathcal{A}}^{\text{NIP}}(\lambda) := \left| \Pr \left[\text{V}(c, w, s) = 1 \mid (c, w) \leftarrow \text{T}(1^\lambda), \rho_{\mathcal{A}} \leftarrow \{0, 1\}^z, s \leftarrow \mathcal{A}(c; \rho_{\mathcal{A}}) \right] \right. \\ \left. - \Pr \left[\text{V}(c, w, s) = 1 \mid (c, w) \leftarrow \text{T}(1^\lambda), \rho_{\text{U}} \leftarrow \{0, 1\}^z, s \leftarrow \text{U}(c; \rho_{\text{U}}) \right] \right|.$$

Before we formally define simple reductions, we make a convention about cryptographic proofs. A proof can be presented as a sequence of games \mathbf{G}_i , where typically \mathbf{G}_0 is the original security game and \mathbf{G}_{i+1} results from \mathbf{G}_i by making small changes. In the final game it will be clear that the advantage of an adversary is negligible. If one can show that in every step, changing the game only changes the advantage of the adversary by a negligible amount, the proof is complete. This is shown in one of two




Figure 11: Overview of a typical cryptographic proof, summarized by two games $\mathbf{G}_1, \mathbf{G}_2$, where \mathbf{G}_1 is statistically close to the real game. Here, a reduction \mathcal{R} to the problem NIP is used to interpolate between the games. We will show meta-reductions \mathcal{B}_i , that have a high advantage in \mathbf{G}_1 , whereas every adversary has negligible advantage in \mathbf{G}_2 .

Alg $\mathcal{R}^{\mathcal{A}}(c)$ // Simulating $N\text{-MU-CMA-S}$	
01 $\rho_{\mathcal{R}} \xleftarrow{s} \{0, 1\}^z$	05 $(\mathbf{m}^*, \sigma^*) \leftarrow \mathcal{A}_2^{\text{H}}(St_{\mathcal{A}}, (\text{sk}_i)_{i \in [N] \setminus \{j\}})$
02 $(St_{\mathcal{R}}, \text{par}, (\text{pk}_i)_{i \in [N]}) \leftarrow \mathcal{R}_1(c; \rho_{\mathcal{R}})$	06 return $\mathcal{R}_3(St_{\mathcal{R}}, j, \mathbf{m}^*, \sigma^*)$
03 $(j, St_{\mathcal{A}}) \leftarrow \mathcal{A}_1^{\text{H}}(\text{par}, (\text{pk}_i)_{i \in [N]})$	Oracle $\text{H}(\text{query})$
04 $(St_{\mathcal{R}}, (\text{sk}_i)_{i \in [N] \setminus \{j\}}) \leftarrow \mathcal{R}_2(St_{\mathcal{R}}, j)$	07 $(St_{\mathcal{R}}, h) \leftarrow \mathcal{R}_{RO}(St_{\mathcal{R}}, \text{query})$
	08 return h

Figure 12: Syntax of a simple reduction $\mathcal{R} = (\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3, \mathcal{R}_{RO})$ in an execution with an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, used in the proof of the impossibility result in Section 5. Here, \mathcal{R} simulates the game $N\text{-MU-CMA-S}_{\text{SIG}}^{\mathcal{A}}$ for \mathcal{A} .

ways: Either, one can argue that two subsequent games look statistically close to the adversary, or one uses a reduction that interpolates between the games to show that the advantages are close under some computational assumption. Clearly, we can summarize all the steps into one initial statistical step and one computational step using a reduction \mathcal{R} , as it is presented in Fig. 11. Note that this also captures reductions to search problems, as one can always define the final game to reject everything. The reduction solves the computational problem whenever the difference between the advantages in \mathbf{G}_1 and \mathbf{G}_2 is non-negligible. This means that, when we analyze the advantage of adversaries or meta-reductions, we can focus on \mathbf{G}_1 , as every (even unbounded) adversary has negligible advantage in \mathbf{G}_2 . Hence, in our analysis we only have to deal with the case where \mathcal{R} 's simulation is statistically close to the real game. With this convention in mind, we can now move towards the definition.

Definition 5.6 (Simple Reduction). Let NIP be a non-interactive computational problem and SIG be a signature scheme. A simple (NIP, SIG)-reduction \mathcal{R} is an algorithm against NIP that has one-time black box-access to an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ against the $N\text{-MU-CMA-S}$ security of SIG. In this case, \mathcal{R} can be represented by four algorithms $(\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3, \mathcal{R}_{RO})$, where $\mathcal{R}_2, \mathcal{R}_3, \mathcal{R}_{RO}$ are deterministic polynomial time algorithms and \mathcal{R}_1 is PPT, such that

- $\mathcal{R}_1(c)$ takes as input a NIP challenge c and outputs a state, parameters and public keys $(St_{\mathcal{R}}, \text{par}, (\text{pk}_i)_{i \in [N]})$.
- $\mathcal{R}_2(St_{\mathcal{R}}, j)$ takes as input a state $St_{\mathcal{R}}$ and an index $j \in [N]$ and outputs a new state and secret keys $(St_{\mathcal{R}}, (\text{sk}_i)_{i \in [N] \setminus \{j\}})$.
- $\mathcal{R}_3(St_{\mathcal{R}}, j, \mathbf{m}^*, \sigma^*)$ takes as input a state, an index $j \in [N]$, a message \mathbf{m}^* and a signature σ^* and outputs a NIP solution s .
- $\mathcal{R}_{RO}(St_{\mathcal{R}}, \text{query})$ takes as input a state $St_{\mathcal{R}}$ and a random oracle query query and outputs a new state and a hash value $(St_{\mathcal{R}}, h)$.

The joint execution of \mathcal{R} with adversary \mathcal{A} is formally given in Fig. 12. We say that \mathcal{R} is $(N, \delta_{\mathcal{R}}, L)$ -simple, if \mathcal{R} 's simulation has statistical distance at most $\delta_{\mathcal{R}}$ from the game MU-CMA-S and for all \mathcal{A} as above, it holds that

$$\text{Adv}_{\mathcal{R}^{\mathcal{A}}}^{\text{NIP}}(\lambda) \geq L(\lambda, N, \text{Adv}_{\mathcal{A}, \text{SIG}}^{N\text{-MU-CMA-S}}(\lambda)).$$

Note that in our definition we can assume that \mathcal{R}_1 is the only probabilistic part of the reduction as it can save random coins for $\mathcal{R}_2, \mathcal{R}_3, \mathcal{R}_{RO}$ in the state $St_{\mathcal{R}}$.

OUR IMPOSSIBILITY RESULT. We formalize and prove our impossibility result.

Theorem 5.7 *Let ID be a canonical identification scheme, which is ε_{zk} -HVZK, parameter-verifiable, key-verifiable and key-rerandomizable. Define the signature scheme $SIG := SIG_p[ID, H]$. Then for every $(N, \delta_{\mathcal{R}}, L)$ -simple (NIP, SIG)-reduction $\mathcal{R} = (\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3, \mathcal{R}_{RO})$ there is an algorithm \mathcal{B} such that*

$$\text{Adv}_{\mathcal{B}}^{\text{NIP}}(\lambda) \geq L(\lambda, N, 1) - 2(\delta_{\mathcal{R}} + \varepsilon_{zk}) - 1/N$$

and $\mathbf{T}(\mathcal{B}) \leq N \cdot \mathbf{T}(\mathcal{R}) + N(N-1)\mathbf{T}(\text{VerK}) + \mathbf{T}(\text{VerP}) + \mathbf{T}(\text{RerandK}) + \mathbf{T}(\text{Sig})$.

Proof. Let $\mathcal{R} = (\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3, \mathcal{R}_{RO})$ be a reduction as defined above. To prove our impossibility result, we construct a sequence of adversaries and show that they can win the MU-CMA-S game with high probability. The first few adversaries will be inefficient. However, the final adversary is efficient by rewinding the reduction \mathcal{R} . This is a common way to present meta-reductions, although often there is only one inefficient algorithm [BJLS16]. Our main task is to show that the success probabilities of the reduction do not change significantly when we move from one adversary to the next. The first adversary $\mathcal{A}^* = (\mathcal{A}_1^*, \mathcal{A}_2^*)$, formally presented in Fig. 13, obtains parameters and keys $\text{par}, (\text{pk}_i)_{i \in [N]}$ from the challenger, samples $j^* \xleftarrow{\$} [N]$ and gives it to the challenger. After obtaining all secret keys except sk_{j^*} , \mathcal{A}^* samples a random secret key with bit 0, i.e. a secret key with respect to ID, par and $\text{pk}_{j^*,0}$. Note that this is why \mathcal{A}^* is inefficient. It then signs a random message $\text{m} \xleftarrow{\$} \mathcal{M}$ and returns it. In terms of success probability the following claim is then clear:

Lemma 5.8 $\text{Adv}_{\mathcal{A}^*, \text{SIG}_p[\text{ID}, \text{H}]}^{N\text{-MU-CMA-S}}(\lambda) = 1$.

We will now present and analyze the other adversaries, which are implicitly given as meta-reductions $\mathcal{B}_1, \dots, \mathcal{B}_5$ modeling the adversary and the reduction in their joint execution. That is, they run in the NIP game and use \mathcal{R} as a subroutine. \mathcal{B}_5 will be efficient. A formal description can be found in Fig. 15. The changes can be summarized as follows:

- \mathcal{B}_1 is as $\mathcal{R}^{\mathcal{A}^*}$ except that \mathcal{B}_1 makes the following steps, summarized in the subroutine `Rewind` in Fig. 14: After obtaining $St_{\mathcal{R},1}, \text{par}$ and $(\text{pk}_i)_{i \in [N]}$ from \mathcal{R}_1 it runs \mathcal{R}_2 independently for every $j \in [N]$, stores all secret keys obtained and uses a flag $\text{succ}[j]$ to keep track of those runs in which all secret keys returned by \mathcal{R}_2 were valid. Then it samples a random j^* as \mathcal{A}^* does, continues with the j^* -th run as \mathcal{A}^* and returns whatever \mathcal{R}_3 returns.
- \mathcal{B}_2 additionally checks for an event `bad` between sampling the index j^* and continuing with the j^* -run. The event occurs if $\text{succ}[j^*] = 1$ and $\text{succ}[j] = 0$ for all other $j \neq j^*$, i.e. \mathcal{R} could only return valid secret keys for one index j^* given to \mathcal{R}_2 . If the event holds, \mathcal{B}_2 aborts.
- \mathcal{B}_3 is as \mathcal{B}_2 but additionally brute forces a random secret key for $\text{pk}_{j^*,1}$ and then uses both secret keys $\text{sk}_{j^*,0}, \text{sk}_{j^*,1}$ to compute the signature instead of using the algorithm `Sim`. The computation of the signature with two keys is summarized in Fig. 14.
- \mathcal{B}_4 is as \mathcal{B}_3 , but if `bad` does not occur, it will have received a valid secret key $(b, \text{sk}_{j^*,b})$ for pk_{j^*} from some execution of \mathcal{R}_2 with index $j \neq j^*$. It will use this secret key (rerandomized) to generate the signature instead of a brute forced one. The other key $\text{sk}_{j^*,1-b}$ is still brute forced and `Sim` is still not used.
- \mathcal{B}_5 now uses only the rerandomized $\text{sk}_{j^*,b}$ and the algorithm `Sim` to generate the signature. Note the \mathcal{B}_5 does not brute force any secret key anymore and is efficient. We set $\mathcal{B} := \mathcal{B}_5$.

We will now argue, that the success probability of \mathcal{R} does not significantly change when we change our adversaries.

Lemma 5.9 $\text{Adv}_{\mathcal{R}^{\mathcal{A}^*}}^{\text{NIP}}(\lambda) = \text{Adv}_{\mathcal{B}_1}^{\text{NIP}}(\lambda)$.

Proof. First, note that the output of \mathcal{B}_1 does not depend on the executions of $\mathcal{R}_2(St_{\mathcal{R},1}, j)$ for $j \neq j^*$. That is, only one iteration of the loop in Fig. 14, Line 05 has an influence on the output of \mathcal{R}_3 and hence \mathcal{B}_1 . Considering only this iteration, $\mathcal{R}^{\mathcal{A}^*}$ and \mathcal{B}_1 are exactly the same, where it may be worth mentioning that Line 03 in Fig. 15 and Line 05 in Fig. 13 are equivalent conditions. \square

Alg $\mathcal{A}_1^{\text{H}}(\text{par}, (\text{pk}_i)_{i \in [N]})$ 01 if VerP(par) $\neq 1$: return \perp 02 $j^* \xleftarrow{\$} [N]$ 03 $St := (\text{par}, (\text{pk}_i)_{i \in [N]}, j^*)$ 04 return (j^*, St)	Alg $\mathcal{A}_2^{\text{H}}(St, (\text{sk}_i = (b_i, \text{sk}_{i,b_i}))_{i \in [N] \setminus \{j^*\}})$ 05 if $\exists i \in [N] \setminus \{j^*\} : \text{VerK}(\text{par}, \text{pk}_{i,b_i}, \text{sk}_{i,b_i}) = 0$: return \perp 06 $\text{sk}_0 \xleftarrow{\$} \mathcal{SK}(\text{pk}_{j^*,0})$ 07 $m^* \xleftarrow{\$} \mathcal{M}, \sigma^* \leftarrow \text{Sig}((0, \text{sk}_0), m^*)$ 08 return (m^*, σ^*)
---	---

Figure 13: The optimal (but inefficient) adversary $\mathcal{A}^* = (\mathcal{A}_1^*, \mathcal{A}_2^*)$, winning the game **MU-CMA-S** for the signature scheme $\text{SIG}_p[\text{ID}, \text{H}]$.

