



# Bifurcated Signatures: Folding the Accountability vs. Anonymity Dilemma into a Single Private Signing Scheme

Benoît Libert, Khoa Nguyen, Thomas Peters, Moti Yung

## ► To cite this version:

Benoît Libert, Khoa Nguyen, Thomas Peters, Moti Yung. Bifurcated Signatures: Folding the Accountability vs. Anonymity Dilemma into a Single Private Signing Scheme. Eurocrypt 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Oct 2021, Zagreb (partiellement virtuel), Croatia. pp.1-48, 10.1007/978-3-030-77883-5\_18 . hal-03380687

HAL Id: hal-03380687

<https://hal.inria.fr/hal-03380687>

Submitted on 15 Oct 2021

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Bifurcated Signatures: Folding the Accountability vs. Anonymity Dilemma into a Single Private Signing Scheme

Benoît Libert<sup>1,2</sup>, Khoa Nguyen<sup>3</sup>, Thomas Peters<sup>4</sup>, and Moti Yung<sup>5</sup>

<sup>1</sup> CNRS, Laboratoire LIP, France

<sup>2</sup> ENS de Lyon, Laboratoire LIP (U. Lyon, CNRS, ENSL, Inria, UCBL), France

<sup>3</sup> Nanyang Technological University, SPMS, Singapore

<sup>4</sup> FNRS and UCLouvain (ICTEAM), Belgium

<sup>5</sup> Google and Columbia University, USA

**Abstract.** Over the development of modern cryptography, often, alternative cryptographic schemes are developed to achieve goals that in some important respect are orthogonal. Thus, we have to choose either a scheme which achieves the first goal and not the second, or vice versa. This results in two types of schemes that compete with each other. In the basic area of user privacy, specifically in anonymous (multi-use credentials) signing, such an orthogonality exists between anonymity and accountability. The conceptual contribution of this work is to reverse the above orthogonality by design, which essentially typifies the last 25 years or so, and to suggest an alternative methodology where the opposed properties are carefully folded into a single scheme. The schemes will support both opposing properties simultaneously in a bifurcated fashion, where:

- First, based on rich semantics expressed over the message’s context and content, the user, etc., the relevant property is applied point-wise per message operation depending on a predicate; and
- Secondly, at the same time, the schemes provide what we call “branch-hiding;” namely, the resulting calculated value hides from outsiders which property has actually been locally applied.

Specifically, we precisely define and give the first construction and security proof of a “Bifurcated Anonymous Signature” (BiAS): A scheme which supports either absolute anonymity or anonymity with accountability, based on a specific contextual predicate, while being branch-hiding. This novel signing scheme has numerous applications not easily implementable or not considered before, especially because: (i) the conditional traceability does *not* rely on a trusted authority as it is (non-interactively) encapsulated into signatures; and (ii) signers *know* the predicate value and can make a conscious choice at each signing time.

Technically, we realize BiAS from homomorphic commitments for a general family of predicates that can be represented by bounded-depth circuits. Our construction is generic and can be instantiated in the standard model from lattices and, more efficiently, from bilinear maps. In particular, the signature length is independent of the circuit size when we use commitments with suitable efficiency properties.

**Keywords.** New primitive, privacy, anonymity, accountability, group signatures, conditional traceability, predicate-based privacy.

## 1 Introduction

Properties provided by cryptographic primitives (such as confidentiality and anonymity) generate a natural tension between the requirements of individual users (such as privacy and other rights), and those of society (such as safety and individual accountability). This fact has created a very rigid positioning of cryptosystems: Designs that serve individual needs, and those which serve societal concerns. A classical example of the above rigidity is the scenario of anonymous signing. On the one hand, there are group signatures [23], central privacy tools allowing users to anonymously sign messages in the name of a population of users they belong to. In order to keep users accountable for their actions, group signatures involve a trusted opening authority (OA) which is called upon when needed only, and is endowed with some privileged information allowing it to trace any signature back to its author. This accountability mechanism, therefore, in these cases, revokes the anonymity of that user. On the other hand, ring signatures [54] and related primitives [37, 47, 24, 15] allow users to sign whatever they like in the name of a population while retaining unconditional anonymity.

In light of this over quarter-of-a-century old situation, we claim that for many applications, in fact, group and ring signatures fall short of providing an appropriate tradeoff between anonymity and accountability that would be sufficiently fair for, both, signers and authorities. Privacy-aware signers naturally want to protect their privacy as much as possible. At the same time, authorities aim to ensure that signers of all problematic signatures can be caught. What we argue in this work is that in many real-life situations, the ability to trace a signature or not should actually depend on the content and context of the message, and should be provided programmatically by the primitive rather than being supported by a one sided mechanism which is part of the primitive specification.

Consider the scenario where each signature authenticates an anonymous financial transaction associated with a hidden amount of money and between users in different countries that should only be known to the payer and the payee (this can be done by employing additional cryptographic mechanisms, e.g., the amount and countries are encrypted or committed to, as in the privacy-preserving cryptocurrency system Monero [49]). For money-laundering detection purposes, the authorities would like to make sure that transactions with amounts above a certain threshold between two specific countries can be traced. On the other hand, to satisfy privacy-aware users, the system should also provide absolute anonymity for transactions amounts below the traceability threshold or within the same country, say, or for any other messages of harmless content. Importantly, for privacy reasons as well (e.g., keeping statistics of the financial transactions hidden), the public should not be able to determine whether a given signed transaction, corresponds to a traceable or to an untraceable type.

As another example, imagine that visitors of a digital library are required to register and sign before reviewing specific e-books. The ability/ inability to identify signers should naturally depend on whether the books in question are totally benign (e.g., comics, essays containing controversial but inoffensive political opinions, etc.) or potentially harmful (like chemistry books explaining

how to produce bombs, any form of advocacy of hatred, etc.).

As yet another example, note that service providers often ask users to attest to personal attributes, for example, to guarantee the veracity of answers to questions like “Have you been to one of these countries in the last 6 months?”, “Are you above 18?”, etc. In this case we argue that while suspicious online activities and alert-raising messages should be traceable by some warranted authority, regular well-behaved typical users should not have to reveal their history and information to service providers that verify their signatures.

The above examples motivate the design of a new anonymous signature primitive where the ability to trace a signature back to its source is determined by a predicate that depends on the signed message and the user’s credential, and where the traceability property of a signature is hidden from the general public. Such schemes are highly desirable, so that they can support the above scenarios, since they provide a fair privacy-preserving non-rigid setting which users and authorities both have strong incentives to deploy. However, to our knowledge, a non-rigid conditional setting, and such signatures have never been considered so far. We call this type of schemes that allow this on the fly flexibility “Bifurcated Cryptosystems.”

### 1.1 Our Contributions

We introduce the study of cryptographic primitives that provide tradeoffs between competing requirements like end-to-end privacy and accountability, by suggesting *bifurcated anonymous signatures* (BiAS). In short, BiAS schemes are anonymous signature schemes allowing to bifurcate into absolute anonymity and identity escrow at the signing time, where computing signatures is associated with a predicate  $P$ . They enable unconditional anonymity when the predicate  $P(M, \text{id}, w)$  evaluates to 1 on input of the message  $M$ , the user’s identity  $\text{id}$  and some secret context-dependent piece of information  $w$  which we call a witness. At the same time, signatures should be traceable by an authority whenever  $P(M, \text{id}, w) = 0$ .

As a first major requirement, BiAS must be *branch-hiding*: verifiers as well as the issuing credential authority should be unable to figure out whether a signature is traceable or not.

In our BiAS primitive, whether a signature is traceable or not depends on the content which is being signed and the user’s identity. Since users know the predicate and its value before signing, they know beforehand when they will be subject to tracing and they can make an educated decision as to whether they can afford to sign a specific message or not. At the same time, *only* the tracing authority will be able to learn whether the signatures are traceable or not. The users are assured that if their signature is not traceable, no one, even if the authority’s keys are available, will be able to trace them. In fact, let us emphasize that this is an unconditional anonymity property, which is of high importance in some applications, such as the case of journalists signing an article unfavorable to the local regime, in a place where their life is in danger upon eventual identification.

As a natural second requirement, we pair the notion of branch-hiding with

the notion of *branch-soundness*. Branch-soundness prevents users from generating untraceable signatures when the signer should have been identified or vice versa. Said otherwise, no signers can fool the system even with the help of the authorities and be able to produce a signature of which the traceability does not respect the predicate.

We first give precise syntax and security definitions for the BiAS primitive. The guarantee offered by this notion allows us to extend the notions of traceability, non-frameability, and anonymity, borrowed from ordinary group signatures to our more general predicate-based primitive.

Secondly, we provide a generic BiAS realization where predicates may consist of polynomial-size circuits of a priori bounded depth. As building blocks, our construction relies on the homomorphic equivocal commitment (HEC) primitive defined by Katsumata *et al.* [34], dual-mode non-interactive zero-knowledge (NIZK) arguments [31, 50, 33], and a variant of the  $\mathcal{R}$ -lossy encryption primitive of Boyle *et al.* [18]. Our constructions are instantiable in the standard model for arbitrary polynomial-size Boolean circuits under the Learning-With-Errors (LWE) assumption [53]. For Boolean formulas (equivalently,  $\text{NC}^1$  circuits), more efficient instantiations are possible under falsifiable assumptions in groups endowed with a bilinear map. In both cases, our schemes enjoy the property that the signature size only depends on the maximal circuit depth, and not on its size. The signature size is dominated by  $O((\log N + |w|) \cdot \lambda^c)$  bits (where  $N$  is the group size and  $c$  is a constant) committing to the witness  $w$  and the user’s identity together with NIZK arguments showing that the ciphertexts were properly generated.

## 1.2 Technical Overview

**DEFINING SECURITY.** Our security model puts forth the notions of branch-hiding and branch-soundness for our bifurcated anonymous signature (BiAS) primitive. We advocate, more generally, that it is the first instance of a new fundamental notion of bifurcated cryptosystems (balancing based on a predicate in one scheme, both, user concerns and public safety issues). Further, to capture anonymity we extend the CCA-like notion of unlinkability of signatures which can be now produced from different predicate values. We call the resulting notion anonymity “in the traceable case” which implies the branch-hiding property of BiAS. We augment this anonymity notion with the anonymity “in the non-traceable case” where all the signatures are generated from a predicate value equals to 1, i.e., from the branch leading to unconditional anonymity. A BiAS is, then, called fully anonymous if it fulfills both anonymity notions. Branch-hiding and full anonymity, primarily, take care of privacy of BiAS.

To prevent misuse of the BiAS functionality, we build on two security notions from the Kiayias-Yung model [36] of group signatures. First, the security against *mis-identification attacks* (a.k.a. traceability) which requires that, even if the adversary can introduce users under its control in the group of signers, it cannot produce a signature that traces outside the dishonest coalition. Second, the notion of security against *framing attacks* which implies that honest users can never be falsely accused of having signed messages, even if the whole system conspires

against them. However, extending these security notions is not straightforward or immediate in our model, since we have to detect whether a given signature contradicts one of these properties even, if that signature is untraceable. Indeed, to build a reduction in a security proof, for instance, we have to figure out if a given untraceable signature has been generated honestly by a legitimate signer or if it is a forgery. However, being able to do so, in fact, seems to contradict statistical (unconditional) anonymity.

To circumvent the apparent incompatibility between privacy and security, we rely on our branch-soundness notion. It is a two-stage definition which first defines an extractable mode of the scheme only useful for the sake of proving security: it generates parameters of the scheme allowing to extract the identity and the witness behind *any* valid signatures. Such an extraction allows evaluating the predicate *a posteriori* given any signature. As a second stage, we require that the (real) tracing algorithm can be indistinguishably emulated from the (ideal) extractable mode, even when the authorities' keys are exposed. A BiAS satisfying branch-soundness thus ensures the hardness of "cheating" with the predicate, even for corrupt authorities. The reason is that it implies the infeasibility of producing signatures that: (i) can be traced while the context allows retaining statistical anonymity, i.e.,  $P(M, \text{id}, w) = 1$ , and conversely (ii) cannot be traced while the context allows retaining identity escrow, i.e.,  $P(M, \text{id}, w) = 0$ . We stress that the indistinguishability in branch-soundness cannot be statistical since, otherwise, untraceable signatures would no longer be statistically anonymous. Based on the branch-soundness notion, we can now extend the security notion of [36] from the (ideal) extractable mode of the BiAS.

**UNDERLYING PRIMITIVES.** Our construction for bounded-depth circuits is based on combining a number of primitives. First, it is built on the homomorphic equivocable commitments (HEC) of Katsumata *et al.* [34]. An HEC is a commitment scheme that allows committing to a message  $\mathbf{x}$  using random coins  $R$  in such a way that anyone can publicly evaluate a circuit  $C$  over the commitment  $\text{com}$  to obtain an evaluated commitment  $\text{com}_{ev} = \text{Eval}(C, \text{com}_{ev})$  to  $C(\mathbf{x})$ . Using the pair  $(\mathbf{x}, R)$ , the committer can internally run a private evaluation algorithm over  $(\mathbf{x}, R)$  in order to compute a proof  $\pi$  which will convince a verifier that  $\text{com}_{ev}$  is a commitment to  $C(\mathbf{x})$ . The primitive is instantiatable for all circuits via the fully homomorphic commitments of Gorbunov, Vaikuntanathan and Wichs [29], which in turn, is built on the FHE scheme of Gentry, Sahai and Waters (GSW) [28]. Katsumata *et al.* [34] also gave a construction for log-depth circuits under pairing-related assumptions [34]. In order to combine statistical anonymity in non-tracing mode and extractability in the security proofs, we employ a dual-mode NIZK argument, which either provides statistical zero-knowledge and computational soundness or vice versa, depending on the distribution of the common reference string. In the public verifiability setting, dual-mode NIZK is known to exist under standard assumptions in pairing-friendly groups, as shown by Groth, Ostrovsky and Sahai [31]. Peikert and Shiehian [50] (inspired by the earlier work of Canetti *et al.* [22]) recently gave constructions under the Learning-With-Errors (LWE) assumption. In order to smoothly interact with HEC schemes, the dual-mode

NIZK system makes use of dual-mode commitments [31], where the commitment key can be tuned to give either statistically hiding or extractable commitments.

**CONSTRUCTION.** At a high level, our construction proceeds as follows. When a user joins the group, he generates a fresh public key for a digital signature  $\mathbf{pk}_{\mathbf{id}}$  and obtains from the group manager (GM) a membership certificate  $\mathbf{cert}_{\mathbf{id}}$  consisting of the GM’s signature on the pair  $(\mathbf{id}, \mathbf{pk}_{\mathbf{id}})$ , where  $\mathbf{id}$  is the user’s identity. In order to sign a message w.r.t. the predicate  $P$ , a group member computes an HEC commitment  $\mathbf{com}_{(\mathbf{id}, w)}$  to the witness  $w$  and his identity  $\mathbf{id}$ . At the same time, the signer computes a dual-mode commitment to  $(\mathbf{id}, w)$ , which is configured to be statistically hiding in the real scheme. The signer then considers the message-dependent circuit  $C_M(\cdot, \cdot)$  which evaluates  $C_M(\mathbf{id}, w) = P(M, \mathbf{id}, w)$  on input of  $(\mathbf{id}, w)$ . He runs the private HEC evaluation algorithm to compute a proof  $\pi_{C, M}$  that  $\mathbf{com}_{ev} = \mathbf{Eval}(C_M, \mathbf{com}_{ev})$  is really a commitment to  $C_M(\mathbf{id}, w)$ . It finally computes a public key encryption  $\mathbf{ct}_{\mathbf{id}} \leftarrow \mathbf{Encrypt}(\mathbf{pk}, (1 - C_M(\mathbf{id}, w)) \cdot \mathbf{id})$  of the product  $(1 - C_M(\mathbf{id}, w)) \cdot \mathbf{id}$ , so that  $\mathbf{ct}_{\mathbf{id}}$  encrypts 0 when unconditional anonymity is enabled (i.e., when  $C_M(\mathbf{id}, w) = 1$ ) and  $\mathbf{id}$  otherwise. A dual-mode NIZK argument then allows arguing that all steps were properly carried out.

When the tracing capability is enabled (namely, when  $C_M(\mathbf{id}, w) = 0$ ), we need to rely on a special kind of dual-mode commitment to prove computational anonymity in the CCA sense (i.e., when the adversary has access to a signature opening oracle). Specifically, we need to program the commitment key to make commitments extractable in all signature opening queries and statistically hiding in the challenge phase. This is achieved using tag-based commitments, where each commitment is computed under a tag-dependent commitment key that either provides statistically hiding or extractable commitments, depending on the tag. In order to instantiate these dual-mode tag-based commitments, we use a recent variant [41] of the  $\mathcal{R}$ -lossy encryption of Boyle *et al.* [18], for which we give a DDH-based realization (as well as an LWE-based realization adapted from [41]). Using this  $\mathcal{R}$ -lossy encryption to instantiate the dual-mode commitment component, we can make it statistically hiding in the challenge phase (for a specific tag corresponding to a one-time verification key  $\mathbf{VK}^*$ ) and statistically extractable for all other tags  $\mathbf{VK} \neq \mathbf{VK}^*$ . This allows us to proceed with a sequence of hybrid games where, instead of answering opening queries by decrypting the ciphertext  $\mathbf{ct}_{\mathbf{id}}$ , we can extract the committed  $(\mathbf{id}, w)$ . This can be seen as applying the two-key paradigm of Naor and Yung [48] for CCA2-secure encryption.

The above construction crucially relies on HEC to avoid the signature length from depending on the circuit size. If we were to give up the circuit-size independent property, constructions would be possible from any non-interactive statistically hiding commitment.

**OPEN QUESTIONS.** Naturally, our work, being in a new area (with some new technical and definitional challenges as described above), leaves many open problems which we believe to be interesting. Conceptually, one can ask how bifurcated cryptography applies elsewhere. Technically, the first problem that comes to mind is finding practically efficient instantiations from lattice assumptions (even in the random oracle model) as our LWE-based construction is only meant to be a first

feasibility result. In particular, it would be interesting to provide more efficient lattice-based solutions using, e.g., the Fiat-Shamir-with-abort method [44,14] or the techniques recently suggested in [58]. The second open problem is to determine the extent to which more specific predicate families can be realized more efficiently (with or without random oracles) and under different assumptions. While the work of Katsumata *et al.* [34] implies a pairing-based construction for Boolean formulas, it is only known to achieve circuit-size independence under a  $q$ -type (although falsifiable) assumption. A sufficient condition to avoid relying on variable-size assumptions in the pairing setting would be to build an HEC scheme based on simple assumptions, where the size of partial openings does not depend on the circuit size. Finally, while the rest of the paper will concentrate on BiAS, exploring further primitives providing point-wise predicate-based built-in privacy or confidentiality bifurcated tradeoffs seems like a new area for broader investigations.

### 1.3 Related Work

The rigid anonymity vs. accountability situation indeed generated much discomfort in the community. Group signatures traditionally allow opening authorities to identify the author of any signature. As advocated by Sakai *et al.* [55], it may be desirable to prevent the OA from seeing the entire signature history of all group members. Restricting the power of the opening authorities is a challenging research direction, where a few steps have been taken. For example, traceable signatures [35], group signatures with message-dependent opening [55], and accountable tracing signatures [40] can all be viewed as group signatures with restricted opening authorities. However, the OA can still freely break the anonymity of some subset of signatures without the user’s agreement. These primitives offer more privacy to users than ordinary group signatures, but not at a level that privacy-sensitive users would hope for.

Ring signatures [54,9] confer everlasting anonymity to group members. They depart from group signatures in that signers are not required to register in the system and signers have complete freedom on the list of their ring-mates at each signature. Compared to group signatures, they stand at the opposite extreme of the spectrum as they do not provide any accountability at all. Linkable ring signature [43], traceable ring signatures [25], as well as  $k$ -time anonymous authentication systems [56] only introduce a weak form of accountability in ring signatures as users only lose anonymity to some extent: for example, if they issue two or more signatures for some message, their signatures become linkable but they can still create one controversial signature and disappear from the system without being caught.

Accountable tracing signatures (ATS) [39,40] take a different approach to balance accountability and anonymity, by allowing the two extreme settings of ring signatures and group signatures to co-exist. In ATS schemes, a given user is either always unconditionally anonymous or always traceable, based on a decision made by the authority when the user joins the system. In addition, users are never notified about their traceability status. In our setting and use-cases, the

ability to trace or not should depend point-wise on the message and not only on group members’ identities. Moreover, we deliberately aim at leaving users some control on when and under which circumstances they want to accept traceability.

Accountable ring signatures (ARS) [57,13] provide another kind of tradeoff where anonymity and traceability can live together. Xu and Yung [57] consider a threshold opening mechanism where no single opener is given the entire power. The ARS model of Bootle *et al.* [13] also provides some flexibility in the choice of tracing authorities as signers are allowed to choose which opener they trust without necessarily leaving the full tracing capability to a pre-determined authority. On the other hand, neither of these models [57,13] provides full expressiveness as to which messages can be signed with unconditional anonymity and which ones should always be traceable.

Bangerter *et al.* [5] considered an informal framework allowing to monitor the release of certified data. Their (interactive) model fully trusts the opening authority to only disclose users’ data when specific conditions are met. In contrast, BiAS does not trust the opener when de-anonymization conditions are not fulfilled and even requires unconditional anonymity in this case.

Boyen and Delerablée [16] introduced a variant of group signatures allowing group members to flexibly and expressively choose a subgroup wherein they hide their identity. Our goals are orthogonal to theirs since their model always allows tracing authorities to identify signers whereas we accurately control the conditions under which the signer’s identity can be uncovered.

Garms and Lehmann [26] put forth the concept of convertably linkable signatures (CLS) which are group signatures where a “converter” can blindly relate a bunch of signatures to some randomized pseudonyms. To convert the given signatures into linkable ones, another authority first blinds the signatures in order to mitigate the power of the converter. However, the converter is actually an opening authority that can always trace a given signature as long as it was not blinded. At any time, CLS thus “only” provide computational anonymity.

Attribute-based signatures (ABS) [45] allow a signer to sign a message while simultaneously showing possession of credentials satisfying a public predicate. Policy-based signatures (PBS) [6] are signature schemes where users obtain a policy-based signing key (associated with some predicate  $P$ ) from some authority, which allows them to sign exactly those messages  $M$  for which  $P(M) = 1$ . ABS and PBS address different problems than our BiAS primitive in that they provide fine-grained control over “who can sign” and “which messages can be signed at all”, respectively. As such, they do not give users control over which signatures can be traced (with user knowledge of it). Our BiAS functionality departs from PBS [6] in that predicates are not associated with keys but with signed messages and may vary across signatures. In terms of generic implications, PBS were shown [6] to imply digital signatures, NIZK [10] and CCA-secure encryption [52]. However, they are not known to imply homomorphic equivocable commitments with circuit-size independent verification, and the relationship between PBS and BiAS thus remains unclear. In particular, we do not see any obvious way to obtain BiAS realizations by generically using a PBS.

Functional signatures [17] differ from conditionally traceable signatures in that, in the same spirit as PBS, they accurately control which messages can be signed. In contrast, BiAS controls which signatures can be traced.

In the context of anonymous compact e-cash [20], Camenisch *et al.* [21] were bothered by the anonymity vs. accountability issue in a specific scenario. They considered a conditional anonymity setting which restricts transactions to be untraceable only when they do not exceed a specific amount. In their model, the threshold amount is fixed for each merchant over multiple transactions. If a user performs a number of transactions with total values exceeding a threshold, he can be traced based on public records. On the other hand, if the total amount remains under the threshold, the traceable authority is unable to extract the user’s identity. However, the anonymity remains computational even in that case. In contrast, our BiAS primitive is a generic add-on mechanism, and ensures the statistical anonymity of signers as long as the predicate equals 1, e.g., the amount does not reach a fixed threshold.

In summary, there was a lot of discomfort with the existing dichotomy between anonymity and accountability. This has produced a large number of interesting and useful cases and solutions. However, a bifurcated solution with choice at the user’s hand folded into a single scheme, where the choice is driven by a local predicate which further remains undetected by others, has not been considered. This new BiAS system, which gives the user the best possible (i.e., unconditional) anonymity when permitted, in fact, constitutes the most private and the most versatile solution for the problem of balancing user’s vs. society’s needs within anonymous signing scenarios.

## 2 Preliminaries

### 2.1 $\mathcal{R}$ -Lossy Public-Key Encryption

Boyle *et al.* [18] formalized the notion of  $\mathcal{R}$ -lossy encryption. The primitive is a tag-based encryption scheme [38] where the tag space  $\mathcal{T}$  is partitioned into *injective* tags and *lossy* tags. When ciphertexts are generated under an injective tag, the decryption algorithm recovers the underlying plaintext. On lossy tags, the ciphertext statistically hides the underlying plaintext. In  $\mathcal{R}$ -lossy PKE schemes, the tag space is partitioned according to a binary relation  $\mathcal{R} \subseteq \mathcal{K} \times \mathcal{T}$ . The key generation algorithm inputs an initialization value  $K \in \mathcal{K}$  and partitions  $\mathcal{T}$  in such a way that injective tags  $t \in \mathcal{T}$  are exactly those for which  $(K, t) \in \mathcal{R}$  (i.e., all tags  $t$  for which  $(K, t) \notin \mathcal{R}$  are lossy).

Libert *et al.* [41] considered a flavor of  $\mathcal{R}$ -lossy PKE schemes with two distinct key generation algorithms and equivocable lossy ciphertexts. For our purposes, we do not need to equivocate lossy ciphertexts but we still need two distinct key generation algorithms. Looking ahead, our proof of anonymity (in Lemma 4), requires to switch from a setting where all tags are lossy to a setting where only one tag is lossy. Also, proving other security notions requires to move from the “all lossy” setting to the “all injective” setting in the proof of Theorem 1.

