Lattice-based Signatures with Tight Adaptive Corruptions and More

Jiaxin Pan[⋆1] and Benedikt Wagner[⋆⋆2]

[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.

⋆ 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).

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. [8,9]). Many works (e.g. [16,12,26]) also consider a relaxed tightness notion, called "*almost tight*", where L depends at most linearly on the security parameter λ. We do 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 [32] (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$^{\mathsf{corr}}$ security in [6,31]. We note that there is a weaker notion of multi-user security considered in [41,49] (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 [10] and Canetti-Krawczyk [14] 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. [6] 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 [31,43,38] for constructing more efficient AKE protocols, and the notion is also used to prove the tight security of real-world protocols [22,20]. 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, [11,26]) or tight MU-CMA security (for instance, [41,49]). None of them will lead to a tightly MU-CMA-Corr secure scheme. Furthermore, Bader et al.[7] 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 [5,6,1,31,21,35] 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. [35] identified a gap in the security proof of the compact and tightly MU-CMA-Corr-secure scheme in [6] and closed this gap by following the blueprint of the pairing-based HIBE scheme in [42].

We highlight that the tight lattice-based signature schemes in [2,11,12] and isogeny-based scheme in [24] 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 [5,6] and [1, 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 [55] 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 [1, 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 [21] a generic construction was proposed by transforming a lossy identification (ID) scheme [2] to a tightly MU-CMA-Corr secure signature scheme using the sequential OR proof technique [4,25]. 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 [31,35] crucially rely on number-theoretic assumptions and the underlying algebraic structure. More precisely, [31] requires

the Decisional Diffie-Hellman (DDH) assumption and a proof system for the equality of discrete logarithms, and the compact scheme in [35] 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 Learing 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 (see our full version) 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 [44] and Ducas et al. [23], 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	$\lvert\mathsf{sk}\rvert$	$\lvert\mathsf{pk}\rvert$	$\lvert\sigma\rvert$
GPV [29]	SIS	N	T	M	mz
Lyu [44]	SIS	$Q_N/\mathsf{Adv}_{\mathcal{A}}$	mn	$n^2z + M$	$\omega(\log\lambda) + mz$
DDLL [23]	SIS	$Q_N/\mathsf{Adv}_{\mathcal{A}}$	M	$mn + M$	$m + n + mz$
AFLT [2]	RLWE	N	$2nz$	nz	$3nz$
KLS [40]	MLWE	N	$2knz$	k^2nz	$3knz$
Ours (Fig. 6)	LWE	λ	$1 + T$	$4M$	$(4n + 2m)z$
Ours (full version)	LWE	1	$1 + T$	$2M$	$(2n^2 + 2nm)z$
Ours (full version)	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 $q = \mathsf{poly}(n)$ 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. [21] which used the *sequential* OR proofs of Abe et al. [4] and Fischlin et al. [25]. The abstraction is a generic transformation from dual-mode commitment to lossy ID, and the existing lossy ID schemes [34,15,2,1] are concrete instantiations of our transformation. More importantly, this yield 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 [24,54], 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. [18], 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 [15]. We establish this impossibility result using meta-reduction techniques [17,7,1]. We note that our impossibility result does not apply to more generic but less efficient OR-proof-based tight construction in [6], since they use the OR-proof ideas in a different manner.

Our result is very different to the previous impossibility results [17,39,37,7] about tight signatures, and it enriches our understanding on constructing tight signature schemes. More precisely, Bader et al. [7] 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. [7] summarized results [17,39,37]. 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.

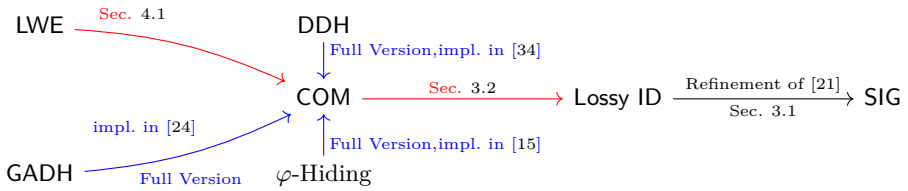

Fig. 1. Overview of our construction. All implications are tight. New implications are marked with red, and new implications that implicitly exist in previous work are marked with blue. The assumption GADH is a generic assumption about group actions, capturing isogeny-based assumptions.

Our Starting Point: The Diemert et al. (DGJL) Approach [21]. 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 [19]). 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 can be tightly satisfied by the random self-reducibility of DDH- and φ-Hiding-based schemes in [15,34], 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 [2] and Module-LWE [40] 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 [46]. However, for plain LWE assumption the number of samples does not influence security [52,50,13,27], 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 lossiness[3].

[3] A 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).

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 [52] 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 [44].

Moreover, we show that many well-known lossy ID schemes [34,15,2,1,24] are obtained from dual-mode commitment schemes. In particular, we give a new analysis of the isogeny-based scheme in [24] 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 [21]. 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 [21], we allow the lossy keys to be correlated, which is necessary for the analysis of the isogney-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.

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 [45] and the known transformation [53].

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 $\mathbf{T}(\mathsf{A})$ 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 $\nu : \mathbb{N} \to \mathbb{R}$ is negligible in its input λ if $\nu \in \lambda^{-\omega(1)}$. The term $\mathsf{negl}(\lambda)$ always denotes a negligible function. If a function ν is at least $1 - \mathsf{negl}(\lambda)$, 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 $y \leftarrow \mathsf{A}(x)$ means that the variable y is assigned to the output of algorithm A on input x. Sometimes we make the randomness used by an algorithm explicit by writing $y = \mathsf{A}(x; r)$ if $r \in \{0,1\}^*$ is $\mathsf{A}'s$ randomness. If we want to state that y is a possible output of A on input x, we write $y \in \mathsf{A}(x)$. 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 $\mathbf{G}^{\mathcal{A}}_{\Pi}(1^\lambda) \Rightarrow b$ 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 $\mathsf{Adv}^{\mathsf{G}}_{\mathcal{A},\Pi}(\lambda)$. 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 $[n] := \{1, \ldots, n\}$ we denote the set of the first n natural numbers.

We present the standard background on lattices in the full version.

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 1 (Dual-Mode Commitment Scheme). *A dual-mode* $(\varepsilon_{\mathsf{b}}, \varepsilon_{\mathsf{t}}, N)$*-commitment scheme is a tuple of PPT algorithms* $\mathsf{CMT} = (\mathsf{Setup}, \mathsf{TSetup}, \mathsf{Gen}, \mathsf{TGen}, \mathsf{Com}, \mathsf{TCom}, \mathsf{Open}, \mathsf{TCol})$ *with the following syntax:*

- $\mathsf{Setup}(1^\lambda)$ *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.*
- $\mathsf{Gen}(\mathsf{par}, 1^N)$ *outputs* N *commitment keys* $\mathsf{ck}_1, \ldots, \mathsf{ck}_N \in \mathcal{K}$*.*
- $\mathsf{Com}(\mathsf{ck}, \mathsf{m})$ *outputs a commitment* $\mathsf{c} \in \mathcal{C}$ *and a decommitment* $\mathsf{dc} \in \mathcal{D}$*.*
- $\mathsf{Open}(\mathsf{ck}, \mathsf{m}, \mathsf{dc}, \mathsf{c})$ *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.*
- $\mathsf{TGen}(\mathsf{par})$ *outputs a commitment key* $\mathsf{ck} \in \mathcal{K}$ *and a trapdoor* $\mathsf{td} \in \mathcal{T}$*.*
- $\mathsf{TCom}(\mathsf{ck}, \mathsf{td})$ *outputs a commitment* $\mathsf{c} \in \mathcal{C}$ *and a state* St*.*
- $\mathsf{TCol}(\mathsf{ck}, \mathsf{td}, St, \mathsf{m})$ *outputs* $\mathsf{dc}' \in \mathcal{D}$*.*

We say that CMT *is* ρ*-complete if for all* $\mathsf{par} \in \mathsf{TSetup}(1^\lambda), (\mathsf{ck}, \mathsf{td}) \in \mathsf{TGen}(\mathsf{par}), \mathsf{m} \in \mathcal{M}$ *we have that* $\Pr\left[\mathsf{Open}(\mathsf{ck}, \mathsf{m}, \mathsf{dc}, \mathsf{c}) = 1 \mid (\mathsf{c}, \mathsf{dc}) \leftarrow \mathsf{Com}(\mathsf{ck}, \mathsf{m})\right] \geq \rho$*.*
Finally, the following security properties should hold:

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

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

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

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

  and

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

- $(\varepsilon_{\mathsf{b}}, N)$***-Statistically Binding:*** *The following probability is at most* ε_{b}*:*

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

  where the probability is taken over

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

Signature Schemes. We define the standard notion of signature schemes and their security.