Alg Rewind$^{\mathcal{R}}(c)$ 01 $\rho_{\mathcal{R}} \xleftarrow{\$} \{0, 1\}^z$ 02 $(St_{\mathcal{R},1}, \text{par}, (\text{pk}_i)_{i \in [N]}) \leftarrow \mathcal{R}_1(c; \rho_{\mathcal{R}})$ 03 if VerP(par) $\neq 1$: return \perp 04 $\text{succ} := [0, \dots, 0]$ 05 for $j \in [N]$: 06 $(St_{\mathcal{R},2,j}, (\text{sk}_i)_{i \in [N] \setminus \{j\}}) \leftarrow \mathcal{R}_2(St_{\mathcal{R},1}, j)$ 07 $\text{succ}[j] := 1$ 08 for $i \in [N] \setminus \{j\}$: 09 let $\text{sk}_i = (b_i, \text{sk}_{i,b_i})$ 10 if $\text{VerK}(\text{par}, \text{pk}_{i,b_i}, \text{sk}_{i,b_i}) = 0$: $\text{succ}[j] := 0$ 11 if $\text{succ}[j] = 1$: 12 for $i \in [N] \setminus \{j\} : \text{sk}[i] := \text{sk}_i$ 13 return $(\text{par}, \text{succ}[\cdot], \text{sk}[\cdot], (St_{\mathcal{R},2,j})_{j \in [N]})$	Alg FakeSign$((\text{sk}_0, \text{sk}_1), m)$ 14 $(\text{cmt}_0, St_0) \leftarrow \text{P}_1(\text{sk}_0)$ 15 $(\text{cmt}_1, St_1) \leftarrow \text{P}_1(\text{sk}_1)$ 16 $\text{ch}_0 \xleftarrow{\$} \text{ChSet}$ 17 $\text{ch} \leftarrow \text{H}(\text{pk}, \text{cmt}_0, \text{cmt}_1, m)$ 18 $\text{ch}_1 := \text{ch}_0 \oplus \text{ch}$ 19 $\text{rs}_0 \leftarrow \text{P}_2(\text{sk}_0, \text{ch}_0, St_0)$ 20 $\text{rs}_1 \leftarrow \text{P}_2(\text{sk}_1, \text{ch}_1, St_1)$ 21 if $\text{rs}_0 = \perp \vee \text{rs}_1 = \perp$: 22 return \perp 23 $\sigma := (\text{cmt}_0, \text{cmt}_1, \text{ch}_0, \text{ch}_1, \text{rs}_0, \text{rs}_1)$ 24 return σ
--	---

Figure 14: Subroutines Rewind and FakeSign, used in algorithms \mathcal{B}_i given in Fig. 15.

Lemma 5.10 $\left| \text{Adv}_{\mathcal{B}_1}^{\text{NIP}}(\lambda) - \text{Adv}_{\mathcal{B}_2}^{\text{NIP}}(\lambda) \right| \leq 1/N$.

Proof. Note that \mathcal{B}_1 and \mathcal{B}_2 only differ if the event **bad** occurs, which implies that $\text{succ}[j^*] = 1$ and $\text{succ}[j] = 0$ for all other $j \neq j^*$. Further, the set of possible j^* satisfying this condition is either empty or has one element. This means that

$$\left| \text{Adv}_{\mathcal{B}_1}^{\text{NIP}}(\lambda) - \text{Adv}_{\mathcal{B}_2}^{\text{NIP}}(\lambda) \right| \leq \Pr[\text{bad}] \leq \frac{1}{N}.$$

□

Lemma 5.11 $\left| \text{Adv}_{\mathcal{B}_2}^{\text{NIP}}(\lambda) - \text{Adv}_{\mathcal{B}_3}^{\text{NIP}}(\lambda) \right| \leq \delta_{\mathcal{R}} + \varepsilon_{\text{zk}}$.

Proof. Both \mathcal{B}_2 and \mathcal{B}_3 return \perp if $\text{VerP}(\text{par}) \neq 1$ (see Fig. 14) and also if **bad** = 1. Therefore we can assume that $\text{par} \in \text{ISetup}(1^\lambda)$. As we assume that the reduction \mathcal{R} simulated a game of statistical distance $\delta_{\mathcal{R}}$ to the real game, the public key $\text{pk}_{j^*,1}$ is of statistical distance at most $\delta_{\mathcal{R}}$ to an honest key. Hence with probability at least $1 - \delta_{\mathcal{R}}$, \mathcal{B}_3 will be able to successfully sample a random sk_1 such that $(\text{pk}_{j^*,1}, \text{sk}_1) \in \text{IGen}(\text{par})$. Note that in this case what \mathcal{R} sees is $(m^*, \sigma^*, \text{query})$ where $m^* \xleftarrow{\$} \mathcal{M}, \sigma^* \leftarrow \text{FakeSign}((\text{sk}_0, \text{sk}_1), m^*)$ and $\text{query} = ((\text{pk}_{j^*,0}, \text{pk}_{j^*,1}), \text{cmt}_0, \text{cmt}_1, m^*)$ is the random oracle query (observable by \mathcal{R}_{RO}) that occurs during the run of FakeSig. Similarly, in the execution of \mathcal{B}_2 it sees $(m^*, \sigma^*, \text{query})$ where $m^* \xleftarrow{\$} \mathcal{M}, \sigma^* \leftarrow \text{Sig}((0, \text{sk}_0), m^*)$ and $\text{query} = ((\text{pk}_{j^*,0}, \text{pk}_{j^*,1}), \text{cmt}_0, \text{cmt}_1, m^*)$. Note that the difference is only the way how the transcript $(\text{cmt}_1, \text{ch}_1, \text{rs}_1)$, which is part of σ^* , is generated. Further, query can be efficiently computed without knowing how that transcript was generated. Hence by ε_{zk} -HVZK, these have statistical distance at most ε_{zk} , which implies that \mathcal{R}_3 's final output in the execution of \mathcal{B}_3 is distributed as the same output in \mathcal{B}_2 , except with probability at most ε_{zk} . Note that this is the step where the entire argument fails for sequential OR proofs, as the additional value query that \mathcal{R} observes would have an order that allows \mathcal{R} to distinguish (Recall that a sequential OR proof makes two random oracle queries during signing). □

Lemma 5.12 $\text{Adv}_{\mathcal{B}_3}^{\text{NIP}}(\lambda) = \text{Adv}_{\mathcal{B}_4}^{\text{NIP}}(\lambda)$.

<p>Alg $\mathcal{B}_1(c), \mathcal{B}_2(c)$</p> <pre> 01 (par, succ[·], sk[·], (St_{R,2,j})_{j∈[N]}) ← Rewind^R(c) 02 j* ←^s [N] 03 if succ[j*] ≠ 1 : return 0 04 if ∀j ∈ [N] \ {j*} : succ[j] = 0 : 05 bad := 1, return ⊥ 06 sk₀ ←^s SK(pk_{j*,0}) 07 m* ←^s M, σ* ← Sig((0, sk₀), m*) 08 return R₃(St_{R,2,j*}, j*, m*, σ*) Alg $\mathcal{B}_5(c)$ 09 (par, succ[·], sk[·], (St_{R,2,j})_{j∈[N]}) ← Rewind^R(c) 10 j* ←^s [N] 11 if succ[j*] ≠ 1 : return 0 12 if ∀j ∈ [N] \ {j*} : succ[j] = 0 : 13 bad := 1, return ⊥ 14 let sk[j*] = (b, sk̄) 15 sk_b ← RerandK(par, pk_{j*,b}, sk̄) 16 m* ←^s M, σ* ← Sig(sk_b, m*) 17 return R₃(St_{R,2,j*}, j*, m*, σ*) </pre>	<p>Alg $\mathcal{B}_3(c), \mathcal{B}_4(c)$</p> <pre> 18 (par, succ[·], sk[·], (St_{R,2,j})_{j∈[N]}) ← Rewind^R(c) 19 j* ←^s [N] 20 if succ[j*] ≠ 1 : return 0 21 if ∀j ∈ [N] \ {j*} : succ[j] = 0 : 22 bad := 1, return ⊥ 23 sk₀ ←^s SK(pk_{j*,0}) 24 sk₁ ←^s SK(pk_{j*,1}) 25 let sk[j*] = (b, sk̄) 26 sk_b ← RerandK(par, pk_{j*,b}, sk̄) 27 m* ←^s M 28 σ* ← FakeSign((sk₀, sk₁), m*) 29 return R₃(St_{R,2,j*}, j*, m*, σ*) </pre>
---	---

Figure 15: The (inefficient) algorithms $\mathcal{B}_1, \dots, \mathcal{B}_4$ and the efficient algorithm \mathcal{B}_5 used in the proof of Thm. 5.7. The subroutines Rewind, FakeSign are given in Fig. 14.

Proof. The only difference between \mathcal{B}_3 uses sk_b sampled uniformly random from $\text{SK}(\text{pk}_{j^*,b})$ to generate the signature via FakeSign and \mathcal{B}_4 uses $\text{sk}_b \leftarrow \text{RerandK}(\text{par}, \text{pk}_{j^*,b}, \bar{\text{sk}})$. If bad does not occur, then there will be some $j \neq j^*$, such that $\text{succ}[j] = 1$. Fix the largest such j , then $\text{sk}[j^*] = (b, \bar{\text{sk}})$ is defined and by definition of succ and key-verifiability we have that $(\text{pk}_{j^*,b}, \bar{\text{sk}}) \in \text{IGen}(\text{par})$. By our assumption that ID is key-rerandomizable, we then know that these keys are distributed sk_b as used in \mathcal{B}_4 is distributed uniformly over $\text{SK}(\text{pk}_{j^*,b})$, which proves the claim. \square

Lemma 5.13 $\left| \text{Adv}_{\mathcal{B}_4}^{\text{NIP}}(\lambda) - \text{Adv}_{\mathcal{B}_5}^{\text{NIP}}(\lambda) \right| \leq \delta_{\mathcal{R}} + \varepsilon_{\text{zk}}$.

Proof. The proof is exactly the same as for Lem. 5.11, applying ε_{zk} -HVZK to $(\text{pk}_{j^*,1-b}, \text{sk}_{1-b})$. \square

In summary, combining all claims we obtain that

$$\text{Adv}_{\mathcal{B}_5}^{\text{NIP}}(\lambda) \geq \text{Adv}_{\mathcal{R}^{A^*}}^{\text{NIP}}(\lambda) - 2(\delta_{\mathcal{R}} + \varepsilon_{\text{zk}}) - 1/N,$$

and \mathcal{B}_5 is efficient, which proves Thm. 5.7. \square

Acknowledgments We thank the anonymous reviewers from Asiacrypt 2021 for suggesting more in-depth discussions on the difficulties in constructing tightly secure lattice-based signature schemes.

References

- [ABP19] Michel Abdalla, Fabrice Benhamouda, and David Pointcheval. On the tightness of forward-secure signature reductions. *Journal of Cryptology*, 32(1):84–150, January 2019. (Cited on page 2, 3, 4, 5.)
- [ADMP20] Navid Alamati, Luca De Feo, Hart Montgomery, and Sikhar Patranabis. Cryptographic group actions and applications. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part II*, volume 12492 of *LNCS*, pages 411–439. Springer, Heidelberg, December 2020. (Cited on page 37.)