**Definition 1.** Let  $\mathcal{R} \subseteq \mathcal{K}_\lambda \times \mathcal{T}_\lambda$  be an efficiently computable binary relation. An  $\mathcal{R}$ -lossy PKE scheme with efficient opening is a 5-uple of PPT algorithms  $(\text{Par-Gen}, \text{Keygen}, \text{LKeygen}, \text{Encrypt}, \text{Decrypt})$  such that:

**Parameter generation:** On input of a security parameter  $\lambda$ , a desired length of initialization values  $L \in \text{poly}(\lambda)$  and a lower bound  $B \in \text{poly}(\lambda)$  on the message length,  $\text{Par-Gen}(1^\lambda, 1^L, 1^B)$  outputs public parameters  $\Gamma$  that specify a tag space  $\mathcal{T}$ , a space of initialization values  $\mathcal{K}$ , a public key space  $\mathcal{PK}$  and a secret key space  $\mathcal{SK}$ .

**Key generation:** For an initialization value  $K \in \mathcal{K}$  and public parameters  $\Gamma$ , algorithm  $\text{Keygen}(\Gamma, K)$  outputs an injective public key  $\text{pk} \in \mathcal{PK}$  and a decryption key  $\text{sk} \in \mathcal{SK}$ . The public key specifies a ciphertext space  $\text{CtSp}$  and a randomness space  $R^{\text{RLE}}$ .

**Lossy Key generation:** Given an initialization value  $K \in \mathcal{K}$  and public parameters  $\Gamma$ , the lossy key generation algorithm  $\text{LKeygen}(\Gamma, K)$  outputs a lossy public key  $\text{pk} \in \mathcal{PK}$  and a lossy secret key  $\text{sk} \in \mathcal{SK}$ .

**Decryption under injective tags:** For any  $\Gamma \leftarrow \text{Par-Gen}(1^\lambda, 1^L, 1^B)$ , any  $K \in \mathcal{K}$ , any  $t \in \mathcal{T}$  such that  $(K, t) \in \mathcal{R}$ , and any  $\text{Msg} \in \text{MsgSp}$ , we have

$$\Pr[\exists r \in R^{\text{RLE}} : \text{Decrypt}(\text{sk}, t, \text{Encrypt}(\text{pk}, t, \text{Msg}; r)) \neq \text{Msg}] < \nu(\lambda) ,$$

for some negligible function  $\nu(\lambda)$ , where  $(\text{pk}, \text{sk}) \leftarrow \text{Keygen}(\Gamma, K)$  and the probability is taken over the randomness of  $\text{Keygen}$ .

**Indistinguishability:** For any  $\Gamma \leftarrow \text{Par-Gen}(1^\lambda, 1^L, 1^B)$ , the key generation algorithms  $\text{LKeygen}$  and  $\text{Keygen}$  satisfy the following:

- (i) For any  $K \in \mathcal{K}$ , the distributions  $D_{\text{inj}} = \{\text{pk} \mid (\text{pk}, \text{sk}) \leftarrow \text{Keygen}(\Gamma, K)\}$  and  $D_{\text{loss}} = \{\text{pk} \mid (\text{pk}, \text{sk}) \leftarrow \text{LKeygen}(\Gamma, K)\}$  are computationally indistinguishable.
- (ii) For any initialization values  $K, K' \in \mathcal{K}$ , the two distributions  $\{\text{pk} \mid (\text{pk}, \text{sk}) \leftarrow \text{LKeygen}(\Gamma, K)\}$  and  $\{\text{pk} \mid (\text{pk}, \text{sk}) \leftarrow \text{LKeygen}(\Gamma, K')\}$  are statistically indistinguishable.

**Lossiness:** For any  $\Gamma \leftarrow \text{Par-Gen}(1^\lambda, 1^L, 1^B)$ , any initialization value  $K \in \mathcal{K}$  and tag  $t \in \mathcal{T}$  such that  $(K, t) \notin \mathcal{R}$ , any  $(\text{pk}, \text{sk}) \leftarrow \text{Keygen}(\Gamma, K)$ , and any  $\text{Msg}_0, \text{Msg}_1 \in \text{MsgSp}$ , the following distributions are statistically close:

$$\{C \mid C \leftarrow \text{Encrypt}(\text{pk}, t, \text{Msg}_0)\} \approx_s \{C \mid C \leftarrow \text{Encrypt}(\text{pk}, t, \text{Msg}_1)\}.$$

For any  $(\text{pk}, \text{sk}) \leftarrow \text{LKeygen}(\Gamma, K)$ , the above holds for any tag  $t$ .

We will use an  $\mathcal{R}$ -lossy encryption scheme for the inequality relation.

**Definition 2.** Let  $\mathcal{K} = \{0, 1\}^L$  and  $\mathcal{T} = \{0, 1\}^L \setminus \{\mathbf{0}^L\}$ , for some  $L \in \text{poly}(\lambda)$ . The **inequality relation**  $\mathcal{R}_{\text{NEQ}} : \mathcal{K} \times \mathcal{T} \rightarrow \{0, 1\}$  is the relation for which we have  $\mathcal{R}_{\text{NEQ}}(K, t) = 1$  if and only if  $K \neq t$ .

We note that, since we exclude the all-zeroes string  $\mathbf{0}^L$  from  $\mathcal{T}$ , running Keygen on input of the initialization value  $K = \mathbf{0}^L$  produces a key  $\mathbf{pk}$  that is injective for *all* tags. In contrast LKeygen produces keys that are lossy for any tag and any initialization value.

We now give an  $\mathcal{R}_{\text{NEQ}}$ -Lossy PKE realization under the Decision Diffie-Hellman (DDH) assumption (MDDH generalization is obvious). In Supplementary Material E.4, we also provide a construction from the LWE assumption.

**An  $\mathcal{R}_{\text{NEQ}}$ -Lossy PKE Scheme from DDH.** The construction below is inspired from [42], which is itself inspired from Groth-Sahai commitments [32] to scalars. We first recall the definition of the DDH problem.

**Definition 3.** *In a cyclic group  $\mathbb{G}$  of prime order  $p$ , the **Decision Diffie-Hellman Problem** (DDH) in  $\mathbb{G}$ , is to distinguish the distributions  $(g, g^a, g^b, g^{ab})$  and  $(g^a, g^b, g^c)$ , with  $a, b, c \leftarrow \mathbb{Z}_p$ . The **Decision Diffie-Hellman assumption** is the intractability of DDH for any PPT algorithm  $\mathcal{D}$ .*

**Par-Gen**( $1^\lambda, 1^L, 1^B$ ): Define  $\mathcal{K} = \{0, 1\}^L$ , and  $\mathcal{T} = \{0, 1\}^L \setminus \{\mathbf{0}^L\}$ . Define public parameters as  $\Gamma = (1^\lambda, 1^L, 1^B)$ .

**Keygen**( $\Gamma, K$ ): On input of  $\Gamma$  and  $K \in \mathcal{K}$ , generate a key pair as follows.

1. Choose a cyclic group  $\mathbb{G}$  or prime order  $p > 2^\lambda$  with a generator  $g \leftarrow U(\mathbb{G})$ . Choose  $\alpha \leftarrow U(\mathbb{Z}_p)$  and compute  $h = g^\alpha$ . Define  $\mathbf{g}_0 = (g, h) \in \mathbb{G}^2$  and  $\mathbf{g} = \mathbf{g}_0^\beta \cdot (1, g) \in \mathbb{G}^2$ , where  $\beta \leftarrow U(\mathbb{Z}_p)$ .
2. Pick  $\gamma \leftarrow U(\mathbb{Z}_p)$  and compute  $\mathbf{u} = \mathbf{g}_0^\gamma \cdot \mathbf{g}^{-K} \in \mathbb{G}^2$ , where  $K \in \{0, 1\}^L$  is interpreted as an element of  $\mathbb{Z}_p$ .

Define  $R^{\text{RLE}} = \mathbb{Z}_p^B$  and output  $\mathbf{sk} = (\alpha, K)$  as well as  $\mathbf{pk} := (\mathbb{G}, \mathbf{g}_0, \mathbf{g}, \mathbf{u})$ .

**LKeygen**( $\Gamma, K$ ): This algorithm is identical to Keygen with the difference that, at step 1, it computes  $\mathbf{g}$  as  $\mathbf{g} = \mathbf{g}_0^\beta \in \mathbb{G}^2$ , where  $\beta \leftarrow U(\mathbb{Z}_p)$ . It defines  $R^{\text{RLE}} = \mathbb{Z}_p^B$  and outputs  $\mathbf{sk} = (\alpha, K)$  as well as  $\mathbf{pk} := (\mathbb{G}, \mathbf{g}_0, \mathbf{g}, \mathbf{u})$ .

**Encrypt**( $\mathbf{pk}, t, \text{Msg}$ ): To encrypt  $\text{Msg} \in \{0, 1\}^B$ , interpret the tag  $t \in \mathcal{T}$  as an element of  $\mathbb{Z}_p$ . For each index  $i \in [B]$ , pick  $r_i \leftarrow U(\mathbb{Z}_p^*)$  and compute  $\mathbf{ct}_i = (\mathbf{u} \cdot \mathbf{g}^t)^{\text{Msg}[i]} \cdot \mathbf{g}_0^{r_i} \in \mathbb{G}^2$ . Then, output  $\mathbf{ct} = (\mathbf{ct}_1, \dots, \mathbf{ct}_B) \in \mathbb{G}^{2B}$ .

**Decrypt**( $\mathbf{sk}, t, \mathbf{ct}$ ): Given  $\mathbf{sk} = (\alpha, K)$  and  $t \in \{0, 1\}^L$ , interpret  $t$  as an element of  $\mathbb{Z}_p$ . If  $t = K$ , return  $\perp$ . Otherwise, for each  $i \in [B]$ , do the following:

1. Parse  $\mathbf{ct}_i$  as  $(\mathbf{ct}_{i,1}, \mathbf{ct}_{i,2}) \in \mathbb{G}^2$
2. Set  $\text{Msg}[i] = 0$  if  $\mathbf{ct}_{i,2} = \mathbf{ct}_{i,1}^\alpha$  and  $\text{Msg}[i] = 1$  otherwise.

Output  $\text{Msg} \in \{0, 1\}^B$ .

The proof of Lemma 1 is straightforward. The first indistinguishability property follows immediately from the semantic security of ElGamal and the observation that Keygen and LKeygen only differ in the distribution of  $\mathbf{g}$ . The second indistinguishability property follows from the fact that, for any  $K \in \mathcal{K}$  and any public key  $\mathbf{pk}$  generated by LKeygen,  $\mathbf{u} \in \mathbb{G}^2$  is uniformly distributed in the subspace

spanned by  $\mathbf{g}_0$ . The same holds for any ciphertext encrypted under an injective key for the lossy tag  $t = K$  or a lossy key for any tag. The construction readily extends to rely on the  $k$ -linear assumption [11] for  $k > 1$ .

**Lemma 1.** *The above construction is an  $\mathcal{R}_{\text{NEQ}}$ -lossy public-key encryption scheme under the DDH assumption.*

## 2.2 Homomorphic Equivocal Commitments

We now recall the definition of homomorphic equivocal commitment, as formalized by Katsumata *et al.* [34].

**Definition 4.** *A HEC scheme with message space  $\mathcal{X}$ , randomness space  $R^{\text{HEC}}$  and randomness distribution  $\mathcal{D}^{\text{HEC}}$  over  $R^{\text{HEC}}$  for a circuit class  $\mathcal{C} = \{C : \mathcal{X} \rightarrow \{0, 1\}\}$  is a tuple of PPT algorithms  $\text{HEC} = (\text{Setup}, \text{Commit}, \text{Open}, \text{Eval}^{\text{in}}, \text{Eval}^{\text{out}}, \text{Verify})$  with the following specifications:*

**Setup**( $1^\lambda$ ): *Inputs a security parameter  $1^\lambda$  and outputs public parameters  $\text{pp}$ , an evaluation key  $\text{ek}$  and a master secret key  $\text{msk}$ .*

**Commit**( $\text{pp}, \mathbf{x}, r$ ): *Takes as input public parameters  $\text{pp}$ , a message  $\mathbf{x} \in \mathcal{X}$  and randomness  $r \in R^{\text{HEC}}$ . It outputs a commitment  $\text{com}$ . When  $r$  is omitted from the notation  $\text{Commit}(\text{pp}, \mathbf{x})$ , we mean that  $r$  is sampled from  $\mathcal{D}^{\text{HEC}}$ .*

**Open**( $\text{msk}, \mathbf{x}, r, \mathbf{x}'$ ): *Takes as input a master secret key  $\text{msk}$ , messages  $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$ , and randomness  $r \in R^{\text{HEC}}$ . It outputs fake randomness  $r' \in R^{\text{HEC}}$ .*

**Eval**<sup>in</sup>( $\text{ek}, C, \mathbf{x}, r$ ): *The inner evaluation algorithm inputs a key  $\text{ek}$ , a circuit  $C \in \mathcal{C}$ , a message  $\mathbf{x}$  and randomness  $r \in R^{\text{HEC}}$ . It outputs a proof  $\pi$ .*

**Eval**<sup>out</sup>( $\text{ek}, C, \text{com}$ ): *The outer evaluation algorithm is a deterministic algorithm that inputs an evaluation key  $\text{ek}$ , a circuit  $C \in \mathcal{C}$  and a commitment  $\text{com}$ . It outputs an evaluated commitment  $\text{com}_{\text{ev}}$ .*

**Verify**( $\text{pp}, \text{com}_{\text{ev}}, z, \pi$ ): *The verification algorithm takes as input public parameters  $\text{pp}$ , an evaluated commitment  $\text{com}_{\text{ev}}$ , a message  $z \in \{0, 1\}$ , and a proof  $\pi$ . It outputs 0 or 1.*

In addition, it should satisfy the following properties:

**Evaluation correctness:** For all  $\lambda \in \mathbb{N}$ , all  $(\text{pp}, \text{ek}, \text{msk}) \leftarrow \text{HEC}.\text{Setup}(1^\lambda)$ , any input  $\mathbf{x} \in \mathcal{X}$ , any randomness  $r \in R^{\text{HEC}}$ , and any circuit  $C \in \mathcal{C}$ , if  $\text{com} = \text{HEC}.\text{Commit}(\text{pp}, \mathbf{x}; r)$ ,  $\pi \leftarrow \text{HEC}.\text{Eval}^{\text{in}}(\text{msk}, C, \mathbf{x}, r)$ , and

$$\text{com}_{\text{ev}} = \text{HEC}.\text{Eval}^{\text{out}}(\text{ek}, C, \text{com}),$$

then  $\Pr[\text{HEC}.\text{Verify}(\text{pp}, \text{com}_{\text{ev}}, C(\mathbf{x}), \pi) = 1] \geq 1 - \nu(\lambda)$ , for some function  $\nu(\lambda) \in \text{negl}(\lambda)$ .

**Distributional equivalence of Open:** For all  $\lambda \in \mathbb{N}$ , any  $(\text{pp}, \text{ek}, \text{msk}) \leftarrow \text{HEC}.\text{Setup}(1^\lambda)$ , any  $\mathbf{x}, \bar{\mathbf{x}} \in \mathcal{X}$ , randomness  $r, r' \leftarrow \mathcal{D}^{\text{HEC}}$ , the distributions  $\{(\text{pp}, \text{ek}, \text{msk}, \mathbf{x}, r, \text{com}) \mid r \leftarrow \mathcal{D}^{\text{HEC}}, \text{com} = \text{HEC}.\text{Commit}(\text{pp}, \mathbf{x}; r)\}$  and

$$\begin{aligned} & \{(\text{pp}, \text{ek}, \text{msk}, \mathbf{x}, r' = \text{HEC}.\text{Open}(\text{msk}, \bar{\mathbf{x}}, \bar{r}, \mathbf{x}), \text{com}') \mid \\ & \quad \bar{r} \leftarrow \mathcal{D}^{\text{HEC}}, \text{com}' = \text{HEC}.\text{Commit}(\text{pp}, \bar{\mathbf{x}}; \bar{r})\} \end{aligned}$$

are statistically close.

**Computational binding on evaluated commitments:** For any PPT adversary  $\mathcal{A}$ , we have

$$\begin{aligned} \Pr \left[ \text{HEC.Verify}(\text{pp}, \text{com}_{ev}, z^*, \pi^*) = 1 \wedge z^* \neq C(\mathbf{x}) \mid \right. \\ (\text{pp}, \text{ek}, \text{msk}) \leftarrow \text{HEC.Setup}(1^\lambda), \\ (\mathbf{x}, r, C, z^*, \pi^*) \leftarrow \mathcal{A}(\text{pp}, \text{ek}), \\ \text{com} = \text{HEC.Commit}(\text{pp}, \mathbf{x}; r), \\ \left. \text{com}_{ev} = \text{HEC.Eval}^{out}(\text{ek}, C, \text{com}) \right] \in \text{negl}(\lambda) \end{aligned}$$

**Efficient committing:** There exists a polynomial  $\text{poly}(\lambda)$  such that, for all  $(\text{pp}, \text{ek}, \text{msk}) \leftarrow \text{HEC.Setup}(1^\lambda)$ ,  $\mathbf{x} \in \mathcal{X}$  and  $r \in R^{\text{HEC}}$ , the running time of  $\text{com} = \text{HEC.Commit}(\text{pp}, \mathbf{x}; r)$  is bounded by  $|\mathbf{x}| \cdot \text{poly}(\lambda)$ .

**Efficient verification:** There exists a polynomial  $\text{poly}(\lambda)$  such that, for all  $(\text{pp}, \text{ek}, \text{msk}) \leftarrow \text{HEC.Setup}(1^\lambda)$  and any  $\mathbf{x} \in \mathcal{X}$ ,  $r \in R^{\text{HEC}}$ ,  $C \in \mathcal{C}$  and  $z \in \{0, 1\}$ , if  $\text{com} = \text{HEC.Commit}(\text{pp}, \mathbf{x}; r)$ ,  $\pi \leftarrow \text{HEC.Eval}^{in}(\text{msk}, C, \mathbf{x}, r)$  and  $\text{com}_{ev} = \text{HEC.Eval}^{out}(\text{ek}, C, \text{com})$ , then  $|\pi|, |\text{com}_{ev}| \leq \text{poly}(\lambda)$  and the running time of  $\text{HEC.Verify}(\text{pp}, \text{com}_{ev}, z, \pi)$  is at most  $\text{poly}(\lambda)$  (which does not depend on  $|C|$ ).

**Context hiding:** There exists a PPT simulator  $\text{HEC.ProofSim}$  such that, for all  $\lambda \in \mathbb{N}$ ,  $(\text{pp}, \text{ek}, \text{msk}) \leftarrow \text{HEC.Setup}(1^\lambda)$ ,  $\mathbf{x} \in \mathcal{X}$ ,  $C \in \mathcal{C}$ ,  $r \in R^{\text{HEC}}$  and  $\text{com} = \text{HEC.Commit}(\text{pp}, \mathbf{x}; r)$ , the distribution  $\{\pi \leftarrow \text{HEC.Eval}^{in}(\text{msk}, C, \mathbf{x}, R)\}$  is statistically indistinguishable from  $\{\pi' \leftarrow \text{HEC.ProofSim}(\text{msk}, \text{com}, C, C(\mathbf{x}))\}$ .

We note that the distributional equivalence of **Open** implies that the commitment is statistically hiding. Here, we only need the statistically hiding property and we do not rely on equivocation. We do not explicitly rely on the context hiding property either since partial openings  $\pi$  produced by **Open** will be part of witnesses in a NIZK argument. On the other hand, we will exploit the efficient verification property to achieve circuit-size independence in terms of signature size.<sup>6</sup>

### 3 Bifurcated Anonymous Signatures

This section formalizes the primitive of bifurcated anonymous signature (BiAS). BiAS is the first general signature primitive reconciling statistical anonymity and accountability in front of dishonest authorities in a single scheme. Nevertheless, our model and syntax are inspired by those in the context of dynamic group signatures given by Kiayias and Yung [36] —who extended the work of Bellare, Micciancio and Warinschi [7] on static group signatures.

#### 3.1 Syntax

Like in [36], we consider dynamically growing groups. The syntax includes an interactive protocol which allows users to be enrolled as new members of the

<sup>6</sup> As pointed out in [34, Remark 3.3], it is straightforward to build an HEC without the context-hiding and efficient verification properties, using any statistically hiding commitment.

group at any time. Analogously to the similar model of Bellare, Shi and Zhang [8], the Kiayias-Yung (KY) model assumes an interactive *join* protocol whereby a prospective user becomes a group member by interacting with the group manager responsible to issuing credentials. This protocol provides the user with a membership certificate,  $\text{cert}_i$ , and a membership secret,  $\text{sec}_i$ .

In the syntax below, we define a space  $\mathcal{ID}$  of user identifiers as well as a space of witnesses  $\mathcal{W}$ . For any message-identity-witness triple  $(M, \text{id}, w) \in \mathcal{M} \times \mathcal{ID} \times \mathcal{W}$ , we adopt the convention that  $P(M, \text{id}, w) = 0$  whenever the user of identity  $\text{id}$  was only allowed to sign  $M$  using the witness  $w$  while being subject to tracing, by the opening authority. In contrast, having  $P(M, \text{id}, w) = 1$  allows the user to use the witness  $w$  to generate a signature on  $M$  while retaining unconditional anonymity. For generality and more applicability, a BiAS defines a family  $\mathcal{P}$  of authorized *public* predicates  $P$  which are needed to verify signatures.

**Definition 5 (Bifurcated Anonymous Signature).** *A bifurcated anonymous signature (BiAS) scheme for a predicate family  $\mathcal{P}$  consists of the following algorithms or protocols.*

**Setup**( $1^\lambda, 1^N$ ): given a security parameter  $\lambda$  and a maximal number of group members  $N \in \text{poly}(\lambda) \cap \mathbb{N}$ , this algorithm is run by a trusted party to generate a group public key  $\mathcal{Y}$  associated to the predicate family  $\mathcal{P}_\lambda$ , which specifies a message space  $\mathcal{M}$ , a space of user identifiers  $\mathcal{ID}$ , and a witness space  $\mathcal{W}$ . It also outputs the group manager's private key  $\mathcal{S}_{GM}$  and the opening authority's private key  $\mathcal{S}_{OA}$ . Each key is given to the appropriate authority while  $\mathcal{Y}$  is made public. The algorithm also initializes a public state  $St$  comprising a set data structure  $St_{\text{users}} = \emptyset$  and a string data structure  $St_{\text{trans}} = \epsilon$ . From now on we assume  $\lambda$  is implicit.

**Join:** is an interactive protocol between the group manager  $GM$  and a user  $\mathcal{U}$  who gets a unique identifier  $\text{id} \in \mathcal{ID}$  (both are responsible for enforcing the uniqueness of the identifier). The protocol involves two interactive Turing machines  $J_{\text{user}}$  and  $J_{GM}$  that both take  $\mathcal{Y}$  as input. The execution, denoted as  $[J_{\text{user}}(\mathcal{Y}), J_{GM}(\mathcal{Y}, \mathcal{S}_{GM}, St)]$ , ends with user  $\mathcal{U}_{\text{id}} := \mathcal{U}$  obtaining a membership secret  $\text{sec}_{\text{id}}$ , that no one else knows, and a membership certificate  $\text{cert}_{\text{id}}$ . If the protocol is successful, the group manager updates the public state  $St$  by setting  $St_{\text{users}} := St_{\text{users}} \cup \{\text{id}\}$  as well as  $St_{\text{trans}} := St_{\text{trans}} \parallel \langle \text{id}, \text{transcript}_{\text{id}} \rangle$ .

**Sign**( $\text{id}, \text{cert}_{\text{id}}, \text{sec}_{\text{id}}, M, w, P$ ): given an identifier  $\text{id} \in \mathcal{ID}$ , a membership certificate  $\text{cert}_{\text{id}}$ , a membership secret  $\text{sec}_{\text{id}}$ , a message  $M \in \mathcal{M}$ , a witness  $w \in \mathcal{W}$ , and a predicate  $P \in \mathcal{P}$ , this probabilistic algorithm outputs a signature  $\sigma$ .

**Verify**( $\mathcal{Y}, M, \sigma, P$ ): given a message  $M \in \mathcal{M}$ , a signature  $\sigma$ , a predicate  $P \in \mathcal{P}$  and a group public key  $\mathcal{Y}$ , this deterministic algorithm returns either 0 or 1.

**Open**( $\mathcal{Y}, \mathcal{S}_{OA}, M, \sigma, P, St$ ): takes as input the opening authority's private key  $\mathcal{S}_{OA}$ , a message  $M \in \mathcal{M}$ , a signature  $\sigma$  w.r.t.  $\mathcal{Y}$  and a predicate  $P \in \mathcal{P}$  as well as the public state  $St$ . It outputs  $\text{id} \in St_{\text{users}} \cup \{\perp\}$ , which is the identity of a group member or a symbol indicating anonymity.

Correctness basically requires that, if all parties *honestly* run the protocols, all algorithms are correct with respect to their specification described as above.

*Correctness of Bifurcated Anonymous Signatures.* Following the Kiayias-Yung terminology [36], we say that a public state  $St$  is *valid* if it can be reached from  $St = (\emptyset, \epsilon)$  by a Turing machine having oracle access to  $J_{GM}$ . Also, a state  $St'$  is said to *extend* another state  $St$  if it is within reach from  $St$ . Moreover, we write  $cert_{id} \Leftarrow_{\mathcal{Y}} sec_{id}$  to mean that there exists coin tosses  $\varpi$  for  $J_{GM}$  and  $J_{user}$  such that, for some valid public state  $St'$ , the execution of the interactive protocol  $[J_{user}(\mathcal{Y}), J_{GM}(\mathcal{Y}, \mathcal{S}_{GM}, St')](\varpi)$  provides  $J_{user}$  with  $\langle id, cert_{id}, cert_{id} \rangle$ .