Definition 2 (Digital Signature Scheme). *A signature scheme is a tuple of PPT algorithms* $\mathsf{SIG} = (\mathsf{Setup}, \mathsf{Gen}, \mathsf{Sig}, \mathsf{Ver})$*, where*

- $\mathsf{Setup}(1^\lambda)$ *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.*
- $\mathsf{Gen}(\mathsf{par})$ *outputs public and secret key* $(\mathsf{pk}, \mathsf{sk}) \in \mathcal{K}_p \times \mathcal{K}_s$*.*
- $\mathsf{Sig}(\mathsf{sk}, \mathsf{m})$ *returns a signature* $\sigma \in \mathcal{S}$*.*
- $\mathsf{Ver}(\mathsf{pk}, \mathsf{m}, \sigma)$ *is deterministic and returns* $b \in \{0, 1\}$*.*

We say that SIG *is* ρ*-complete, if for all* $\mathsf{par} \in \mathsf{Setup}(1^\lambda)$*, all* $(\mathsf{pk}, \mathsf{sk}) \in \mathsf{Gen}(\mathsf{par})$*, all* $\mathsf{m} \in \mathcal{M}$ *we have* $\Pr\left[\mathsf{Ver}(\mathsf{pk}, \mathsf{m}, \sigma) = 1 \mid \sigma \leftarrow \mathsf{Sig}(\mathsf{sk}, \mathsf{m})\right] \geq \rho$*.*

Definition 3 (Multi-User Security). *Consider a signature scheme* $\mathsf{SIG} = (\mathsf{Setup}, \mathsf{Gen}, \mathsf{Sig}, \mathsf{Ver})$*, let* $N \in \mathbb{N}$ *be a natural number and consider the game* N-**MU-CMA-Corr** *given in Fig. 2. We say that* SIG *is* N-MU-CMA-Corr *secure, if for every PPT adversary* $\mathcal{A}$ *the following advantage is negligible in* λ*:*

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

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

Game N-**MU-CMA-Corr**$^{\mathcal{A}}_{\mathsf{SIG}}(\lambda)$

01 $\mathsf{par} \leftarrow \mathsf{Setup}(1^\lambda)$
02 **for** $i \in [N] : (\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow \mathsf{Gen}(\mathsf{par})$
03 $\mathrm{O} := (\textsc{Sig}, \textsc{Key})$
04 $(i^*, \mathsf{m}^*, \sigma^*) \leftarrow \mathcal{A}^{\mathrm{O}}(\mathsf{par}, (\mathsf{pk}_i)_{i=1}^N)$
05 **if** $i^* \in \mathcal{L}_{id}$: **return** 0
06 **if** $\exists \sigma : (i^*, \mathsf{m}^*, \sigma) \in \mathcal{L}_m$: **return** 0
07 **return** $\mathsf{Ver}(\mathsf{pk}_{i^*}, \mathsf{m}^*, \sigma^*)$

Oracle $\textsc{Key}(i)$

08 $\mathcal{L}_{id} := \mathcal{L}_{id} \cup \{i\}$
09 **return** sk_i

Oracle $\textsc{Sig}(i, \mathsf{m})$

10 $\sigma \leftarrow \mathsf{Sig}(\mathsf{sk}_i, \mathsf{m})$
11 $\mathcal{L}_m := \mathcal{L}_m \cup \{(i, \mathsf{m}, \sigma)\}$
12 **return** σ

Fig. 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**.

IDENTIFICATION SCHEMES. Here, we introduce identification schemes and their properties, where we extend the notions of [2,40] to the multi-user setting.

Definition 4 (Canonical Identification Scheme). *A canonical identification scheme* ID *is defined as a tuple of PPT algorithms* $\mathsf{ID} := (\mathsf{ISetup}, \mathsf{IGen}, \mathsf{P} := (\mathsf{P}_1, \mathsf{P}_2), \mathsf{V})$*, with the following properties:*

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

– $\mathsf{V}(\mathsf{pk}, \mathsf{cmt}, \mathsf{ch}, \mathsf{rs})$ *is deterministic and outputs* $b \in \{0, 1\}$.

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

$\mathbf{Alg}\ \mathsf{Tran}(\mathsf{pk}, \mathsf{sk}, \mathsf{ch})$
01 $(\mathsf{cmt}, St) \leftarrow \mathsf{P}_1(\mathsf{sk}), \mathsf{rs} \leftarrow \mathsf{P}_2(\mathsf{sk}, \mathsf{ch}, St)$
02 **if** $\mathsf{rs} = \perp$: $(\mathsf{cmt}, \mathsf{ch}) = (\perp, \perp)$
03 **return** $(\mathsf{cmt}, \mathsf{ch}, \mathsf{rs})$

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

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

From now on, without loss of generality, we assume that V *accepts an honestly generated transcript if and only if* $\mathsf{P}_2(\mathsf{sk}, \mathsf{ch}, St) \neq \perp$*. 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 $\mathsf{ID} = (\mathsf{ISetup}, \mathsf{IGen}, \mathsf{P} = (\mathsf{P}_1, \mathsf{P}_2), \mathsf{V})$ be a canonical identification scheme.

Definition 5 (Special Honest Verifier Zero-Knowledge). *We say that* ID *is* $\varepsilon_{\mathsf{zk}}$*-special honest verifier zero-knowledge (HVZK) if there is a PPT algorithm* Sim *such that for all* $\mathsf{par} \in \mathsf{ISetup}(1^\lambda)$*, all* $(\mathsf{pk}, \mathsf{sk}) \in \mathsf{IGen}(\mathsf{par})$ *the following distributions have statistical distance at most* $\varepsilon_{\mathsf{zk}}$*:*

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

and

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

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

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

Definition 7 (Multi-Key Lossiness). *Let* N *be a natural number. We say that* ID *is* $(\varepsilon_{\mathsf{mkl}}, N)$*-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* $\mathsf{pk}_1, \ldots, \mathsf{pk}_N$ *such that the following holds:*

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

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

– *The following inequality holds:*

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

where we take the expectation, maximum and probability over

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

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 $\varepsilon_{\mathsf{mkl}}$.

Note that N-multi-key lossiness for $N=1$ is just lossiness as defined in [2].

Remark 1 (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 [21] 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 the full version.

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 $\mathsf{ID}:=(\mathsf{ISetup},\mathsf{IGen},\mathsf{P}:=(\mathsf{P}_1,\mathsf{P}_2),\mathsf{V})$ 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 $\mathsf{SIG}_s[\mathsf{ID},\mathsf{H},\ell]$ with random oracle

$\mathsf{H} : \{0,1\}^* \to \mathsf{ChSet}$ 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 [2,25], in a sense that we combine the sequential OR proof from [25] with the lossy identification framework and the repetition as in [2]. Completeness is straight-forward.

Alg $\mathsf{Gen}(\mathsf{par})$
01 $(\mathsf{pk}_0, \mathsf{sk}_0) \leftarrow \mathsf{IGen}(\mathsf{par})$
02 $(\mathsf{pk}_1, \mathsf{sk}_1) \leftarrow \mathsf{IGen}(\mathsf{par})$
03 $b \xleftarrow{\$} \{0,1\}, \mathsf{sk} := (b, \mathsf{sk}_b)$
04 $\mathsf{pk} := (\mathsf{pk}_0, \mathsf{pk}_1)$
05 **return** $(\mathsf{pk}, \mathsf{sk})$

Alg $\mathsf{Ver}(\mathsf{pk}, \mathsf{m}, \sigma)$
06 $\mathsf{ch}_1 \leftarrow \mathsf{H}(0, \mathsf{pk}, \mathsf{cmt}_0, \mathsf{m})$
07 $\mathsf{ch}_0 \leftarrow \mathsf{H}(1, \mathsf{pk}, \mathsf{cmt}_1, \mathsf{m})$
08 $v_0 \leftarrow \mathsf{V}(\mathsf{pk}_0, \mathsf{cmt}_0, \mathsf{ch}_0, \mathsf{rs}_0)$
09 $v_1 \leftarrow \mathsf{V}(\mathsf{pk}_1, \mathsf{cmt}_1, \mathsf{ch}_1, \mathsf{rs}_1)$
10 **return** $(v_0 \wedge v_1)$