- [AFLT12] Michel Abdalla, Pierre-Alain Fouque, Vadim Lyubashevsky, and Mehdi Tibouchi. Tightly-secure signatures from lossy identification schemes. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 572–590. Springer, Heidelberg, April 2012. (Cited on page 2, 3, 4, 5, 7, 9, 18, 29, 31, 33, 34.)
- [AGO11] Masayuki Abe, Jens Groth, and Miyako Ohkubo. Separating short structure-preserving signatures from non-interactive assumptions. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 628–646. Springer, Heidelberg, December 2011. (Cited on page 19.)
- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th ACM STOC*, pages 99–108. ACM Press, May 1996. (Cited on page 28.)
- [AOS02] Masayuki Abe, Miyako Ohkubo, and Koutarou Suzuki. 1-out-of-n signatures from a variety of keys. In Yuliang Zheng, editor, *ASIACRYPT 2002*, volume 2501 of *LNCS*, pages 415–432. Springer, Heidelberg, December 2002. (Cited on page 3.)
- [Bad14] Christoph Bader. Efficient signatures with tight real world security in the random-oracle model. In Dimitris Gritzalis, Aggelos Kiayias, and Ioannis G. Askoxylakis, editors, *CANS 14*, volume 8813 of *LNCS*, pages 370–383. Springer, Heidelberg, October 2014. (Cited on page 2.)
- [BBM00] Mihir Bellare, Alexandra Boldyreva, and Silvio Micali. Public-key encryption in a multi-user setting: Security proofs and improvements. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 259–274. Springer, Heidelberg, May 2000. (Cited on page 1.)
- [BHJ⁺15] Christoph Bader, Dennis Hofheinz, Tibor Jäger, Eike Kiltz, and Yong Li. Tightly-secure authenticated key exchange. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part I*, volume 9014 of *LNCS*, pages 629–658. Springer, Heidelberg, March 2015. (Cited on page 2, 4.)
- [BJS16] Christoph Bader, Tibor Jäger, Yong Li, and Sven Schäge. On the impossibility of tight cryptographic reductions. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 273–304. Springer, Heidelberg, May 2016. (Cited on page 2, 4, 19, 21.)
- [BKKP15] Olivier Blazy, Saqib A. Kakvi, Eike Kiltz, and Jiaxin Pan. Tightly-secure signatures from chameleon hash functions. In Jonathan Katz, editor, *PKC 2015*, volume 9020 of *LNCS*, pages 256–279. Springer, Heidelberg, March / April 2015. (Cited on page 2.)
- [BL16] Xavier Boyen and Qinyi Li. Towards tightly secure lattice short signature and id-based encryption. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 404–434. Springer, Heidelberg, December 2016. (Cited on page 1, 2.)
- [BLP⁺13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 575–584. ACM Press, June 2013. (Cited on page 5, 28.)
- [BR94] Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In Douglas R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 232–249. Springer, Heidelberg, August 1994. (Cited on page 2.)
- [BR09] Mihir Bellare and Thomas Ristenpart. Simulation without the artificial abort: Simplified proof and improved concrete security for Waters' IBE scheme. In Antoine Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 407–424. Springer, Heidelberg, April 2009. (Cited on page 1.)
- [CDS94] Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In Yvo Desmedt, editor, *CRYPTO'94*, volume 839 of *LNCS*, pages 174–187. Springer, Heidelberg, August 1994. (Cited on page 4.)

- [CK01] Ran Canetti and Hugo Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 453–474. Springer, Heidelberg, May 2001. (Cited on page 2.)
- [CLM⁺18] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: An efficient post-quantum commutative group action. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part III*, volume 11274 of *LNCS*, pages 395–427. Springer, Heidelberg, December 2018. (Cited on page 37.)
- [Cor02] Jean-Sébastien Coron. Optimal security proofs for PSS and other signature schemes. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 272–287. Springer, Heidelberg, April / May 2002. (Cited on page 4.)
- [CP93] David Chaum and Torben P. Pedersen. Wallet databases with observers. In Ernest F. Brickell, editor, *CRYPTO'92*, volume 740 of *LNCS*, pages 89–105. Springer, Heidelberg, August 1993. (Cited on page 4, 5, 18, 19, 31.)
- [CW13] Jie Chen and Hoeteck Wee. Fully, (almost) tightly secure IBE and dual system groups. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 435–460. Springer, Heidelberg, August 2013. (Cited on page 1.)
- [Dam10] Ivan Damgård. On Σ -protocols. <https://cs.au.dk/~ivan/Sigma.pdf>, 2010. (Cited on page 4.)
- [DDLL13] Léo Ducas, Alain Durmus, Tancrede Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal Gaussians. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 40–56. Springer, Heidelberg, August 2013. (Cited on page 3.)
- [DG21] Hannah Davis and Felix Günther. Tighter proofs for the SIGMA and TLS 1.3 key exchange protocols. *ACNS 2021*, 2021. (Cited on page 2.)
- [DGJL21] Denis Diemert, Kai Gellert, Tibor Jager, and Lin Lyu. More efficient digital signatures with tight multi-user security. In Juan Garay, editor, *PKC 2021, Part II*, volume 12711 of *LNCS*, pages 1–31. Springer, Heidelberg, May 2021. (Cited on page 2, 3, 4, 5, 9.)
- [DJ20] Denis Diemert and Tibor Jager. On the tight security of TLS 1.3: Theoretically-sound cryptographic parameters for real-world deployments. *Journal of Cryptology*, 2020. (Cited on page 2.)
- [EHK⁺13] Alex Escala, Gottfried Herold, Eike Kiltz, Carla Ràfols, and Jorge Villar. An algebraic framework for Diffie-Hellman assumptions. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 129–147. Springer, Heidelberg, August 2013. (Cited on page 31.)
- [EKP20] Ali El Kaafarani, Shuichi Katsumata, and Federico Pintore. Lossy CSI-FiSh: Efficient signature scheme with tight reduction to decisional CSIDH-512. In Aggelos Kiyias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part II*, volume 12111 of *LNCS*, pages 157–186. Springer, Heidelberg, May 2020. (Cited on page 2, 4, 5, 17, 37.)
- [FHJ20] Marc Fischlin, Patrick Harasser, and Christian Janson. Signatures from sequential-OR proofs. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 212–244. Springer, Heidelberg, May 2020. (Cited on page 3, 9.)
- [FS87] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO'86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987. (Cited on page 29.)
- [GHKP18] Romain Gay, Dennis Hofheinz, Lisa Kohl, and Jiaxin Pan. More efficient (almost) tightly secure structure-preserving signatures. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 230–258. Springer, Heidelberg, April / May 2018. (Cited on page 1, 2.)

- [GJ18] Kristian Gjøsteen and Tibor Jager. Practical and tightly-secure digital signatures and authenticated key exchange. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 95–125. Springer, Heidelberg, August 2018. (Cited on page 2, 3.)
- [GMPW20] Nicholas Genise, Daniele Micciancio, Chris Peikert, and Michael Walter. Improved discrete gaussian and subgaussian analysis for lattice cryptography. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 623–651. Springer, Heidelberg, May 2020. (Cited on page 5.)
- [GMR88] Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal on Computing*, 17(2):281–308, April 1988. (Cited on page 2.)
- [GPV07] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. Cryptology ePrint Archive, Report 2007/432, 2007. <https://eprint.iacr.org/2007/432>. (Cited on page 28.)
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008. (Cited on page 3, 28.)
- [GQ90] Louis C. Guillou and Jean-Jacques Quisquater. A “paradoxical” indentity-based signature scheme resulting from zero-knowledge. In Shafi Goldwasser, editor, *CRYPTO’88*, volume 403 of *LNCS*, pages 216–231. Springer, Heidelberg, August 1990. (Cited on page 4, 5, 18, 31.)
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92. Springer, Heidelberg, August 2013. (Cited on page 17.)
- [GVW15] Sergey Gorbunov, Vinod Vaikuntanathan, and Daniel Wichs. Leveled fully homomorphic signatures from standard lattices. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 469–477. ACM Press, June 2015. (Cited on page 17.)
- [HJK12] Dennis Hofheinz, Tibor Jager, and Edward Knapp. Waters signatures with optimal security reduction. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 66–83. Springer, Heidelberg, May 2012. (Cited on page 4.)
- [HJK⁺21] Shuai Han, Tibor Jager, Eike Kiltz, Shengli Liu, Jiaxin Pan, Doreen Riepel, and Sven Schäge. Authenticated key exchange and signatures with tight security in the standard model. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part IV*, volume 12828 of *LNCS*, pages 670–700, Virtual Event, August 2021. Springer, Heidelberg. (Cited on page 2, 3.)
- [HL10] Carmit Hazay and Yehuda Lindell. *Efficient Secure Two-Party Protocols - Techniques and Constructions*. ISC. Springer, Heidelberg, 2010. (Cited on page 13.)
- [JKRS21] Tibor Jager, Eike Kiltz, Doreen Riepel, and Sven Schäge. Tightly-secure authenticated key exchange, revisited. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part I*, volume 12696 of *LNCS*, pages 117–146. Springer, Heidelberg, October 2021. (Cited on page 2.)
- [KK12] Saqib A. Kakvi and Eike Kiltz. Optimal security proofs for full domain hash, revisited. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 537–553. Springer, Heidelberg, April 2012. (Cited on page 4.)
- [KLS18] Eike Kiltz, Vadim Lyubashevsky, and Christian Schaffner. A concrete treatment of Fiat-Shamir signatures in the quantum random-oracle model. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 552–586. Springer, Heidelberg, April / May 2018. (Cited on page 3, 5, 7.)

- [KMP16] Eike Kiltz, Daniel Masny, and Jiaxin Pan. Optimal security proofs for signatures from identification schemes. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 33–61. Springer, Heidelberg, August 2016. (Cited on page 2, 18, 31.)
- [LLGW20] Xiangyu Liu, Shengli Liu, Dawu Gu, and Jian Weng. Two-pass authenticated key exchange with explicit authentication and tight security. In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part II*, volume 12492 of *LNCS*, pages 785–814. Springer, Heidelberg, December 2020. (Cited on page 2.)
- [LM18] Vadim Lyubashevsky and Daniele Micciancio. Asymptotically efficient lattice-based digital signatures. *Journal of Cryptology*, 31(3):774–797, July 2018. (Cited on page 6.)
- [LP19] Roman Langrehr and Jiaxin Pan. Tightly secure hierarchical identity-based encryption. In Dongdai Lin and Kazue Sako, editors, *PKC 2019, Part I*, volume 11442 of *LNCS*, pages 436–465. Springer, Heidelberg, April 2019. (Cited on page 2.)
- [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 1–23. Springer, Heidelberg, May / June 2010. (Cited on page 5.)
- [Lyu12] Vadim Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 738–755. Springer, Heidelberg, April 2012. (Cited on page 3, 5.)
- [MP11] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. Cryptology ePrint Archive, Report 2011/501, 2011. <https://eprint.iacr.org/2011/501>. (Cited on page 28.)
- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Heidelberg, April 2012. (Cited on page 28.)
- [MR04] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. In *45th FOCS*, pages 372–381. IEEE Computer Society Press, October 2004. (Cited on page 28.)
- [Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 333–342. ACM Press, May / June 2009. (Cited on page 5, 28.)
- [PR20] Jiaxin Pan and Magnus Ringerud. Signatures with tight multi-user security from search assumptions. In Liqun Chen, Ninghui Li, Kaitai Liang, and Steve A. Schneider, editors, *ESORICS 2020, Part II*, volume 12309 of *LNCS*, pages 485–504. Springer, Heidelberg, September 2020. (Cited on page 2.)
- [PVW08] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 554–571. Springer, Heidelberg, August 2008. (Cited on page 14.)
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005. (Cited on page 5, 14, 28.)
- [SPW07] Ron Steinfeld, Josef Pieprzyk, and Huaxiong Wang. How to strengthen any weakly unforgeable signature into a strongly unforgeable signature. In Masayuki Abe, editor, *CT-RSA 2007*, volume 4377 of *LNCS*, pages 357–371. Springer, Heidelberg, February 2007. (Cited on page 6.)
- [Sto12] Anton Stolbunov. *Cryptographic schemes based on isogenies*. PhD thesis, Norwegian University of Science and Technology, 2012. (Cited on page 4.)

- [Unr15] Dominique Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 755–784. Springer, Heidelberg, April 2015. (Cited on page 2.)

Supplementary Material

A Lattice Background

An m -dimensional lattice Λ is a discrete additive subgroup of \mathbb{R}^m . Let $\mathbf{c} \in \mathbb{R}^m$ be a vector. Then we denote the discrete Gaussian distribution with parameter $s > 0$ over the coset $\mathbf{c} + \Lambda$ by $D_{\mathbf{c}+\Lambda,s}$, which is the distribution proportional to $\rho_s(\mathbf{x}) := \exp(-\pi\|\mathbf{x}\|^2/s^2)$ restricted to the coset $\mathbf{c} + \Lambda$. For an $n \times m$ matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $m > n$ we can define m -dimensional lattices and cosets:

$$\Lambda_q^\perp(\mathbf{A}) := \{\mathbf{z} \in \mathbb{Z}^m : \mathbf{A}\mathbf{z} = \mathbf{0} \pmod{q}\}, \quad \Lambda_{\mathbf{u}}^\perp(\mathbf{A}) := \{\mathbf{z} \in \mathbb{Z}^m : \mathbf{A}\mathbf{z} = \mathbf{u} \pmod{q}\}.$$

The following lemmas are well-known [Ajt96, MR04, Reg05, GPV07, GPV08] and can be obtained by using Lemmas 5.1, 5.2 and 5.3 in [GPV07] and Lemma 4.4 in [MR04].