**Definition 6 (Correctness).** A BiAS scheme is correct if the following conditions are all satisfied for any  $(St, \mathcal{Y}, \mathcal{S}_{GM}, \mathcal{S}_{OA}) \leftarrow \text{Setup}(1^\lambda, 1^N)$ :

- (1) In a valid state  $St$ ,  $|St_{users}| = |St_{trans}|$  always holds and two distinct entries of  $St_{trans}$  always contain certificates with distinct id.
- (2) If  $[J_{user}(\mathcal{Y}), J_{GM}(\mathcal{Y}, \mathcal{S}_{GM}, St)]$  is run by two honest parties following the protocol and  $\langle id, cert_{id}, sec_{id} \rangle$  is obtained by  $J_{user}$ , then we have  $cert_{id} \Leftarrow_{\mathcal{Y}} sec_{id}$ .
- (3) For each  $\langle id, cert_{id}, sec_{id} \rangle$  such that  $cert_{id} \Leftarrow_{\mathcal{Y}} sec_{id}$ , any message  $M \in \mathcal{M}$ , any witness  $w \in \mathcal{W}$  and any predicate  $P \in \mathcal{P}$ , we have

$$\text{Verify}(\mathcal{Y}, M, \text{Sign}(id, cert_{id}, sec_{id}, M, w, P), P) = 1.$$

- (4) For any  $M \notin \mathcal{M}$  or any  $P \notin \mathcal{P}$ , and any  $\sigma$ , we have  $\text{Verify}(\mathcal{Y}, M, \sigma, P) = 0$ .
- (5)  $\text{Open}(\mathcal{Y}, \mathcal{S}_{OA}, M, \sigma, P, St) \in \mathcal{ID} \cup \{\perp\}$  as long as  $\text{Verify}(\mathcal{Y}, M, \sigma, P) = 1$ .
- (6) For any outcome  $\langle id, cert_{id}, sec_{id} \rangle$  of  $[J_{user}(\cdot, \cdot), J_{GM}(\cdot, St, \cdot, \cdot)]$ , for some valid  $St$ , any predicate  $P \in \mathcal{P}$ , any message  $M \in \mathcal{M}$ , any witness  $w \in \mathcal{W}$  and  $\sigma \leftarrow \text{Sign}(id, cert_{id}, sec_{id}, M, w, P)$ , with overwhelming probability:
  - (a) if  $P(M, id, w) = 0$ , then  $\text{Open}(\mathcal{Y}, \mathcal{S}_{OA}, M, \sigma, P, St) = id$ ;
  - (b) if  $P(M, id, w) = 1$ , then  $\text{Open}(\mathcal{Y}, \mathcal{S}_{OA}, M, \sigma, P, St) = \perp$ .

We formalize security properties via experiments where the adversary interacts with a stateful interface  $\mathcal{I}$  that maintains the following variables:

- $\text{state}_{\mathcal{I}}$ : is a data structure representing the state of the interface as the adversary invokes the various oracles available in the attack games. It is initialized as  $\text{state}_{\mathcal{I}} = (St, \mathcal{Y}, \mathcal{S}_{GM}, \mathcal{S}_{OA}) \leftarrow \text{Setup}(1^\lambda, 1^N)$ . It includes the (initially empty) set  $St_{users}$  of group members and a dynamically growing database  $St_{trans}$  storing the transcripts of previously executed join protocols.
- $n = |St_{users}| \leq N$  denotes the current cardinality of the group.
- $\text{Sigs}$ : is a database of honestly generated signatures created by the signing oracle. Each entry consists of a tuple  $(id, M, w, \sigma, P)$  indicating that message  $M$  was signed by user  $id$  with respect to the witness  $w$  and the predicate  $P$ .
- $U^a$ : is an initially empty set of users that are introduced by the adversary in the system in an execution of the join protocol.
- $U^b$ : is an initially empty set of honest users introduced in the system by the adversary acting as a dishonest group manager. For these users, the adversary obtains the transcript of  $[J_{user}, J_{GM}]$  but not the user's membership secret.

In attack games, adversaries are granted access to the following oracles:

- $Q_{\text{pub}}$ ,  $Q_{\text{keyGM}}$  and  $Q_{\text{keyOA}}$ : when these oracles are invoked, the interface looks up  $\text{state}_{\mathcal{I}}$  and returns the group public key  $\mathcal{Y}$ , the GM’s private key  $\mathcal{S}_{\text{GM}}$  and the opening authority’s private key  $\mathcal{S}_{\text{OA}}$  respectively. Once the oracle  $Q_{\text{keyGM}}$  (resp.  $Q_{\text{keyOA}}$ ) is invoked, it updates the initially empty key state  $St_{\text{GM}} \leftarrow \{\mathcal{S}_{\text{GM}}\}$  (resp.  $St_{\text{OA}} \leftarrow \{\mathcal{S}_{\text{OA}}\}$ ).
- $Q_{\text{a-join}}$ : allows the adversary to introduce users under its control in the group. On behalf of the GM, the interface runs  $J_{\text{GM}}$  in interaction with the  $J_{\text{user}}$ -executing adversary who plays the role of the prospective user in the join protocol. At the beginning of  $J_{\text{user}}$ , the user chooses an identifier  $\text{id}$  and the interface aborts if  $\text{id}$  was previously assigned to a different user in  $U^a$ . If this protocol successfully ends, the interface updates  $St$  by inserting the new user  $\text{id}$  in both sets  $St_{\text{users}}$  and  $U^a$ . It also sets  $St_{\text{trans}} := St_{\text{trans}} \parallel \langle \text{id}, \text{transcript}_{\text{id}} \rangle$ .
- $Q_{\text{b-join}}$ : allows the adversary, acting as a corrupted group manager, to introduce new honest group members of its choice. The interface triggers an execution of  $[J_{\text{user}}, J_{\text{GM}}]$  and runs  $J_{\text{user}}$  in interaction with the adversary who runs  $J_{\text{GM}}$ . If the protocol successfully completes, the interface adds user  $\text{id}$  to  $St_{\text{users}}$  and  $U^b$  and sets  $St_{\text{trans}} := St_{\text{trans}} \parallel \langle \text{id}, \text{transcript}_{\text{id}} \rangle$ . It stores the membership certificate  $\text{cert}_{\text{id}}$  and the membership secret  $\text{sec}_{\text{id}}$  in a *private* part of  $\text{state}_{\mathcal{I}}$ .
- $Q_{\text{sig}}$ : given a tuple  $(M, w, P)$  and an identifier  $\text{id}$ , the interface returns  $\perp$  if  $\text{id} \notin U^b$ . Otherwise, the private area of  $\text{state}_{\mathcal{I}}$  must contain a certificate  $\text{cert}_{\text{id}}$  and a membership secret  $\text{sec}_{\text{id}}$ . The interface outputs a signature  $\sigma$  on behalf of user  $\text{id}$  and also updates  $\text{Sigs} \leftarrow \text{Sigs} \parallel \langle \text{id}, M, w, \sigma, P \rangle$ .
- $Q_{\text{open}}$ : when this oracle is invoked on input of a valid triple  $(M, \sigma, P)$ , the interface runs algorithm  $\text{Open}$  using the current state  $St$ . When  $S$  is a set of tuples of the form  $(M, \sigma, P)$ ,  $Q_{\text{open}}^S$  denotes a restricted oracle that only applies the opening algorithm to tuples  $(M, \sigma, P)$  which are not in  $S$ .
- $Q_{\text{read}}$  and  $Q_{\text{write}}$ : are used by the adversary to read and write the content of  $St$ . At each invocation,  $Q_{\text{read}}$  outputs the state  $St$  of the interface. By using  $Q_{\text{write}}$ , the adversary can modify  $St$  at will as long as it does not invalidate  $St$ : for example, the adversary is allowed to create dummy users as long as it does not re-use already existing certificates.

In the random oracle model we implicitly assume that all the BiAS algorithms and protocols have access to the random oracle.

### 3.2 Branch-Hiding and Privacy

*Branch-Hiding.* The notion of branch-hiding captures the infeasibility, even for a corrupt group manager, to decide whether a user signs a message  $M$  for a given predicate  $P$  while enabling traceability or not. In particular,  $P(M, \text{id}, w)$  remains computationally hidden. Said otherwise, signatures do not betray any potential intent of a user to remain untraceable or accept traceability. The formal description is given in Section B as we require a stronger privacy notion for BiAS.

*Full Anonymity.* The notion of anonymity is formalized via two games parametrized by a bit  $d$  and involving a two-stage adversary. The first stage is called *play*

stage and allows the adversary  $\mathcal{A}$  to modify  $\text{state}_{\mathcal{I}}$  via  $Q_{\text{write}}$ -queries and open arbitrary signatures by probing  $Q_{\text{open}}$ . When the play stage ends,  $\mathcal{A}$  chooses a message-predicate pair  $(M^*, P^*)$  as well as two 4-ules  $(\text{id}_0^*, w_0^*, \text{sec}_0^*, \text{cert}_0^*)$  and  $(\text{id}_1^*, w_1^*, \text{sec}_1^*, \text{cert}_1^*)$ , both containing a valid membership certificate and a corresponding membership secret. Then, depending on  $d \in \{0, 1\}$ , the adversary is given a challenge signature  $\sigma^*$  computed using  $(\text{id}_d^*, w_d^*, \text{sec}_d^*, \text{cert}_d^*)$  with the task of eventually guessing the bit  $d \in \{0, 1\}$ . Before doing so, it is allowed further oracle queries throughout the second stage, called **guess** stage, but is restricted not to query  $Q_{\text{open}}$  for  $(M^*, \sigma^*, P^*)$ . We note that the adversary is allowed to choose  $(\text{id}_0^*, \text{sec}_0^*, \text{cert}_0^*)$  and  $(\text{id}_1^*, \text{sec}_1^*, \text{cert}_1^*)$  such that  $P^*(M^*, \text{id}_0^*, w_0^*) \neq P^*(M^*, \text{id}_1^*, w_1^*)$ . Our definition of anonymity thus reflects the inability of a verifier to distinguish signatures that are traceable from those that are not. To strengthen the model, the definition even allows the adversary to corrupt the opening authority as long as  $P^*(M^*, \text{id}_0^*, w_0^*) = 1 = P^*(M^*, \text{id}_1^*, w_1^*)$ . In such a non-traceable case, we require that the indistinguishability is statistically independent of the bit  $d$ . We elaborate more on this adversarial complexity just after.

**Definition 7.** *A BiAS is fully anonymous if it satisfies the next conditions:*

**Traceable case** *For any PPT adversary  $\mathcal{A}$  the following advantage is negligible.*

$$\mathbf{Adv}_{\mathcal{A}, N}^{\text{anon}}(\lambda) := |\Pr[\mathbf{Exp}_{\mathcal{A}, N}^{\text{anon-}1}(\lambda) = 1] - \Pr[\mathbf{Exp}_{\mathcal{A}, N}^{\text{anon-}0}(\lambda) = 1]|$$

**Non-traceable case** *For any (unbounded) adversary involved in  $\mathbf{Exp}_{\mathcal{A}, N}^{\text{anon-}d}$  and  $\mathbf{Exp}_{\mathcal{A}, N}^{\text{anon-ntr-}d}$  (defined in Fig. 1), the following advantage is negligible.*

$$\mathbf{Adv}_{\mathcal{A}, N}^{\text{anon-ntr}}(\lambda) := |\Pr[\mathbf{Exp}_{\mathcal{A}, N}^{\text{anon-ntr-}1}(\lambda) = 1] - \Pr[\mathbf{Exp}_{\mathcal{A}, N}^{\text{anon-ntr-}0}(\lambda) = 1]|$$

```

1 stateI :=  $(St, \mathcal{Y}, \mathcal{S}_{\text{GM}}, \mathcal{S}_{\text{OA}}) \leftarrow \text{Setup}(1^\lambda, 1^N)$ ;
2  $(aux, M^*, w_0^*, w_1^*, (\text{id}_0^*, \text{sec}_0^*, \text{cert}_0^*), (\text{id}_1^*, \text{sec}_1^*, \text{cert}_1^*), P^*)$ 
    $\leftarrow \mathcal{A}(\text{play}; Q_{\text{pub}}, Q_{\text{keyGM}}, Q_{\text{open}}, Q_{\text{read}}, Q_{\text{write}}, Q_{\text{keyOA}})$ ;
3 if  $\neg(\text{cert}_{\text{id}_0^*} \Rightarrow_y \text{sec}_{\text{id}_0^*}) \vee \neg(\text{cert}_{\text{id}_1^*} \Rightarrow_y \text{sec}_{\text{id}_1^*})$  then return 0;
4 if  $(M^* \notin \mathcal{M}) \vee (w_0^*, w_1^* \notin \mathcal{W}) \vee (P^* \notin \mathcal{P})$  then return 0;
5  $\sigma_d \leftarrow \text{Sign}(\mathcal{Y}, \text{id}_d^*, \text{cert}_d^*, \text{sec}_d^*, M^*, w_d^*, P^*)$  ;
6  $d' \leftarrow \mathcal{A}(\text{guess}; \sigma_d, aux, Q_{\text{pub}}, Q_{\text{keyGM}}, Q_{\text{open}}^{\neg\{(M^*, \sigma_d, P^*)\}}, Q_{\text{read}}, Q_{\text{write}}, Q_{\text{keyOA}})$ ;
7 if  $\lceil (P^*(M^*, \text{id}_0^*, w_0^*) = 1) \wedge (P^*(M^*, \text{id}_1^*, w_1^*) = 1) \rceil$  then return  $\lceil d' \rceil$ ;
8 if  $\lceil (St_{\text{OA}} = \emptyset) \rceil$  then return  $\lceil d' \rceil$ ;
9 return 0;

```

**Fig. 1:** Experiment  $\mathbf{Exp}_{\mathcal{A}, N}^{\text{anon-}d}(\lambda)$  (*resp.*  $\mathbf{Exp}_{\mathcal{A}, N}^{\text{anon-ntr-}d}(\lambda)$ ) excluding the dotted (*resp.* solid) box.

The anonymity definition has two parts: a first one that captures the (CCA) unlinkability against all entities but the OA regardless of the predicate value; and a second one which captures the unlinkability even against the OA when the predicate evaluates to 1. Clearly, the first case can never be statistical if the predicate is not constantly equal to 1. However, while the requirement of the second case could only have been computational, we stress that having two cases has nothing to do with the running time of the adversary.

The reason we are requiring statistical anonymity is because we advocate the need to enhance the privacy branch in the context of anonymous signatures. When the predicate equals 1, the signer should have full confidence in his anonymity. Allowing computational anonymity leaves room for a potential backdoor in the system, which could be exploitable in an unexpected way in some applications.

In Section C of the Supplementary Material, we suggest an even stronger notion of anonymity, called unsubversive anonymity, in the non-traceable case. This notion allows for adversarially-generated authorities' keys. Since it only seems achievable in the random oracle model, we do not include it in the general BiAS model and leave the design of a BiAS achieving it for future research.

### 3.3 Branch-Soundness and Security

Defining strong unforgeability-related notions requires being able to check whether an adversary fools the underlying predicate value embedded into signatures. However, checking such a relation needs extracting meaningful information even from (statistically) non-traceable signatures. To circumvent this apparent conflicting requirements we first define the branch-soundness notion which sets an indistinguishable extractable mode even if all the keys are exposed. It also captures the inability of any efficient adversary to produce valid signatures in the extractable mode that contradict the openability of signatures in the real mode. Equipped with a setting where identities and witnesses are extractable “all the time” we can turn to other security notions.

*Branch-Soundness.* To be able to extract  $(\text{id}, w)$  from *any* valid signatures, we introduce an indistinguishable setting allowing such extractions for the purpose of testing the underlying predicate value  $P(M, \text{id}, w)$ . As long as signatures are traceable, we require  $\text{id}$  to be consistent with the outcome of  $\text{Open}$ .

**Definition 8.** *A BiAS scheme satisfies the branch-soundness property if there is a pair of efficient algorithms with the following specifications:*

**SimSetup**( $1^\lambda, 1^N$ ): *given a security parameter  $\lambda$  and a maximal number of users  $N \in \text{poly}(\lambda) \cap \mathbb{N}$ , this algorithm generates a group public key  $\mathcal{Y}$ , the group manager's secret key  $\mathcal{S}_{\text{GM}}$ , the opening authority's secret key  $\mathcal{S}_{\text{OA}}$  as well as an extraction trapdoor  $\tau_{\text{ext}}$ . The algorithm also initializes a public state  $St = (St_{\text{users}}, St_{\text{trans}}) := (\emptyset, \epsilon)$  as in Setup;*

**Extract**( $\mathcal{Y}, \tau_{\text{ext}}, M, \sigma, P, St$ ): *inputs a valid message-signature pair  $(M, \sigma)$  w.r.t.  $\mathcal{Y}$  and a predicate  $P \in \mathcal{P}$ , the extraction trapdoor  $\tau_{\text{ext}}$  as well as the public state  $St$ . It outputs an identity  $\text{id} \in \mathcal{ID}$  and a witness  $w \in \mathcal{W}$ .*

In addition, these algorithms must satisfy the following notions.

**Extractable correctness:** for any  $(St, \mathcal{Y}, \mathcal{S}_{GM}, \mathcal{S}_{OA}, \tau_{ext}) \leftarrow \text{SimSetup}(1^\lambda, 1^N)$ , for any outcome  $\langle id, \text{cert}_{id}, \text{sec}_{id} \rangle$  such that  $\text{cert}_{id} \Leftarrow_{\mathcal{Y}} \text{sec}_{id}$ , any message  $M$ , any witness  $w$ , and any predicate  $P \in \mathcal{P}$ : if  $\sigma \leftarrow \text{Sign}(id, \text{cert}_{id}, \text{sec}_{id}, M, w, P)$  and  $(id', w') \leftarrow \text{Extract}(\mathcal{Y}, \tau_{ext}, M, \sigma, P, St)$ , then  $(id, w) = (id', w')$  with overwhelming probability.

**Extractable soundness:** For any PPT adversary  $\mathcal{A}$  involved in the experiments defined in Fig. 2, the following advantage function must be negligible:

$$\mathbf{Adv}_{\mathcal{A}}^{\text{ext-s}}(\lambda) = \left| \Pr[\mathbf{Exp}_{\mathcal{A}}^{\text{real}}(\lambda) = 1] - \Pr[\mathbf{Exp}_{\mathcal{A}}^{\text{ext}}(\lambda) = 1] \right|.$$

```

1  $(St, \mathcal{Y}, \mathcal{S}_{GM}, \mathcal{S}_{OA}) \leftarrow \text{Setup}(1^\lambda, 1^N)$   $\boxed{(St, \mathcal{Y}, \mathcal{S}_{GM}, \mathcal{S}_{OA}, \tau_{ext}) \leftarrow \text{SimSetup}(1^\lambda, 1^N)}$ ;
2  $(M, \sigma, P, st) \leftarrow \mathcal{A}(St, \mathcal{Y}, \mathcal{S}_{GM}, \mathcal{S}_{OA})$ ;
3 if  $\text{Verify}(\mathcal{Y}, M, \sigma, P) = 0$  then return 0;
4  $\boxed{\text{id}^* \leftarrow \text{Open}(\mathcal{S}_{OA}, \mathcal{Y}, M, \sigma, P, St)}$   $\boxed{(\text{id}, w) \leftarrow \text{Extract}(\tau_{ext}, \mathcal{Y}, M, \sigma, P, St)}$ ;
5 if  $\boxed{P(\text{id}, M, w) = 1}$  then  $\boxed{\text{id}^* \leftarrow \perp}$  else  $\boxed{\text{id}^* \leftarrow \text{id}}$ ;
6 return  $\mathcal{A}(st, \text{id}^*)$ ;

```

**Fig. 2:** Experiment  $\mathbf{Exp}_{\mathcal{A}}^{\text{real}}(\lambda)$  (resp.  $\mathbf{Exp}_{\mathcal{A}}^{\text{ext}}(\lambda)$ ) excluding the dotted (resp. solid) boxes.

In the random oracle model, Item 2 of Fig. 2 is modified as follows:

```

2'  $\boxed{(M, \sigma, P, st) \leftarrow \mathcal{A}^{H_0}(St, \mathcal{Y}, \mathcal{S}_{GM}, \mathcal{S}_{OA})}$   $\boxed{(M, \sigma, P, st) \leftarrow \mathcal{A}^{H_1}(St, \mathcal{Y}, \mathcal{S}_{GM}, \mathcal{S}_{OA})}$ ;

```

Here,  $H_0$  and  $H_1$  are random oracles which privately evaluate and return the digests of given inputs. In the real setup, the BiAS algorithms have access to  $H_0$  whereas, in the extractable setup, they have access to  $H_1$ .

We stress that all secret keys but the extraction trapdoor are given to the distinguisher/adversary. This is necessary because we need the extractable properties even in presence of dishonest authorities. In the extractable setting, we require **Extract** to output a potential identifier  $\text{id} \in \mathcal{ID}$  and a witness  $w \in \mathcal{W}$  with overwhelming probability, even on adversarially-chosen verifying signatures and when both authorities are corrupted. This extractable mode makes it possible to compute the predicate in a meaningful way. Further, the extractable soundness property implies the hardness of computing a valid signature that traces to some user  $\text{id}$  for some predicate  $P$  although this predicate would have allowed user  $\text{id}$  to sign the message with statistical anonymity. While **Extract** is consistent with **Open**, we still do not have the complementary property of the hardness of computing a valid signature that cannot be traced although the tracing operation should have been possible. Indeed, if **Open** identifies a signature as non-traceable, we still have no clue about the meaning of the identity-witness pair produced by **Extract** on adversarially generated valid signatures that are not honestly generated (as otherwise, extractable-correctness implies the match with the actual pair)

*Security Against Misidentification Attacks (a.k.a. traceability).* In a misidentification attack, the adversary can corrupt the opening authority using the  $Q_{\text{keyOA}}$  oracle and introduce malicious users in  $U^a$  via  $Q_{\text{a-join}}$ -queries. It aims at producing a valid signature  $\sigma^*$  that does not open to any adversarially-controlled user.

**Definition 9.** A BiAS scheme is secure against misidentification attacks if it is branch-sound and, for any PPT adversary  $\mathcal{A}$  involved in experiment  $\mathbf{Exp}_{\mathcal{A},N}^{\text{mis-id}}$  (as defined in Fig.3), we have:  $\mathbf{Adv}_{\mathcal{A},N}^{\text{mis-id}}(\lambda) = \Pr[\mathbf{Exp}_{\mathcal{A},N}^{\text{mis-id}}(\lambda) = 1] \in \text{negl}(\lambda)$ .

```

1 stateT := (St, Y, SGM, SOA); (St, Y, SGM, SOA, τext) ← SimSetup(1λ, 1N) ;
2 (M, σ, P) ← A(Qpub, Qa-join, Qread, Qwrite, QkeyOA);
3 if Verify(Y, M, σ, P) = 0 then return 0;
4 (id, w) ← Extract(τext, Y, M, σ, P, St);
5 if id ∈ ID \ Ua then return 1;
6 return 0;

```

Fig. 3: Experiment  $\mathbf{Exp}_{\mathcal{A},N}^{\text{mis-id}}(\lambda)$ .

The winning condition is also checkable without the extractor if we rather define  $\text{id}^* \leftarrow \text{Open}(S_{OA}, Y, M, \sigma, P, St)$  in the experiment, as long as  $\text{id}^* \neq \perp$  in the winning condition of line 5. In that case, the analogue security with the real setup is implied by the extractable soundness property. Nevertheless, in the extractable mode the definition also captures the unforgeability of anonymous signatures, i.e. those which would have made  $\text{Open}$  to return  $\text{id}^* = \perp$  at line 5, if the extracted id does not correspond to a corrupt user when the group manager remains honest.

*Non-Frameability.* Framing attacks consider the case where the entire system is colluding against some honest user. The adversary can corrupt the group manager as well as the opening authority (via oracles  $Q_{\text{keyGM}}$  and  $Q_{\text{keyOA}}$ , respectively). It can also introduce honest group members (via  $Q_{\text{b-join}}$ -queries), observe the system while these users sign messages and create dummy users using  $Q_{\text{write}}$ . The adversary aims at framing an honest group member. Moreover, the adversary is also deemed successful if it is able to create a non-traceable valid signature which could have been created by an honest user but who never computed it: even a corrupted group manager is unable to compute a non-traceable signature using the identity of an honest user. For example, if the predicate of a BiAS only allows some users to compute perfectly anonymous signatures, it is infeasible to compute such signatures without corrupting at least one of these users. The definition follows the indistinguishable approach of security against misidentification attacks.