Alg $\mathsf{Sig}(\mathsf{sk}, \mathsf{m})$
11 $ctr := 0$
12 **while** $ctr \leq \ell \wedge (\mathsf{rs}_0 = \perp \vee \mathsf{rs}_1 = \perp)$:
13 $\quad ctr := ctr + 1$
14 $\quad (\mathsf{cmt}_b, St_b) \leftarrow \mathsf{P}_1(\mathsf{sk}_b)$
15 $\quad \mathsf{ch}_{1-b} \leftarrow \mathsf{H}(b, \mathsf{pk}, \mathsf{cmt}_b, \mathsf{m})$
16 $\quad (\mathsf{cmt}_{1-b}, \mathsf{rs}_{1-b}) \leftarrow \mathsf{Sim}(\mathsf{pk}_{1-b}, \mathsf{ch}_{1-b})$
17 $\quad \mathsf{ch}_b \leftarrow \mathsf{H}(1-b, \mathsf{pk}, \mathsf{cmt}_{1-b}, \mathsf{m})$
18 $\quad \mathsf{rs}_b \leftarrow \mathsf{P}_2(\mathsf{sk}_b, \mathsf{ch}_b, St_b)$
19 **if** $\mathsf{rs}_0 = \perp \vee \mathsf{rs}_1 = \perp$: **return** $\perp$
20 **return** $\sigma := (\mathsf{cmt}_0, \mathsf{cmt}_1, \mathsf{rs}_0, \mathsf{rs}_1)$

Fig. 3. The signature scheme $\mathsf{SIG}_s[\mathsf{ID}, \mathsf{H}, \ell] = (\mathsf{Setup}, \mathsf{Gen}, \mathsf{Sig}, \mathsf{Ver})$ for a canonical identification scheme $\mathsf{ID} := (\mathsf{ISetup}, \mathsf{IGen}, \mathsf{P} := (\mathsf{P}_1, \mathsf{P}_2), \mathsf{V})$ with HVZK simulator Sim, where $\mathsf{Setup} := \mathsf{ISetup}$.

Theorem 1. *Let* ID *be a canonical identification scheme. If* ID *is* $\varepsilon_{\mathsf{zk}}$*-uHVZK and* $(\varepsilon_{\mathsf{mkl}}, N)$*-multi-key lossy for negligible* $\varepsilon_{\mathsf{zk}}, \varepsilon_{\mathsf{mkl}}$*, then* $\mathsf{SIG}_s[\mathsf{ID}, \mathsf{H}, \ell]$ *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*

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

Due to space limitation and its similarities with [21] we postpone the proof to the full version.

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 the full version.

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 commitment schemes [36]. 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 $\mathsf{CMT} = (\mathsf{Setup}, \mathsf{TSetup}, \mathsf{Gen}, \mathsf{TGen}, \mathsf{Com}, \mathsf{TCom}, \mathsf{Open}, \mathsf{TCol})$ be a dual-mode commitment with message space $\mathcal{M}$. We construct a canonical identification scheme $\mathsf{ID}[\mathsf{CMT}]$ in Fig. 4.

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.

Alg $\mathsf{LIGen}(1^N)$
01 $\mathsf{par} \leftarrow \mathsf{Setup}(1^\lambda)$
02 $(\mathsf{ck}_1, \ldots, \mathsf{ck}_N) \leftarrow \mathsf{Gen}(\mathsf{par}, 1^N)$
03 **for** $i \in [N]$: $\mathsf{pk}_i := \mathsf{ck}_i$
04 **return** $(\mathsf{par}, \mathsf{pk}_1, \ldots, \mathsf{pk}_N)$

Alg $\mathsf{P}_1(\mathsf{sk} = \mathsf{td})$
05 $(\mathsf{c}, St) \leftarrow \overline{\mathsf{TCom}}(\mathsf{pk}, \mathsf{sk})$
06 **return** $(\mathsf{cmt} := \mathsf{c}, St)$

Alg $\mathsf{P}_2(\mathsf{sk}, \mathsf{ch}, St)$
07 $\mathsf{dc} \leftarrow \overline{\mathsf{TCol}}(\mathsf{pk}, \mathsf{sk}, St, \mathsf{ch})$
08 **return** dc

Alg $\mathsf{V}(\mathsf{pk}, \mathsf{cmt}, \mathsf{ch}, \mathsf{rs})$
09 $\mathsf{c} := \mathsf{cmt}, \mathsf{m} := \mathsf{ch}, \mathsf{dc} := \mathsf{rs}$
10 **return** $\mathsf{Open}(\mathsf{pk}, \mathsf{m}, \mathsf{dc}, \mathsf{c})$

Alg $\mathsf{Sim}(\mathsf{pk}, \mathsf{ch})$
11 $(\mathsf{c}, \mathsf{dc}) \leftarrow \overline{\mathsf{Com}}(\mathsf{pk}, \mathsf{ch})$
12 **return** $(\mathsf{cmt} := \mathsf{c}, \mathsf{rs} := \mathsf{dc})$

Fig. 4. The identification scheme $\mathsf{ID}[\mathsf{CMT}] = (\mathsf{ISetup} := \mathsf{TSetup}, \mathsf{IGen} := \mathsf{TGen}, \mathsf{P}, \mathsf{V})$ with challenge set $\mathsf{ChSet} := \mathcal{M}$ and related algorithms $\mathsf{Sim}, \mathsf{LIGen}$ for a given dual-mode commitment $\mathsf{CMT} = (\mathsf{Setup}, \mathsf{TSetup}, \mathsf{Gen}, \mathsf{TGen}, \mathsf{Com}, \mathsf{TCom}, \mathsf{Open}, \mathsf{TCol})$ with message space $\mathcal{M}$.

Lemma 1 (uHVZK and Completeness). *If* CMT *is a* ρ*-complete dual-mode* $(\varepsilon_{\mathsf{bind}}, \varepsilon_{\mathsf{trap}}, N)$*-commitment scheme, then* $\mathsf{ID}[\mathsf{CMT}]$ *is* $\varepsilon_{\mathsf{zk}}$*-uHVZK and* ρ'*-complete, where* $\varepsilon_{\mathsf{zk}} \leq \varepsilon_{\mathsf{trap}}$ *and* $\rho' \geq \rho - \varepsilon_{\mathsf{trap}}$.

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

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

and

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

and the former is exactly the distribution output by Sim on input $\mathsf{ch} = \mathsf{m}$, 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 $\rho - \varepsilon_{\mathsf{trap}}$, which finishes the proof. □

Lemma 2 (Multi-Key Lossiness). *If* CMT *is a dual-mode* $(\varepsilon_{\mathsf{bind}}, \varepsilon_{\mathsf{trap}}, N)$*-commitment scheme, then* $\mathsf{ID}[\mathsf{CMT}]$ *is* $(\varepsilon_{\mathsf{mkl}}, N)$*-multi-key lossy, where*

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

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*

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

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

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

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

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

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

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

Define the random variable W as

$$\mathsf{W} := \max_{i \in [N]} \max_{\mathsf{cmt} \in \mathsf{CmtSet}} \Pr_{\mathsf{ch} \xleftarrow{\$} \mathsf{ChSet}} [\exists \mathsf{rs} \in \mathsf{RspSet} : \mathsf{V}(\mathsf{pk}_i, \mathsf{cmt}, \mathsf{ch}, \mathsf{rs}) = 1].$$

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

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

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

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

□

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 $\mathsf{CMT}_{\mathsf{LWE}}$ based on the LWE assumption is presented in Fig. 5. It is inspired by the classical lattice cryptosystem by Regev [52] and its extension to multiple bits from [51]. 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 [47][4]) to work, we need to ensure that

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

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.

Lemma 3 (Completeness, Trapdoor Property). *The scheme* $\mathsf{CMT}_{\mathsf{LWE}}$ *is* ρ*-complete and satisfies the* ε_{t}*-trapdoor property with* $\rho \geq 1 - \mathsf{negl}(\lambda)$ *and* $\varepsilon_{\mathsf{t}} \leq \mathsf{negl}(\lambda)$.

Proof. Let $(\mathsf{ck} = \mathbf{A}, \mathsf{td} = \mathbf{T}_{\mathbf{A}}) \leftarrow \mathsf{TGen}(\mathsf{par})$. 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 comleteness.

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

$$\mathbf{A}\mathbf{z} = \mathbf{y} = \mathbf{u} - \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rceil \cdot \mathbf{m} \end{bmatrix} \Longrightarrow \mathbf{A}\mathbf{z} + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rceil \cdot \mathbf{m} \end{bmatrix} = \mathbf{u},$$

and with overwhelming probability (see [48,28]) $\|\mathbf{z}\| \leq s \cdot \sqrt{m}$ (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 $\mathsf{TCom}, \mathsf{TCol}$, hence we only have to show that for every $\mathbf{m}$ the distributions

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

[4] For the exact statements we use, we refer to the full version of our paper.