Lemma A.1 *Let $n, m \in \mathbb{N}$, $q \in \mathbb{P}$ at least polynomial in n , $m \geq 2n \log q$. Consider any $\omega(\sqrt{\log m})$ function and $s \geq \omega(\sqrt{\log m})$. Then for all but a negligible (in n) fraction of all $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ the following distribution is statistically close to uniform over \mathbb{Z}_q^n : $\{\mathbf{A}\mathbf{e} \mid \mathbf{e} \leftarrow D_{\mathbb{Z}^m,s}\}$. Furthermore, the conditional distribution of $\mathbf{e} \leftarrow D_{\mathbb{Z}^m,s}$ given $\mathbf{u} = \mathbf{A}\mathbf{e} \pmod{q}$ is exactly $D_{\Lambda_{\mathbf{u}}^\perp(\mathbf{A}),s}$.*

Lemma A.2 *Let $n \in \mathbb{N}$, $q \in \mathbb{P}$ and $m \geq 2n \log q$. Consider any $\omega(\sqrt{\log m})$ function and $s \geq \omega(\sqrt{\log m})$. Then for all but an at most q^{-n} fraction of all $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and any vector $\mathbf{u} \in \mathbb{Z}_q^n$, we have $\Pr[\|\mathbf{x}\| > s\sqrt{m} \mid \mathbf{x} \leftarrow D_{\Lambda_{\mathbf{u}}^\perp(\mathbf{A}),s}] \leq 2^{-m+1}$.*

We let \mathbf{G} be the fixed gadget matrix defined in [MP12]. Let $n, m, q \in \mathbb{N}$, $m \geq n \lceil \log q \rceil$ and $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$. A matrix $\mathbf{R} \in \mathbb{Z}^{(m-n \lceil \log q \rceil) \times n \lceil \log q \rceil}$ is a trapdoor for \mathbf{A} if $\mathbf{A}[-\mathbf{R}^t \mid \mathbf{I}_{n \lceil \log q \rceil}]^t = \mathbf{G}$. We summarize the results in [MP12] in the next lemma, where we obtained the precise statements from Section 5 of [MP12] and we use the statistical instantiation of trapdoors. The constant C_0 is taken such that for Gaussian matrices \mathbf{X} we have $s_1(\mathbf{X}) \leq C_0 \cdot s \cdot (\sqrt{m} + \sqrt{n})$ except with negligible probability (see Lemma 2.9 in [MP11]). The constant C_1 satisfies $C_1 = \sqrt{s_1(\Sigma_{\mathbf{G}}) + 2} \leq 3$ (see [MP12]).

Lemma A.3 *There are PPT algorithms GenTrap , SampleD and constants $C_0 > 0$, $C_1 \leq 3$ such that for $n, q, m \in \mathbb{N}$, $q \geq 2$, $m \geq 3n \log q$, $w := n \lceil \log q \rceil$ and any $\omega(\sqrt{\log n})$ function the following holds with overwhelming probability over all random choices:*

- For any $s \geq \omega(\sqrt{\log n})$ the algorithm $\text{GenTrap}(1^n, 1^m, s, q)$ outputs matrices $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{R} \in \mathbb{Z}^{(m-w) \times w}$ such that \mathbf{A} is statistically close to uniform, \mathbf{R} is a trapdoor for \mathbf{A} with entries sampled from $D_{\mathbb{Z},s}$ and $s_1(\mathbf{R}) \leq s \cdot C_0 \cdot (\sqrt{m-w} + \sqrt{w})$.
- For any matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ with trapdoor \mathbf{R} , for any $\mathbf{u} \in \mathbb{Z}_q^n$ and any $s \geq C_1 \cdot \sqrt{s_1(\mathbf{R})^2 + 1} \cdot \omega(\sqrt{\log n})$, the distribution $\{\mathbf{z} \mid \mathbf{z} \leftarrow \text{SampleD}(\mathbf{A}, \mathbf{R}, \mathbf{u}, s)\}$ is statistically close to $D_{\Lambda_{\mathbf{u}}^\perp(\mathbf{A}),s}$.

Finally, we introduce the LWE assumption.

Definition A.4 (Learning With Errors (LWE) Assumption). Let $\lambda, n = n(\lambda) \in \mathbb{N}$, $q = q(n)$ be prime number and $\chi = \chi(n)$ be a distribution over \mathbb{Z} . We say that the $\text{LWE}_{n,q,\chi}$ assumption holds, if for every PPT algorithm \mathcal{B} and every polynomial $m = \text{poly}(n)$ the following advantage is negligible in λ :

$$\text{Adv}_{\mathcal{B}}^{\text{LWE}_{n,q,\chi}}(\lambda) := \left| \Pr[\mathcal{B}(\mathbf{A}, \mathbf{b}) = 1 \mid \mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}, \mathbf{b} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}] - \Pr[\mathcal{B}(\mathbf{A}, \mathbf{A}^t \mathbf{s} + \mathbf{e}) = 1 \mid \mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n \times m}, \mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n, \mathbf{e} \leftarrow \chi^m] \right|.$$

The hardness of LWE for discrete Gaussian error distributions of parameter αq with $\alpha q \geq 2\sqrt{n}$ based on the worst-case hardness of lattice approximation problems is shown by [Reg05, Pei09, BLP⁺13]. We omit the number m of samples from the advantage, as according to these results, it does not play a role in the hardness of the problem.

B Multi-User Signatures without Corruptions

Let ID be a multi-key lossy identification scheme with challenge set ChSet , $H : \{0, 1\}^* \rightarrow \text{ChSet}$ be a random oracle, and $\ell \in \mathbb{N}$ be a natural number. Here we show that multi-user security without corruptions can be tightly achieved via the Fiat-Shamir transformation $\text{SIG}[ID, H, \ell]$ [FS87, AFLT12], presented in Fig. 16. Note that compared to the presentation in [AFLT12] we added the public key as input of the random oracle. This minor detail will be helpful in the proof of Lem. B.1. For a proof of completeness of $\text{SIG}[ID, H, \ell]$, see [AFLT12].

<u>Alg Sig(sk, m)</u>	<u>Alg Ver(pk, m, σ)</u>
01 $ctr := 0$	09 let $\sigma = (\text{cmt}, \text{rs})$
02 while $ctr \leq \ell \wedge \text{rs}_0 = \perp$:	10 $\text{ch} \leftarrow H(\text{pk}, \text{cmt}, \text{m})$
03 $ctr := ctr + 1$	11 $v \leftarrow V(\text{pk}, \text{cmt}, \text{ch}, \text{rs})$
04 $(\text{cmt}, St) \leftarrow P_1(\text{sk})$	12 return v
05 $\text{ch} \leftarrow H(\text{pk}, \text{cmt}, \text{m})$	
06 $\text{rs} \leftarrow P_2(\text{sk}, \text{ch}, St)$	
07 if $\text{rs} = \perp$: return \perp	
08 return $\sigma := (\text{cmt}, \text{rs})$	

Figure 16: The signature scheme $\text{SIG}[ID, H, \ell] = (\text{Setup}, \text{Gen}, \text{Sig}, \text{Ver})$ for a canonical identification scheme $ID := (I\text{Setup}, I\text{Gen}, P := (P_1, P_2), V)$, where $\text{Setup} := I\text{Setup}$, $\text{Gen} = I\text{Gen}$.

Lemma B.1 *Let ID be a canonical identification scheme with commitment space of min-entropy $\beta = \beta(\lambda)$. If ID is ε_{zk} -HVZK and $(\varepsilon_{\text{mkl}}, N)$ -multi-key lossy for negligible $\varepsilon_{\text{zk}}, \varepsilon_{\text{mkl}}$, then $\text{SIG}[ID, H, \ell]$ is N -MU-CMA secure, with a tight reduction. More precisely, for any adversary \mathcal{A} making at most Q_S signing queries and Q_H hash queries (including the indirect ones induced by signing queries), there exists an adversary \mathcal{D} such that $\mathbf{T}(\mathcal{D}) \approx \mathbf{T}(\mathcal{A})$ and*

$$\text{Adv}_{\mathcal{A}, \text{SIG}[ID, H, \ell]}^{N\text{-MU-CMA}}(\lambda) \leq \text{Adv}_{\mathcal{D}, ID}^{N\text{-keydist}}(\lambda) + Q_S \cdot \frac{Q_H + 1}{2^\beta} + \ell \cdot Q_S \cdot \varepsilon_{\text{zk}} + (Q_H + 1) \cdot \varepsilon_{\text{mkl}}.$$

Proof. The proof follows the proof for the similar statement for single-user security in [AFLT12], but using multi-key lossiness instead of lossiness. We only give a sketch here. In every game i , we denote the advantage of an adversary \mathcal{A} as $\text{Adv}_i(\mathcal{A}) := \Pr[\mathbf{G}_i^{\mathcal{A}} \Rightarrow 1]$. \mathbf{G}_0 is the original game MU-CMA.

Game \mathbf{G}_1 : In game \mathbf{G}_1 we require the adversary to query the values $\text{pk}_{i^*}, \text{cmt}^*, \text{m}^*$ from the forgery as a random oracle query. For every adversary \mathcal{A} winning \mathbf{G}_0 , we can easily construct an adversary \mathcal{B} winning \mathbf{G}_1 with the same advantage: \mathcal{B} runs \mathcal{A} , but makes the required query before finally forwarding \mathcal{A} 's forgery. This will only increase Q_H by one.

Game \mathbf{G}_2 : In game \mathbf{G}_2 the game holds a variable `bad`, which is initially set to false. In any signing query for message m , the game checks whether the random oracle value $H(\text{pk}_i, \text{cmt}, \text{m})$ is already defined. If so, it sets the variable `bad` to true and chooses the challenge `ch` freshly at random instead of using the output of $H(\text{pk}_i, \text{cmt}, \text{m})$. The games differ only if `bad` is set, hence

$$|\text{Adv}_1(\mathcal{B}) - \text{Adv}_2(\mathcal{B})| \leq \Pr[\text{bad}] \leq Q_S \cdot (Q_H + 1) \cdot 2^{-\beta}.$$

Game \mathbf{G}_3 : \mathbf{G}_3 does not longer use the variable `bad`, which was only used for the analysis, so the advantage does not change.

Game \mathbf{G}_4 : \mathbf{G}_4 uses the HVZK simulator and program the random oracle to answer signature queries. Because of our changes in \mathbf{G}_2 and as the public key is included in the query this causes no contradictory

programming. We can apply the statistical HVZK property for each iteration in each signature query and get

$$|\text{Adv}_3(\mathcal{B}) - \text{Adv}_4(\mathcal{B})| \leq \ell \cdot Q_S \cdot \varepsilon_{zk}.$$

As the secret keys are no longer needed, we can apply multi-key lossiness:

Game \mathbf{G}_5 : In game \mathbf{G}_5 we change the parameter and key generation to

$$(\text{par}, \text{pk}_1, \dots, \text{pk}_N) \leftarrow \text{LIGen}(1^N).$$

An easy reduction \mathcal{D} shows:

$$|\text{Adv}_4(\mathcal{B}) - \text{Adv}_5(\mathcal{B})| \leq \text{Adv}_{\mathcal{D}, \text{ID}}^{N\text{-keydist}}(\lambda).$$

Finally, $\text{Adv}_5(\mathcal{B})$ can be bounded by

$$\text{Adv}_5(\mathcal{B}) \leq (Q_H + 1) \cdot \varepsilon_{\text{mkl}}.$$

Informally, this is because even an unbounded adversary \mathcal{B} can only try only $(Q_H + 1)$ values `cmt` to get a challenge `ch` from the random oracle, for which \mathcal{B} needs to find a response. We omit a formal analysis here, but a similar one can be found in the proof of Thm. 3.2. \square

C Explaining well-known Lossy ID in Our Framework

Let us show how to apply our transformation in Section 3.2 to obtain well-known lossy identification schemes. First, we present the dual-mode commitments that lead to the Chaum-Pedersen [CP93] and Guillou-Quisquater [GQ90] scheme via our transformation. The commitments are given in Figs. 17 and 18. The proof of their properties is a simple reformulation of HVZK and lossiness proofs, see e.g. [KMP16]. Note that we presented them with fixed $N = 1$, but it is easy to see that they tightly support every polynomial number N of keys, due to their random self-reducibility [KMP16]. Later, we also show schemes based on the decisional short discrete logarithm problem (DSDL) and the subset sum problem. Applying our transformation in Section 3.2 leads to the lossy identification schemes in [AFLT12]. The proof of their security is a reformulation of the lossiness proofs in [AFLT12]. For concrete parameters, see [AFLT12]. We highlight that for multi-key lossiness, i.e. $N > 1$ we obtain non-tight reductions. However, they still show the applicability of our transformation from dual-mode commitments to lossy identification.

Remark C.1 (On the Generality of our Approach). One may wonder if this can always be accomplished. To be more precise, if we construct a commitment scheme from a given lossy identification scheme by reversing the above transformation, one may ask if it will be dual-mode and statistically binding. Unfortunately this is not the case, as we explain below.

Note that all previously known schemes use the same idea to show lossiness: For a random lossy key, with high probability there is only one challenge for a given commitment such that there exists a valid response. This can be seen as an example of our transformation, as this is exactly the binding property of the underlying commitment. However, this may not be a necessary property for lossiness, as one may also show that with high probability there are only a constant number of such challenges. This is sufficient for lossiness, but is not useful for showing that the commitment is statistical binding. Hence, a direct converse of our transformation may not be correct. Instead, one could show a weaker binding property of the commitment scheme obtained from the lossy identification scheme: Using a lossy key as commitment key, the pre-image size of a commitment can be bounded by $\varepsilon_1 \cdot |\mathcal{M}|$, where \mathcal{M} is the message space of the resulting commitment scheme and ε_1 is the bound on lossiness, see Def. 2.7.