**Definition 10.** A BiAS scheme is secure against framing attacks if it satisfies branch-soundness and, for any PPT adversary  $\mathcal{A}$  involved in experiment  $\mathbf{Exp}_{\mathcal{A},N}^{\text{fra}}$  (as defined in Fig.4), we have:  $\mathbf{Adv}_{\mathcal{A},N}^{\text{fra}}(\lambda) = \Pr[\mathbf{Exp}_{\mathcal{A},N}^{\text{fra}}(\lambda) = 1] \in \text{negl}(\lambda)$ .

```

1 state $\mathcal{I}$  :=  $(St, \mathcal{Y}, \mathcal{S}_{GM}, \mathcal{S}_{OA})$ ;  $(St, \mathcal{Y}, \mathcal{S}_{GM}, \mathcal{S}_{OA}, \tau_{ext}) \leftarrow \text{SimSetup}(1^\lambda, 1^N)$  ;
2  $(M^*, \sigma^*, P^*) \leftarrow \mathcal{A}(Q_{pub}, Q_{keyGM}, Q_{keyOA}, Q_{b-join}, Q_{sig}, Q_{read}, Q_{write})$ ;
3 if  $\text{Verify}(\mathcal{Y}, M^*, \sigma^*, P^*) = 0$  then return 0;
4  $(id, w) \leftarrow \text{Extract}(\tau_{ext}, \mathcal{Y}, M^*, \sigma^*, P^*, St)$ ;
5 if  $(id \in U^b) \wedge (id, M^*, w, \sigma^*, P^*) \notin \text{Sigs}$  then return 1;
6 return 0;

```

**Fig. 4:** Experiment  $\text{Exp}_{\mathcal{A}, N}^{\text{fra}}(\lambda)$

Let  $\text{id}^* = \text{Open}(\mathcal{S}_{OA}, \mathcal{Y}, M^*, \sigma^*, P^*, St)$  in the framing experiment. Then, we can derive two winning conditions depending on whether  $\text{id}^* \in \mathcal{ID}$  or  $\text{id}^* = \perp$ . In the former case, the branch-soundness tells us that  $\text{id}^* = \text{id}$ . This traceable case is thus the analogue of the usual framing attack of KY in group signature transposed to our BiAS primitive. In the latter case, the signature  $\sigma^*$  of a successful adversary is deemed non-traceable, but it would have been created on behalf of an honest signer with identifier  $\text{id}$  who never produced it. This further justifies the need of all these security notions as we now have the complementary property discussed after Definition 8: a branch-sound BiAS scheme whose extracting algorithm returns independent identity-witness pairs given non-honest valid signatures cannot be secure against framing attacks.

Finally, we note that a signature does not only authenticate the message  $M$ , but it also binds the predicate value as well as the hidden  $(\text{id}, w)$  to  $M$ . The framing resistance also guarantees that the signature itself is not malleable as the winning condition is akin to the “strong”-unforgeability notion of standard signatures. This requirement is actually necessary since, in order to achieve anonymity in the “CCA sense”, we need to prevent signatures from being malleable.

*Discussion on the witness.* Our model does not assume any property of the witness. At first glance, it may seem strange to apparently let the users choose their witnesses arbitrarily at the signature generation time. This syntactic choice makes BiAS more flexible to be combined with other building blocks. For instance, the witness  $w$  may already be committed in an external commitment, i.e. outside our syntax, and bound by the application. Additional zero-knowledge proofs between  $w$ , the context, and the BiAS scheme are of course possible, which might prevent the user from choosing  $w$  freely.

In a money-laundering prevention application, a signer has no incentive in authenticating a transaction for a big amount of money  $w$  if he does not want to pay such an amount. Therefore, even if  $P(M, \text{id}, w)$  may vary when  $w$  varies at each transaction, the context prevents the user with identity  $\text{id}$  from fixing  $w$  in an arbitrary way. We thus leave it to the applications to define their own rules on the  $w$ ’s and the desire and the way to keep their level of secrecy.

## 4 Generic Construction

We provide a generic construction of BiAS for an arbitrary predicate family  $\mathcal{P} : \{0, 1\}^* \rightarrow \{0, 1\}$ . Our construction relies on the following building blocks:

- An  $\mathcal{R}_{\text{NEQ}}$ -lossy PKE scheme  $\Pi^{\text{RLE}} = (\text{Par-Gen}, \text{Keygen}, \text{Encrypt}, \text{Decrypt})$ ;
- An ordinary lossy PKE scheme  $\Pi^{\text{lpke}} = (\text{Keygen}, \text{LKeygen}, \text{Encrypt}, \text{Decrypt})$  where the message space has size at least  $N$  and forms an additive group;
- A digital signature scheme  $\Pi^{\text{sig}} = (\text{Kg}, \text{Sign}, \text{Verify})$  with signature space  $\mathcal{S}$  and public key space  $\mathcal{VK}$ ;
- A one-time signature  $\Pi^{\text{ots}} = (\text{Kg}, \text{Sign}, \text{Verify})$ ;
- A homomorphic equivocal commitment scheme  $\text{HEC} = (\text{Setup}, \text{Commit}, \text{Open}, \text{Eval}^{\text{in}}, \text{Eval}^{\text{out}}, \text{Verify})$ , where  $\text{Commit}$  samples its random coins from a distribution  $\mathcal{D}^{\text{HEC}}$  over a randomness space  $R^{\text{HEC}}$ ;
- A dual-mode statistical NIZK argument system  $\text{NIZK} = (\text{Setup}, \text{ExtSetup}, \text{Prove}, \text{Verify}, \text{Sim}, \text{Extract})$ , as defined in Supplementary Material A.

Since an  $\mathcal{R}_{\text{NEQ}}$ -lossy PKE scheme implies a standard lossy PKE scheme, the only ingredients we need are an  $\mathcal{R}_{\text{NEQ}}$ -lossy PKE system, a digital signature, a homomorphic equivocable commitment and a dual-mode NIZK argument.

For our purposes, it is sufficient to use an HEC scheme without the context-hiding property since we combine it with NIZK arguments where its partial openings serve as witnesses.<sup>7</sup> By using an HEC with the efficient verification property, we can make the signature length independent of the circuit size. Katsumata *et al.* [34] gave such a pairing-based HEC construction under a  $q$ -type assumption for NC<sup>1</sup> circuits. In the lattice setting, the fully homomorphic commitments of Gorbunov *et al.* [29] provide efficient verification (for bounded-depth circuits) under the Short Integer Solution [3] assumption, as recalled in Supplementary Material E.3. At the expense of a signature length depending on the circuit size, the construction can be simplified to use any statistically hiding commitment instead of an efficiently verifiable HEC. However, we aim at avoiding the circuit-size dependency.

Intuitively, the construction encrypts the group member’s identity  $\text{id}$  and the witness  $w$  using an HEC *and* simultaneously encrypts them into  $\text{ct}_{(\text{id}, w)}$  using the  $\mathcal{R}_{\text{NEQ}}$ -lossy PKE scheme, which realizes either a statistically hiding or an extractable commitment to  $(\text{id}, w)$ . Our proofs of anonymity require statistically-hiding commitments (as in the real scheme). In our proofs of security against mis-identification attacks and framing attacks, we will switch  $\text{ct}_{(\text{id}, w)}$  to its extractable mode because we need to be able to extract the underlying  $w$  and  $\text{id}$ .

In the signing algorithm, the group member next computes an evaluated HEC commitment  $\text{com}_{ev}$  of the predicate evaluation  $C_M(\text{id}, w)$  by homomorphically computing over  $\text{com}_{(\text{id}, w)}$  (note that  $\text{com}_{ev}$  need not be included in the signature since the verifier can recompute it from  $\text{com}_{(\text{id}, w)}$ ). Then, the signer computes a ciphertext  $\text{ct}_{\text{id}}$  that verifiably encrypts a product  $(1 - C_M(\text{id}, w)) \cdot \text{id}$  of his

---

<sup>7</sup> A context-hiding construction can still improve the efficiency by outputting partial openings in the clear in each signature.

identity  $\text{id}$  and the logical NOT of  $C_M(\text{id}, w)$ . When the predicate evaluates to  $C_M(\text{id}, w) = 0$ ,  $\text{ct}_{\text{id}}$  is distributed as a lossy encryption<sup>8</sup> of  $\text{id}$ . When  $C_M(\text{id}, w) = 1$ ,  $\text{ct}_{\text{id}}$  is completely independent of the signer's identity as it encrypts  $\mathbf{0}^{|\text{id}|}$ .

**Setup**( $1^\lambda, 1^N, \mathcal{P}_d$ ): Given a security parameter  $\lambda$ , a predicate family  $\mathcal{P}_d$  modeled by circuits of depth  $d = d(\lambda)$  and the maximal number of group members  $N = 2^\ell \in \text{poly}(\lambda)$ , do the following.

1. Generate a key pair  $(\text{pk}_{\text{sig}}, \text{sk}_{\text{sig}}) \leftarrow \Pi^{\text{sig}}.\text{Kg}(1^\lambda)$  for the signature scheme. We assume that each public key has bitlength  $\ell_{\text{sig}} \in \text{poly}(\lambda)$ .
2. Run  $(\text{pp}, \text{ek}, \text{msk}) \leftarrow \text{HEC}.\text{Setup}(1^\lambda)$  to generate parameters for the homomorphic equivocal commitment, together with an evaluation key  $\text{ek}$  and a master key  $\text{msk}$ .
3. Choose a one-time signature scheme  $\Pi^{\text{ots}} = (\text{Kg}, \text{Sign}, \text{Verify})$  with verification key space  $\{0, 1\}^L$ , for some  $L \in \text{poly}(\lambda)$ .
4. Choose public parameters  $\Gamma \leftarrow \Pi^{\text{RLE}}.\text{Par-Gen}(1^\lambda, 1^L, 1^B)$  for an  $\mathcal{R}_{\text{NEQ}}$ -lossy PKE scheme with tag space  $\mathcal{K} = \mathcal{T} = \{0, 1\}^L$  and message length  $B = \ell + \ell_w$ , where  $\ell_w \in \text{poly}(\lambda)$  is the bitlength of witnesses from the witness space  $\mathcal{W} = \{0, 1\}^{\ell_w}$ . Then, generate lossy keys  $(\text{pk}_{\text{RLE}}, \text{sk}_{\text{RLE}}) \leftarrow \Pi^{\text{RLE}}.\text{LKeygen}(\Gamma, \mathbf{0}^L)$  for the initialization value  $K = \mathbf{0}^L$ .
5. Generate an injective key pair  $(\text{pk}_e, \text{sk}_e) \leftarrow \Pi^{\text{lpke}}.\text{Keygen}(1^\lambda)$  for the standard lossy PKE scheme.
6. Generate a common reference string  $\rho$  from  $(\rho, \zeta) \leftarrow \text{NIZK}.\text{Setup}(1^\lambda)$  for a dual-mode NIZK argument in its statistical ZK mode.

The algorithm outputs  $(\mathcal{Y}, \mathcal{S}_{\text{GM}}, \mathcal{S}_{\text{OA}})$ , where the group public key is as

$$\mathcal{Y} := (\rho, \text{pk}_{\text{sig}}, (\text{pp}, \text{ek}), (\Gamma, \text{pk}_{\text{RLE}}), \text{pk}_e),$$

the opening authority's private key is  $\mathcal{S}_{\text{OA}} := \text{sk}_e$  and the private key of the group manager consists of  $\mathcal{S}_{\text{GM}} := \text{sk}_{\text{sig}}$ .  $\mathcal{Y}$  implicitly initializes  $St$ .

**Join**( $(\text{GM}, \mathcal{U}_{\text{id}})$ ): the group manager and the prospective user  $\mathcal{U}_{\text{id}}$  run the following interactive protocol  $[\text{J}_{\text{user}}(\lambda, \mathcal{Y}), \text{J}_{\text{GM}}(\lambda, St, \mathcal{Y}, \mathcal{S}_{\text{GM}})]$ :

1. User  $\mathcal{U}_{\text{id}}$  generates a key pair  $(\text{sk}_{\text{id}}, \text{pk}_{\text{id}}) \leftarrow \Pi^{\text{sig}}.\text{Kg}(1^\lambda)$  and sends the public key  $\text{pk}_{\text{id}}$  together with his identity  $\text{id} \in \{0, 1\}^\ell \setminus \{\mathbf{0}^\ell\}$  and an ordinary signature  $\text{sig}_{\text{id}} \leftarrow \Pi^{\text{sig}}.\text{Sign}(\text{usk}[\text{id}], (\text{id}, \text{pk}_{\text{id}}))$  to GM.
2.  $\text{J}_{\text{GM}}$  verifies that:  $\text{id} \neq \mathbf{0}^\ell$ ;  $\text{id}$  was not previously used by a registered user;  $\text{sig}_{\text{id}}$  is a valid signature on  $(\text{id}, \text{pk}_{\text{id}})$  w.r.t.  $\text{upk}[\text{id}]$ . It aborts if this is not the case. Otherwise, it computes  $\text{cert}_{\text{id}} \leftarrow \Pi^{\text{sig}}.\text{Sign}(\text{sk}_{\text{sig}}, (\text{id}, \text{pk}_{\text{id}}))$  as a signature on the message  $(\text{id}, \text{pk}_{\text{id}})$ . The membership certificate  $\text{cert}_{\text{id}}$  is sent to  $\mathcal{U}_{\text{id}}$ . Then,  $\text{J}_{\text{user}}$  verifies that  $\Pi^{\text{sig}}.\text{Verify}(\text{pk}_{\text{sig}}, (\text{id}, \text{pk}_{\text{id}}), \text{cert}_{\text{id}}) = 1$ . If this condition is not satisfied,  $\text{J}_{\text{user}}$  aborts. Otherwise,  $\text{J}_{\text{user}}$  defines the membership certificate as  $\text{cert}_{\text{id}}$ . The membership secret  $\text{sec}_{\text{id}}$  is defined to be  $\text{sec}_{\text{id}} = \text{sk}_{\text{id}}$ .  $\text{J}_{\text{GM}}$  stores  $\text{transcript}_{\text{id}} = (\text{id}, \text{pk}_{\text{id}}, \text{cert}_{\text{id}}, \text{upk}[\text{id}], \text{sig}_{\text{id}})$  in the database  $St_{\text{trans}}$  of joining transcripts.

<sup>8</sup> It is possible to compute  $\text{ct}_{\text{id}}$  using an ordinary (i.e., non-lossy) PKE scheme but it requires to rely on the simulation-soundness of NIZK in the proof of Lemma 4.

**Sign**( $\text{id}, \text{cert}_{\text{id}}, \text{sec}_{\text{id}}, M, w, P$ ): To sign a message  $M \in \{0, 1\}^{\ell_m}$  using the witness  $w = w[1] \dots w[\ell_w] \in \{0, 1\}^{\ell_w}$  w.r.t. the predicate  $P \in \mathcal{P}_d$ , let  $C_M : \{0, 1\}^{\ell_w} \times \{0, 1\}^\ell \rightarrow \{0, 1\}$  be the message-dependent Boolean circuit of depth  $\leq d$  that evaluates  $P(M, \text{id}, w)$  on input of  $(w[1], \dots, w[\ell_w], \text{id}[1], \dots, \text{id}[\ell])$ .

1. Generate a one-time signature key pair  $(\text{VK}, \text{SK}) \leftarrow \Pi^{\text{ots}}.\text{Kg}(1^\lambda)$ .
2. Choose  $r_{\text{id}, w} \leftarrow R^{\text{RLE}}$  in the randomness space of  $\Pi^{\text{RLE}}$  and encrypt the identity-witness pair  $(\text{id}, w) \in \{0, 1\}^{\ell + \ell_w}$  as an  $\mathcal{R}_{\text{NEQ}}$ -lossy encryption

$$\text{ct}_{(\text{id}, w)} = \Pi^{\text{RLE}}.\text{Encrypt}(\text{pk}_{\text{RLE}}, \text{VK}, (\text{id}, w); r_w) \quad (1)$$

under the tag  $\text{VK} \in \{0, 1\}^L$ .

3. Sample random coins  $r^{\text{hec}} \leftarrow \mathcal{D}^{\text{HEC}}$  and compute a commitment

$$\text{com}_{(\text{id}, w)} = \text{HEC}.\text{Commit}(\text{pp}, \text{ek}, (\text{id}, w); r^{\text{hec}}). \quad (2)$$

4. Using the homomorphic evaluation algorithm of  $\text{HEC}$ , compute

$$\begin{aligned} \pi_{C, M} &\leftarrow \text{HEC}.\text{Eval}^{\text{in}}(\text{ek}, C_M, (\text{id}, w), r^{\text{hec}}) \\ \text{com}_{ev} &= \text{HEC}.\text{Eval}^{\text{out}}(\text{ek}, C_M, \text{com}_{(\text{id}, w)}). \end{aligned}$$

5. Choose  $r^{\text{lpke}} \leftarrow R^{\text{lpke}}$  and compute

$$\text{ct}_{\text{id}} = \Pi^{\text{lpke}}.\text{Encrypt}(\text{pk}_e, (1 - c_{ev}) \cdot \text{id}; r^{\text{lpke}}), \quad (3)$$

where  $c_{ev} = C_M(w_1, \dots, w_{\ell_w}, \text{id}_1, \dots, \text{id}_\ell) \in \{0, 1\}$ .

6. Generate  $\sigma \leftarrow \Pi^{\text{sig}}.\text{Sign}(\text{sk}_{\text{id}}, (M, P, \text{ct}_{(\text{id}, w)}))$  as a signature on the message  $(M, P, \text{ct}_{(\text{id}, w)})$ .
7. Generate a NIZK argument  $\pi \leftarrow \text{NIZK}.\text{Prove}(\rho, \mathbf{x}, \mathbf{w})$  for the statement  $\mathbf{x}$  that there exists a witnesses  $\mathbf{w}$  comprised of  $(\text{id}, w) \in \{0, 1\}^{\ell + \ell_w}$ ,  $(\text{pk}_{\text{id}}, \text{cert}_{\text{id}}, \sigma) \in \mathcal{VK} \times \mathcal{S} \times \mathcal{S}$ ,  $r_{\text{id}, w} \in R^{\text{RLE}}$ ,  $r^{\text{hec}} \in R^{\text{HEC}}$ ,  $r^{\text{lpke}} \in R^{\text{lpke}}$ ,  $c_{ev} \in \{0, 1\}$  and  $\pi_{C, M}$ , which satisfy the relations (1)-(3) as well as

$$\begin{aligned} \Pi^{\text{sig}}.\text{Verify}(\text{pk}_{\text{sig}}, (\text{id}, \text{pk}_{\text{id}}), \text{cert}_{\text{id}}) &= 1 \\ \Pi^{\text{sig}}.\text{Verify}(\text{pk}_{\text{id}}, (M, P, \text{ct}_{(\text{id}, w)}), \sigma) &= 1 \\ \text{HEC}.\text{Verify}(\text{pp}, \text{com}_{ev}, c_{ev}, \pi_{C, M}) &= 1. \end{aligned} \quad (4)$$

8. Compute  $\text{sig} \leftarrow \Pi^{\text{ots}}.\text{Sign}(\text{SK}, (\text{ct}_{(\text{id}, w)}, \text{com}_{(\text{id}, w)}, \text{ct}_{\text{id}}, \pi))$ .

Return the signature

$$\Sigma = (\text{VK}, (\text{ct}_{(\text{id}, w)}, \text{com}_{(\text{id}, w)}, \text{ct}_{\text{id}}, \pi), \text{sig}) \quad (5)$$

**Verify**( $\mathcal{Y}, M, \Sigma, P$ ): Parse  $\Sigma$  as above. Return 1 if and only if: (i)  $\text{sig}$  is a valid one-time signature on  $(\text{ct}_{(\text{id}, w)}, \text{com}_{(\text{id}, w)}, C_{\text{id}}, \pi)$  for the verification key  $\text{VK}$ ; (ii) The NIZK argument  $\pi$  properly verifies for the commitment  $\text{com}_{ev}$  publicly obtained as  $\text{com}_{ev} = \text{HEC}.\text{Eval}^{\text{out}}(\text{ek}, C_M, \text{com}_{(\text{id}, w)})$ .

**Open**( $\mathcal{Y}, \mathcal{S}_{\text{OA}}, M, \Sigma, P, St$ ): Given the opener’s secret key  $\mathcal{S}_{\text{OA}} := \text{sk}_e$ , parse the signature  $\Sigma$  as in (5). Compute  $t_{\text{id}} = \Pi^{\text{lpke}}.\text{Decrypt}(\text{sk}_e, \text{ct}_{\text{id}})$ . If  $t_{\text{id}} = \mathbf{0}^\ell$ , return  $\perp$ . Otherwise, check if the string  $t_{\text{id}} \in \{0, 1\}^\ell$  appears in a record  $(t_{\text{id}}, \text{transcript}_{\text{id}} = (t_{\text{id}}, \text{pk}_{\text{id}}, \text{cert}_{\text{id}}, \text{upk}[\text{id}], \text{sig}_{\text{id}}))$  of  $St_{\text{trans}}$ . If it does, output  $\text{id} = t_{\text{id}} \in \{0, 1\}^\ell$  (and, optionally,  $\text{upk}[\text{id}]$ ). Otherwise, output  $\perp$ .

In Supplementary Material E and F, we provide details on instantiations from lattices and bilinear maps. The lattice-based construction is only a feasibility result based on generic NIZK for NP statements [50]. In the case of NC<sup>1</sup> circuits, the scheme can be instantiated with Groth-Sahai proofs [32] to provide much shorter signatures than using the Groth-Ostrovsky-Sahai techniques [31].

#### 4.1 Branch-Soundness and Security

To prove security under our definitions, we first consider the following **SimSetup** and **Extract** algorithms associated to our BiAS construction.

**SimSetup**( $1^\lambda, 1^N, \mathcal{P}_d$ ): This algorithm is exactly as **Setup**( $1^\lambda, 1^N, \mathcal{P}_d$ ) except that steps 4 and 6 are modified in the following way:

4. Choose public parameters  $\Gamma \leftarrow \Pi^{\text{RLE}}.\text{Par-Gen}(1^\lambda, 1^L, 1^B)$  for an  $\mathcal{R}_{\text{NEQ}}$ -lossy PKE scheme with tag space  $\mathcal{K} = \mathcal{T} = \{0, 1\}^L$  and message length  $B = \ell + \ell_w$ , where  $\ell_w \in \text{poly}(\lambda)$  is the bitlength of witnesses from the witness space  $\mathcal{W} = \{0, 1\}^{\ell_w}$ . Then, generate **injective** keys  $(\text{pk}_{\text{RLE}}, \text{sk}_{\text{RLE}}) \leftarrow \Pi^{\text{RLE}}.\text{Keygen}(\Gamma, \mathbf{0}^L)$  for the initialization value  $K = \mathbf{0}^L$ .
6. Generate a common reference string  $\rho$  from  $(\rho, \xi) \leftarrow \text{NIZK}.\text{ExtSetup}(1^\lambda)$  for an extractable (and thus statistically sound) NIZK proof system.

The algorithm returns the same output as **Setup**, together with an extraction trapdoor  $\tau_{\text{ext}} = (\text{sk}_{\text{RLE}}, \xi)$ , where  $\xi$  is the extraction trapdoor of **NIZK**.

**Extract**( $\mathcal{Y}, \tau_{\text{ext}}, M, \Sigma, P, St$ ): Write  $\Sigma$  as  $(\text{VK}, (\text{ct}_{(\text{id}, w)}, \text{com}_{(\text{id}, w)}, \text{ct}_{\text{id}}, \boldsymbol{\pi}), \text{sig})$  and return  $\perp$  if its components do not parse properly. Otherwise, use  $\text{sk}_{\text{RLE}}$  to decrypt the  $\mathcal{R}_{\text{NEQ}}$ -lossy PKE ciphertexts  $\text{ct}_{(\text{id}, w)}$  (recall that the **NEQ** relations makes all tags injective on a public key produced by **Keygen** for the initialization value  $K = \mathbf{0}^\ell$ ). If any decryption fails, return  $\perp$ . Otherwise, output  $w \in \{0, 1\}^{\ell_w}$  and  $\text{id} \in \{0, 1\}^\ell$ .

The security properties of the NIZK argument system ensure that the common reference strings  $\rho$  produced by **NIZK**.**Setup** and **NIZK**.**ExtSetup** are computationally indistinguishable. Moreover, in the  $\mathcal{R}_{\text{NEQ}}$ -lossy PKE scheme, the public keys produced by **LKeygen** and **Keygen** are computationally indistinguishable as well.

Next, we will show that this extractable BiAS satisfies the extractable soundness notion unless the adversary can break the (statistical) soundness of the proof  $\boldsymbol{\pi}$  included in a valid signature  $\Sigma$ .

**Theorem 1.** *The scheme satisfies branch-soundness if: (i)  $\Pi^{\text{RLE}}$  is a secure  $\mathcal{R}_{\text{NEQ}}$ -lossy PKE scheme; (ii) NIZK is a dual-mode NIZK argument system (i.e., its statistically sound and statistically ZK modes are computationally indistinguishable); (iii) HEC is computationally binding for evaluated commitments.*

*Proof.* To prove the result, we consider a sequence of games. In each game, we call  $W_i$  the event that the challenger outputs 1.

**Game 0:** This is the real experiment  $\mathbf{Exp}_{\mathcal{A}}^{\text{real}}(\lambda)$ , where the adversary  $\mathcal{A}$  is given  $(\mathcal{Y}, \mathcal{S}_{\text{GM}}, \mathcal{S}_{\text{OA}})$ , where  $(St, \mathcal{Y}, \mathcal{S}_{\text{GM}}, \mathcal{S}_{\text{OA}}) \leftarrow \text{Setup}(1^\lambda, 1^N)$ . The adversary outputs a tuple  $(M, \Sigma, P, st)$ , where  $\Sigma = (\mathbf{VK}, (\mathbf{ct}_{(\text{id}, w)}, \mathbf{com}_{(\text{id}, w)}, \mathbf{ct}_{\text{id}}, \boldsymbol{\pi}), sig)$ . If  $\Sigma$  does not verify, the challenger outputs 0. Otherwise, it runs  $\text{Open}$  to obtain  $\text{id}^* \in \{0, 1\}^\ell$  and feeds  $\mathcal{A}$  with  $\text{id}^*$ . Then, the challenger outputs whatever  $\mathcal{A}$  outputs. By definition,  $\Pr[W_0] = \Pr[\mathbf{Exp}_{\mathcal{A}}^{\text{real}}(\lambda)] = 1$ .