Alg $\mathsf{TGen}(\mathsf{par})$
01 $(\mathbf{A}, \mathbf{T_A}) \leftarrow \mathsf{GenTrap}(1^{n+k}, 1^m, s_0, q)$
02 $\mathsf{ck} := \mathbf{A} \in \mathbb{Z}_q^{(n+k)\times m}, \mathsf{td} := \mathbf{T_A}$
03 **return** $(\mathsf{ck}, \mathsf{td})$

Alg $\mathsf{Gen}(\mathsf{par}, 1^N)$
04 $\bar{\mathbf{S}} \xleftarrow{\$} \mathbb{Z}_q^{n\times k}$
05 **for** $i \in [N]$:
06 $\quad \bar{\mathbf{A}}_i \xleftarrow{\$} \mathbb{Z}_q^{n\times m}, \bar{\mathbf{E}}_i \leftarrow D_{\mathbb{Z},\alpha q}^{m\times k}$
07 $\quad \mathsf{ck}_i := \mathbf{A}_i := \begin{bmatrix} \bar{\mathbf{A}}_i \\ \bar{\mathbf{S}}^t\bar{\mathbf{A}}_i + \bar{\mathbf{E}}_i^t \end{bmatrix}$
08 **return** $(\mathsf{ck}_1, \ldots, \mathsf{ck}_N)$

Alg $\mathsf{Com}(\mathsf{ck}, \mathbf{m})$
09 $\mathbf{z} \leftarrow D_{\mathbb{Z}^m,s}, \mathbf{u} := \mathbf{Az} + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rceil \cdot \mathbf{m} \end{bmatrix}$
10 **if** $\|\mathbf{z}\| > s \cdot \sqrt{m}$: **return** $\perp$
11 **return** $(\mathsf{c} := \mathbf{u}, \mathsf{dc} := \mathbf{z})$

Alg $\mathsf{TCom}(\mathsf{ck}, \mathsf{td})$
12 $\mathbf{u} \xleftarrow{\$} \mathbb{Z}_q^{n+k}$
13 **return** $(\mathbf{u}, St := \mathbf{u})$

Alg $\mathsf{TCol}(\mathsf{ck}, \mathsf{td}, St, \mathbf{m})$
14 $\mathbf{y} := \mathbf{u} - \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rceil \cdot \mathbf{m} \end{bmatrix}$
15 $\mathbf{z} \leftarrow \mathsf{SampleD}(\mathbf{A}, \mathbf{T_A}, \mathbf{y}, s)$
16 **if** $\|\mathbf{z}\| > s \cdot \sqrt{m}$: **return** $\perp$
17 **return** $\mathbf{z}$

Alg $\mathsf{Open}(\mathsf{ck}, \mathbf{m}, \mathbf{z}, \mathbf{u})$
18 **if** $\|\mathbf{z}\| > s \cdot \sqrt{m}$: **return** 0
19 **if** $\mathbf{Az} + \begin{bmatrix} \mathbf{0} \\ \lfloor q/2 \rceil \cdot \mathbf{m} \end{bmatrix} \neq \mathbf{u}$:
20 $\quad$ **return** 0
21 **return** 1

Fig. 5. The dual-mode commitment $\mathsf{CMT}_{\mathsf{LWE}} = (\mathsf{Setup}, \mathsf{TSetup}, \mathsf{Gen}, \mathsf{TGen}, \mathsf{Com}, \mathsf{TCom}, \mathsf{Open}, \mathsf{TCol})$ with message space $\mathcal{M} = \{0,1\}^k$, where $\mathsf{Setup} = \mathsf{TSetup}$ sets parameters par as in the text.

and

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

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

$$\mathcal{D}_1' := \left\{ (\mathbf{y}, \mathbf{z}) \middle| \mathbf{y} \xleftarrow{\$} \mathbb{Z}_q^{n+k}, \mathbf{z} \leftarrow \mathsf{SampleD}(\mathbf{A}, \mathbf{T_A}, \mathbf{y}, s) \right\}$$

and

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

Standard lattice trapdoor techniques (see [48,28]) imply that these are statistically close, which finishes the proof. □

Lemma 4 (Key Indistinguishability). *Let $N = \mathsf{poly}(\lambda)$ be a natural number. Then $\mathsf{CMT}_{\mathsf{LWE}}$ satisfies key indistinguishability, under the $\mathsf{LWE}_{n,q,D_{\mathbb{Z},\alpha q}}$ 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*

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

Due to space limitations, we postpone the proof to the full version.

Lemma 5 (Binding Property). *For any $N = \mathsf{poly}(\lambda)$ the scheme $\mathsf{CMT}_{\mathsf{LWE}}$ is $(\varepsilon_{\mathsf{b}}, N)$-statistically binding, with $\varepsilon_{\mathsf{b}} \leq \mathsf{negl}(\lambda)$.*

Proof. Consider the random experiment

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

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

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

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

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

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

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

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. □

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

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

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 2 (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. 2 becomes useless and lossiness is

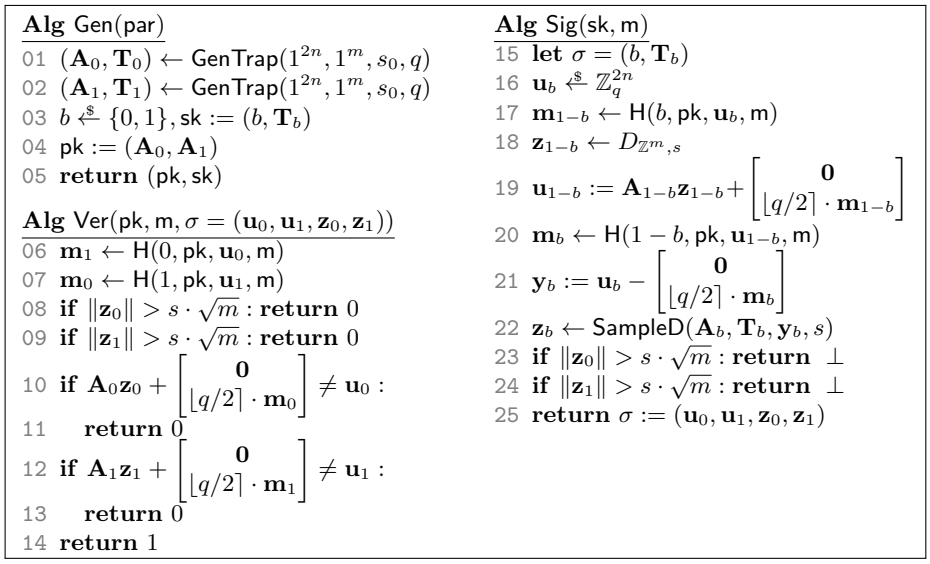

Fig. 6. The signature scheme $\mathsf{SIG}_s[\mathsf{ID}[\mathsf{CMT}_{\mathsf{LWE}}], \mathsf{H}, 1] = (\mathsf{Setup}, \mathsf{Gen}, \mathsf{Sig}, \mathsf{Ver})$, where Setup sets parameters as in Section 4.1.

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. 2. The second variant is to use commitments resulting from [30,33] 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 $\mathbf{C} := \mathbf{AR} + \mathbf{x}^t \otimes \mathbf{G}$. 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 the full version.

We will now instantiate our generic construction in Section 3 with the dual-mode commitment scheme $\mathsf{CMT}_{\mathsf{LWE}}$. 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. 6. Note that signatures contain a linear number of elements from $\mathbb{Z}_q$. The signature schemes based on the completely tight dual-mode commitments mentioned above are formally presented in the full version.

4.2 Instantiation based on Isogenies

We show how to instantiate our approach in the isogeny setting. In [24] 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 can also show that the previously known lossy ID schemes [34,15,2,41] are concrete instantiations of our transformation in Section 3.2. Due to space limitations, we postpone these results to the full version.