C.1 Dual-Mode Commitment based on DDH

Let $\mathbb{G} = \langle g \rangle$ be some cyclic group of order $p \in \mathbb{P}$, which is output by some group generation algorithm $\text{GGen}(1^\lambda)$. We assume that p is a λ -bit prime. We will represent elements of \mathbb{G} implicitly via their exponents in the following way [EHK⁺13]. The representation of a group element $h = g^a \in \mathbb{G}$ is $[a]$. We can extend this notation to matrices, vectors and tuples by applying this isomorphism in every component. For example, if $\mathbf{A} = (a_{ij})_{i,j} \in \mathbb{Z}_p^{m \times n}$, then we have $[\mathbf{A}] = (g^{a_{ij}})_{i,j} \in \mathbb{G}^{m \times n}$. Let $\ell > k$ be integers. Informally, for a distribution $\mathcal{D}_{\ell,k}$ of full-rank matrices from $\mathbb{Z}_p^{\ell \times k}$ we say that the Matrix-Diffie-Hellman Assumption holds relative to GGen and $\mathcal{D}_{\ell,k}$ if no PPT algorithm can distinguish $(\mathbb{G}, g, [\mathbf{A}], [\mathbf{Aw}])$ from $(\mathbb{G}, g, [\mathbf{A}], [\mathbf{u}])$, where $\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}$, $\mathbf{w} \leftarrow_{\mathcal{S}} \mathbb{Z}_p^k$, $\mathbf{u} \leftarrow_{\mathcal{S}} \mathbb{Z}_p^\ell$. The scheme CMT_{DDH} given in Fig. 17 is a dual-mode commitment scheme based on this assumption.

C.2 Dual-Mode Commitment based on φ -Hiding

For the scheme based on the φ -Hiding assumption we follow the notation in [KMP16]. Denote the set of ℓ -bit primes by \mathbb{P}_ℓ . Let $n = p \cdot q$ be the product of two distinct $\lambda/2$ -bit primes $p, q \in \mathbb{P}_{\lambda/2}$. We denote the set of all such (n, p, q) by RSA_λ and for a relation \mathcal{R} on n, p, q we define the subset

$$\text{RSA}_\lambda[\mathcal{R}] := \{(n, p, q) \in \text{RSA}_\lambda \mid \mathcal{R}(n, p, q) = 1\}.$$

Denote Euler's totient function of n as $\varphi(n) := (p-1)(q-1)$ and the integers modulo n as \mathbb{Z}_n . The subset of invertible elements modulo n is denoted as \mathbb{Z}_n^\times . For a constant $0 < c < 1/4$ we define the distributions

$$\begin{aligned} \mathcal{I}_\lambda &:= \{(n, e) \mid e \leftarrow_{\mathcal{S}} \mathbb{P}_{c\lambda}, (n, p, q) \leftarrow_{\mathcal{S}} \text{RSA}_\lambda[\text{gcd}(e, \varphi(n) = 1)]\}, \\ \mathcal{L}_\lambda &:= \left\{ (n, e) \mid e \leftarrow_{\mathcal{S}} \mathbb{P}_{c\lambda}, (n, p, q) \leftarrow_{\mathcal{S}} \text{RSA}_\lambda \left[\begin{array}{l} p = 1 \pmod{e}, \\ p \neq 1 \pmod{e^2}, \\ q \neq 1 \pmod{e} \end{array} \right] \right\}. \end{aligned}$$

Alg Setup (1^λ) 01 $\mathbb{G} \leftarrow \text{GGen}(1^\lambda)$ 02 $\mathbf{A} \leftarrow \mathcal{D}_{\ell,k}$ 03 $\text{par} := (\mathbb{G}, [\mathbf{A}])$ 04 return par	Alg Com ($\text{ck}, m \in \mathbb{Z}_p$) 10 $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_p^k$ 11 $\mathbf{c} := [\mathbf{A}\mathbf{s} - m\mathbf{t}]$ 12 return ($\mathbf{c}, \text{dc} := \mathbf{s}$)	Alg TCom (ck, td) 18 $\mathbf{r} \xleftarrow{\$} \mathbb{Z}_p^k$ 19 $\mathbf{c} := [\mathbf{A}\mathbf{r}], St := \mathbf{r}$ 20 return (\mathbf{c}, St)
Alg Gen (par) 05 $\text{ck} := [\mathbf{t}] \xleftarrow{\$} \mathbb{G}^\ell$ 06 return ck	Alg Open ($\text{ck}, m, \text{dc}, \mathbf{c}$) 13 let $\mathbf{c} = [\mathbf{u}]$ 14 let $\text{dc} = \mathbf{s}$ 15 if $\mathbf{u} \neq [\mathbf{A}\mathbf{s} - m\mathbf{t}]$ 16 return 0 17 return 1	Alg TCol ($\text{ck}, \text{td}, St, m$) 21 let $St = \mathbf{r}, \text{td} = \mathbf{x}$ 22 return $m\mathbf{x} + \mathbf{r}$
Alg TGen (par) 07 $\mathbf{x} \xleftarrow{\$} \mathbb{Z}_p^k, [\mathbf{t}] := [\mathbf{A}\mathbf{x}]$ 08 $\text{ck} := [\mathbf{t}], \text{td} := \mathbf{x}$ 09 return (ck, td)		

Figure 17: Dual-mode commitment $\text{CMT}_{\text{MDDH}} = (\text{Setup}, \text{TSetup} = \text{Setup}, \text{Gen}, \text{TGen}, \text{Com}, \text{TCom}, \text{Open}, \text{TCol})$ with message space $\mathcal{M} := \mathbb{Z}_p$. Applying the transformation to an identification scheme we obtain $\text{ID}[\text{CMT}_{\text{MDDH}}]$, which is the (generalized) Chaum-Pedersen scheme. For simplicity, we just presented algorithm Gen with fixed $N = 1$.

Informally, the φ -Hiding assumption is that no PPT algorithm can distinguish these two distributions. The scheme CMT_φ given in Fig. 18 is a dual-mode commitment scheme based on this assumption.

Alg Setup (1^λ) 01 $\text{par} := (n, e) \xleftarrow{\$} \mathcal{L}_\lambda$ 02 return par	Alg Com ($\text{ck}, m \in \mathbb{Z}_e$) 10 $s \xleftarrow{\$} \mathbb{Z}_n^\times$ 11 $\mathbf{c} := R := s^e X^{-m}$ 12 return ($\mathbf{c}, \text{dc} := s$)	Alg TCom (ck, td) 21 $r \xleftarrow{\$} \mathbb{Z}_n^\times$ 22 $R := r^e$ 23 $\mathbf{c} := R, St := r$ 24 return (\mathbf{c}, St)
Alg TSetup (1^λ) 03 $\text{par} := (n, e) \xleftarrow{\$} \mathcal{I}_\lambda$ 04 return par	Alg Open ($\text{ck}, m, \text{dc}, \mathbf{c}$) 13 let $\mathbf{c} = R, \text{dc} = s$ 14 if $R \notin \mathbb{Z}_n^\times$: 15 return 0 16 if $s \notin \mathbb{Z}_n^\times$: 17 return 0 18 if $RX^m \neq s^e$ 19 return 0 20 return 1	Alg TCol ($\text{ck}, \text{td}, St, m$) 25 let $St = r, \text{td} = x$ 26 return $x^m r$
Alg Gen (par) 05 $\text{ck} := X \xleftarrow{\$} \mathbb{Z}_n^\times$ 06 return ck		
Alg TGen (par) 07 $x \xleftarrow{\$} \mathbb{Z}_n^\times, X := x^e$ 08 $\text{ck} := X, \text{td} := x$ 09 return (ck, td)		

Figure 18: Dual-mode commitment $\text{CMT}_\varphi = (\text{Setup}, \text{TSetup}, \text{Gen}, \text{TGen}, \text{Com}, \text{TCom}, \text{Open}, \text{TCol})$ with message space $\mathcal{M} := \mathbb{Z}_e$. All computations are done modulo n . Applying the transformation to an identification scheme we obtain $\text{ID}[\text{CMT}_\varphi]$, which is the Guillou-Quisquater identification scheme. For simplicity, we just presented algorithm Gen with fixed $N = 1$.

C.3 Dual-Mode Commitment based on DSDL

Let $\mathbb{G} = \langle g \rangle$ be a group of prime order $q \in \mathbb{P}$, which is output by some group generation algorithm $\text{GGen}(1^\lambda)$. Informally, the decisional short discrete logarithm assumption with parameter $c \in \mathbb{N}$ relative to GGen is that no PPT algorithm can distinguish between the distributions

$$\{(g, h) \mid h \xleftarrow{\$} \mathbb{G}\} \text{ and } \{(g, g^x) \mid x \xleftarrow{\$} \{0, \dots, 2^c - 1\}\}.$$

The dual-mode commitment scheme based on this assumption makes use of additional parameters $k, k' \in \mathbb{N}$ and is presented in Fig. 19.

<p>Alg Setup(1^λ)</p> 01 $\mathbb{G} \leftarrow \text{GGen}(1^\lambda)$ 02 return $\text{par} := \mathbb{G}$	<p>Alg Com(ck, m)</p> 11 $s \xleftarrow{\$} \{B_l, \dots, B_r\}$ 12 $c := R := g^s h^{-m}$ 13 $\text{coin} \xleftarrow{\$} \{0, 1\}^{k'}$ 14 if $\text{coin} = 0^{k'}$: 15 return \perp 16 return $(c, \text{dc} := s)$	<p>Alg TCom(ck, td)</p> 23 $r \xleftarrow{\$} \{0, \dots, B_r\}$ 24 $R := g^r$ 25 $c := R, St := r$ 26 return (c, St)
<p>Alg TSetup(1^λ)</p> 03 $\mathbb{G} \leftarrow \text{GGen}(1^\lambda)$ 04 return $\text{par} := \mathbb{G}$	<p>Alg Open($\text{ck}, m, \text{dc}, c$)</p> 17 let $c = R, \text{dc} = s$ 18 if $s \notin \{B_l, \dots, B_r\}$: 19 return 0 20 if $R \cdot h^m \neq g^s$: 21 return 0 22 return 1	<p>Alg TCol($\text{ck}, \text{td}, St, m$)</p> 27 let $St = r, \text{td} = x$ 28 $s = mx + r$ 29 if $s \notin \{B_l, \dots, B_r\}$: 30 return \perp 31 return s
<p>Alg Gen(par)</p> 05 $\text{ck} := h \xleftarrow{\$} \mathbb{G}$ 06 return ck		
<p>Alg TGen(par)</p> 07 $x \xleftarrow{\$} \{0, \dots, 2^c - 1\}$ 08 $h := g^x$ 09 $\text{ck} := h, \text{td} := x$ 10 return (ck, td)		

Figure 19: Dual-mode commitment $\text{CMT}_{\text{DSDL}} = (\text{Setup}, \text{TSetup}, \text{Gen}, \text{TGen}, \text{Com}, \text{TCom}, \text{Open}, \text{TCol})$ with message space $\mathcal{M} := \{0, \dots, 2^k - 1\}$ and $B_l := 2^{k+c}, B_r := 2^{k+k'+c} - 1$. Applying the transformation to an identification scheme we obtain $\text{ID}[\text{CMT}_{\text{DSDL}}]$, which is the DSDL-based lossy identification scheme in [AFLT12]. We just presented algorithm Gen with fixed $N = 1$.

C.4 Dual-Mode Commitment based on Subset Sum

Informally, the (decisional) subset sum assumption with parameters $n, M \approx 2^n \in \mathbb{N}$ states that no PPT algorithm can distinguish between the distributions

$$\{(\mathbf{a}, t) \mid \mathbf{a} \xleftarrow{\$} \mathbb{Z}_M^n, t \xleftarrow{\$} \mathbb{Z}_M\} \text{ and } \{(\mathbf{a}, \sum_{i=1}^n a_i s_i) \mid \mathbf{a} \xleftarrow{\$} \mathbb{Z}_M^n, \mathbf{s} \xleftarrow{\$} \{0, 1\}^n\}.$$

For the dual-mode commitment scheme based on this assumption we need

$$M > (2\lambda n + 1)^n \cdot 3^{2\lambda}$$

and a parameter $\rho \approx (1 - 1/n)^n$. The scheme is given in Fig. 20.