**Game 1:** This game is identical to Game 0 except that, at step 4 of the  $\text{Setup}$  algorithm, the challenger computes  $(\mathbf{pk}_{\text{RLE}}, \mathbf{sk}_{\text{RLE}}) \leftarrow \Pi^{\text{RLE}}.\text{Keygen}(\Gamma, \mathbf{0}^L)$  instead of  $(\mathbf{pk}_{\text{RLE}}, \mathbf{sk}_{\text{RLE}}) \leftarrow \Pi^{\text{RLE}}.\text{LKeygen}(\Gamma, \mathbf{0}^L)$ . By the first indistinguishability property of  $\Pi^{\text{RLE}}$ , we have  $|\Pr[W_1] - \Pr[W_0]| \in \text{negl}(\lambda)$ .

**Game 2:** This game is like Game 1 except that, at step 6 of the  $\text{Setup}$  algorithm, the challenger generates  $(\rho, \xi) \leftarrow \text{NIZK}.\text{ExtSetup}(1^\lambda)$  instead of  $(\rho, \zeta) \leftarrow \text{NIZK}.\text{Setup}(1^\lambda)$  and keeps the extraction trapdoor  $\tau_{\text{ext}} = (\mathbf{sk}_{\text{RLE}}, \xi)$  to itself. By the dual-mode property of  $\text{NIZK}$ , the CRSes produced by  $\text{NIZK}.\text{Setup}$  and  $\text{NIZK}.\text{ExtSetup}$  have computationally indistinguishable distributions, thus ensuring that  $|\Pr[W_2] - \Pr[W_1]| \in \text{negl}(\lambda)$  for any PPT adversary  $\mathcal{A}$ .

**Game 3:** In this game, the challenger makes use of the trapdoor  $\tau_{\text{ext}} = (\mathbf{sk}_{\text{RLE}}, \xi)$ . When  $\mathcal{A}$  outputs a tuple  $(M, \Sigma, P, st)$ , the challenger parses the signature  $\Sigma$  as  $(\mathbf{VK}, (\mathbf{ct}_{(\text{id}, w)}, \mathbf{com}_{(\text{id}, w)}, \mathbf{ct}_{\text{id}}, \boldsymbol{\pi}), sig)$  and uses  $\mathbf{sk}_{\text{RLE}}$  to extract  $(\text{id}^\dagger, w^\dagger)$ . From the NIZK proof  $\boldsymbol{\pi}$ , it uses  $\xi$  to extract the witnesses  $(\text{id}, w) \in \{0, 1\}^{\ell_w + \ell}$ ,  $(\mathbf{pk}_{\text{id}}, \mathbf{cert}_{\text{id}}, \sigma) \in \mathcal{VK} \times \mathcal{S} \times \mathcal{S}$ ,  $r_{\text{id}, w} \in R^{\text{RLE}}$ ,  $r^{\text{hec}} \in R^{\text{HEC}}$ ,  $r^{\text{lpke}} \in R^{\text{lpke}}$ ,  $c_{ev} \in \{0, 1\}$  and  $\pi_{C, M}$ . Then, the challenger halts and outputs a random bit if  $c_{ev} \neq C_M(w_1, \dots, w_{\ell_w}, \text{id}[1], \dots, \text{id}[\ell])$ .

We claim that  $|\Pr[W_3] - \Pr[W_2]| \in \text{negl}(\lambda)$  as the two games only differ when  $\mathcal{A}$  breaks the computational binding property of  $\text{HEC}$  for evaluated commitments. Indeed, by the statistical soundness of  $\text{NIZK}$  on a CRS  $\rho$  produced by  $\text{NIZK}.\text{ExtSetup}$ , we have  $(\text{id}, w) = (\text{id}^\dagger, w^\dagger)$  and extracted witnesses satisfy the relations (1)-(3). In particular, we have  $\mathbf{com}_{(\text{id}, w)} = \text{HEC}.\text{Commit}(\mathbf{pp}, \mathbf{ek}, (\text{id}, w); r^{\text{hec}})$  and the extracted  $c_{ev} \in \{0, 1\}$ ,  $\pi_{C, M}$  satisfy  $\text{HEC}.\text{Verify}(\mathbf{pp}, \mathbf{com}_{ev}, c_{ev}, \pi_{C, M}) = 1$ , where  $\mathbf{com}_{ev} = \text{HEC}.\text{Eval}^{\text{out}}(\mathbf{ek}, C_M, \mathbf{com}_{(\text{id}, w)})$ . It is easy to see that Game<sub>3</sub> only differs from Game<sub>2</sub> when the extracted  $\pi_{C, M}$  differs from

$$\bar{\pi}_{C, M} \leftarrow \text{HEC}.\text{Eval}^{\text{in}}(\mathbf{ek}, C_M, (\text{id}, w), r^{\text{hec}}),$$

which is the value that would satisfy  $\text{HEC}.\text{Verify}(\mathbf{pp}, \mathbf{com}_{ev}, C_M(w, \text{id}), \bar{\pi}_{C, M}) = 1$ . Hence, if  $|\Pr[W_3] - \Pr[W_2]|$  is noticeable, the challenger can break the binding property of  $\text{HEC}$  by outputting  $((\text{id}, w), r^{\text{hec}}, C_M, c_{ev}, \pi_{C, M})$ .

**Game 4:** This game is identical to Game 3 with the difference that, after having extracted  $(\text{id}, w)$ , the challenger computes  $C_M(w, \text{id}) \in \{0, 1\}$ , which is identical to the extracted  $c_{ev} \in \{0, 1\}$  unless the failure event of Game 3 occurs. If  $C_M(w, \text{id}) = 0$ , it overwrites  $\text{id}^* \leftarrow \text{Open}(\mathcal{S}_{\text{OA}}, \mathcal{Y}, M, \sigma, P, St)$

with  $\text{id}^* = \text{id}$ , which was extracted from  $\text{ct}_{(\text{id}, w)}$ . If  $C_M(w, \text{id}) = 1$ , it sets  $\text{id}^* = \perp$ . In both cases, it feeds  $\mathcal{A}$  with  $\text{id}^*$  and returns whatever  $\mathcal{A}$  outputs in reaction. This change does not modify the output distribution of  $\mathcal{A}$  because, as long as  $c_{ev} = C_M(w, \text{id})$ , the statistical soundness of  $\pi$  ensures that  $\text{ct}_{\text{id}} = \Pi^{\text{lpke}}.\text{Encrypt}(\text{pk}_e, (1 - C_M(w, \text{id})) \cdot \text{id}; r^{\text{lpke}})$ , where  $r^{\text{lpke}}$  and  $\text{id}$  are extracted from  $\pi$ . Hence, unless  $\mathcal{A}$  breaks the statistical soundness of  $\pi$ , Game 4 eventually returns  $\text{id} = \perp$  or  $\text{id} = \text{id}^*$  to  $\mathcal{A}$  whenever Game 3 does.

**Game 5:** This game is like Game 4 but we remove the restriction introduced in Game 3. Namely, the challenger does no longer replace  $\mathcal{A}$ 's output by a random bit when the witnesses  $c_{ev} \in \{0, 1\}$ ,  $(\text{id}, w) \in \{0, 1\}^{\ell_w + \ell}$  extracted from  $\pi$  are such that  $c_{ev} \neq C_M(w_1, \dots, w_{\ell_w}, \text{id}[1], \dots, \text{id}[\ell])$ . The same arguments as those in the transition between the first two games show that  $|\Pr[W_5] - \Pr[W_4]| \in \text{negl}(\lambda)$  so long as  $\text{HEC}$  is computationally binding.

We conclude the proof by noting that Game 5 is identical to  $\text{Exp}_{\mathcal{A}}^{\text{ext}}(\lambda)$ , so that we have  $|\Pr[\text{Exp}_{\mathcal{A}}^{\text{real}}(\lambda) = 1] - \Pr[\text{Exp}_{\mathcal{A}}^{\text{ext}}(\lambda) = 1]| = |\Pr[W_0] - \Pr[W_5]|$ .  $\square$

*Security Against Mis-Identification and Framing Attacks.*

**Lemma 2.** *The scheme is secure against misidentification attacks if: (i)  $\Pi^{\text{sig}}$  is existentially unforgeable under chosen-message attacks; (ii) The NIZK argument is computationally sound.* (The proof is in Supplementary Material D.1.)

**Lemma 3.** *The scheme is secure against framing attacks provided: (i)  $\Pi^{\text{sig}}$  is strongly unforgeable under chosen-message attacks; (ii) The NIZK argument is computationally sound.* (The proof is given in Supplementary Material D.2.)

## 4.2 Branch-Hiding and Privacy

The branch-hiding property follows from the full anonymity of our scheme.

**Theorem 2.** *The scheme provides full anonymity if:  $\Pi^{\text{RLE}}$  and  $\Pi^{\text{lpke}}$  are secure  $\mathcal{R}_{\text{NEQ}}$ -lossy PKE and standard lossy PKE schemes, respectively; (ii) NIZK is a computationally sound NIZK argument; (iii)  $\Pi^{\text{ots}}$  is strongly unforgeable.*

To prove Theorem 2, we separately consider the tracing and non-tracing modes. Lemma 4 first considers the former case where the adversary does not corrupt the opening authority. Lemma 5 shows that even an unbounded adversary is unable to distinguish group members' signatures in non-tracing mode.

**Lemma 4.** *The scheme provides anonymity in tracing mode assuming that: (i)  $\Pi^{\text{RLE}}$  is a secure  $\mathcal{R}_{\text{NEQ}}$ -lossy PKE scheme; (ii)  $\Pi^{\text{lpke}}$  is a standard lossy PKE scheme; (ii) The NIZK argument system provides soundness; (iii)  $\Pi^{\text{ots}}$  is strongly unforgeable.* (The proof is given in Supplementary Material D.3.)

**Lemma 5.** *The scheme provides statistical anonymity in non-tracing mode.*

*Proof.* Recall that experiment  $\mathbf{Exp}_{\mathcal{A}}^{\text{anon-ntr-}d}(\lambda)$  allows the adversary to obtain a challenge for the non-tracing mode. Namely, it is allowed to corrupt the opening authority and obtain  $\mathcal{S}_{\text{OA}}$  as long as, in the challenge phase, it chooses a pair  $(M^*, P^*)$  and two tuples  $(\text{id}_0^*, w_0^*, \text{sec}_0^*, \text{cert}_0^*)$  and  $(\text{id}_1^*, w_1^*, \text{sec}_1^*, \text{cert}_1^*)$ , such that  $P^*(M^*, \text{id}_0^*, w_0^*) = P^*(M^*, \text{id}_1^*, w_1^*) = 1$ . In this scenario, we will prove that, even after having obtained  $\mathcal{S}_{\text{OA}}$ , an unbounded adversary  $\mathcal{A}$  remains unable to infer anything about the bit  $d \in \{0, 1\}$  used by the challenger to compute the signature  $\Sigma^* = (\text{VK}^*, (\text{ct}_{(\text{id}, w)}^*, \text{com}_{(\text{id}, w)}^*, \text{ct}_{\text{id}}^*, \pi^*), \text{sig}^*)$  using  $(\text{id}_d^*, w_d^*, \text{sec}_d^*, \text{cert}_d^*)$ .

To this end, we consider two statistically indistinguishable games. The first one is the real game whereas the second one appeals to the statistical honest-verifier zero-knowledge simulator of the argument system.

**Game<sup>(d)</sup> 0:** This is the real game, which is as in the proof of Lemma 4.

**Game<sup>(d)</sup> 1:** This game is like **Game<sup>(d)</sup> 0** except that, in the challenge signature  $\Sigma^*$ , we use the simulation trapdoor  $\zeta$  generated from  $\text{NIZK}.\text{Setup}$  and the statistical NIZK simulator  $\text{NIZK}.\text{Sim}$  to generate  $\pi^*$ . Owing to the statistical ZK property of NIZK, the simulated  $\pi^*$  is statistically close to a real  $\pi^*$  that would be generated using the witnesses. Moreover, it is statistically independent of the witnesses used to compute  $\text{ct}_{(\text{id}, w)}^*$ ,  $\text{com}_{(\text{id}, w)}^*$  and  $\text{ct}_{\text{id}}^*$ .

In **Game<sup>(d)</sup> 1**, we note that, when  $C_{M^*}(w_1^*, \dots, w_{\ell_w}^*, \text{id}_1^*, \dots, \text{id}_{\ell}^*) = 1$ , the ciphertext  $\text{ct}_{\text{id}}^*$  is of the form  $\text{ct}_{\text{id}}^* = \Pi^{\text{lpke}}.\text{Encrypt}(\text{pk}_e, \mathbf{0}^{\ell}; r^{\text{lpke}})$ , where  $r^{\text{lpke}} \leftarrow R^{\text{lpke}}$ , so that  $\text{ct}_{\text{id}}^*$  is independent of  $d \in \{0, 1\}$  although  $\text{pk}_e$  is an injective public key. Moreover,  $\text{ct}_{(\text{id}, w)}^*$ ,  $\text{com}_{(\text{id}, w)}^*$  statistically hide the underlying pair  $(\text{id}, w)$  since, by definition, the homomorphic equivocable commitment  $\text{com}_{(\text{id}, w)}^*$  is statistically hiding and the  $\mathcal{R}_{\text{NEQ}}$ -lossy encryption  $\text{ct}_{(\text{id}, w)}^*$  is computed under a lossy key produced by  $\text{LKeygen}$ .  $\square$

## Acknowledgements

Part of this research was funded by the French ANR ALAMBIC project (ANR-16-CE39-0006). This work was also supported in part by the European Union PROMETHEUS project (Horizon 2020 Research and Innovation Program, grant 780701). Khoa Nguyen was supported in part by the Gopalakrishnan - NTU PPF 2018, by A\*STAR, Singapore under research grant SERC A19E3b0099, and by Vietnam National University HoChiMinh City (VNU-HCM) under grant number NCM2019-18-01. Thomas Peters is a senior research associate of the Belgian Fund for Scientific Research (F.R.S.-FNRS).

## References

1. M. Abe, G. Fuchsbauer, J. Groth, K. Haralambiev, and M. Ohkubo. Structure-preserving signatures and commitments to group elements. In *Crypto*, 2010.
2. S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (H)IBE in the standard model. In *Eurocrypt*, 2010.
3. M. Ajtai. Generating hard instances of lattice problems. In *STOC*, 1996.

4. W. Banaszczyk. New bounds in some transference theorems in the geometry of number. *Math. Ann.*, 296:625–635, 1993.
5. E. Bangerter, J. Camenisch, and A. Lysyanskaya. A cryptographic framework for the controlled release of certified data. In *Security Protocols*, 2004.
6. M. Bellare and G. Fuchsbauer. Policy-based signatures. In *PKC*, 2014.
7. M. Bellare, D. Micciancio, and B. Warinschi. Foundations of group signatures: Formal definitions, simplified requirements, and a construction based on general assumptions. In *Eurocrypt*, 2003.
8. M. Bellare, H. Shi, and C. Zhang. Foundations of group signatures: The case of dynamic groups. In *CT-RSA*, 2005.
9. A. Bender, J. Katz, and R. Morselli. Ring Signatures: Stronger Definitions, and Constructions without Random Oracles. *J. Cryptology*, 22(1):114–138, 2009.
10. M. Blum, M. Feldman, and S. Micali. Non-interactive zero-knowledge and its applications. In *STOC*, 1988.
11. D. Boneh, X. Boyen, and H. Shacham. Short group signatures. In *Crypto*, 2004.
12. D. Boneh, C. Gentry, S. Gorbunov, S. Halevi, V. Nikolaenko, G. Segev, V. Vaikuntanathan, and D. Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *Eurocrypt*, 2014.
13. J. Bootle, A. Cerulli, P. Chaidos, E. Ghadafi, J. Groth, and C. Petit. Short accountable ring signatures based on DDH. In *ESORICS*, 2015.
14. J. Bootle, V. Lyubashevsky, and G. Seiler. Algebraic techniques for short(er) exact lattice-based zero-knowledge proofs. In *Crypto*, 2019.
15. X. Boyen. Mesh signatures. In *Eurocrypt*, 2007.
16. X. Boyen and C. Delerabée. Expressive subgroup signatures. In *SCN*, 2008.
17. E. Boyle, S. Goldwasser, and I. Ivan. Functional signatures and pseudorandom functions. In *PKC*, 2014.
18. E. Boyle, G. Segev, and D. Wichs. Fully leakage-resilient signatures. In *Eurocrypt*, 2011.
19. Z. Brakerski and V. Vaikuntanathan. Lattice-based FHE as secure as PKE. In *ITCS*, 2014.
20. J. Camenisch, S. Hohenberger, and A. Lysyanskaya. Compact e-cash. In *Eurocrypt*, 2005.
21. J. Camenisch, S. Hohenberger, and A. Lysyanskaya. Balancing accountability and privacy using e-cash. In *SCN*, 2006.
22. R. Canetti, Y. Chen, J. Holmgren, A. Lombardi, G. Rothblum, R. Rothblum, and D. Wichs. Fiat-Shamir: From practice to theory. In *STOC*, 2019.
23. D. Chaum and E. Van Heyst. Group signatures. In *Eurocrypt*, 1991.
24. Y. Dodis, A. Kiayias, A. Nicolosi, and V. Shoup. Anonymous identification in ad-hoc groups. In *Eurocrypt*, 2004.
25. E. Fujisaki and K. Suzuki. Traceable ring signature. In *PKC*, 2007.
26. L. Garms and A. Lehmann. Group signatures with selective linkability. In *PKC*, 2019.
27. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, 2008.
28. C. Gentry, A. Sahai, and B. Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *Crypto*, 2013.
29. S. Gorbunov, V. Vaikuntanathan, and D. Wichs. Leveled fully homomorphic signatures from standard lattices. In *STOC*, 2015.
30. J. Groth. Fully anonymous group signatures without random oracles. In *Asiacrypt*, 2007.

31. J. Groth, R. Ostrovsky, and A. Sahai. Perfect non-interactive zero-knowledge for NP. In *Eurocrypt*, 2006.
32. J. Groth and A. Sahai. Efficient non-interactive proof systems for bilinear groups. In *Eurocrypt*, 2008.
33. D. Hofheinz and B. Ursu. Dual-mode NIZKs from obfuscation. In *Asiacrypt*, 2019.
34. S. Katsumata, R. Nishimaki, S. Yamada, and T. Yamakawa. Exploring constructions of compact nizks from various assumptions. In *Crypto*, 2019.
35. A. Kiayias, Y. Tsiounis, and M. Yung. Traceable signatures. In *Eurocrypt*, 2004.
36. A. Kiayias and M. Yung. Secure scalable group signature with dynamic joins and separable authorities. *Int. Journal of Security and Networks*, 1(1):24–45, 2006.
37. J. Kilian and E. Petrank. Identity escrow. In *Crypto*, 1998.
38. E. Kiltz. Chosen-ciphertext security from tag-based encryption. In *TCC*, 2006.
39. M. Kohlweiss and I. Miers. Accountable tracing signatures. *IACR Cryptology ePrint Archive*, 2014:824, 2014.
40. M. Kohlweiss and I. Miers. Accountable metadata-hiding escrow: A group signature case study. In *PoPETs*, 2015.
41. B. Libert, K. Nguyen, A. Passelègue, and R. Titiu. Simulation-sound arguments for LWE and applications to KDM-CCA2 security. In *Asiacrypt*, 2020.
42. B. Libert and M. Yung. Non-interactive CCA-secure threshold cryptosystems with adaptive security: New framework and constructions. In *TCC*, 2012.
43. J. Liu, V. Wei, and D. Wong. Linkable spontaneous anonymous group signature for ad hoc groups (extended abstract). In *ACISP*, 2004.
44. V. Lyubashevsky. Fiat-Shamir with Aborts: Applications to Lattice and Factoring-Based Signatures. In *Asiacrypt*, 2009.
45. H. Maji, M. Prabhakaran, and M. Rosulek. Attribute-based signatures. In *CT-RSA*, 2011.
46. D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Eurocrypt*, 2012.
47. M. Naor. Deniable ring authentication. In *Crypto*, 2002.
48. M. Naor and M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC*, 1990.
49. S. Noether and A. Mackenzie. Ring confidential transactions. *Ledger*, 1:1–18, 2016.
50. C. Peikert and S. Shiehian. Non-interactive zero knowledge for NP from (plain) Learning With Errors. In *Crypto*, 2019.
51. C. Peikert and B. Waters. Lossy trapdoor functions and their applications. In *STOC*, 2008.
52. C. Rackoff and D. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In *Crypto*, 1991.
53. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, 2005.
54. R. Rivest, A. Shamir, and Y. Tauman. How to leak a secret. In *Asiacrypt*, 2001.
55. Y. Sakai, K. Emura, G. Hanaoka, Y. Kawai, T. Matsuda, and K. Omote. Group signatures with message-dependent opening. In *Pairing*, 2012.
56. I. Teranishi, J. Furukawa, and K. Sako. k-times anonymous authentication (extended abstract). In *Asiacrypt*, 2004.
57. S. Xu and M. Yung. Accountable ring signatures: A smart card approach. In *CARDIS*, 2004.
58. R. Yang, M.-H. Au, Z. Zhang, Q. Xu, Z. Yu, and W. Whyte. Efficient lattice-based zero-knowledge arguments with standard soundness: Construction and applications. In *Crypto*, 2019.

# Supplementary Material

## A Definition of Dual-Mode NIZK

A dual-mode NIZK argument [31,50,33] is a NIZK system that operates either in hiding or binding mode. In the former, it satisfies statistical zero-knowledge and computational soundness. In the latter, it satisfies computational zero-knowledge, statistical soundness and, optionally, statistical extractability when the common reference string contains an extraction trapdoor. The common reference strings generated by the two modes are computationally indistinguishable.

In Definition 11, we use the term “argument” in the statistical ZK setting and “proof” in the statistical soundness setting.

**Definition 11.** Let  $\mathcal{L} \subseteq \{0,1\}^n$ , where  $n = n(\lambda)$ , be an NP language associated with an NP relation  $\mathcal{R}$ . A dual-mode statistical NIZK system for  $\mathcal{L}$  is a tuple of efficient algorithms  $\text{NIZK} = (\text{Setup}, \text{ExtSetup}, \text{Prove}, \text{Verify}, \text{Sim}, \text{Extract})$  with the following properties.

- $\text{Setup}(1^\lambda) \rightarrow (\rho, \zeta)$ : On input security parameter  $\lambda$ , outputs a common reference string  $\text{crs}$  and a simulation trapdoor  $\zeta$ .
- $\text{ExtSetup}(1^\lambda) \rightarrow (\rho, \xi)$ : On input security parameter  $\lambda$ , outputs a common reference string  $\text{crs}$  and an extraction trapdoor  $\xi$ .
- $\text{Prove}(\rho, \mathbf{x}, \mathbf{w}) \rightarrow \boldsymbol{\pi}$ : On input  $\rho$ , a statement  $\mathbf{x}$  and a witness  $\mathbf{w}$ , outputs a proof/argument  $\boldsymbol{\pi}$ .
- $\text{Verify}(\rho, \mathbf{x}, \boldsymbol{\pi}) \rightarrow \{0,1\}$ : On input  $\rho$ , a statement  $\mathbf{x}$  and a purported proof or argument  $\boldsymbol{\pi}$ , outputs either 1 or 0.
- $\text{Sim}(\rho, \zeta, \mathbf{x}) \rightarrow \{\boldsymbol{\pi}^*, \perp\}$ : On input  $\rho$ , a simulation trapdoor  $\zeta$  and statement  $\mathbf{x}$ , outputs a simulated argument  $\boldsymbol{\pi}^*$  or symbol  $\perp$  indicating failure.
- $\text{Extract}(\rho, \xi, \mathbf{x}, \boldsymbol{\pi})$ : On input  $\rho$ , an extraction trapdoor  $\xi$ , a statement  $\mathbf{x}$  and a proof  $\boldsymbol{\pi}$ , outputs a witness  $\mathbf{w}$ .

**CRS indistinguishability.** The distribution of the common reference strings generated via  $\text{Setup}(1^\lambda)$  and  $\text{ExtSetup}(1^\lambda)$  are computationally indistinguishable.

**Perfect completeness in both modes.** For any  $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$ , we have

$$\Pr \left[ \text{Verify}(\rho, \mathbf{x}, \boldsymbol{\pi}) = 1 \mid (\rho, \zeta) \leftarrow \text{Setup}(1^\lambda), \boldsymbol{\pi} \leftarrow \text{Prove}(\rho, \mathbf{x}, \mathbf{w}) \right] = 1,$$

$$\Pr \left[ \text{Verify}(\rho, \mathbf{x}, \boldsymbol{\pi}) = 1 \mid (\rho, \xi) \leftarrow \text{ExtSetup}(1^\lambda), \boldsymbol{\pi} \leftarrow \text{Prove}(\rho, \mathbf{x}, \mathbf{w}) \right] = 1,$$

**Statistical soundness in binding mode.** For every adversary  $\mathcal{A}$ ,

$$\Pr \left[ \text{Verify}(\rho, \mathbf{x}, \boldsymbol{\pi}) = 1 \wedge \mathbf{x} \notin \mathcal{L} \mid \begin{array}{l} (\rho, \xi) \leftarrow \text{ExtSetup}(1^\lambda), \\ (\mathbf{x}, \boldsymbol{\pi}) \leftarrow \mathcal{A}(\rho) \end{array} \right] = \text{negl}(\lambda).$$

**Statistical extractability in binding mode.** For every adversary  $\mathcal{A}$ ,

$$\Pr \left[ \begin{array}{l} \text{Verify}(\rho, \mathbf{x}, \boldsymbol{\pi}) = 1 \quad \wedge \\ (\mathbf{x}, \text{Extract}(\rho, \xi, \mathbf{x}, \boldsymbol{\pi})) \notin \mathcal{R} \end{array} \mid \begin{array}{l} (\rho, \xi) \leftarrow \text{ExtSetup}(1^\lambda), \\ (\mathbf{x}, \boldsymbol{\pi}) \leftarrow \mathcal{A}(\rho) \end{array} \right] = \text{negl}(\lambda).$$

**Statistical zero-knowledge in hiding mode.** For any  $(\rho, \zeta) \leftarrow \text{Setup}(1^\lambda)$ , and for all  $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$ , the distributions  $\{\text{Prove}(\rho, \mathbf{x}, \mathbf{w})\}$  and  $\{\text{Sim}(\rho, \zeta, \mathbf{x})\}$  are statistically indistinguishable.