5 Impossibility Result for Parallel OR Proofs

In this section, we consider a canonical identification scheme $\mathsf{ID} = (\mathsf{ISetup}, \mathsf{IGen}, \mathsf{P} := (\mathsf{P}_1, \mathsf{P}_2), \mathsf{V})$ with challenge set ChSet and a random oracle $\mathsf{H} : \{0,1\}^* \to \mathsf{ChSet}$. 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 $\mathsf{SIG}_p[\mathsf{ID}, \mathsf{H}]$ defined in Fig. 7 is unlikely. For

Alg $\mathsf{Gen}(\mathsf{par})$
01 $(\mathsf{pk}_0, \mathsf{sk}_0) \leftarrow \mathsf{IGen}(\mathsf{par})$
02 $(\mathsf{pk}_1, \mathsf{sk}_1) \leftarrow \mathsf{IGen}(\mathsf{par})$
03 $b \xleftarrow{\$} \{0,1\}, \mathsf{sk} := (b, \mathsf{sk}_b)$
04 **return** $(\mathsf{pk} := (\mathsf{pk}_0, \mathsf{pk}_1), \mathsf{sk})$

Alg $\mathsf{Ver}(\mathsf{pk}, \mathsf{m}, \sigma)$
05 $\mathsf{ch} \leftarrow \mathsf{H}(\mathsf{pk}, \mathsf{cmt}_0, \mathsf{cmt}_1, \mathsf{m})$
06 **if** $\mathsf{ch}_0 \oplus \mathsf{ch}_1 \neq \mathsf{ch}$: **return** 0
07 $v_0 \leftarrow \mathsf{V}(\mathsf{pk}_0, \mathsf{cmt}_0, \mathsf{ch}_0, \mathsf{rs}_0)$
08 $v_1 \leftarrow \mathsf{V}(\mathsf{pk}_1, \mathsf{cmt}_1, \mathsf{ch}_1, \mathsf{rs}_1)$
09 **return** $(v_0 \wedge v_1)$

Alg $\mathsf{Sig}(\mathsf{sk}, \mathsf{m})$
10 $(\mathsf{cmt}_b, St_b) \leftarrow \mathsf{P}_1(\mathsf{par}, \mathsf{sk}_b)$
11 $\mathsf{ch}_{1-b} \xleftarrow{\$} \mathsf{ChSet}$
12 $(\mathsf{cmt}_{1-b}, \mathsf{rs}_{1-b}) \leftarrow \mathsf{Sim}(\mathsf{pk}_{1-b}, \mathsf{ch}_{1-b})$
13 $\mathsf{ch} \leftarrow \mathsf{H}(\mathsf{pk}, \mathsf{cmt}_0, \mathsf{cmt}_1, \mathsf{m})$
14 $\mathsf{ch}_b := \mathsf{ch} \oplus \mathsf{ch}_{1-b}$
15 $\mathsf{rs}_b \leftarrow \mathsf{P}_2(\mathsf{sk}_b, \mathsf{ch}_b, St_b)$
16 **if** $\mathsf{rs}_0 = \perp \vee \mathsf{rs}_1 = \perp$: **return** $\perp$
17 $\sigma := (\mathsf{cmt}_0, \mathsf{cmt}_1, \mathsf{ch}_0, \mathsf{ch}_1, \mathsf{rs}_0, \mathsf{rs}_1)$
18 **return** σ

Fig. 7. The signature scheme $\mathsf{SIG}_p[\mathsf{ID}, \mathsf{H}] = (\mathsf{Setup}, \mathsf{Gen}, \mathsf{Sig}, \mathsf{Ver})$ for a canonical identification scheme $\mathsf{ID} := (\mathsf{ISetup}, \mathsf{IGen}, \mathsf{P} := (\mathsf{P}_1, \mathsf{P}_2), \mathsf{V})$ with HVZK simulator Sim, where $\mathsf{Setup} := \mathsf{ISetup}$.

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 Notions 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**$^{\mathcal{A}}_{\mathsf{SIG}}(\lambda)$

$\mathsf{par} \leftarrow \mathsf{Setup}(1^\lambda)$
for $i \in [N] : (\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow \mathsf{Gen}(\mathsf{par})$ // Assume pk_i's pairwise distinct
$(j, St_{\mathcal{A}}) \leftarrow \mathcal{A}_1(\mathsf{par}, (\mathsf{pk}_i)_{i\in[N]})$
if $j \notin [N]$: **return** 0
$(\mathsf{m}^*, \sigma^*) \leftarrow \mathcal{A}_2(St_{\mathcal{A}}, (\mathsf{sk}_i)_{i\in[N]\setminus\{j\}})$
return $\mathsf{Ver}(\mathsf{pk}_j, \mathsf{m}^*, \sigma^*)$

Fig. 8. Game **MU-CMA-S** for a signature scheme $\mathsf{SIG} = (\mathsf{Setup}, \mathsf{Gen}, \mathsf{Sig}, \mathsf{Ver})$, used in the proof of the impossibility result in Section 5. We assume that the keys $\mathsf{pk}_1, \ldots, \mathsf{pk}_N$ are pairwise distinct.

Definition 8 (Static Multi-User Security). *Let* $\mathsf{SIG} = (\mathsf{Setup}, \mathsf{Gen}, \mathsf{Sig}, \mathsf{Ver})$ *be a signature scheme and* $N \in \mathbb{N}$ *be a natural number. Consider the game* **MU-CMA-S** *given in Fig. 8. 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*:*

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

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 [7]. 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 [7] 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 $\mathcal{SK}(\mathsf{pk})$. More formally $\mathcal{SK}(\mathsf{pk}) := \{\mathsf{sk} | (\mathsf{pk}, \mathsf{sk}) \in \mathsf{IGen}(\mathsf{par})\}$.

Definition 9 (Verifiability). *Let* $\mathsf{ID} = (\mathsf{ISetup}, \mathsf{IGen}, \mathsf{P}, \mathsf{V})$ *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*:*

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

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

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

Definition 10 (Key-Rerandomization). *Let* $\mathsf{ID} = (\mathsf{ISetup}, \mathsf{IGen}, \mathsf{P}, \mathsf{V})$ *be a canonical identification scheme. We say that* ID *is key-rerandomizable if there is a PPT algorithm* RerandK *such that for all* $\mathsf{par} \in \mathsf{ISetup}(1^\lambda)$ *and all* $(\mathsf{pk}, \mathsf{sk}) \in \mathsf{IGen}(\mathsf{par})$ *the key* $\mathsf{sk}' \leftarrow \mathsf{RerandK}(\mathsf{par}, \mathsf{pk}, \mathsf{sk})$ *is distributed uniformly over* $\mathcal{SK}(\mathsf{pk})$*.*

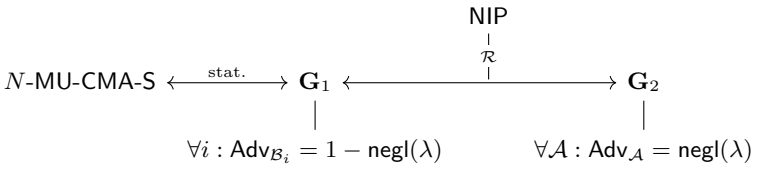

Fig. 9. Overview of a typical cryptographic proof, summarized by two games $\mathbf{G}_1, \mathbf{G}_2$, where $\mathbf{G}_1$ is statistically close to the real game. Here, a reduction $\mathcal{R}$ to the problem NIP is used to interpolate between the games. We will show meta-reductions $\mathcal{B}_i$, that have a high advantage in $\mathbf{G}_1$, whereas every adversary has negligible advantage in $\mathbf{G}_2$.

We note that these properties are quite natural and are satisfied for example by the Chaum-Pedersen (CP) lossy identification scheme [15], which is easy to see.

Example 1. 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 undelying problem, where we follow the notation in [3,7].

Definition 11 (Non-Interactive Problem). *A non-interactive computational problem is a triple of algorithms* $\mathsf{NIP} = (\mathsf{T}, \mathsf{V}, \mathsf{U})$*, where*

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

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

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 ways: Either, one can argue that two subsequent games look statistically close to the adversary, or one uses a reduction

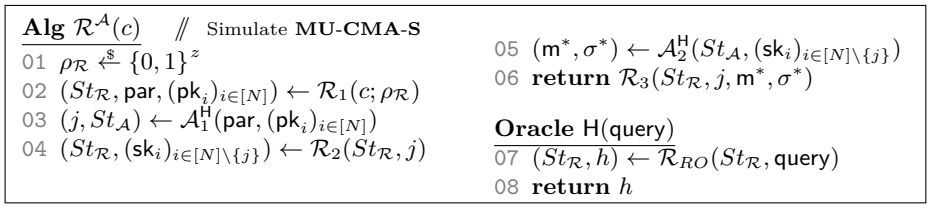

Fig. 10. Syntax of a simple reduction $\mathcal{R} = (\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3, \mathcal{R}_{RO})$ in an execution with an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, used in the proof of the impossibility result in Section 5. Here, $\mathcal{R}$ simulates the game N-**MU-CMA-S**$^{\mathcal{A}}_{\mathsf{SIG}}$ for $\mathcal{A}$.