<p>Alg Gen(par)</p> <p>01 $\mathbf{a} \xleftarrow{\\$} \mathbb{Z}_M^n, \mathbf{t} \xleftarrow{\\$} \mathbb{Z}_M^n$</p> <p>02 return $\text{ck} := (\mathbf{a}, \mathbf{t})$</p> <p>Alg TGen(par)</p> <p>03 $\mathbf{X} \xleftarrow{\\$} \{0, 1\}^{n \times \lambda}, \mathbf{a} \xleftarrow{\\$} \mathbb{Z}_M^n$</p> <p>04 $\mathbf{t} := \mathbf{a}^t \mathbf{X} \pmod{M}$</p> <p>05 $\text{ck} := (\mathbf{a}, \mathbf{t}), \text{td} := \mathbf{X}$</p> <p>06 return (ck, td)</p> <p>Alg Open(ck, m, dc, c)</p> <p>07 let $c = R, \text{dc} = \mathbf{s}$</p> <p>08 if $\mathbf{s} \notin \{B_l, \dots, B_r\}^n$: return 0</p> <p>09 if $\mathbf{a}^t \mathbf{s} - \mathbf{t}^t \mathbf{m} \pmod{M} \neq R$:</p> <p>10 return 0</p> <p>11 return 1</p>	<p>Alg Com(ck, m)</p> <p>12 $\text{dc} := \mathbf{s} \xleftarrow{\\$} \{B_l, \dots, B_r\}^n$</p> <p>13 $\mathbf{c} := R := \mathbf{a}^t \mathbf{s} - \mathbf{t}^t \mathbf{m} \pmod{M}$</p> <p>14 With prob. $1 - \rho$: return \perp</p> <p>15 return (\mathbf{c}, dc)</p> <p>Alg TCom(ck, td)</p> <p>16 $\mathbf{y} \xleftarrow{\\$} \{-\lambda n, \dots, \lambda n\}^n$</p> <p>17 $R := \mathbf{a}^t \mathbf{y} \pmod{M}$</p> <p>18 return $(\mathbf{c} := R, St := \mathbf{y})$</p> <p>Alg TCol(ck, td, St, m)</p> <p>19 let $St := \mathbf{y}, \mathbf{X}$</p> <p>20 $\mathbf{s} = \mathbf{X} \mathbf{m} + \mathbf{y}$</p> <p>21 if $\mathbf{s} \notin \{B_l, \dots, B_r\}^n$: return \perp</p> <p>22 return \mathbf{s}</p>
---	--

Figure 20: Dual-mode commitment $\text{CMT}_\Sigma = (\text{Setup}, \text{TSetup}, \text{Gen}, \text{TGen}, \text{Com}, \text{TCom}, \text{Open}, \text{TCol})$ with message space $\mathcal{M} := \{0, 1\}^\lambda$ and $B_l := -\lambda n + \lambda, B_r := \lambda n + \lambda$. Algorithms Setup and TSetup return \perp . Applying the transformation to an identification scheme we obtain $\text{ID}[\text{CMT}_\Sigma]$, which is the Subset Sum-based lossy identification scheme in [AFLT12]. We just presented algorithm Gen with fixed $N = 1$.

D Descriptions of Tight Variants based on LWE

Here we formally define completely tight variants of our LWE-based signature scheme. First, the dual-mode commitment schemes CMT'_{LWE} and $\text{CMT}''_{\text{LWE}}$ are presented in Figs. 21 and 22. Applying our transformation from Sections 3.1 and 3.2, we obtain signature schemes $\text{SIG}_s[\text{ID}[\text{CMT}'_{\text{LWE}}], \text{H}, 1]$ and $\text{SIG}_s[\text{ID}[\text{CMT}''_{\text{LWE}}], \text{H}, 1]$ with tight N -MU-CMA-Corr security. These signature schemes are presented in Figs. 23 and 24.

<p>Alg TGen(par)</p> <p>01 $(\mathbf{A}, \mathbf{T}_A) \leftarrow \text{GenTrap}(1^{n+1}, 1^m, s_0, q)$</p> <p>02 $\text{ck} := \mathbf{A} \in \mathbb{Z}_q^{(n+1) \times m}, \text{td} := \mathbf{T}_A$</p> <p>03 return (ck, td)</p> <p>Alg Gen(par, 1^N)</p> <p>04 $\bar{s} \xleftarrow{s} \mathbb{Z}_q^n$</p> <p>05 for $i \in [N]$:</p> <p>06 $\bar{\mathbf{A}}_i \xleftarrow{s} \mathbb{Z}_q^{n \times m}, \bar{\mathbf{e}}_i \leftarrow D_{\mathbb{Z}^m, \alpha q}$</p> <p>07 $\text{ck}_i := \mathbf{A}_i := \begin{bmatrix} \bar{\mathbf{A}}_i \\ \bar{s}^t \bar{\mathbf{A}}_i + \bar{\mathbf{e}}_i^t \end{bmatrix}$</p> <p>08 return $(\text{ck}_1, \dots, \text{ck}_N)$</p> <p>Alg Com(ck, m)</p> <p>09 $\mathbf{Z} = [\mathbf{z}_1 \mid \dots \mid \mathbf{z}_k] \leftarrow D_{\mathbb{Z}, s}^{m \times k}$</p> <p>10 $\mathbf{U} := \mathbf{AZ} + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m}^t \end{bmatrix}$</p> <p>11 if $\exists i \in [k] : \ \mathbf{z}_i\ > s \cdot \sqrt{m}$: return \perp</p> <p>12 return (c := \mathbf{U}, dc := \mathbf{Z})</p>	<p>Alg TCom(ck, td)</p> <p>13 $\mathbf{U} \xleftarrow{s} \mathbb{Z}_q^{(n+1) \times k}$</p> <p>14 return ($\mathbf{U}, St := \mathbf{U}$)</p> <p>Alg TCol(ck, td, St, \mathbf{m})</p> <p>15 $[\mathbf{y}_1 \mid \dots \mid \mathbf{y}_k]$</p> <p style="text-align: right;">$:= \mathbf{U} - \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m}^t \end{bmatrix}$</p> <p>16 for $i \in [k]$:</p> <p>17 $\mathbf{z}_i \leftarrow \text{SampleD}(\mathbf{A}, \mathbf{T}_A, \mathbf{y}_i, s)$</p> <p>18 if $\ \mathbf{z}_i\ > s \cdot \sqrt{m}$: return \perp</p> <p>19 return $\mathbf{Z} := [\mathbf{z}_1 \mid \dots \mid \mathbf{z}_k]$</p> <p>Alg Open(ck, m, \mathbf{Z}, \mathbf{U})</p> <p>20 if $\exists i \in [k] : \ \mathbf{z}_i\ > s \cdot \sqrt{m}$:</p> <p>21 return 0</p> <p>22 if $\mathbf{AZ} + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m}^t \end{bmatrix} \neq \mathbf{U}$:</p> <p>23 return 0</p> <p>24 return 1</p>
--	--

Figure 21: The dual-mode commitment $\text{CMT}'_{\text{LWE}} = (\text{Setup}, \text{TSetup}, \text{Gen}, \text{TGen}, \text{Com}, \text{TCom}, \text{Open}, \text{TCol})$ with message space $\mathcal{M} = \{0, 1\}^k$, where $\text{Setup} = \text{TSetup}$ sets parameters par as in the text. Note that CMT'_{LWE} is CMT_{LWE} with $k = 1$ applied to each message bit individually.

<p>Alg TGen(par)</p> <p>01 $(\mathbf{A}, \mathbf{T}_A) \leftarrow \text{GenTrap}(1^n, 1^m, s_0, q)$</p> <p>02 $\text{ck} := \mathbf{A} \in \mathbb{Z}_q^{n \times m}, \text{td} := \mathbf{T}_A$</p> <p>03 return (ck, td)</p> <p>Alg Gen(par, 1^N)</p> <p>04 $\bar{s} \xleftarrow{s} \mathbb{Z}_q^{n-1}$</p> <p>05 for $i \in [N]$:</p> <p>06 $\bar{\mathbf{A}}_i \xleftarrow{s} \mathbb{Z}_q^{(n-1) \times m}, \bar{\mathbf{e}}_i \leftarrow D_{\mathbb{Z}^m, \alpha q}$</p> <p>07 $\text{ck}_i := \mathbf{A}_i := \begin{bmatrix} \bar{\mathbf{A}}_i \\ \bar{s}^t \bar{\mathbf{A}}_i + \bar{\mathbf{e}}_i^t \end{bmatrix}$</p> <p>08 return $(\text{ck}_1, \dots, \text{ck}_N)$</p> <p>Alg Com(ck, m)</p> <p>09 $\mathbf{Z} = [\mathbf{z}_1 \mid \dots \mid \mathbf{z}_{kn \log q}] \leftarrow D_{\mathbb{Z}, s}^{m \times kn \log q}$</p> <p>10 $\mathbf{U} := \mathbf{AZ} + \mathbf{m}^t \otimes \mathbf{G}$</p> <p>11 if $\exists i \in [kn \log q] : \ \mathbf{z}_i\ > s \cdot \sqrt{m}$:</p> <p>12 return \perp</p> <p>13 return (c := \mathbf{U}, dc := \mathbf{Z})</p>	<p>Alg TCom(ck, td)</p> <p>14 $\mathbf{U} \xleftarrow{s} \mathbb{Z}_q^{n \times kn \log q}$</p> <p>15 return ($\mathbf{U}, St := \mathbf{U}$)</p> <p>Alg TCol(ck, td, St, \mathbf{m})</p> <p>16 $[\mathbf{y}_1 \mid \dots \mid \mathbf{y}_{kn \log q}] := \mathbf{U} - \mathbf{m}^t \otimes \mathbf{G}$</p> <p>17 for $i \in [kn \log q]$:</p> <p>18 $\mathbf{z}_i \leftarrow \text{SampleD}(\mathbf{A}, \mathbf{T}_A, \mathbf{y}_i, s)$</p> <p>19 if $\ \mathbf{z}_i\ > s \cdot \sqrt{m}$: return \perp</p> <p>20 return $\mathbf{Z} := [\mathbf{z}_1 \mid \dots \mid \mathbf{z}_k]$</p> <p>Alg Open(ck, m, \mathbf{Z}, \mathbf{U})</p> <p>21 if $\exists i \in [kn \log q] : \ \mathbf{z}_i\ > s \cdot \sqrt{m}$:</p> <p>22 return 0</p> <p>23 if $\mathbf{AZ} + \mathbf{m}^t \otimes \mathbf{G} \neq \mathbf{U}$:</p> <p>24 return 0</p> <p>25 return 1</p>
---	---

Figure 22: The dual-mode commitment $\text{CMT}''_{\text{LWE}} = (\text{Setup}, \text{TSetup}, \text{Gen}, \text{TGen}, \text{Com}, \text{TCom}, \text{Open}, \text{TCol})$ with message space $\mathcal{M} = \{0, 1\}^k$, where $\text{Setup} = \text{TSetup}$ sets parameters par as in the text.

<p>Alg Gen(par)</p> <pre> 01 $(\mathbf{A}_0, \mathbf{T}_0) \leftarrow \text{GenTrap}(1^{n+1}, 1^m, s_0, q)$ 02 $(\mathbf{A}_1, \mathbf{T}_1) \leftarrow \text{GenTrap}(1^{n+1}, 1^m, s_0, q)$ 03 $b \xleftarrow{s} \{0, 1\}, \text{sk} := (b, \mathbf{T}_b)$ 04 $\text{pk} := (\mathbf{A}_0, \mathbf{A}_1)$ 05 return (pk, sk) </pre> <p>Alg Ver(pk, m, $\sigma = (\mathbf{U}_0, \mathbf{U}_1, \mathbf{Z}_0, \mathbf{Z}_1)$)</p> <pre> 06 $\mathbf{m}_1 \leftarrow \text{H}(0, \text{pk}, \mathbf{U}_0, \mathbf{m})$ 07 $\mathbf{m}_0 \leftarrow \text{H}(1, \text{pk}, \mathbf{U}_1, \mathbf{m})$ 08 let $\mathbf{Z}_0 = [\mathbf{z}_{0,1} \mid \cdots \mid \mathbf{z}_{0,k}]$ 09 let $\mathbf{Z}_1 = [\mathbf{z}_{1,1} \mid \cdots \mid \mathbf{z}_{1,k}]$ 10 if $\exists i \in [k] : \ \mathbf{z}_{0,i}\ > s \cdot \sqrt{m} : \text{return } 0$ 11 if $\exists i \in [k] : \ \mathbf{z}_{1,i}\ > s \cdot \sqrt{m} : \text{return } 0$ 12 if $\mathbf{A}_0 \mathbf{Z}_0 + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m}_0^t \end{bmatrix} \neq \mathbf{U}_0 :$ 13 return 0 14 if $\mathbf{A}_1 \mathbf{Z}_1 + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m}_1^t \end{bmatrix} \neq \mathbf{U}_1 :$ 15 return 0 16 return 1 </pre>	<p>Alg Sig(sk, m)</p> <pre> 17 let $\sigma = (b, \mathbf{T}_b)$ 18 $\mathbf{U}_b \xleftarrow{s} \mathbb{Z}_q^{(n+1) \times k}$ 19 $\mathbf{m}_{1-b} \leftarrow \text{H}(b, \text{pk}, \mathbf{U}_b, \mathbf{m})$ 20 $\mathbf{Z}_{1-b} = [\mathbf{z}_{1-b,1} \mid \cdots \mid \mathbf{z}_{1-b,k}]$ $\leftarrow D_{\mathbb{Z},s}^{m \times k}$ 21 \mathbf{U}_{1-b} $:= \mathbf{A}_{1-b} \mathbf{Z}_{1-b} + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m}_{1-b}^t \end{bmatrix}$ 22 $\mathbf{m}_b \leftarrow \text{H}(1-b, \text{pk}, \mathbf{U}_{1-b}, \mathbf{m})$ 23 $[\mathbf{y}_1 \mid \cdots \mid \mathbf{y}_k]$ $:= \mathbf{U}_b - \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rfloor \cdot \mathbf{m}_b^t \end{bmatrix}$ 24 for $i \in [k] :$ 25 $\mathbf{z}_{b,i} \leftarrow \text{SampleD}(\mathbf{A}_b, \mathbf{T}_b, \mathbf{y}_i, s)$ 26 $\mathbf{Z}_b := [\mathbf{z}_{b,1} \mid \cdots \mid \mathbf{z}_{b,k}]$ 27 if $\exists i \in [k] : \ \mathbf{z}_{0,i}\ > s \cdot \sqrt{m} :$ 28 return 0 29 if $\exists i \in [k] : \ \mathbf{z}_{1,i}\ > s \cdot \sqrt{m} :$ 30 return 0 31 return $\sigma := (\mathbf{U}_0, \mathbf{U}_1, \mathbf{Z}_0, \mathbf{Z}_1)$ </pre>
--	--

Figure 23: The signature scheme $\text{SIG}_s[\text{ID}[\text{CMT}'_{\text{LWE}}, \text{H}, 1] = (\text{Setup}, \text{Gen}, \text{Sig}, \text{Ver})$, where Setup sets parameters as in Section 4.1.