## B Branch-Hiding: Formal Definition

**Definition 12.** A BiAS is branch-hiding if, for any PPT adversary  $\mathcal{A}$  involved in experiment  $\text{Exp}_{\mathcal{A}, N}^{\text{bh-d}}$ , the following advantage function is negligible:

$$\mathbf{Adv}_{\mathcal{A}, N}^{\text{bh}}(\lambda) := \left| \Pr \left[ \text{Exp}_{\mathcal{A}, N}^{\text{bh-1}}(\lambda) = 1 \right] - \Pr \left[ \text{Exp}_{\mathcal{A}, N}^{\text{bh-0}}(\lambda) = 1 \right] \right|.$$

```

1 stateT := (St, Y, SGM, SOA) ← Setup(1λ, 1N);
2 (aux, M*, w0*, w1*, id*, sec*, cert*, P*) ← A(play; Qpub, QkeyGM, Qopen, Qwrite);
3 if (certid* ≠Y secid*) ∨ (M* ∉ M) ∨ (w0*, w1* ∉ W) ∨ (P* ∉ P) then return 0;
4 σd ← Sign(Y, id*, cert*, sec*, M*, wd*, P*);
5 d' ← A(guess; σd, aux, Qpub, QkeyGM, Qopen¬{(M*, σd, P*)}, Qread, Qwrite);
6 return d';

```

**Fig. 5:** Experiment  $\text{Exp}_{\mathcal{A}, N}^{\text{bh-d}}(\lambda)$

This notion also ensures the computational confidentiality of the witness. Nonetheless, if  $P(M, id, w)$  is independent of the witness  $w \in W$  so that  $P(M, id, \star)$  is constant, the BiAS cannot hide this value. However, it is the purpose of the construction of a BiAS to define the appropriate predicate set  $P$  for useful applications. Still, the witness remains hidden even in that case.

In the case of anonymous signature, the branch-hiding property of the bifurcated cryptosystem is implied by the notion of *full anonymity*, Definition 7.

## C Stronger Anonymity Notion

The non-traceable case of the definition of anonymity is formalized to avoid running a BiAS where a hidden trapdoor would allow to break the anonymity. Since the adversary may be unbounded it could retrieve a trapdoor in the parameters if there is one and gain more chance to identify signers while the probability of this to happen should remain negligible.

Nevertheless, the definition assumes an honest generation of these parameters: The first line in  $\text{Exp}$ . 1 runs  $\text{Setup}$ . In practice, users do not know whether the parameters have been faithfully generated from the proper algorithm. This is a problem if an adversary is able to produce valid-looking  $(\mathcal{Y}, \mathcal{S}_{\text{GM}}, \mathcal{S}_{\text{OA}})$  which are not in the range of  $\text{Setup}$ , or if  $\text{Setup}$  contains a secret mode producing subverted keys that can only be activated by an adversary —such as a malevolent designer. Here, we strengthen the anonymity definition to capture and prevent such kind of subversion attacks. We call this notion *unsubversive anonymity*.

*Unsubversive Anonymity* To model unsubversive anonymity, we consider a two-stage adversary  $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$  where  $\mathcal{A}_1$  is a PPT algorithm that produces computationally (weakly) indistinguishable keys  $(\mathcal{Y}, \mathcal{S}_{\text{GM}}, \mathcal{S}_{\text{OA}})$  from those generated by  $\text{Setup}$ . We could have allowed  $\mathcal{A}_1$  to generate an additional state  $\text{st}$  to transfer it as auxiliary input to  $\mathcal{A}_2$ . However, there is no need to complicate the description since  $\mathcal{A}_2$  might be unbounded. The second-stage adversary plays similar experiments that of anonymity in the non-traceable experiments but based on the  $\mathcal{A}_1$ 's chosen keys.

**Definition 13.** *A BiAS scheme is unsubversively anonymous if, for any two-stage adversary  $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$  where  $\mathcal{A}_1$  is a PPT algorithm and  $\mathcal{A}_2$  any (unbounded) algorithm, the following conditions hold:*

**Weak Indistinguishability of Keys** *For any PPT distinguisher  $\mathcal{D}$  the following distance is non-noticeable.*

$$|\Pr[1 \leftarrow \mathcal{D}^{\text{Setup}}(1^\lambda, 1^N)] - \Pr[1 \leftarrow \mathcal{D}^{\mathcal{A}_1}(1^\lambda, 1^N)]|$$

**Unsubverion case** *The following advantage is negligible.*

$$\mathbf{Adv}_{\mathcal{A}, N}^{\text{anon-subv}}(\lambda) := \left| \Pr\left[\mathbf{Exp}_{\mathcal{A}, N}^{\text{anon-subv-1}}(\lambda) = 1\right] - \Pr\left[\mathbf{Exp}_{\mathcal{A}, N}^{\text{anon-subv-0}}(\lambda) = 1\right] \right|,$$

*where the experiments  $\mathbf{Exp}_{\mathcal{A}, N}^{\text{anon-subv-}d}$  are defined in Exp. 6.*

*In the random oracle model,  $\text{Setup}$ ,  $\mathcal{A}_1$  and  $\mathcal{A}_2$ , as well as all the BiAS algorithm have access to (the same)  $H_0$ .*

Relying on a weak indistinguishability of the key generations (instead of a “usual” indistinguishability) allows modeling more threats. Indeed, it suffices that we do *not notice* that  $\mathcal{A}_1$  produced the keys using another algorithm than  $\text{Setup}$  to capture potential anonymity subversion. The distribution of subverted keys does not have to be *negligibly* closed to the distribution of honestly generated keys. Note: by taking  $\mathcal{A}_1 = \text{Setup}$  we see that an unsubversively anonymous BiAS is also anonymous in the non-traceable case.

```

1 stateτ := (St, Y, SGM, SOA) ← A1(1λ, 1N);
2 (aux, M*, w0*, w1*, (id0*, sec0*, cert0*), (id1*, sec1*, cert1*), P*)
   ← A2(play; Y, SGM, SOA, Qread, Qwrite);
3 if  $\neg(\text{cert}_{\text{id}_0^*} \Leftarrow y \text{ sec}_{\text{id}_0^*}) \vee \neg(\text{cert}_{\text{id}_1^*} \Leftarrow y \text{ sec}_{\text{id}_1^*})$  then return 0;
4 if  $(M^* \notin \mathcal{M}) \vee (w_0^*, w_1^* \notin \mathcal{W}) \vee (P^* \notin \mathcal{P})$  then return 0;
5  $\sigma_d \leftarrow \text{Sign}(Y, \text{id}_d^*, \text{cert}_d^*, \text{sec}_d^*, w_d^*, M^*, P^*)$  ;
6  $d' \leftarrow A_2(\text{guess}; \sigma_d, aux, Q_{\text{pub}}, Q_{\text{read}}, Q_{\text{write}})$ ;
7 if  $(P^*(M^*, \text{id}_0^*, w_0^*) = 1) \wedge (P^*(M^*, \text{id}_1^*, w_1^*) = 1)$  then return d';
8 return 0;

```

Fig. 6: Experiment  $\text{Exp}_{(\mathcal{A}_1, \mathcal{A}_2), N}^{\text{anon-subv-}d}(\lambda)$ .

*Compatibility of the definitions* The reason why unsubversive anonymity and branch-soundness (Definition 8) are compatible lies in two observations: (1) it is hard for an efficient  $\mathcal{A}_1$  in the anonymity definition above to produce  $\text{SimSetup}$  fulfilling extractable soundness, and (2) In the random oracle model,  $\mathcal{A}_1$  shares (the same)  $H_0$  with the honest BiAS algorithms, so with those of the real mode of the experiment in **Exp. 2** of the branch-soundness definition, and so does  $\mathcal{A}_2$ , while in the extractable mode, the BiAS algorithms have access to possibly different but computationally indistinguishable random oracle  $H_1$ .

To be more pragmatic, if the (real) key generation contains a proof that the keys have been produced faithfully,  $\mathcal{A}_1$  is forced to adopt an honest behavior except if it is able to break the soundness of the proof. In the extractable setting, we can simulate a false proof and possibly also rely on a pseudo-random oracle  $H_1$ . However, note that  $H_1$  does not exist in the experiment for  $\mathcal{A}_2$ .

## D Deferred Proofs for the Generic Construction

### D.1 Proof of Lemma 2

*Proof.* The proof is based on standard arguments. Let us assume an adversary  $\mathcal{A}$  for which experiment  $\text{Exp}_{\mathcal{A}, N}^{\text{mis-id}}(\lambda)$  returns 1 at line 5 (because the Extract algorithm reveals  $\text{id} \in \mathcal{ID} \setminus U^a$ ) with non-negligible probability. We prove that  $\mathcal{A}$  implies an algorithm  $\mathcal{B}$  attacking one of the underlying components with noticeable probability.

Recall that, in  $\text{Exp}_{\mathcal{A}, N}^{\text{mis-id}}(\lambda)$ ,  $\text{SimSetup}$  generates injective keys for the  $\mathcal{R}$ -lossy encryption  $(\mathsf{pk}_{\text{RLE}}, \mathsf{sk}_{\text{RLE}}) \leftarrow \Pi^{\text{RLE}}.\text{Keygen}(\mathcal{F}, \mathbf{0}^L)$  for  $K = \mathbf{0}^L$ . This implies that, in each signature  $\Sigma = (\mathsf{VK}, (\mathsf{ct}_{(\text{id}, w)}, \mathsf{com}_{(\text{id}, w)}, \mathsf{ct}_{\text{id}}, \boldsymbol{\pi}), \mathsf{sig})$ ,  $\mathsf{ct}_{(\text{id}, w)}$  is computed as a statistically extractable commitment to  $w \in \{0, 1\}^{\ell_w}$  and  $\text{id} \in \{0, 1\}^{\ell}$ . Moreover,  $\text{SimSetup}$  generate a common reference string  $\rho$  from  $(\rho, \xi) \leftarrow \text{NIZK}.\text{ExtSetup}(1^\lambda)$  for an extractable NIZK proof system, which allows the challenger to compute witnesses from any verifying NIZK argument (with overwhelming probability) thanks to the extraction trapdoor  $\xi$  contained in  $\tau_{\text{ext}} = (\mathsf{sk}_{\text{RLE}}, \xi)$ . In the adversary's

forgery  $\Sigma^* = (\text{VK}^*, (\text{ct}_{(\text{id}, w)}^*, \text{com}_{(\text{id}, w)}^*, \text{ct}_{\text{id}}^*, \boldsymbol{\pi}^*), \text{sig}^*)$ , the Extract algorithm of  $\text{Exp}_{\mathcal{A}, N}^{\text{mis-id}}(\lambda)$  can thus use the secret key  $\text{sk}_e$  as an extraction trapdoor to recover  $\{w_i^*\}_{i=1}^{\ell_w}$  and  $\{\text{id}_i^*\}_{i=1}^{\ell}$  from  $\text{ct}_{(\text{id}, w)}^*$ .

We now describe a PPT reduction  $\mathcal{B}$  that uses  $\mathcal{A}$  to forge a signature for  $\Pi^{\text{sig}}$ . Specifically,  $\mathcal{B}$  takes as input a public key  $\text{pk}_{\text{sig}}$  obtain from its existential unforgeability (EUF-CMA) challenger. Then, it runs steps 2-6 of the **SimSetup** algorithm to produce a group public key  $\mathcal{Y} = (\rho, \text{pk}_{\text{sig}}, (\text{pp}, \text{ek}), (\Gamma, \text{pk}_{\text{RLE}}), \text{pk}_e)$  together with  $\tau_{\text{ext}} = (\text{sk}_{\text{RLE}}, \xi)$  and provides  $\mathcal{A}$  with  $\mathcal{Y}$ .

**Queries.** The reduction  $\mathcal{B}$  starts interacting with the adversary  $\mathcal{A}$  and handles  $\mathcal{A}$ 's queries to the  $\mathcal{Q}_{\text{a-join}}$  oracle with the help of its EUF-CMA challenger. When  $\mathcal{A}$  triggers an execution of the joining protocol, it chooses a pair  $(\text{id}, \text{pk}_{\text{id}})$ . If these do not parse properly or if  $\text{id}$  is not a fresh identifier with a valid signature  $\text{sig}_{\text{id}}$ ,  $\mathcal{B}$  returns  $\perp$ . Otherwise,  $\mathcal{B}$  invokes its EUF-CMA challenger to obtain a signature  $\text{cert}_{\text{id}} \leftarrow \Pi^{\text{sig}}.\text{Sign}(\text{sk}_{\text{sig}}, (\text{id}, \text{pk}_{\text{id}}))$  on the message  $(\text{id}, \text{pk}_{\text{id}})$ . Then,  $\mathcal{B}$  returns  $\text{cert}_{\text{id}}$  to  $\mathcal{A}$  before updating  $St_{\text{trans}}$  as the real group manager would.

**Forgery.** When the adversary  $\mathcal{A}$  halts, it comes up with a signature  $\Sigma^* = (\text{VK}^*, (\text{ct}_{(\text{id}, w)}^*, \text{com}_{(\text{id}, w)}^*, \text{ct}_{\text{id}}^*, \boldsymbol{\pi}^*), \text{sig}^*)$  on some message  $M^*$  and for some predicate  $P^*$ . At this point,  $\mathcal{B}$  uses the extraction trapdoor  $\tau_{\text{ext}}$  produced by the **SimSetup** algorithm to extract  $\text{ct}_{(\text{id}, w)}^*$  and obtain strings  $\text{id}^* \in \{0, 1\}^\ell$ ,  $w^* \in \{0, 1\}^{\ell_w}$  as well as  $(\text{pk}_{\text{id}}^*, \text{cert}_{\text{id}}^*, \sigma^*) \in \mathcal{VK} \times \mathcal{S} \times \mathcal{S}$  among other witnesses. The statistical soundness of  $\boldsymbol{\pi}^*$  ensures that, with overwhelming probability, the extracted pair  $(\text{pk}_{\text{id}}^*, \text{cert}_{\text{id}}^*, \sigma^*)$  satisfies  $\Pi^{\text{sig}}.\text{Verify}(\text{pk}_{\text{sig}}, (\text{id}^*, \text{pk}_{\text{id}}^*), \text{cert}_{\text{id}}^*) = 1$ . Since  $\mathcal{A}$  is assumed to win the misidentification experiment with noticeable probability  $\varepsilon = \text{Adv}_{\mathcal{A}, N}^{\text{mis-id}}(\lambda)$ , we must have  $\text{id}^* \in \mathcal{ID} \setminus U^a$ , which implies that  $\text{cert}_{\text{id}}^*$  is a non-trivial signature on a pair  $(\text{id}^*, \text{pk}_{\text{id}}^*)$  that  $\mathcal{B}$  did not submit to its signing oracle. This allows  $\mathcal{B}$  to output  $(\text{id}^*, \text{pk}_{\text{id}}^*)$ , which allows it to win its EUF-CMA security game with advantage  $\varepsilon - \nu(\lambda)$ , for some negligible function  $\nu(\lambda)$  accounting for the soundness error of  $\boldsymbol{\pi}^*$ .  $\square$

## D.2 Proof of Lemma 3

*Proof.* Let us assume that an adversary  $\mathcal{A}$  can win experiment  $\text{Exp}_{\mathcal{A}, N}^{\text{fra}}(\lambda)$  with noticeable advantage, meaning that the experiment returns 1 at line 5 because the Extract algorithm reveals  $(\text{id}, w)$  such that  $\text{id} \in U^b$  and the adversary's forgery  $\Sigma^*$  was never generated by the signing oracle for  $(M^*, P')$  using the pair  $(\text{id}', w')$ . We prove that  $\mathcal{A}$  implies a forger  $\mathcal{B}$  with noticeable probability against the signature scheme  $\Pi^{\text{sig}}$ .

Recall that, in experiment  $\text{Exp}_{\mathcal{A}, N}^{\text{fra}}(\lambda)$ , **SimSetup** generates injective keys  $(\text{pk}_{\text{RLE}}, \text{sk}_{\text{RLE}}) \leftarrow \Pi^{\text{RLE}}.\text{Keygen}(\Gamma, \mathbf{0}^L)$  for the initialization value  $K = \mathbf{0}^L$ . In each signature  $\Sigma = (\text{VK}, (\text{ct}_{(\text{id}, w)}, \text{com}_{(\text{id}, w)}, \text{ct}_{\text{id}}, \boldsymbol{\pi}), \text{sig})$ , the commitment  $\text{ct}_{(\text{id}, w)}$  is thus computed as a statistically extractable commitment to  $(\text{id}, w) \in \{0, 1\}^{\ell + \ell_w}$ . Moreover, **SimSetup** generates a CRS  $\rho$  from  $(\rho, \xi) \leftarrow \text{NIZK}.\text{ExtSetup}(1^\lambda)$  for an extractable NIZK proof system, which allows the challenger to recover witnesses from any verifying NIZK argument (with overwhelming probability) using the

extraction trapdoor  $\xi$  contained in  $\tau_{\text{ext}} = (\text{sk}_{\text{RLE}}, \xi)$ . In the adversary's forgery  $\Sigma^* = (\text{VK}^*, (\text{ct}_{(\text{id}, w)}^*, \text{com}_{(\text{id}, w)}^*, \text{ct}_{\text{id}}^*, \pi^*), \text{sig}^*)$ , the Extract algorithm can thus use  $\text{sk}_e$  to recover  $\text{id}^* \in \{0, 1\}^\ell$ ,  $w^* \in \{0, 1\}^{\ell_w}$  as well as  $(\text{pk}_{\text{id}}, \text{cert}_{\text{id}}^*, \sigma^*) \in \mathcal{VK} \times \mathcal{S} \times \mathcal{S}$  together with other witnesses which are extracted from  $\pi^*$ .

- If the pair  $(\text{id}^*, \text{pk}_{\text{id}}^*)$  was never used by the reduction when simulating the  $Q_{\text{b-join}}$ -oracle, it means that  $\text{id}^*$  was introduced in  $U^b$  when, at step 1 of the Join protocol, it sent a signature  $\text{sig}_{\text{id}^*}$  on a pair  $(\text{id}^*, \text{pk}_{\text{id}})$  with  $\text{pk}_{\text{id}} \neq \text{pk}_{\text{id}^*}$ . This means that  $\mathcal{A}$  was able to break the unforgeability of  $\Pi^{\text{sig}}$  for the long-term public key  $\text{upk}[\text{id}^*]$  that user  $\text{id}^*$  created for itself.
- If the pair  $(\text{id}^*, \text{pk}_{\text{id}}^*)$  was used in an execution of the Join protocol triggered by a  $Q_{\text{b-join}}$ -query, the tuple  $(M^*, P^*, \text{ct}_{(\text{id}, w)}^*)$  was never signed by  $\mathcal{B}$  when answering a  $Q_{\text{sig}}$ -query. Indeed, at line 5 of  $\text{Exp}_{\mathcal{A}, N}^{\text{fra}}(\lambda)$ , the adversary's winning conditions  $(\text{id}^*, M^*, w^*, \sigma^*, P^*) \notin \text{Sigs}$  implies a breach in the strong unforgeability of  $\Pi^{\text{sig}}$  for the public key  $\text{pk}_{\text{id}}^*$ .

If  $\mathcal{A}$ 's advantage is  $\varepsilon = \text{Exp}_{\mathcal{A}, N}^{\text{fra}}(\lambda)$ , we can thus have the reduction  $\mathcal{B}$  break the strong unforgeability of  $\Pi^{\text{sig}}$  with advantage  $\varepsilon/(2 \cdot Q_S) - \nu(\lambda)$ , where  $Q_S$  is the number of  $Q_{\text{sig}}$ -queries and  $\nu(\lambda)$  is a negligible function accounting for the soundness error of  $\pi^*$ .  $\square$

### D.3 Proof of Lemma 4

*Proof.* We prove the result via a sequence of games which are computationally indistinguishable. The first is identical to the real anonymity experiment  $\text{Exp}_{\mathcal{A}, N}^{\text{anon-}d}(\lambda)$  which is parameterized by a bit  $d \in \{0, 1\}$  that dictates the challenger's choice in the challenge phase. The last game is the same regardless of whether  $d = 0$  or  $d = 1$ . It follows that, under the stated assumptions, no PPT adversary can distinguish  $\text{Exp}_{\mathcal{A}}^{\text{anon-}0}$  from  $\text{Exp}_{\mathcal{A}}^{\text{anon-}1}$  with noticeable advantage.

**Game**<sup>(d)</sup> 0: This is the real anonymity experiment  $\text{Exp}_{\mathcal{A}}^{\text{anon-}d}(\lambda)$  as described in Definition 7. More precisely, the challenger starts by running the algorithm  $\text{Setup}(1^\lambda, 1^N, \mathcal{P}_d)$  to obtain  $(\mathcal{Y}, \mathcal{S}_{\text{GM}} = \text{sk}_{\text{sig}}, \mathcal{S}_{\text{OA}} = \text{sk}_e)$  along with state information  $\mathcal{St}$ . The challenger gives the public parameters  $\mathcal{Y}$  and the group manager key  $\mathcal{S}_{\text{GM}}$  to the adversary  $\mathcal{A}$ . At each opening query involving a signature  $\Sigma = (\text{VK}, (\text{ct}_{(\text{id}, w)}, \text{com}_{(\text{id}, w)}, \text{ct}_{\text{id}}, \pi), \text{sig})$ , the challenger uses  $\mathcal{S}_{\text{OA}} = \text{sk}_e$  to decrypt  $\text{ct}_{\text{id}}$  as in the real Open algorithm. If  $\text{ct}_{\text{id}}$  decrypts to  $\mathbf{0}^\ell$ , the challenger outputs  $\perp$  to indicate that  $\Sigma$  was not traceable.

At the challenge phase, the adversary  $\mathcal{A}$  chooses a message-predicate pair  $(M^*, P^*)$  as well as two tuples  $(\text{id}_0^*, w_0^*, \text{sec}_0^*, \text{cert}_0^*)$  and  $(\text{id}_1^*, w_1^*, \text{sec}_1^*, \text{cert}_1^*)$ , each of which contains a valid membership certificate and a membership secret. Then, the challenger uses its secret coin  $d \leftarrow \{0, 1\}$  to compute a challenge signature  $\Sigma^* = (\text{VK}^*, (\text{ct}_{(\text{id}, w)}^*, \text{com}_{(\text{id}, w)}^*, \text{ct}_{\text{id}}^*, \pi^*), \text{sig}^*)$  using  $(\text{id}_d^*, w_d^*, \text{sec}_d^*, \text{cert}_d^*)$ . Note that  $\text{ct}_{\text{id}}^*$  is a PKE encryption of  $\text{id}_d^* \in \{0, 1\}^\ell$ .

When  $\mathcal{A}$  halts, it outputs a bit  $d' \in \{0, 1\}$  which is also the output of the experiment.

**Game<sup>(d)</sup> 1:** In this experiment, we modify Game<sup>(d)</sup> 0 as follows. At the very beginning of the game, the challenger generates the one-time signature key pair  $(\text{VK}^*, \text{SK}^*)$  that will be used in the challenge phase. During the game, if  $\mathcal{A}$  queries the opening of a valid signature  $\Sigma = (\text{VK}, (\text{ct}_{(\text{id}, w)}, \text{com}_{(\text{id}, w)}, \text{ct}_{\text{id}}, \boldsymbol{\pi}), \text{sig})$  where  $\text{VK} = \text{VK}^*$ , the challenger returns a random bit and aborts. If we call  $F_1$  this event, it is easy to see that it would contradict the strong unforgeability of the one-time signature  $\Pi^{\text{ots}}$ . Indeed, before the challenge phase  $\text{VK}^*$  is independent of  $\mathcal{A}$ 's view and the probability that  $\text{VK}^*$  accidentally appears in  $\mathcal{A}$ 's queries is negligible. After the challenge phase  $\Sigma^*$ , if  $\mathcal{A}$  queries the opening of a valid signature  $\Sigma = (\text{VK}, (\text{ct}_{(\text{id}, w)}, \text{com}_{(\text{id}, w)}, \text{ct}_{\text{id}}, \boldsymbol{\pi}), \text{sig})$  such that  $\text{VK} = \text{VK}^*$ , then  $\text{sig}$  is a fake one-time signature, which defeats the strong unforgeability of  $\Pi^{\text{ots}}$ . Therefore the probability  $\Pr[F_1]$  that the challenger aborts in this experiment is negligible. From this point forward, we thus assume that valid signature opening queries do not involve  $\text{VK}^*$ .

**Game<sup>(d)</sup> 2:** In this game, we modify step 4 of the **Setup** algorithm in the following way at the beginning of the game. Instead of generating  $(\text{pk}_{\text{RLE}}, \text{sk}_{\text{RLE}}) \leftarrow \Pi^{\text{RLE}}.\text{LKeygen}(\Gamma, \mathbf{0}^L)$  for the initialization value  $K = \mathbf{0}^L$ , we now generate them as  $(\text{pk}_{\text{RLE}}, \text{sk}_{\text{RLE}}) \leftarrow \Pi^{\text{RLE}}.\text{LKeygen}(\Gamma, \text{VK}^*)$  for the initialization value  $K = \text{VK}^*$ . By the second indistinguishability property of the  $\mathcal{R}_{\text{NEQ}}$ -lossy PKE scheme, this modification does not impact the adversary's view since the distribution of  $\text{pk}_{\text{RLE}}$  remains statistically unchanged.