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. 9. 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 12 (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-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* $(St_{\mathcal{R}}, \mathsf{par}, (\mathsf{pk}_i)_{i\in[N]})$*.*
- $\mathcal{R}_2(St_{\mathcal{R}}, j)$ *takes as input a state* $St_{\mathcal{R}}$ *and an index* $j \in [N]$ *and outputs a new state and secret keys* $(St_{\mathcal{R}}, (\mathsf{sk}_i)_{i\in[N]\setminus\{j\}})$*.*
- $\mathcal{R}_3(St_{\mathcal{R}}, j, \mathsf{m}^*, \sigma^*)$ *takes as input a state, an index* $j \in [N]$*, a message* m^* *and a signature* σ^* *and outputs a* NIP *solution* s*.*
- $\mathcal{R}_{RO}(St_{\mathcal{R}}, \mathsf{query})$ *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. 10. 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*

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

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

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

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

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 $\mathsf{MU\text{-}CMA\text{-}S}$ 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 [7]. 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. 11, obtains parameters and keys $\mathsf{par}, (\mathsf{pk}_i)_{i\in[N]}$ 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 $\mathsf{ID}, \mathsf{par}$ and $\mathsf{pk}_{j^*,0}$. Note that this is why $\mathcal{A}^*$ is inefficient. It then signs a random message $\mathsf{m} \xleftarrow{\$} \mathcal{M}$ and returns it. In terms of success probability the following claim is then clear:

Lemma 6. $\mathsf{Adv}^{N\text{-}\mathsf{MU\text{-}CMA\text{-}S}}_{\mathcal{A}^*,\mathsf{SIG}_p[\mathsf{ID},\mathsf{H}]}(\lambda) = 1$.

We will now present and analyze the other adversaries, which are implicitly given as meta-reductions $\mathcal{B}_1, \ldots \mathcal{B}_5$ 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. 13. 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. 12: After obtaining $St_{\mathcal{R},1}, \mathsf{par}$ and $(\mathsf{pk}_i)_{i\in[N]}$ from $\mathcal{R}_1$ it runs $\mathcal{R}_2$ independently for every $j \in [N]$, stores all secret keys obtained and uses a flag $\mathsf{succ}[j]$ 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 $\mathsf{succ}[j^*] = 1$ and $\mathsf{succ}[j] = 0$ 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 $\mathsf{pk}_{j^*,1}$ and then uses both secret keys $\mathsf{sk}_{j^*,0}, \mathsf{sk}_{j^*,1}$ to compute the signature instead of using the algorithm Sim. The computation of the signature with two keys is summarized in Fig. 12.
- $\mathcal{B}_4$ is as $\mathcal{B}_3$, but if bad does not occur, it will have received a valid secret key $(b, \mathsf{sk}_{j^*,b})$ for pk_{j^*} from some execution of $\mathcal{R}_2$ with index $j \neq j^*$. It will use

Alg $\mathcal{A}_1^{*\mathsf{H}}(\mathsf{par}, (\mathsf{pk}_i)_{i\in[N]})$
01 **if** $\mathsf{VerP}(\mathsf{par}) \neq 1$:
02 **return** $\perp$
03 $j^* \xleftarrow{\$} [N]$
04 $St := (\mathsf{par}, (\mathsf{pk}_i)_{i\in[N]}, j^*)$
05 **return** (j^*, St)

Alg $\mathcal{A}_2^{*\mathsf{H}}(St, (\mathsf{sk}_i = (b_i, \mathsf{sk}_{i,b_i}))_{i\in[N]\setminus\{j^*\}})$
06 **if** $\exists i \in [N] \setminus \{j^*\} : \mathsf{VerK}(\mathsf{par}, \mathsf{pk}_{i,b_i}, \mathsf{sk}_{i,b_i}) = 0$:
07 **return** $\perp$
08 $\mathsf{sk}_0 \xleftarrow{\$} \mathcal{SK}(\mathsf{pk}_{j^*,0})$
09 $\mathsf{m}^* \xleftarrow{\$} \mathcal{M}, \sigma^* \leftarrow \mathsf{Sig}((0, \mathsf{sk}_0), \mathsf{m}^*)$
10 **return** (m^*, σ^*)

Fig. 11. The optimal (but inefficient) adversary $\mathcal{A}^* = (\mathcal{A}_1^*, \mathcal{A}_2^*)$, winning the game **MU-CMA-S** for the signature scheme $\mathsf{SIG}_p[\mathsf{ID}, \mathsf{H}]$.

this secret key (rerandomized) to generate the signature instead of a brute forced one. The other key $\mathsf{sk}_{j^*,1-b}$ is still brute forced and Sim is still not used.
– $\mathcal{B}_5$ now uses only the rerandomized $\mathsf{sk}_{j^*,b}$ 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 $\mathcal{B} := \mathcal{B}_5$.

Alg $\mathsf{Rewind}^{\mathcal{R}}(c)$
01 $\rho_{\mathcal{R}} \xleftarrow{\$} \{0,1\}^z$
02 $(St_{\mathcal{R},1}, \mathsf{par}, (\mathsf{pk}_i)_{i\in[N]}) \leftarrow \mathcal{R}_1(c; \rho_{\mathcal{R}})$
03 **if** $\mathsf{VerP}(\mathsf{par}) \neq 1$: **return** $\perp$
04 $\mathsf{succ} := [0, \ldots, 0]$
05 **for** $j \in [N]$:
06 $(St_{\mathcal{R},2,j}, (\mathsf{sk}_i)_{i\in[N]\setminus\{j\}}) \leftarrow \mathcal{R}_2(St_{\mathcal{R},1}, j)$
07 $\mathsf{succ}[j] := 1$
08 **for** $i \in [N] \setminus \{j\}$:
09 **let** $\mathsf{sk}_i = (b_i, \mathsf{sk}_{i,b_i})$
10 **if** $\mathsf{VerK}(\mathsf{par}, \mathsf{pk}_{i,b_i}, \mathsf{sk}_{i,b_i}) = 0$:
11 $\mathsf{succ}[j] := 0$
12 **if** $\mathsf{succ}[j] = 1$:
13 **for** $i \in [N] \setminus \{j\} : \mathsf{sk}[i] := \mathsf{sk}_i$
14 **return** $(\mathsf{par}, \mathsf{succ}[\cdot], \mathsf{sk}[\cdot], (St_{\mathcal{R},2,j})_{j\in[N]})$

Alg $\mathsf{FakeSign}((\mathsf{sk}_0, \mathsf{sk}_1), \mathsf{m})$
15 $(\mathsf{cmt}_0, St_0) \leftarrow \mathsf{P}_1(\mathsf{sk}_0)$
16 $(\mathsf{cmt}_1, St_1) \leftarrow \mathsf{P}_1(\mathsf{sk}_1)$
17 $\mathsf{ch}_0 \xleftarrow{\$} \mathsf{ChSet}$
18 $\mathsf{ch} \leftarrow \mathsf{H}(\mathsf{pk}, \mathsf{cmt}_0, \mathsf{cmt}_1, \mathsf{m})$
19 $\mathsf{ch}_1 := \mathsf{ch}_0 \oplus \mathsf{ch}$
20 $\mathsf{rs}_0 \leftarrow \mathsf{P}_2(\mathsf{sk}_0, \mathsf{ch}_0, St_0)$
21 $\mathsf{rs}_1 \leftarrow \mathsf{P}_2(\mathsf{sk}_1, \mathsf{ch}_1, St_1)$
22 **if** $\mathsf{rs}_0 = \perp \vee \mathsf{rs}_1 = \perp$:
23 **return** $\perp$
24 $\sigma := (\mathsf{cmt}_0, \mathsf{cmt}_1, \mathsf{ch}_0, \mathsf{ch}_1, \mathsf{rs}_0, \mathsf{rs}_1)$
25 **return** σ

Fig. 12. Subroutines Rewind and $\mathsf{FakeSign}$, used in algorithms $\mathcal{B}_i$ given in Fig. 13.

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

Lemma 7. $\mathsf{Adv}^{\mathsf{NIP}}_{\mathcal{R}^{\mathcal{A}^*}}(\lambda) = \mathsf{Adv}^{\mathsf{NIP}}_{\mathcal{B}_1}(\lambda)$.