<p>Alg Gen(par)</p> <pre> 01 $(\mathbf{A}_0, \mathbf{T}_0) \leftarrow \text{GenTrap}(1^n, 1^m, s_0, q)$ 02 $(\mathbf{A}_1, \mathbf{T}_1) \leftarrow \text{GenTrap}(1^n, 1^m, s_0, q)$ 03 $b \xleftarrow{s} \{0, 1\}, \text{sk} := (b, \mathbf{T}_b)$ 04 $\text{pk} := (\mathbf{A}_0, \mathbf{A}_1)$ 05 return (pk, sk) </pre> <p>Alg Ver(pk, m, $\sigma = (\mathbf{U}_0, \mathbf{U}_1, \mathbf{Z}_0, \mathbf{Z}_1)$)</p> <pre> 06 $\mathbf{m}_1 \leftarrow \text{H}(0, \text{pk}, \mathbf{U}_0, \mathbf{m})$ 07 $\mathbf{m}_0 \leftarrow \text{H}(1, \text{pk}, \mathbf{U}_1, \mathbf{m})$ 08 let $\mathbf{Z}_0 = [\mathbf{z}_{0,1} \mid \cdots \mid \mathbf{z}_{0,kn \log q}]$ 09 let $\mathbf{Z}_1 = [\mathbf{z}_{1,1} \mid \cdots \mid \mathbf{z}_{1,kn \log q}]$ 10 if $\exists i \in [kn \log q] : \ \mathbf{z}_{0,i}\ > s \cdot \sqrt{m} :$ 11 return 0 12 if $\exists i \in [kn \log q] : \ \mathbf{z}_{1,i}\ > s \cdot \sqrt{m} :$ 13 return 0 14 if $\mathbf{A}_0 \mathbf{Z}_0 + \mathbf{m}_0^t \otimes \mathbf{G} \neq \mathbf{U}_0 : \text{return } 0$ 15 if $\mathbf{A}_1 \mathbf{Z}_1 + \mathbf{m}_1^t \otimes \mathbf{G} \neq \mathbf{U}_1 : \text{return } 0$ 16 return 1 </pre>	<p>Alg Sig(sk, m)</p> <pre> 17 let $\sigma = (b, \mathbf{T}_b)$ 18 $\mathbf{U}_b \xleftarrow{s} \mathbb{Z}_q^{n \times kn \log q}$ 19 $\mathbf{m}_{1-b} \leftarrow \text{H}(b, \text{pk}, \mathbf{U}_b, \mathbf{m})$ 20 $\mathbf{Z}_{1-b} = [\mathbf{z}_{1-b,1} \mid \cdots \mid \mathbf{z}_{1-b,kn \log q}]$ $\leftarrow D_{\mathbb{Z},s}^{m \times kn \log q}$ 21 $\mathbf{U}_{1-b} := \mathbf{A}_{1-b} \mathbf{Z}_{1-b} + \mathbf{m}_{1-b}^t \otimes \mathbf{G}$ 22 $\mathbf{m}_b \leftarrow \text{H}(1-b, \text{pk}, \mathbf{U}_{1-b}, \mathbf{m})$ 23 $[\mathbf{y}_1 \mid \cdots \mid \mathbf{y}_{kn \log q}] := \mathbf{U}_b - \mathbf{m}_b^t \otimes \mathbf{G}$ 24 for $i \in [kn \log q] :$ 25 $\mathbf{z}_{b,i} \leftarrow \text{SampleD}(\mathbf{A}_b, \mathbf{T}_b, \mathbf{y}_i, s)$ 26 $\mathbf{Z}_b := [\mathbf{z}_{b,1} \mid \cdots \mid \mathbf{z}_{b,kn \log q}]$ 27 if $\exists i \in [kn \log q] : \ \mathbf{z}_{0,i}\ > s \cdot \sqrt{m} :$ 28 return 0 29 if $\exists i \in [kn \log q] : \ \mathbf{z}_{1,i}\ > s \cdot \sqrt{m} :$ 30 return 0 31 return $\sigma := (\mathbf{U}_0, \mathbf{U}_1, \mathbf{Z}_0, \mathbf{Z}_1)$ </pre>
--	--

Figure 24: The signature scheme $\text{SIG}_s[\text{ID}[\text{CMT}''_{\text{LWE}}, \text{H}, 1] = (\text{Setup}, \text{Gen}, \text{Sig}, \text{Ver})$, where Setup sets parameters as in Section 4.1.

E Instantiation based on Secure Group Actions

We will now show that the lossy identification scheme in [EKP20] is essentially obtained from a dual-mode commitment scheme. More importantly, our analysis shows that the weak algebraic structure of a group action is enough to obtain tight multi-user security with adaptive corruptions. We use the fact that our definition allows lossy keys to be correlated, and we can tightly switch to a set of N lossy correlated keys.

The lossy identification scheme in [EKP20] can be interpreted generically by using a specific group acting on a set and a computational assumption related to it as a black box. Similar to that, [ADMP20] define and use cryptographic group actions generically to construct a variety of primitives. We recall the algebraic definition of a group action.

Definition E.1 (Group Action). Let \mathbb{G} be a group with identity element $1_{\mathbb{G}}$ and \mathbb{S} be a set. A map $\star : \mathbb{G} \times \mathbb{S} \rightarrow \mathbb{S}$ is a group action if for all $g, h \in \mathbb{G}$ and all $X \in \mathbb{S}$ the following hold:

$$\begin{aligned} 1_{\mathbb{G}} \star X &= X, \\ (gh) \star X &= g \star (h \star X). \end{aligned}$$

Further, we say that the group acts freely and transitively on the set if for all $X \in \mathbb{S}$ the map

$$\mathbb{G} \longrightarrow \mathbb{S}, \quad g' \longmapsto g' \star X$$

is a bijection.

Let $\mathbb{A} := (\mathbb{G}, g, n, \mathbb{S}, \star) \leftarrow \text{GAGen}(1^\lambda)$ be a PPT algorithm that outputs the description of a cyclic group $\mathbb{G} = \langle g \rangle$ of order n , the description of a set \mathbb{S} , and an efficiently computable group action $\star : \mathbb{G} \times \mathbb{S} \rightarrow \mathbb{S}$. We assume that the group acts freely and transitively on \mathbb{S} and that membership and equality testing on \mathbb{G} and \mathbb{S} are efficient. Finally, we make use of the following assumption.

Definition E.2 (Group Action Diffie-Hellman). We say that the GADH assumption holds relative to GAGen , if for every PPT algorithm \mathcal{B} the following advantage is negligible in λ :

$$\begin{aligned} \text{Adv}_{\mathcal{B}, \text{GAGen}}^{\text{GADH}}(\lambda) := & \left| \Pr \left[\mathcal{B}(\mathbb{A}, E, H, g^a \star E, g^a \star H) = 1 \mid \begin{array}{l} \mathbb{A} \leftarrow \text{GAGen}(1^\lambda), \\ E, H \xleftarrow{\$} \mathbb{S}, a \xleftarrow{\$} \mathbb{Z}_n \end{array} \right] \right. \\ & \left. - \Pr \left[\mathcal{B}(\mathbb{A}, E, H, g^a \star E, g^b \star H) = 1 \mid \begin{array}{l} \mathbb{A} \leftarrow \text{GAGen}(1^\lambda), \\ E, H \xleftarrow{\$} \mathbb{S}, a, b \xleftarrow{\$} \mathbb{Z}_n \end{array} \right] \right|. \end{aligned}$$

Note that if \star is exponentiation in a cyclic prime order group, then this is exactly the DDH assumption. However, the more interesting case is if \mathbb{G} is a class group and \mathbb{S} a set of supersingular elliptic curves, see [CLM⁺18, EKP20]. Also, note that this assumption is implied by assuming that the group action is weakly pseudorandom, as defined in [ADMP20].

We will now construct a (multi-key) dual-mode commitment CMT_{ga} based on the generic assumption GADH with tight security. Our scheme is presented in Fig. 25. It makes use of a parallel repetition parameter $t \in \omega(\log \lambda)$ that ensures that the message space is large enough to apply Lem. 3.4. The scheme CMT_{ga} is ρ -complete with $\rho = 1$, which follows easily by inspection.

Remark E.3 (Almost Tight Variants). We note that [EKP20] present different variations of their scheme using a multi-instance variant of the underlying assumption, thereby showing a tradeoff between almost tightness and communication costs. Our analysis can also be combined with this tradeoff. However, this tradeoff can only increase the challenge space to polynomial size and thus can not remove parallel repetition entirely. Furthermore, the technical new part of our analysis is tight multi-key indistinguishability, which is why we want to keep the presentation clean and simple and focus completely tight variant of their scheme.

Lemma E.4 (Trapdoor Property). *The scheme CMT_{ga} perfectly satisfies the ε_t -trapdoor property, i.e. $\varepsilon_t = 0$.*

Proof. Let $(\mathbb{A}, E, H) \in \text{TSetup}(1^\lambda), ((E', H'), a) \in \text{TGen}((\mathbb{A}, E, H))$. Let $m \in \{0, 1\}^\lambda$ be a message. Using a standard hybrid argument, it is sufficient to consider a fixed $i \in [t]$ and the distributions of

<p>Alg Setup(1^λ)</p> <pre> 01 $\mathbb{A} \leftarrow \text{GAGen}(1^\lambda)$ 02 $E, H \xleftarrow{\\$} \mathbb{S}$ 03 return par := (\mathbb{A}, E, H) Alg Gen(par, 1^N) 04 $a_0, b_0 \xleftarrow{\\$} \mathbb{Z}_n$ 05 $E'_0 := g^{a_0} \star E, H'_0 := g^{b_0} \star H$ 06 for $i \in [N]$: 07 $d_i \xleftarrow{\\$} \mathbb{Z}_n$ 08 $E'_i := g^{d_i} \star E'_0, H'_i := g^{d_i} \star H'_0$ 09 $\text{ck}_i := (E'_i, H'_i)$ 10 return ($\text{ck}_1, \dots, \text{ck}_N$) Alg TGen(par) 11 $a \xleftarrow{\\$} \mathbb{Z}_n$ 12 $E' := g^a \star E, H' := g^a \star H$ 13 $\text{ck} := (E', H'), \text{td} := a$ 14 return (ck, td) Alg TCol($\text{ck}, \text{td}, St, m$) 15 let $St = (r_1, \dots, r_t), \text{td} = a$ 16 for $i \in [t]$: $s_i := r_i - m_i \cdot a$ 17 return $\text{dc} := (s_1, \dots, s_t)$ </pre>	<p>Alg Com(ck, m)</p> <pre> 18 for $i \in [t]$: 19 $s_i \xleftarrow{\\$} \mathbb{Z}_n$ 20 if $m_i = 0$: 21 $R_{i,0} := g^{s_i} \star E, R_{i,1} := g^{s_i} \star H$ 22 if $m_i = 1$: 23 $R_{i,0} := g^{s_i} \star E', R_{i,1} := g^{s_i} \star H'$ 24 $\mathbf{c} := ((R_{1,0}, R_{1,1}), \dots, (R_{t,0}, R_{t,1}))$ 25 $\text{dc} := (s_1, \dots, s_t)$ 26 return (\mathbf{c}, dc) Alg Open($\text{ck}, m, \text{dc}, \mathbf{c}$) 27 let $\text{ck} = (E', H')$ 28 let $\text{dc} = (s_1, \dots, s_t)$ 29 let $\mathbf{c} = ((R_{1,0}, R_{1,1}), \dots, (R_{t,0}, R_{t,1}))$ 30 for $i \in [t]$: 31 if $m_i = 0$: 32 if $g^{s_i} \star E \neq R_{i,0} \vee g^{s_i} \star H \neq R_{i,1}$: 33 return 0 34 if $m_i = 1$: 35 if $g^{s_i} \star E' \neq R_{i,0} \vee g^{s_i} \star H' \neq R_{i,1}$: 36 return 0 37 return 1 Alg TCom(ck, td) 38 for $i \in [t]$: 39 $r_i \xleftarrow{\\$} \mathbb{Z}_n$ 40 $R_{i,0} := g^{r_i} \star E, R_{i,1} := g^{r_i} \star H$ 41 $\mathbf{c} := ((R_{1,0}, R_{1,1}), \dots, (R_{t,0}, R_{t,1}))$ 42 $St := (r_1, \dots, r_t)$ 43 return (\mathbf{c}, St) </pre>
--	--

Figure 25: Dual-mode commitment $\text{CMT}_{ga} = (\text{Setup}, \text{TSetup} = \text{Setup}, \text{Gen}, \text{TGen}, \text{Com}, \text{TCom}, \text{Open}, \text{TCol})$ with message space $\mathcal{M} := \{0, 1\}^t$.