**Game<sup>(d)</sup> 3:** We now modify the distribution of  $\text{pk}_{\text{RLE}}$  at step 4 of the **Setup** algorithm. Instead of generating lossy keys  $(\text{pk}_{\text{RLE}}, \text{sk}_{\text{RLE}}) \leftarrow \Pi^{\text{RLE}}.\text{LKeygen}(\Gamma, \text{VK}^*)$ , we now generate injective keys  $(\text{pk}_{\text{RLE}}, \text{sk}_{\text{RLE}}) \leftarrow \Pi^{\text{RLE}}.\text{Keygen}(\Gamma, \text{VK}^*)$  for the initialization value  $K = \text{VK}^*$ . By the first indistinguishability property of the  $\mathcal{R}_{\text{NEQ}}$ -lossy PKE scheme, the two distributions of  $\text{pk}_{\text{RLE}}$  are computationally indistinguishable. This implies that **Game<sup>(d)</sup> 4** is computationally indistinguishable from **Game<sup>(d)</sup> 3** from the adversary's view.

**Game<sup>(d)</sup> 4:** Here, we modify the signature opening oracle. Namely, at each valid opening query  $\Sigma = (\text{VK}, (\text{ct}_{(\text{id}, w)}, \text{com}_{(\text{id}, w)}, \text{ct}_{\text{id}}, \boldsymbol{\pi}), \text{sig})$  (note that we have  $\text{VK} \neq \text{VK}^*$  due to the change introduced in **Game<sup>(d)</sup> 1**), the challenger uses the secret key  $\text{sk}_{\text{RLE}}$  of the  $\mathcal{R}_{\text{NEQ}}$ -lossy PKE scheme to extract  $(\text{id}, w)$  from  $\text{ct}_{(\text{id}, w)}$ . Having extracted  $\text{id} \in \{0, 1\}^\ell$  and  $w \in \{0, 1\}^{\ell_w}$ ,  $\mathcal{B}$  can evaluate  $c_{ev} = C_M(w_1, \dots, w_{\ell_w}, \text{id}_1, \dots, \text{id}_\ell) \in \{0, 1\}$ . If  $c_{ev} = 0$ ,  $\mathcal{B}$  returns  $\text{id} \in \{0, 1\}^\ell$ . Otherwise (i.e., if  $c_{ev} = 1$ ), it returns  $\perp$  to indicate that  $\Sigma$  was not a traceable signature.

Unless  $\mathcal{A}$  can break the computational soundness of the NIZK argument system used to compute  $\boldsymbol{\pi}$ , its view remains the same as in **Game<sup>(d)</sup> 3** since  $\boldsymbol{\pi}$  ensures that  $\text{ct}_{(\text{id}, w)}$  and  $\text{ct}_{\text{id}}$  contain the same string when  $c_{ev} = 0$ . Since the challenger knows both  $\text{sk}_{\text{RLE}}$  and  $\text{sk}_e$ , it can efficiently detect when a breach in the soundness of the argument system occurs, so that we can easily turn

it into a soundness attacker if the adversary can somehow create an unsound argument  $\pi$ . Assuming the soundness of NIZK arguments  $\pi$ , **Game**<sup>(d)</sup> 4 is thus indistinguishable from **Game**<sup>(d)</sup> 3.

In **Game**<sup>(d)</sup> 4, we observe that the OA's secret key  $\mathcal{S}_{\text{OA}} = \mathbf{sk}_e$  is no longer used, which will allow us to switch the public of  $\Pi^{\text{lpke}}$  into a lossy key in the next transition.

**Game**<sup>(d)</sup> 5: We modify again the distribution of  $\mathcal{Y}$ . At Step 5 of the **Setup** algorithm, the challenger now computes  $(\mathbf{pk}_e, \mathbf{sk}_e) \leftarrow \Pi^{\text{lpke}}.\text{LKeygen}(1^\lambda)$  as a lossy key (where  $\mathbf{sk}_e$  may be  $\perp$ ). Under the key indistinguishability property of the lossy PKE scheme  $\Pi^{\text{lpke}}$ , this change should remain unnoticed to  $\mathcal{A}$ .

**Game**<sup>(d)</sup> 6: This game is identical to **Game**<sup>(d)</sup> 5 but, in the challenge signature  $\Sigma^*$ , the NIZK argument  $\pi^*$  is generated using the NIZK simulator. Since the NIZK argument system is statistically zero-knowledge, this modification implies that the simulated  $\pi^*$  is statistically close to a real argument that would be generated using the witnesses as in **Game**<sup>(d)</sup> 5.

In **Game**<sup>(d)</sup> 6, we claim that the adversary's view is statistically independent of the challenger's bit  $d \in \{0, 1\}$ . Indeed,  $\mathbf{ct}_{(\text{id}, w)}^*$  statistically hides the underlying pair  $(\text{id}^*, w^*)$  since it is computed for the tag  $\text{VK}^*$  under which  $\Pi^{\text{RLE}}$  provides lossy encryptions. Moreover,  $\mathbf{com}_{(\text{id}, w)}^*$  is a statistically hiding commitment by design. Finally,  $\mathbf{ct}_{\text{id}}^*$  is now computed under a lossy public key  $\mathbf{pk}_e$  for  $\Pi^{\text{lpke}}$  while  $\pi^*$  is simulated without using the witnesses.  $\square$

## E Material for Lattice-Based Instantiations

### E.1 Lattices

LATTICES AND GAUSSIAN DISTRIBUTIONS. A lattice  $L$  is the set of integer linear combinations of linearly independent basis vectors  $(\mathbf{b}_i)_{i \leq n}$  living in  $\mathbb{R}^m$ . We work with  $q$ -ary lattices, for some prime  $q$ .

**Definition 14.** Let  $m \geq n \geq 1$ , a prime  $q \geq 2$  and  $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$  and  $\mathbf{u} \in \mathbb{Z}_q^n$ , define  $\Lambda_q(\mathbf{A}) := \{\mathbf{e} \in \mathbb{Z}^m \mid \exists \mathbf{s} \in \mathbb{Z}_q^n \text{ s.t. } \mathbf{A}^T \cdot \mathbf{s} = \mathbf{e} \bmod q\}$  as well as

$$\begin{aligned} \Lambda_q^\perp(\mathbf{A}) &:= \{\mathbf{e} \in \mathbb{Z}^m \mid \mathbf{A} \cdot \mathbf{e} = \mathbf{0}^n \bmod q\}, \\ \Lambda_q^\mathbf{u}(\mathbf{A}) &:= \{\mathbf{e} \in \mathbb{Z}^m \mid \mathbf{A} \cdot \mathbf{e} = \mathbf{u} \bmod q\} \end{aligned}$$

For any arbitrary  $\mathbf{t} \in \Lambda_q^\mathbf{u}(\mathbf{A})$ , we also define the shifted lattice  $\Lambda_q^\mathbf{u}(\mathbf{A}) = \Lambda_q^\perp(\mathbf{A}) + \mathbf{t}$ .

For a lattice  $L$ , a vector  $\mathbf{c} \in \mathbb{R}^m$  and a real number  $\sigma > 0$ , define the function  $\rho_{\sigma, \mathbf{c}}(\mathbf{x}) = \exp(-\pi \|\mathbf{x} - \mathbf{c}\|^2 / \sigma^2)$ . The discrete Gaussian distribution of support  $L$ , center  $\mathbf{c}$  and parameter  $\sigma$  is defined as  $D_{L, \sigma, \mathbf{c}}(\mathbf{y}) = \rho_{\sigma, \mathbf{c}}(\mathbf{y}) / \rho_{\sigma, \mathbf{c}}(L)$  for any  $\mathbf{y} \in L$ , where  $\rho_{\sigma, \mathbf{c}}(L) = \sum_{\mathbf{x} \in L} \rho_{\sigma, \mathbf{c}}(\mathbf{x})$ . We denote by  $D_{L, \sigma}(\mathbf{y})$  the distribution centered in  $\mathbf{c} = \mathbf{0}^m$  and exploit the fact that samples from  $D_{L, \sigma}$  are short w.h.p.

**Lemma 6** ([4, Le. 1.5]). *For any lattice  $L \subseteq \mathbb{R}^m$  and positive real number  $\sigma > 0$ , we have  $\Pr_{\mathbf{b} \leftarrow D_{L,\sigma}}[\|\mathbf{b}\| \leq \sqrt{m}\sigma] \geq 1 - 2^{-\Omega(m)}$ .*

**GADGET MATRIX.** In [46], Micciancio and Peikert described a trapdoor mechanism for LWE. Their technique uses a “gadget” matrix  $\mathbf{G} \in \mathbb{Z}_q^{n \times m}$  for which anyone can publicly sample short vectors  $\mathbf{x} \in \mathbb{Z}^m$  such that  $\mathbf{G} \cdot \mathbf{x} = \mathbf{0}$ .

As in [46], we call  $\mathbf{R} \in \mathbb{Z}^{m \times m}$  a  $\mathbf{G}$ -trapdoor for a matrix  $\mathbf{A} \in \mathbb{Z}_q^{2m \times n}$  if  $\mathbf{A} \cdot [\mathbf{R}^\top | \mathbf{I}_m]^\top = \mathbf{H} \cdot \mathbf{G}$  for some invertible matrix  $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$  which is referred to as the trapdoor tag. If  $\mathbf{H} = \mathbf{0}$ , then  $\mathbf{R}$  is called a “punctured” trapdoor for  $\mathbf{A}$ .

**Lemma 7** ([46, Section 5]). *Assume that  $m \geq 2n \log q$ . There exists a PPT algorithm  $\text{GenTrap}$  that takes as inputs matrices  $\bar{\mathbf{A}} \in \mathbb{Z}_q^{n \times m}$ ,  $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$  and outputs matrices  $\mathbf{R} \in \{-1, 1\}^{m \times m}$  and  $\mathbf{A} = [\bar{\mathbf{A}} | -\bar{\mathbf{A}} \cdot \mathbf{R} + \mathbf{H} \cdot \mathbf{G}] \in \mathbb{Z}_q^{n \times 2m}$  such that if  $\mathbf{H} \in \mathbb{Z}_q^{n \times n}$  is invertible, then  $\mathbf{R}$  is a  $\mathbf{G}$ -trapdoor for  $\mathbf{A}$  with tag  $\mathbf{H}$ ; and if  $\mathbf{H} = \mathbf{0}$ , then  $\mathbf{R}$  is a punctured trapdoor.*

Further, in case of a  $\mathbf{G}$ -trapdoor, one can efficiently compute from  $\mathbf{A}, \mathbf{R}$  and  $\mathbf{H}$  a basis  $(\mathbf{b}_i)_{i \leq 2m}$  of  $\Lambda^\perp(\mathbf{A})$  such that  $\max_i \|\mathbf{b}_i\| \leq O(m^{3/2})$ .

We recall the Learning With Errors problem [53]. Note that we make the number of samples  $m$  explicit in our definition.

**Definition 15.** *Let  $\lambda \in \mathbb{N}$  be a security parameter and let integers  $n = n(\lambda)$ ,  $m = m(\lambda)$ ,  $q = q(\lambda)$ . Let  $\chi = \chi(\lambda)$  be an efficiently samplable distribution over  $\mathbb{Z}_q$ . The  $\text{LWE}_{n,m,q,\chi}$  assumption posits that the following distance is a negligible function for any PPT algorithm  $\mathcal{A}$ :*

$$\begin{aligned} \mathbf{Adv}_{n,m,q,\chi}^{\mathcal{A}, \text{LWE}}(\lambda) := & \left| \Pr[\mathcal{A}(1^\lambda, \mathbf{A}, \mathbf{u}) = 1 | \mathbf{A} \leftarrow U(\mathbb{Z}_q^{n \times m}), \mathbf{u} \leftarrow U(\mathbb{Z}_q^m)] \right. \\ & \left. - \Pr[\mathcal{A}(1^\lambda, \mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e}) = 1 | \mathbf{A} \leftarrow U(\mathbb{Z}_q^{m \times n}), \mathbf{s} \leftarrow U(\mathbb{Z}_q^n), \mathbf{e} \leftarrow \chi^m] \right|. \end{aligned}$$

**Definition 16.** *Let  $m, n, q, \beta$  be functions of a parameter  $1^\lambda$ . The Short Integer Solution problem  $\text{SIS}_{m,q,\beta}$  is as follows: Given  $\mathbf{A} \leftarrow U(\mathbb{Z}_q^{n \times n})$ , find a non-zero  $\mathbf{x} \in \mathbb{Z}^m$  such that  $\mathbf{A} \cdot \mathbf{x} = \mathbf{0}^n$  and  $0 < \|\mathbf{x}\| \leq \beta$ .*

## E.2 Randomness Extraction

We first recall the Leftover Hash Lemma, as it was stated in [2].

**Lemma 8** ([2]). *Let  $\mathcal{H} = \{h : X \rightarrow Y\}_{h \in \mathcal{H}}$  be a family of universal hash functions and let  $f : X \rightarrow Z$  be a function. Let  $(T_i)_{i \leq k}$  be independent random variables over  $X$ , for some  $k > 0$ . Letting  $\gamma = \max_{i \leq k} \gamma(T_i)$ , we have  $\Delta((h, (h(T_i), f(T_i))_{i \leq k}), (h, (U(Y))^{(i)}, f(T_i))_{i \leq k})) \leq \frac{k}{2} \cdot \sqrt{\gamma \cdot |Y| \cdot |Z|}$ .*

The Leftover Hash Lemma was used by Agrawal *et al.* [2] to re-randomize matrices over  $\mathbb{Z}_q$  by multiplying them with small-norm matrices.

**Lemma 9.** *Let integers  $m, n, k$  such that  $m > 2n \cdot (\log q + \lambda)$ , for some prime  $q > 2$ . Let  $\mathbf{B} \leftarrow U(\mathbb{Z}_q^{n \times m})$ ,  $\tilde{\mathbf{B}} \leftarrow U(\mathbb{Z}_q^{n \times k})$  and  $\mathbf{R} \leftarrow U(\{-1, 1\}^{m \times k})$ . For any matrix  $\mathbf{F} \in \mathbb{Z}_q^{n \times m}$ , the distributions  $(\mathbf{B}, \mathbf{B} \cdot \mathbf{R}, \mathbf{F} \cdot \mathbf{R})$  and  $(\mathbf{B}, \tilde{\mathbf{B}}, \mathbf{F} \cdot \mathbf{R})$  are within statistical distance  $k \cdot 2^{-\lambda}$ .*

### E.3 Fully Homomorphic Commitments and the GSW FHE

This section recalls the leveled homomorphic encryption scheme of Gentry, Sahai and Waters (GSW) [28] and its use as a fully homomorphic commitment scheme. The description makes use of the “gadget matrix” of Micciancio and Peikert [46].

Micciancio and Peikert [46] proved that, for any  $m \geq n\lceil\log q\rceil$ , there exists an efficiently computable matrix  $\mathbf{G} \in \mathbb{Z}_q^{n \times m}$  and an efficiently computable deterministic “*short preimage*” function  $\mathbf{G}^{-1}(\cdot)$  with the following property: on input of a matrix  $\mathbf{M} \in \mathbb{Z}_q^{n \times m'}$ , for any integer  $m' > 0$ , the function  $\mathbf{G}^{-1}(\mathbf{M})$  outputs a binary matrix  $\mathbf{G}^{-1}(\mathbf{M}) \in \{0, 1\}^{m \times m'}$  such that  $\mathbf{G}\mathbf{G}^{-1}(\mathbf{M}) = \mathbf{M}$ . The matrix  $\mathbf{G}$  can be defined as

$$\mathbf{G} = \left[ \mathbf{I}_n \otimes (1, 2, 4, \dots, 2^{\lceil\log q\rceil})^\top \mid \mathbf{0}^{n \times \lceil\log q\rceil} \right] \in \mathbb{Z}_q^{n \times m},$$

where  $m = 2n\lceil\log q\rceil$ , while  $\mathbf{G}^{-1}(\cdot)$  is the entry-wise binary decomposition function whose outputs are padded with zeroes to reach the desired dimension.

**Keygen**( $1^\lambda, 1^d$ ): On input of a security parameter  $\lambda$  and a maximal circuit depth  $d$ , the key generation algorithm samples the following elements:

- A uniformly random matrix  $\bar{\mathbf{A}} \leftarrow U(\mathbb{Z}_q^{(n-1) \times m})$  with  $n \in \text{poly}(\lambda)$  and  $m = 2n\lceil\log q\rceil$ ;
- A uniformly random vector  $\mathbf{s} \leftarrow U(\mathbb{Z}_q^{n-1})$ ;
- A small-norm  $\mathbf{e} \leftarrow \chi^m$  sampled from the error distribution  $\chi$ .

It defines the secret key  $\mathbf{sk} := \mathbf{t} = \begin{pmatrix} -\mathbf{s} \\ 1 \end{pmatrix} \in \mathbb{Z}_q^n$  and the public key

$$\mathbf{pk} := \mathbf{A} = \begin{pmatrix} \bar{\mathbf{A}} \\ \mathbf{s}^\top \cdot \bar{\mathbf{A}} + \mathbf{e}^\top \end{pmatrix} \in \mathbb{Z}_q^{n \times m}$$

**Encrypt**( $\mathbf{pk}, \mu$ ): In order to encrypt a bit  $\mu \in \{0, 1\}$ , the encryption algorithm samples a uniform matrix  $\mathbf{R} \leftarrow U(\{0, 1\}^{m \times m})$  and computes

$$\mathbf{C} = \mathbf{A} \cdot \mathbf{R} + \mu \cdot \mathbf{G} \in \mathbb{Z}_q^{n \times m},$$

where  $\mathbf{G} \in \mathbb{Z}_q^{n \times m}$  is the gadget matrix of [46].

**Decrypt**( $\mathbf{sk}, \mathbf{C}$ ): On input of  $\mathbf{C} \in \mathbb{Z}_q^{n \times m}$  and  $\mathbf{sk} := \mathbf{t} \in \mathbb{Z}_q^n$ , do the following:

1. Define  $\mathbf{w} = (0, \dots, 0, \lfloor q/2 \rfloor)^\top \in \mathbb{Z}_q^n$  and compute

$$\mathbf{v} = \mathbf{C} \cdot \mathbf{G}^{-1}(\mathbf{w}) \in \mathbb{Z}_q^n.$$

2. If  $|\mathbf{t}^\top \mathbf{v}|$  is close to 0, output  $\mu = 0$ . If  $|\mathbf{t}^\top \mathbf{v}|$  is close to  $\lfloor q/2 \rfloor$ , output  $\mu = 1$ .

**Eval**( $\mathbf{pk}, \mathbf{C}_1, \mathbf{C}_2, op$ ): Given two ciphertexts  $\mathbf{C}_1, \mathbf{C}_2 \in \mathbb{Z}_q^{n \times m}$  and an operation  $op \in \{+, \times\}$  to be evaluated over  $\mathbf{C}_1$  and  $\mathbf{C}_2$ , do the following:

- If  $op = “+”$ , output  $\mathbf{C}^+ := \mathbf{C}_1 + \mathbf{C}_2 \in \mathbb{Z}_q^{n \times m}$ .
- If  $op = “\times”$ , output  $\mathbf{C}^\times := \mathbf{C}_1 \cdot \mathbf{G}^{-1}(\mathbf{C}_2)$ .

If two fresh ciphertexts  $\mathbf{C}_1 = \mathbf{A} \cdot \mathbf{R}_1 + \mu_1 \cdot \mathbf{G}$  and  $\mathbf{C}_2 = \mathbf{A} \cdot \mathbf{R}_2 + \mu_2 \cdot \mathbf{G}$  encrypt  $\mu_1 \in \{0, 1\}$  and  $\mu_2 \in \{0, 1\}$ , respectively, then

$$\begin{aligned}\mathbf{C}^\times &= \mathbf{C}_1 \cdot \mathbf{G}^{-1}(\mathbf{C}_2) = \mathbf{A} \cdot (\mathbf{R}_1 \cdot \mathbf{G}^{-1}(\mathbf{C}_2) + \mu_1 \cdot \mathbf{R}_2) + \mu_1 \cdot \mu_2 \cdot \mathbf{G} \\ &= \mathbf{A} \cdot \mathbf{R}^\times + \mu_1 \cdot \mu_2 \cdot \mathbf{G}\end{aligned}$$

is indeed an encryption of  $\mu_1 \cdot \mu_2 \in \{0, 1\}$  with the small-norm matrix

$$\mathbf{R}^\times = \mathbf{R}_1 \cdot \mathbf{G}^{-1}(\mathbf{C}_2) + \mu_1 \cdot \mathbf{R}_2.$$

A Boolean NAND gate can be evaluated by computing

$$\mathbf{C}^{\text{NAND}} = \mathbf{G} - \mathbf{C}_1 \cdot \mathbf{G}^{-1}(\mathbf{C}_2) = \mathbf{A}(-\mathbf{R}^\times) + (\mu_1 \text{ NAND } \mu_2) \cdot \mathbf{G},$$

which allows evaluating any Boolean circuit of *a priori* bounded depth.

Gorbunov, Vaikuntanathan and Wichs [29] observed that, when the public matrix  $\mathbf{pk} = \mathbf{A}$  is replaced by a truly random matrix  $\mathbf{A} \sim U(\mathbb{Z}_q^{n \times m})$ , GSW implies fully homomorphic statistically hiding commitments. They proved the resulting scheme to be computationally binding on evaluated commitments under the Short Integer Solution assumption [3].

Their observation relies on the following lemma, which provides an HEC scheme with the efficient verification property.

**Lemma 10 (Adapted from [28,12,19]).** *Let  $C : \{0, 1\}^L \rightarrow \{0, 1\}$  be a NAND Boolean circuit of depth  $d$ . Let  $\mathbf{A}_i = \mathbf{A} \cdot \mathbf{R}_i + k_i \cdot \mathbf{G} \in \mathbb{Z}_q^{n \times m}$  with  $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ ,  $\mathbf{R}_i \in \{-1, 1\}^{m \times m}$  and  $k_i \in \{0, 1\}$ , for  $i \leq L$ . There exist deterministic algorithms  $\text{Eval}_{\text{CT}}^{\text{pub}}$  and  $\text{Eval}_{\text{CT}}^{\text{priv}}$  with running time  $\text{poly}(4^d, L, m, n, \log q)$  that satisfy:  $\text{Eval}_{\text{CT}}^{\text{pub}}(C, (\mathbf{A}_i)_i) = \mathbf{A} \cdot \text{Eval}_{\text{CT}}^{\text{priv}}(C, ((\mathbf{R}_i, k_i))_i) + C(k_1, \dots, k_L) \cdot \mathbf{G}$ , and  $\|\text{Eval}_{\text{CT}}^{\text{priv}}(C, (\mathbf{R}_i, k_i)_i)\| \leq 4^d \cdot O(m^{3/2})$ .*

After having committed to a string  $k_1 \dots k_L \in \{0, 1\}^L$  via matrices  $\{\mathbf{A}_i\}_{i=1}^L$ , the committer can compute  $\text{HEC}.\text{Eval}^{\text{in}}$ , which outputs  $\mathbf{R} \leftarrow \text{Eval}_{\text{CT}}^{\text{priv}}(C, ((\mathbf{R}_i, k_i))_i)$ . Anyone can run the  $\text{HEC}.\text{Eval}^{\text{out}}$  algorithm, which computes  $\text{com}_{ev}$  via  $\text{Eval}_{\text{CT}}^{\text{pub}}$ .

#### E.4 An $\mathcal{R}_{\text{NEQ}}$ -Lossy PKE Scheme from LWE

We now give an  $\mathcal{R}_{\text{NEQ}}$ -lossy PKE construction, which is obtained by adapting a scheme from [41] and removing its equivocation algorithms in lossy mode. In order to obtain the second indistinguishability property in the statistical sense, the construction uses  $\Theta(\lambda)$  matrices  $\mathbf{A}_i \in \mathbb{Z}_q^{n \times m}$  in the public key. If we relax the second indistinguishability property to hold in the computational sense (which is sufficient for our application), it is possible to build a more efficient scheme with  $O(1)$  public matrices using lossy trapdoor functions [51].

We rely on the homomorphic properties of GSW encodings via Lemma 10. We note that for  $t, K \in \{0, 1\}^L$ , checking whether  $t \neq K$  can be done using a Boolean circuit  $C_t : \{0, 1\}^L \rightarrow \{0, 1\}$  of depth  $d = O(\log L)$ , where  $t$  is hard-wired. Our construction goes as follows.

**Par-Gen**( $1^\lambda$ ): Given a security parameter  $\lambda \in \mathbb{N}$ , let  $B = \text{poly}(\lambda)$  the length of messages. Choose a prime modulus  $q = \text{poly}(\lambda)$ ; dimensions  $n = B + \Omega(\lambda)$  and  $m = 2n\lceil \log q \rceil + O(\lambda)$ . Define the tag space as  $\mathcal{T} = \{0, 1\}^L \setminus \{\mathbf{0}^L\}$  where  $L = \Theta(\lambda)$ . Define the initialization value space  $\mathcal{K} = \{0, 1\}^L$  and Gaussian parameters  $\sigma = 4^d \cdot O(m^2)$  and  $\alpha \in (0, 1)$  such that  $4^d m^4 \alpha q \sigma < q/4$ , for some fixed  $d = O(\log L)$ . Define public parameters as  $\Gamma = (L, B, q, n, m, \alpha, \sigma)$ .