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. 12, 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. 13 and Line 06 in Fig. 11 are equivalent conditions. □

Lemma 8. $\left|\mathsf{Adv}^{\mathsf{NIP}}_{\mathcal{B}_1}(\lambda) - \mathsf{Adv}^{\mathsf{NIP}}_{\mathcal{B}_2}(\lambda)\right| \leq 1/N.$

Proof. Note that $\mathcal{B}_1$ and $\mathcal{B}_2$ only differ if the event bad occurs, which implies that $\mathsf{succ}[j^*] = 1$ and $\mathsf{succ}[j] = 0$ 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

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

□

Lemma 9. $\left|\mathsf{Adv}^{\mathsf{NIP}}_{\mathcal{B}_2}(\lambda) - \mathsf{Adv}^{\mathsf{NIP}}_{\mathcal{B}_3}(\lambda)\right| \leq \delta_{\mathcal{R}} + \varepsilon_{\mathsf{zk}}.$

Proof. Both $\mathcal{B}_2$ and $\mathcal{B}_3$ return $\perp$ if $\mathsf{VerP}(\mathsf{par}) \neq 1$ (see Fig. 12) and also if $\mathsf{bad} = 1$. Therefore we can assume that $\mathsf{par} \in \mathsf{ISetup}(1^\lambda)$. As we assume that the reduction $\mathcal{R}$ simulated a game of statistical distance $\delta_{\mathcal{R}}$ to the real game, the public key $\mathsf{pk}_{j^*,1}$ 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 $(\mathsf{pk}_{j^*,1}, \mathsf{sk}_1) \in \mathsf{IGen}(\mathsf{par})$. Note that in this case what $\mathcal{R}$ sees is $(\mathsf{m}^*, \sigma^*, \mathsf{query})$ where $\mathsf{m}^* \xleftarrow{\$} \mathcal{M}, \sigma^* \leftarrow \mathsf{FakeSign}((\mathsf{sk}_0, \mathsf{sk}_1), \mathsf{m}^*)$ and $\mathsf{query} = ((\mathsf{pk}_{j^*,0}, \mathsf{pk}_{j^*,1}), \mathsf{cmt}_0, \mathsf{cmt}_1, \mathsf{m}^*)$ is the random oracle query (observable by $\mathcal{R}_{RO}$) that occurs during the run of $\mathsf{FakeSig}$. Similarly, in the execution of $\mathcal{B}_2$ it sees $(\mathsf{m}^*, \sigma^*, \mathsf{query})$ where $\mathsf{m}^* \xleftarrow{\$} \mathcal{M}, \sigma^* \leftarrow \mathsf{Sig}((0, \mathsf{sk}_0), \mathsf{m}^*)$ and $\mathsf{query} = ((\mathsf{pk}_{j^*,0}, \mathsf{pk}_{j^*,1}), \mathsf{cmt}_0, \mathsf{cmt}_1, \mathsf{m}^*)$. Note that the difference is only the way how the transcript $(\mathsf{cmt}_1, \mathsf{ch}_1, \mathsf{rs}_1)$, which is part of σ^*, is generated. Further, query can be efficiently computed without knowing how that transcript was generated. Hence by $\varepsilon_{\mathsf{zk}}$-HVZK, these have statistical distance at most $\varepsilon_{\mathsf{zk}}$, 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 $\varepsilon_{\mathsf{zk}}$. 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 10. $\mathsf{Adv}^{\mathsf{NIP}}_{\mathcal{B}_3}(\lambda) = \mathsf{Adv}^{\mathsf{NIP}}_{\mathcal{B}_4}(\lambda).$

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

Lemma 11. $\left|\mathsf{Adv}^{\mathsf{NIP}}_{\mathcal{B}_4}(\lambda) - \mathsf{Adv}^{\mathsf{NIP}}_{\mathcal{B}_5}(\lambda)\right| \leq \delta_{\mathcal{R}} + \varepsilon_{\mathsf{zk}}.$

Proof. The proof is exactly the same as for Lem. 9, applying $\varepsilon_{\mathsf{zk}}$-HVZK to $(\mathsf{pk}_{j^*,1-b}, \mathsf{sk}_{1-b})$. □

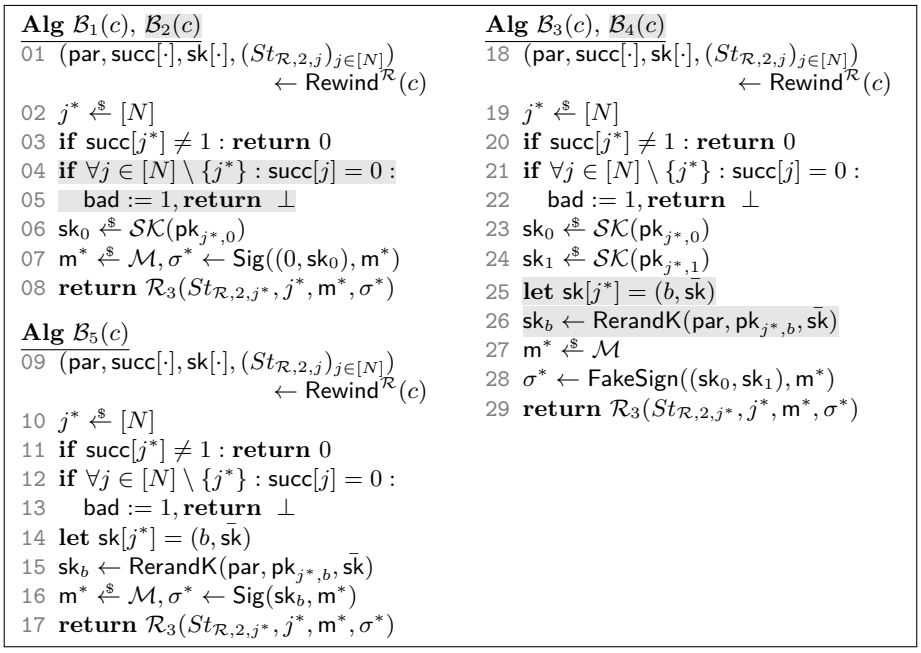

Fig. 13. The (inefficient) algorithms $\mathcal{B}_1, \ldots, \mathcal{B}_4$ and the efficient algorithm $\mathcal{B}_5$ used in the proof of Thm. 2. The subroutines $\mathsf{Rewind}, \mathsf{FakeSign}$ are given in Fig. 12.

In summary, combining all claims we obtain that

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

and $\mathcal{B}_5$ is efficient, which proves Thm. 2. □

References

1. Abdalla, M., Benhamouda, F., Pointcheval, D.: On the tightness of forward-secure signature reductions. Journal of Cryptology 32(1), 84–150 (Jan 2019)
2. Abdalla, M., Fouque, P.A., Lyubashevsky, V., Tibouchi, M.: Tightly-secure signatures from lossy identification schemes. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 572–590. Springer, Heidelberg (Apr 2012)
3. Abe, M., Groth, J., Ohkubo, M.: Separating short structure-preserving signatures from non-interactive assumptions. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 628–646. Springer, Heidelberg (Dec 2011)
4. Abe, M., Ohkubo, M., Suzuki, K.: 1-out-of-n signatures from a variety of keys. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 415–432. Springer, Heidelberg (Dec 2002)
5. Bader, C.: Efficient signatures with tight real world security in the random-oracle model. In: Gritzalis, D., Kiayias, A., Askoxylakis, I.G. (eds.) CANS 14. LNCS, vol. 8813, pp. 370–383. Springer, Heidelberg (Oct 2014)

6. Bader, C., Hofheinz, D., Jager, T., Kiltz, E., Li, Y.: Tightly-secure authenticated key exchange. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part I. LNCS, vol. 9014, pp. 629–658. Springer, Heidelberg (Mar 2015)
7. Bader, C., Jager, T., Li, Y., Schäge, S.: On the impossibility of tight cryptographic reductions. In: Fischlin, M., Coron, J.S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 273–304. Springer, Heidelberg (May 2016)
8. Bellare, M., Boldyreva, A., Micali, S.: Public-key encryption in a multi-user setting: Security proofs and improvements. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 259–274. Springer, Heidelberg (May 2000)
9. Bellare, M., Ristenpart, T.: Simulation without the artificial abort: Simplified proof and improved concrete security for Waters' IBE scheme. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 407–424. Springer, Heidelberg (Apr 2009)
10. Bellare, M., Rogaway, P.: Entity authentication and key distribution. In: Stinson, D.R. (ed.) CRYPTO'93. LNCS, vol. 773, pp. 232–249. Springer, Heidelberg (Aug 1994)
11. Blazy, O., Kakvi, S.A., Kiltz, E., Pan, J.: Tightly-secure signatures from chameleon hash functions. In: Katz, J. (ed.) PKC 2015. LNCS, vol. 9020, pp. 256–279. Springer, Heidelberg (Mar / Apr 2015)
12. Boyen, X., Li, Q.: Towards tightly secure lattice short signature and id-based encryption. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016, Part II. LNCS, vol. 10032, pp. 404–434. Springer, Heidelberg (Dec 2016)
13. Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: Boneh, D., Roughgarden, T., Feigenbaum, J. (eds.) 45th ACM STOC. pp. 575–584. ACM Press (Jun 2013)
14. Canetti, R., Krawczyk, H.: Analysis of key-exchange protocols and their use for building secure channels. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 453–474. Springer, Heidelberg (May 2001)
15. Chaum, D., Pedersen, T.P.: Wallet databases with observers. In: Brickell, E.F. (ed.) CRYPTO'92. LNCS, vol. 740, pp. 89–105. Springer, Heidelberg (Aug 1993)
16. Chen, J., Wee, H.: Fully, (almost) tightly secure IBE and dual system groups. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 435–460. Springer, Heidelberg (Aug 2013)
17. Coron, J.S.: Optimal security proofs for PSS and other signature schemes. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 272–287. Springer, Heidelberg (Apr / May 2002)
18. Cramer, R., Damgård, I., Schoenmakers, B.: Proofs of partial knowledge and simplified design of witness hiding protocols. In: Desmedt, Y. (ed.) CRYPTO'94. LNCS, vol. 839, pp. 174–187. Springer, Heidelberg (Aug 1994)
19. Damgård, I.: On Σ-protocols. https://cs.au.dk/ ivan/Sigma.pdf (2010)
20. Davis, H., Günther, F.: Tighter proofs for the SIGMA and TLS 1.3 key exchange protocols. ACNS 2021 (2021), `https://eprint.iacr.org/2020/1029`
21. Diemert, D., Gellert, K., Jager, T., Lyu, L.: More efficient digital signatures with tight multi-user security. In: Garay, J. (ed.) PKC 2021, Part II. LNCS, vol. 12711, pp. 1–31. Springer, Heidelberg (May 2021)
22. Diemert, D., Jager, T.: On the tight security of TLS 1.3: Theoretically-sound cryptographic parameters for real-world deployments. Journal of Cryptology (2020), `https://eprint.iacr.org/2020/726`
23. Ducas, L., Durmus, A., Lepoint, T., Lyubashevsky, V.: Lattice signatures and bimodal Gaussians. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 40–56. Springer, Heidelberg (Aug 2013)