$((R_{i,0}, R_{i,1}), m_i, s_i)$ output by Com or TCom, TCol . Note that as \star is transitive and free, in both cases $(R_{i,0}, R_{i,1})$ is uniformly distributed over \mathbb{S}^2 and s_i is uniformly distributed over \mathbb{Z}_n conditioned on the relation

$$\begin{aligned} & (m_i = 0 \wedge R_{i,0} = g^{s_i} \star E \wedge R_{i,1} = g^{s_i} \star H) \\ & \vee (m_i = 1 \wedge R_{i,0} = g^{s_i} \star E' \wedge R_{i,1} = g^{s_i} \star H'). \end{aligned}$$

□

Lemma E.5 (Key Indistinguishability). *Let $N = \text{poly}(\lambda)$ be a natural number. Then CMT_{ga} satisfies key indistinguishability, under the GADH assumption relative to GAGen , where for every PPT algorithm \mathcal{A} there exists a PPT algorithm \mathcal{B} , such that $\mathbf{T}(\mathcal{B}) \approx \mathbf{T}(\mathcal{A})$ and*

$$\text{Adv}_{\mathcal{A}, \text{CMT}_{ga}}^{N\text{-keydist}}(\lambda) \leq \text{Adv}_{\mathcal{B}, \text{GAGen}}^{\text{GADH}}(\lambda).$$

Proof. We show the statement via a sequence of distributions. For a cleaner presentation, we ignore the parameters (\mathbb{A}, E, H) in these distributions. In all of them, they should be sampled as $\mathbb{A} \leftarrow \text{GAGen}(1^\lambda), E, H \xleftarrow{\$} \mathbb{S}$. The first distribution is given by the algorithms TSetup and TGen , i.e.

$$\mathcal{D}_1 := \left\{ \left((E'_1, H'_1), \dots, (E'_N, H'_N) \right) \left| \begin{array}{l} \text{for } i \in [N] : a_i \xleftarrow{\$} \mathbb{Z}_n, \\ E'_i := g^{a_i} \star E, H'_i := g^{a_i} \star H \end{array} \right. \right\}.$$

In the second distribution, we define $a_i := a_0 + d_i$ for $a_0 \xleftarrow{\$} \mathbb{Z}_n, d_i \xleftarrow{\$} \mathbb{Z}_n$ and $i \in [N]$. More formally:

$$\mathcal{D}_2 := \left\{ ((E'_1, H'_1), \dots, (E'_N, H'_N)) \left| \begin{array}{l} a_0 \xleftarrow{\$} \mathbb{Z}_n, \text{ for } i \in [N] : d_i \xleftarrow{\$} \mathbb{Z}_n, \\ E'_i := g^{a_0+d_i} \star E, H'_i := g^{a_0+d_i} \star H \end{array} \right. \right\}.$$

Clearly, all a_i are still uniformly random and independent over \mathbb{Z}_n . So \mathcal{D}_1 and \mathcal{D}_2 are exactly the same. We can rewrite \mathcal{D}_3 further without changing the distribution as follows.

$$\mathcal{D}_3 := \left\{ ((E'_1, H'_1), \dots, (E'_N, H'_N)) \left| \begin{array}{l} a_0 \xleftarrow{\$} \mathbb{Z}_n, E'_0 := g^{a_0} \star E, H'_0 := g^{a_0} \star H \\ \text{for } i \in [N] : d_i \xleftarrow{\$} \mathbb{Z}_n, \\ E'_i := g^{d_i} \star E'_0, H'_i := g^{d_i} \star H'_0 \end{array} \right. \right\}.$$

Next, we apply the GADH assumption on the elements (E, H, E'_0, H'_0) , to see that the distribution

$$\mathcal{D}_4 := \left\{ ((E'_1, H'_1), \dots, (E'_N, H'_N)) \left| \begin{array}{l} a_0, b_0 \xleftarrow{\$} \mathbb{Z}_n, E'_0 := g^{a_0} \star E, H'_0 := g^{b_0} \star H \\ \text{for } i \in [N] : d_i \xleftarrow{\$} \mathbb{Z}_n, \\ E'_i := g^{d_i} \star E'_0, H'_i := g^{d_i} \star H'_0 \end{array} \right. \right\}$$

is computationally indistinguishable from \mathcal{D}_3 . Finally, note that \mathcal{D}_4 is exactly the distribution output by Gen , which finishes the proof. \square

Lemma E.6 (Binding Property). *For any $N = \text{poly}(\lambda)$ the scheme CMT_{ga} is (ε_b, N) -statistically binding, with $\varepsilon_b \leq 1/n$.*

Proof. Consider the experiment

$$(\mathbb{A}, E, H) \leftarrow \text{Setup}(1^\lambda), ((E'_1, H'_1), \dots, (E'_N, H'_N)) \leftarrow \text{Gen}(\text{par}, 1^N).$$

Recall that in $\text{Gen}(\text{par}, 1^N)$ the exponents a_0, b_0 are sampled uniformly at random from \mathbb{Z}_n . Denote the event that $a_0 = b_0$ by C . Clearly, we have

$$\Pr[\text{C}] = \frac{1}{n}.$$

We have to bound the following probability of the event

$$\mathbf{E} := \left(\exists i \in [N], c \in \mathcal{C}, m \neq m' \in \mathcal{M} : \begin{array}{l} \exists dc \in \mathcal{D} : \text{Open}(\text{ck}_i, m, dc, c) = 1 \\ \wedge \exists dc' \in \mathcal{D} : \text{Open}(\text{ck}_i, m', dc', c) = 1 \end{array} \right).$$

Clearly, we have

$$\Pr[\mathbf{E}] = \Pr[\mathbf{E} \mid \text{C}] \Pr[\text{C}] + \Pr[\mathbf{E} \mid \neg\text{C}] \Pr[\neg\text{C}] \leq \frac{1}{n} + \Pr[\mathbf{E} \mid \neg\text{C}].$$

Thus, it remains to show that \mathbf{E} can not occur if $a_0 \neq b_0$. Assume towards contradiction, that $i \in [N], j \in [t], m_j \neq m'_j, (R_{j,0}, R_{j,1}) \in \mathbb{S}^2$ are such that there exist s_j, s'_j that make Open accept. Without loss of generality, assume that $m_j = 0, m'_j = 1$. Then we have

$$g^{s_j} \star E = R_{j,0} = g^{s'_j} \star E'_i \wedge g^{s_j} \star H = R_{j,0} = g^{s'_j} \star H'_i,$$

which implies

$$E'_i = g^{s_j - s'_j} \star E \wedge H'_i = g^{s_j - s'_j} \star H.$$

As \star acts freely, we have $g^{a_0} = g^{s_j - s'_j} = g^{b_0}$, and thus $a_0 = b_0$, which is a contradiction. \square

For completeness, we apply our approach from Section 3 to CMT_{ga} and obtain a tightly MU-CMA-Corr secure signature scheme based on the GADH assumption. As CMT_{ga} is perfectly complete, $\ell = 1$ repetition in the sequential OR proof is enough. The scheme is presented in Fig. 26.

Alg Setup(1^λ)

```

01  $\mathbb{A} \leftarrow \text{GAGen}(1^\lambda)$ 
02  $E, H \xleftarrow{\$} \mathbb{S}$ 
03 return  $\text{par} := (\mathbb{A}, E, H)$ 

Alg Ver( $\text{pk}, \text{m}, \sigma = (\text{cmt}_0, \text{cmt}_1, \text{rs}_0, \text{rs}_1)$ )
04 let  $\text{cmt}_0 =$ 
05  $((R_{0,1,0}, R_{0,1,1}), \dots, (R_{0,t,0}, R_{0,t,1}))$ 
06 let  $\text{cmt}_1 =$ 
07  $((R_{1,1,0}, R_{1,1,1}), \dots, (R_{1,t,0}, R_{1,t,1}))$ 
08 let  $\text{rs}_0 = (s_{0,1}, \dots, s_{0,t})$ 
09 let  $\text{rs}_1 = (s_{1,1}, \dots, s_{1,t})$ 
10  $\text{ch}_1 \leftarrow \text{H}(0, \text{pk}, \text{cmt}_0, \text{m})$ 
11  $\text{ch}_0 \leftarrow \text{H}(1, \text{pk}, \text{cmt}_1, \text{m})$ 
12 for  $i \in [t]$  :
13   if  $\text{ch}_{0,i} = 0$  :
14     if  $g^{s_{0,i}} \star E \neq R_{0,i,0}$  : return 0
15     if  $g^{s_{0,i}} \star H \neq R_{0,i,1}$  : return 0
16   if  $\text{ch}_{0,i} = 1$  :
17     if  $g^{s_{0,i}} \star E'_0 \neq R_{0,i,0}$  : return 0
18     if  $g^{s_{0,i}} \star H'_0 \neq R_{0,i,1}$  : return 0
19   if  $\text{ch}_{1,i} = 0$  :
20     if  $g^{s_{1,i}} \star E \neq R_{1,i,0}$  : return 0
21     if  $g^{s_{1,i}} \star H \neq R_{1,i,1}$  : return 0
22   if  $\text{ch}_{1,i} = 1$  :
23     if  $g^{s_{1,i}} \star E'_1 \neq R_{1,i,0}$  : return 0
24     if  $g^{s_{1,i}} \star H'_1 \neq R_{1,i,1}$  : return 0
25 return 1

```

Alg Gen(par)

```

26  $a_0 \xleftarrow{\$} \mathbb{Z}_n$ 
27  $E'_0 := g^{a_0} \star E, H'_0 := g^{a_0} \star H$ 
28  $\text{pk}_0 := (E'_0, H'_0), \text{sk}_0 := a_0$ 
29  $a_1 \xleftarrow{\$} \mathbb{Z}_n$ 
30  $E'_1 := g^{a_1} \star E, H'_1 := g^{a_1} \star H$ 
31  $\text{pk}_1 := (E'_1, H'_1), \text{sk}_1 := a_1$ 
32  $b \xleftarrow{\$} \{0, 1\}, \text{sk} := (b, \text{sk}_b)$ 
33  $\text{pk} := (\text{pk}_0, \text{pk}_1)$ 
34 return  $(\text{pk}, \text{sk})$ 

```

Alg Sig(sk, m)

```

35 for  $i \in [t]$  :
36    $r_{b,i} \xleftarrow{\$} \mathbb{Z}_n$ 
37    $R_{b,i,0} := g^{r_{b,i}} \star E, R_{b,i,1} := g^{r_{b,i}} \star H$ 
38  $\text{cmt}_b :=$ 
39  $((R_{b,1,0}, R_{b,1,1}), \dots, (R_{b,t,0}, R_{b,t,1}))$ 
40  $\text{ch}_{1-b} \leftarrow \text{H}(b, \text{pk}, \text{cmt}_b, \text{m})$ 
41 for  $i \in [t]$  :
42    $s_{1-b,i} \xleftarrow{\$} \mathbb{Z}_n$ 
43   if  $\text{ch}_{1-b,i} = 0$  :
44      $R_{1-b,i,0} := g^{s_{1-b,i}} \star E$ 
45      $R_{1-b,i,1} := g^{s_{1-b,i}} \star H$ 
46   if  $\text{ch}_{1-b,i} = 1$  :
47      $R_{1-b,i,0} := g^{s_{1-b,i}} \star E'_{1-b}$ 
48      $R_{1-b,i,1} := g^{s_{1-b,i}} \star H'_{1-b}$ 
49  $\text{cmt}_{1-b} := ((R_{1-b,1,0}, R_{1-b,1,1}),$ 
50  $\dots, (R_{1-b,t,0}, R_{1-b,t,1}))$ 
51  $\text{ch}_b \leftarrow \text{H}(1-b, \text{pk}, \text{cmt}_{1-b}, \text{m})$ 
52 for  $i \in [t]$  :  $s_{b,i} := r_{b,i} - \text{ch}_{b,i} \cdot a_b$ 
53  $\text{rs}_0 := (s_{0,1}, \dots, s_{0,t})$ 
54  $\text{rs}_1 := (s_{1,1}, \dots, s_{1,t})$ 
55 return  $\sigma := (\text{cmt}_0, \text{cmt}_1, \text{rs}_0, \text{rs}_1)$ 

```

Figure 26: The signature scheme $\text{SIG}_s[\text{ID}[\text{CMT}_{ga}], \text{H}, 1] = (\text{Setup}, \text{Gen}, \text{Sig}, \text{Ver})$ based on the GADH assumption.