24. El Kaafarani, A., Katsumata, S., Pintore, F.: Lossy CSI-FiSh: Efficient signature scheme with tight reduction to decisional CSIDH-512. In: Kiayias, A., Kohlweiss, M., Wallden, P., Zikas, V. (eds.) PKC 2020, Part II. LNCS, vol. 12111, pp. 157–186. Springer, Heidelberg (May 2020)
25. Fischlin, M., Harasser, P., Janson, C.: Signatures from sequential-OR proofs. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part III. LNCS, vol. 12107, pp. 212–244. Springer, Heidelberg (May 2020)
26. Gay, R., Hofheinz, D., Kohl, L., Pan, J.: More efficient (almost) tightly secure structure-preserving signatures. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part II. LNCS, vol. 10821, pp. 230–258. Springer, Heidelberg (Apr / May 2018)
27. Genise, N., Micciancio, D., Peikert, C., Walter, M.: Improved discrete gaussian and subgaussian analysis for lattice cryptography. In: Kiayias, A., Kohlweiss, M., Wallden, P., Zikas, V. (eds.) PKC 2020, Part I. LNCS, vol. 12110, pp. 623–651. Springer, Heidelberg (May 2020)
28. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. Cryptology ePrint Archive, Report 2007/432 (2007), https://eprint.iacr.org/2007/432
29. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Ladner, R.E., Dwork, C. (eds.) 40th ACM STOC. pp. 197–206. ACM Press (May 2008)
30. Gentry, C., Sahai, A., Waters, B.: Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 75–92. Springer, Heidelberg (Aug 2013)
31. Gjøsteen, K., Jager, T.: Practical and tightly-secure digital signatures and authenticated key exchange. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part II. LNCS, vol. 10992, pp. 95–125. Springer, Heidelberg (Aug 2018)
32. Goldwasser, S., Micali, S., Rivest, R.L.: A digital signature scheme secure against adaptive chosen-message attacks. SIAM Journal on Computing 17(2), 281–308 (Apr 1988)
33. Gorbunov, S., Vaikuntanathan, V., Wichs, D.: Leveled fully homomorphic signatures from standard lattices. In: Servedio, R.A., Rubinfeld, R. (eds.) 47th ACM STOC. pp. 469–477. ACM Press (Jun 2015)
34. Guillou, L.C., Quisquater, J.J.: A "paradoxical" indentity-based signature scheme resulting from zero-knowledge. In: Goldwasser, S. (ed.) CRYPTO'88. LNCS, vol. 403, pp. 216–231. Springer, Heidelberg (Aug 1990)
35. Han, S., Jager, T., Kiltz, E., Liu, S., Pan, J., Riepel, D., Schäge, S.: Authenticated key exchange and signatures with tight security in the standard model. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021, Part IV. LNCS, vol. 12828, pp. 670–700. Springer, Heidelberg, Virtual Event (Aug 2021)
36. Hazay, C., Lindell, Y.: Efficient Secure Two-Party Protocols - Techniques and Constructions. ISC, Springer, Heidelberg (2010)
37. Hofheinz, D., Jager, T., Knapp, E.: Waters signatures with optimal security reduction. In: Fischlin, M., Buchmann, J., Manulis, M. (eds.) PKC 2012. LNCS, vol. 7293, pp. 66–83. Springer, Heidelberg (May 2012)
38. Jager, T., Kiltz, E., Riepel, D., Schäge, S.: Tightly-secure authenticated key exchange, revisited. In: Canteaut, A., Standaert, F.X. (eds.) EUROCRYPT 2021, Part I. LNCS, vol. 12696, pp. 117–146. Springer, Heidelberg (Oct 2021)

39. Kakvi, S.A., Kiltz, E.: Optimal security proofs for full domain hash, revisited. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 537–553. Springer, Heidelberg (Apr 2012)
40. Kiltz, E., Lyubashevsky, V., Schaffner, C.: A concrete treatment of Fiat-Shamir signatures in the quantum random-oracle model. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part III. LNCS, vol. 10822, pp. 552–586. Springer, Heidelberg (Apr / May 2018)
41. Kiltz, E., Masny, D., Pan, J.: Optimal security proofs for signatures from identification schemes. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part II. LNCS, vol. 9815, pp. 33–61. Springer, Heidelberg (Aug 2016)
42. Langrehr, R., Pan, J.: Tightly secure hierarchical identity-based encryption. In: Lin, D., Sako, K. (eds.) PKC 2019, Part I. LNCS, vol. 11442, pp. 436–465. Springer, Heidelberg (Apr 2019)
43. Liu, X., Liu, S., Gu, D., Weng, J.: Two-pass authenticated key exchange with explicit authentication and tight security. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020, Part II. LNCS, vol. 12492, pp. 785–814. Springer, Heidelberg (Dec 2020)
44. Lyubashevsky, V.: Lattice signatures without trapdoors. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 738–755. Springer, Heidelberg (Apr 2012)
45. Lyubashevsky, V., Micciancio, D.: Asymptotically efficient lattice-based digital signatures. Journal of Cryptology 31(3), 774–797 (Jul 2018)
46. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 1–23. Springer, Heidelberg (May / Jun 2010)
47. Micciancio, D., Peikert, C.: Trapdoors for lattices: Simpler, tighter, faster, smaller. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 700–718. Springer, Heidelberg (Apr 2012)
48. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on Gaussian measures. In: 45th FOCS. pp. 372–381. IEEE Computer Society Press (Oct 2004)
49. Pan, J., Ringerud, M.: Signatures with tight multi-user security from search assumptions. In: Chen, L., Li, N., Liang, K., Schneider, S.A. (eds.) ESORICS 2020, Part II. LNCS, vol. 12309, pp. 485–504. Springer, Heidelberg (Sep 2020)
50. Peikert, C.: Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In: Mitzenmacher, M. (ed.) 41st ACM STOC. pp. 333–342. ACM Press (May / Jun 2009)
51. Peikert, C., Vaikuntanathan, V., Waters, B.: A framework for efficient and composable oblivious transfer. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 554–571. Springer, Heidelberg (Aug 2008)
52. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: Gabow, H.N., Fagin, R. (eds.) 37th ACM STOC. pp. 84–93. ACM Press (May 2005)
53. Steinfeld, R., Pieprzyk, J., Wang, H.: How to strengthen any weakly unforgeable signature into a strongly unforgeable signature. In: Abe, M. (ed.) CT-RSA 2007. LNCS, vol. 4377, pp. 357–371. Springer, Heidelberg (Feb 2007)
54. Stolbunov, A.: Cryptographic schemes based on isogenies. Ph.D. thesis, Norwegian University of Science and Technology (2012)
55. Unruh, D.: Non-interactive zero-knowledge proofs in the quantum random oracle model. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part II. LNCS, vol. 9057, pp. 755–784. Springer, Heidelberg (Apr 2015)