**Keygen**( $\Gamma, K$ ): On input of public parameters  $\Gamma$  and an initialization value  $K \in \{0, 1\}^L$ , generate a key pair as follows.

1. Sample random matrices  $\bar{\mathbf{B}} \leftarrow U(\mathbb{Z}_q^{(n-B) \times m})$ ,  $\mathbf{S} \leftarrow U(\mathbb{Z}_q^{(n-B) \times B})$  and a small-norm  $\mathbf{E} \leftarrow \chi^{m \times B}$ , where  $\chi = D_{\mathbb{Z}, \alpha q}$ , to compute

$$\mathbf{A} = \begin{bmatrix} \bar{\mathbf{B}} \\ \mathbf{S}^\top \cdot \bar{\mathbf{B}} + \mathbf{E}^\top \end{bmatrix} \in \mathbb{Z}_q^{n \times m}.$$

2. Parse  $K$  as  $K_1 \dots K_L \in \{0, 1\}^L$ . Letting  $\mathbf{G} \in \mathbb{Z}_q^{n \times m}$  denote the gadget matrix, for each  $i \in [L]$ , compute matrices  $\mathbf{A}_i \in \mathbb{Z}_q^{n \times m}$  as

$$\mathbf{A}_i = \mathbf{A} \cdot \mathbf{R}_i + K_i \cdot \mathbf{G} \quad (6)$$

where  $\mathbf{R}_i \leftarrow U(\{-1, 1\}^{m \times m})$  for all  $i \in [L]$ .

Define  $R^{\text{RLE}} = \{\mathbf{r} \in \mathbb{Z}^{2m} \mid \|\mathbf{r}\| \leq \sigma\sqrt{2m}\}$  and output  $\text{sk} = (K, \mathbf{S})$  as well as

$$\text{pk} := (\mathbf{A}, \{\mathbf{A}_i\}_{i \in [L]}).$$

**LKeygen**( $\Gamma, K$ ): This algorithm proceeds identically to Keygen except that, in step 1, the algorithm generates a statistically uniform matrix  $\mathbf{A} \sim U(\mathbb{Z}_q^{n \times m})$ . Define  $R^{\text{RLE}}$  as in Keygen and output  $\text{sk} = (K, \mathbf{S})$  as well as

$$\text{pk} := (\mathbf{A}, \{\mathbf{A}_i\}_{i \in [L]}).$$

**Encrypt**( $\text{pk}, t, \text{Msg}$ ): To encrypt  $\text{Msg} \in \{0, 1\}^B$  for the tag  $t \in \mathcal{T}$ , conduct the following steps.

1. Let  $C_t : \{0, 1\}^L \rightarrow \{0, 1\}$  be the Boolean circuit of depth  $d = O(\log L)$ , where  $t$  is hard-wired, which inputs a string  $K = K_1 \dots K_L \in \{0, 1\}^L$  and outputs  $C_t(K)$  such that  $C_t(K) = 1 \Leftrightarrow t = K$ .

Using Lemma 10, compute  $\mathbf{A}_{C,t} \leftarrow \text{Eval}_{\text{CT}}^{\text{pub}}(C_t, (\mathbf{A}_i)_i) \in \mathbb{Z}_q^{n \times m}$  such that

$$\mathbf{A}_{C,t} = \mathbf{A} \cdot \mathbf{R}_t + C_t(K) \cdot \mathbf{G},$$

where  $\mathbf{R}_t = \text{Eval}_{\text{CT}}^{\text{priv}}(C_t, (\mathbf{R}_i, K_i)_i) \in \mathbb{Z}^{m \times m}$  s.t.  $\|\mathbf{R}_t\| \leq 4^d \cdot O(m^{3/2})$ .

2. Choose  $\mathbf{r} \leftarrow D_{\mathbb{Z}^{2m}, \sigma}$  and output  $\perp$  if  $\mathbf{r} \notin R^{\text{RLE}}$ . Otherwise, output

$$\mathbf{c} = [\mathbf{A} \mid \mathbf{A}_{C,t}] \cdot \mathbf{r} + \begin{bmatrix} \mathbf{0}^{n-B} \\ \text{Msg} \cdot \lfloor q/2 \rfloor \end{bmatrix} \in \mathbb{Z}_q^n. \quad (7)$$

**Decrypt**( $\text{sk}, t, \mathbf{c}$ ): Given the secret key  $\text{sk} = (K, \mathbf{S})$  and the tag  $t \in \mathcal{T}$ , return  $\perp$  if  $R_{\text{NEQ}}(K, t) = 0$ . Otherwise, compute  $\mathbf{w} = [-\mathbf{S}^\top \mid \mathbf{I}_B] \cdot \mathbf{c} \in \mathbb{Z}_q^B$ . For each  $i \in [B]$ , do the following:

1. If neither  $\mathbf{w}[i]$  nor  $|\mathbf{w}[i] - \lfloor q/2 \rfloor|$  is close to 0, halt and return  $\perp$ .
2. Otherwise, set  $\text{Msg}[i] \in \{0, 1\}$  so as to minimize  $|\mathbf{w}[i] - \text{Msg}[i] \cdot \lfloor q/2 \rfloor|$ .

Return  $\text{Msg} = \text{Msg}[1] \dots \text{Msg}[B]$ .

Lemma 11 states that the construction has the required properties under the LWE assumption. The proof is essentially the same as the proof of [41, Theorem 3.2] but we give it for completeness.

**Lemma 11.** *The above construction is an  $\mathcal{R}_{\text{NEQ}}$ -lossy public-key encryption scheme under the LWE assumption.*

*Proof.* We will prove that the scheme satisfies correct decryption with overwhelming probability in the injective mode. We also prove the first indistinguishability property using the LWE assumption, and the second in the statistical sense.

**Decryption under injective tags.** For any initialization value  $K \in \mathcal{K} = \{0, 1\}^L$ , any tag  $t \in \mathcal{T} = \{0, 1\}^L \setminus \{\mathbf{0}^L\}$  such that  $(K, t) \in \mathcal{R}_{\text{NEQ}}$ , any message  $\text{Msg} \in \{0, 1\}^B$ , and any encryption  $\mathbf{c} \in \mathbb{Z}_q^n$  of  $\text{Msg}$  under the  $\text{pk} = (\mathbf{A}, \{\mathbf{A}_i\}_{i \in [L]})$  and  $t$ , we have the following.

1.  $C_t(K) = 0$ , implying that  $\mathbf{A}_{C,t} = \mathbf{A} \cdot \mathbf{R}_t + C_t(K) \cdot \mathbf{G} = \mathbf{A} \cdot \mathbf{R}_t$ , where  $\|\mathbf{R}_t\| \leq 4^d \cdot O(m^{3/2})$  (by Lemma 10).
- 2.

$$\begin{aligned} \mathbf{w} &= [-\mathbf{S}^\top \mid \mathbf{I}_B] \cdot \mathbf{c} = [-\mathbf{S}^\top \mid \mathbf{I}_B] \cdot \mathbf{A} \cdot [\mathbf{I}_m \mid \mathbf{R}_t] \cdot \mathbf{r} + \text{Msg} \cdot \lfloor q/2 \rfloor \\ &= (-\mathbf{S}^\top \cdot \bar{\mathbf{B}} + \mathbf{S}^\top \cdot \bar{\mathbf{B}} + \mathbf{E}^\top) \cdot [\mathbf{I}_m \mid \mathbf{R}_t] \cdot \mathbf{r} + \text{Msg} \cdot \lfloor q/2 \rfloor \\ &= \mathbf{E}^\top \cdot [\mathbf{I}_m \mid \mathbf{R}_t] \cdot \mathbf{r} + \text{Msg} \cdot \lfloor q/2 \rfloor \in \mathbb{Z}_q^B. \end{aligned}$$

We show that, for any  $\mathbf{r} \in \mathbb{Z}^{2m}$  of norm smaller than  $\|\mathbf{r}\|_\infty \leq \|\mathbf{r}\| \leq \sigma\sqrt{2m}$ , we have  $\|\mathbf{E}^\top [\mathbf{I}_m \mid \mathbf{R}_t] \cdot \mathbf{r}\|_\infty < q/4$  with overwhelming probability over the randomness of  $\text{Keygen}$ , so that the decryption algorithm recovers the message. To prove this, note that our definition of the randomness space  $\mathcal{R}^{\text{RLE}}$  imposes  $\|\mathbf{r}\|_\infty \leq \|\mathbf{r}\| \leq \sigma\sqrt{2m}$ . Besides, we also have

$$\|[\mathbf{I}_m \mid \mathbf{R}_t]\|_\infty \leq 1 + \|\mathbf{R}_t\|_\infty \leq 4^d \cdot m^2$$

and

$$\|\mathbf{E}^\top\|_\infty = \max_{i \in [B]} \sum_{j=1}^m |e_{ij}| \leq \sqrt{m} \cdot \max_{i \in [n_0]} \sqrt{\sum_{j=1}^m e_{ij}^2} \leq m \cdot \alpha q$$

with overwhelming probability when  $\mathbf{E}^\top \leftarrow D_{\mathbb{Z}, \alpha q}^{B \times m}$ . Putting the above altogether, our choice of parameters implies that

$$\|\mathbf{E}^\top\|_\infty \cdot \|[\mathbf{I}_m \mid \mathbf{R}_{F,t}]\|_\infty \cdot \|\mathbf{r}\|_\infty \leq 4^d m^4 \alpha q \sigma < q/4 .$$

**Indistinguishability.** The key generation algorithm  $\mathsf{LKeygen}$  and  $\mathsf{Keygen}$  satisfy the following properties:

- (i) The  $\mathsf{LWE}$  assumption implies that, for any  $K \in \mathcal{K}$ , the two distributions  $D_{\text{inj}} = \{\mathsf{pk} \mid (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Keygen}(\Gamma, K)\}$  and  $D_{\text{loss}} = \{\mathsf{pk} \mid (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{LKeygen}(\Gamma, K)\}$  are computationally indistinguishable. These distributions only differ in the generation of the matrix  $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ . The matrix  $\mathbf{A}$  produced by the  $\mathsf{Keygen}$  algorithm is pseudorandom since, under the  $\mathsf{LWE}_{n-B, m, q, \chi}$  assumption, we can replace  $\mathbf{S}^\top \bar{\mathbf{B}} + \mathbf{E}^\top$  by a uniform matrix  $\mathbf{B} \sim U(\mathbb{Z}_q^{B \times m})$  without the adversary noticing. When using  $\mathsf{LKeygen}$ , the matrix  $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$  is statistically uniform by the properties of the  $\mathsf{GenTrap}$  algorithm (specifically, Lemma 7).
- (ii) For any initialization values  $K, K' \in \mathcal{K}$ , the two distributions  $\{\mathsf{pk} \mid (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{LKeygen}(\Gamma, K)\}$  and  $\{\mathsf{pk} \mid (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{LKeygen}(\Gamma, K')\}$  are statistically indistinguishable since the public matrices  $(\mathbf{A}, \{\mathbf{A}_i\}_{i \in [L]})$  are statistically uniform and independent regardless of which initialization value is used as input by  $\mathsf{LKeygen}$ . Recall that matrix  $\mathbf{A}$  produced by the  $\mathsf{LKeygen}$  algorithm is statistically close to  $U(\mathbb{Z}_q^{n \times m})$  by the properties of the  $\mathsf{GenTrap}$ . As for the matrices  $\mathbf{A}_i = \mathbf{A} \cdot \mathbf{R}_i + K_i \cdot \mathbf{G}$ , for  $K_i \in \{0, 1\}$ , the Left-over Hash Lemma implies that the statistical distance between the distributions  $\{(\mathbf{A}, \mathbf{A} \cdot \mathbf{R}_i) \mid \mathbf{A} \leftarrow U(\mathbb{Z}_q^{n \times m}), \mathbf{R}_i \leftarrow U(\{-1, 1\}^{m \times m})\}$  and  $\{(\mathbf{A}, \mathbf{A}_i) \mid \mathbf{A}, \mathbf{A}_i \leftarrow U(\mathbb{Z}_q^{n \times m})\}$  is smaller than  $m \cdot \sqrt{q^n/2^m} < 2^{-\lambda}$ , where the last inequality is implied by our choice of  $m = 2n \lceil \log q \rceil + O(\lambda)$ .

**Lossiness under lossy tags.** It is enough to prove that the distribution of a ciphertext obtained by encrypting under a lossy tag is statistically close to the uniform distribution on  $\mathbb{Z}_q^n$ .

For any initialization value  $K \in \mathcal{K}$  and tag  $t \in \mathcal{T}$  such that  $(K, t) \notin \mathcal{R}_{\mathsf{NEQ}}$ , any pair  $(\mathsf{pk} = (\mathbf{A}, \{\mathbf{A}_i\}_{i=1}^L), \mathsf{sk} = (K, \mathbf{S})) \leftarrow \mathsf{Keygen}(\Gamma, K)$ , and any message  $\mathsf{Msg} \in \{0, 1\}^B$ , an encryption of  $\mathsf{Msg}$  is generated as

$$\mathbf{c} = [\mathbf{A} \mid \bar{\mathbf{A}}_{C,t}] \cdot \mathbf{r} + \begin{bmatrix} \mathbf{0}^{n-B} \\ \mathsf{Msg} \cdot \lfloor q/2 \rfloor \end{bmatrix} \in \mathbb{Z}_q^n .$$

where  $\mathbf{r} \leftarrow D_{\mathbb{Z}^{2m}, \sigma}$  and  $\bar{\mathbf{A}}_{C,t} = [\mathbf{A} \mid \mathbf{A} \cdot \mathbf{R}_t + \mathbf{G}] \in \mathbb{Z}_q^{n \times 2m}$ . The matrix  $\bar{\mathbf{A}}_{C,t}$  is of this form because  $(K, t) \notin \mathcal{R}_{\mathsf{NEQ}}$ , and hence  $C_t(K) = 1$ . This implies that the columns of  $\bar{\mathbf{A}}_{C,t}$  generate  $\mathbb{Z}_q^n$ . By [46, Lemma 5.3], we know that  $\bar{\mathbf{A}}_{C,t}$  has a trapdoor  $\mathbf{T} \in \mathbb{Z}^{2m \times 2m}$  (namely, a short basis of the lattice  $\Lambda^\perp(\bar{\mathbf{A}}_{C,t})$ ) such that  $\|\tilde{\mathbf{T}}\| \leq (\|\mathbf{R}_t\| + 1) \cdot \sqrt{5}$  and thus  $\|\tilde{\mathbf{T}}\| \leq \sqrt{5} \cdot (4^d \cdot O(m^{3/2}) + 1)$ . Again, by

[27, Lemma 3.1], we know that  $\eta_{2-m}(\Lambda^\perp(\bar{\mathbf{A}}_{C,t})) \leq \|\tilde{\mathbf{T}}\| \cdot O(\sqrt{m})$ . By choosing  $\sigma = 4^d \cdot O(m^2)$ , we have  $\sigma \geq \eta_{2-m}(\Lambda^\perp(\bar{\mathbf{A}}_{C,t}))$ . By applying [27, Lemma 5.2], we conclude that  $\bar{\mathbf{A}}_{C,t} \cdot \mathbf{r}$  is statistically close to the uniform distribution  $U(\mathbb{Z}_q^n)$  when  $\mathbf{r} \leftarrow D_{\mathbb{Z}^{2m}, \sigma}$ .  $\square$

## F Efficiency Estimations for the Pairing-Based Instantiation

In this section, we provide an estimation of the signature size when we instantiate our generic construction in symmetric bilinear groups. The purpose is to offer a quick view of the practical instantiability of the resulting scheme in comparison with the lattice-based instantiation. Our estimation assumes symmetric pairings instead of more efficient asymmetric pairings for an equivalent security guarantee. This allows us to directly use the HEC with efficient verification due to [34], which is based on a falsifiable assumption (namely, the CDHER assumption), that holds in the generic group model over symmetric bilinear groups. In the following, we assume that  $(\mathbb{G}, \mathbb{G}_T, g, e, p)$  is a symmetric bilinear group for a prime  $p$ .

For the other primitives to instantiate, we have more common choices. In order to instantiate the  $\mathcal{R}_{\text{NEQ}}$ -lossy PKE scheme  $\Pi^{\text{RLE}}$ , we consider a variant of the DDH-based system in Section 2. To encrypt  $\text{Msg} \in \{0,1\}^B$  under the Decision Linear assumption (DLIN) [11], we need a public key containing  $\mathbf{u}, \mathbf{g}, \mathbf{g}_0, \mathbf{g}_1 \in \mathbb{G}^3$ , where  $\mathbf{u}$  is of the form  $\mathbf{u} = \mathbf{g}_0^{\beta_0} \cdot \mathbf{g}_1^{\beta_1} \cdot \mathbf{g}^{-K}$ . In lossy keys, we have  $\mathbf{g} \in \text{span}(\mathbf{g}_0, \mathbf{g}_1)$  while injective keys are generated so that  $\mathbf{g} \notin \text{span}(\mathbf{g}_0, \mathbf{g}_1)$ . Algorithm  $\text{Encrypt}(\text{pk}, t, \text{Msg})$  interprets the tag  $t \in \mathcal{T}$  as an element of  $\mathbb{Z}_p$  and, for each  $i \in [B]$ , it picks  $r_i, s_i \leftarrow U(\mathbb{Z}_p^*)$  and computes  $\text{ct}_i = (\mathbf{u} \cdot \mathbf{g}^t)^{\text{Msg}[i]} \mathbf{g}_0^{r_i} \mathbf{g}_1^{s_i} \in \mathbb{G}^3$ . The resulting ciphertext is now given by  $\text{ct} = (\text{ct}_1, \dots, \text{ct}_B) \in \mathbb{G}^{3B}$ .

For simplicity, we instantiate the standard lossy encryption scheme  $\Pi^{\text{lpke}}$  using the above bit-by-bit construction (although more efficient DLIN-based lossy PKE schemes exist) with the only difference that we can set  $t = 0$  since we do not have to deal with tags. Note that the ciphertext for a single bit corresponds to a Groth-Sahai (GS) commitment to a 0/1 scalar over  $\mathbb{Z}_p$  in the DLIN-based instantiation that will be used in our instantiation of the dual-mode NIZK argument system.

The signature scheme  $\Pi^{\text{sig}}$  can be instantiated with the structure-preserving scheme of Abe, Haralambiev and Ohkubo [1] (AHO), which was proven unforgeable under the  $q$ -SFP assumption. The signature is comprised of 7 elements of the source group  $\mathbb{G}$  while the verification algorithm only needs to check 2 pairing-product equations (PPEs). Eventually, as the user's verification key  $\text{pk}_{\text{id}}$  is signed by the GM during the  $\text{Join}$  protocol of our generic construction, we note that  $\text{pk}_{\text{id}}$  consists of  $2n + 10$  group elements when  $\mathbf{M} \in \mathbb{G}^n$ .

We are now ready to give a short description of the BiAS signature in the pairing-based instantiation of the generic construction. In order to sign a message  $M \in \{0,1\}^{\ell_m}$  using the witness  $w = w[1] \dots w[\ell_w] \in \{0,1\}^{\ell_w}$  w.r.t. the predicate

$P \in \mathcal{P}_d$ , the signing algorithm first computes the message-dependent Boolean circuit  $C_M : \{0, 1\}^{\ell_w} \times \{0, 1\}^\ell \rightarrow \{0, 1\}$  of depth  $\leq d$  that evaluates  $P(M, \mathbf{id}, w)$  on input of  $(w[1], \dots, w[\ell_w], \mathbf{id}[1], \dots, \mathbf{id}[\ell])$  and then conducts the following steps, where  $c_{ev} = C_M(w_1, \dots, w_{\ell_w}, \mathbf{id}_1, \dots, \mathbf{id}_\ell) \in \{0, 1\}$ .

1. We omit the details on the one-time signature where  $(\mathbf{VK}, \mathbf{SK}) \leftarrow \Pi^{\text{ots}}.\mathbf{Kg}(1^\lambda)$ , which is instantiable with short signatures under the discrete logarithm assumption (see, e.g., [30]).
2. Choose  $r_i, s_i \leftarrow U(\mathbb{Z}_p)$ , for each  $i = 1$  to  $\ell_w + \ell$ , and encrypt the identity-witness pair  $\mu = (\mathbf{id}, w) \in \{0, 1\}^{\ell + \ell_w}$  as an  $\mathcal{R}_{\text{NEQ}}$ -lossy encryption

$$\mathbf{ct}_{(\mathbf{id}, w)} = \left\{ \mathbf{ct}_i = (\mathbf{u} \cdot \mathbf{g}^{\mathbf{VK}})^{\mu[i]} \cdot \mathbf{g}_0^{r_i} \cdot \mathbf{g}_1^{s_i} \right\}_{i=1}^{\ell_w + \ell} \quad (8)$$

under the tag  $\mathbf{VK} \in \{0, 1\}^L$ .

3. Given the public elements  $\{V_{i,b}\}_{i \in [\ell + \ell_w], b \in \{0, 1\}}$  included in the HEC public parameters  $\mathbf{pp}$  of [34], sample  $r^{\text{hec}} \leftarrow U(\mathbb{Z}_p)$  and compute a commitment

$$\mathbf{com}_{(\mathbf{id}, w)} = g^{r^{\text{hec}}} \cdot \prod_{i \in [\ell + \ell_w]} V_{i, \mu[i]}. \quad (9)$$

4. Using the homomorphic evaluation algorithm of  $\mathbf{HEC}$ , compute

$$\begin{aligned} \pi_{C, M} &= (K_1, K_2, K_3) \leftarrow \mathbf{HEC}.\mathbf{Eval}^{in}(\mathbf{ek}, C_M, (\mathbf{id}, w), r^{\text{hec}}) \\ \mathbf{com}_{ev} &= (Z_0, Z_1) = \mathbf{HEC}.\mathbf{Eval}^{out}(\mathbf{ek}, C_M, \mathbf{com}_{(\mathbf{id}, w)}) \end{aligned}$$

which, for some public values  $A \in \mathbb{G}$  and  $T \in \mathbb{G}_T$ , satisfy

$$e(K_1, Z_{c_{ev}}) = e(K_3, g) \quad e(g, K_2) \cdot e(A, K_1) = T, \quad (10)$$

where  $T$  is efficiently computable from  $\ell$ ,  $\ell_w$  and  $d$ .

5. Choose  $r_{\mathbf{id}}, s_{\mathbf{id}} \leftarrow U(\mathbb{Z}_p)$  and compute

$$\mathbf{ct}_{\mathbf{id}} = \left\{ \mathbf{ct}_{\mathbf{id}, i} = \mathbf{v}^{(1 - c_{ev}) \cdot \mathbf{id}[i]} \cdot \mathbf{f}_0^{r_{\mathbf{id}, i}} \cdot \mathbf{f}_1^{s_{\mathbf{id}, i}} \right\}_{i=1}^\ell \quad (11)$$

where  $c_{ev} = C_M(w_1, \dots, w_{\ell_w}, \mathbf{id}_1, \dots, \mathbf{id}_\ell) \in \{0, 1\}$ .

6. Generate a structure-preserving AHO signature  $\sigma$  on the message  $g^{\mathbf{H}(M, P, \mathbf{ct}_{(\mathbf{id}, w)})}$ , where  $\mathbf{H}$  is a collision-resistant hash function. Note that  $n_{\text{user}} = 1$  and then  $\mathbf{pk}_{\mathbf{id}}$  contains 12 elements.
7. Generate a NIZK argument  $\pi \leftarrow \mathbf{NIZK}.\mathbf{Prove}(\rho, \mathbf{x}, \mathbf{w})$  for the statement  $\mathbf{x}$  that there exists a witness  $\mathbf{w}$  comprised of scalars  $(\mathbf{id}, w) \in \{0, 1\}^{\ell + \ell_w}$ ,  $r_{\mathbf{id}, w} = \{(r_i, s_i)\}_{i=1}^{\ell_w + \ell} \in \mathbb{Z}_p^{2\ell_w + 2\ell}$ ,  $r^{\text{hec}} \in \mathbb{Z}_p$ ,  $r^{\text{lpke}} = \{(r_{\mathbf{id}, i}, s_{\mathbf{id}, i})\}_{i=1}^\ell \in \mathbb{Z}_p^{2\ell}$ ,  $c_{ev} \in \{0, 1\}$ , and group elements  $(g^{\mathbf{id}}, \mathbf{pk}_{\mathbf{id}}, \mathbf{cert}_{\mathbf{id}}, \sigma) \in \mathbb{G} \times \mathbb{G}^{12} \times \mathbb{G}^7 \times \mathbb{G}^7$ ,  $Z_{c_{ev}} \in \mathbb{G}$  and  $\pi_{C, M} = (K_1, K_2, K_3) \in \mathbb{G}^3$ , which together satisfy the relations (8)-(11) as well as

$$\begin{aligned} \mathbf{AHO}.\mathbf{Verify}(\mathbf{pk}_{sig}, (g^{\mathbf{id}}, \mathbf{pk}_{\mathbf{id}}), \mathbf{cert}_{\mathbf{id}}) &= 1 \\ \mathbf{AHO}.\mathbf{Verify}(\mathbf{pk}_{\mathbf{id}}, g^{\mathbf{H}(M, P, \mathbf{ct}_{(\mathbf{id}, w)})}, \sigma) &= 1, \end{aligned} \quad (12)$$

where the GM's public key  $\mathbf{pk}_{sig}$  allows signing  $n_{\text{GM}} = 13$  group elements.

8. Compute  $sig \leftarrow \Pi^{\text{ots}}.\text{Sign}(\text{SK}, (\text{ct}_{(\text{id}, w)}, \text{com}_{(\text{id}, w)}, \text{ct}_{\text{id}}, \boldsymbol{\pi})).$

Return the signature

$$\Sigma = (\text{VK}, (\text{ct}_{(\text{id}, w)}, \text{com}_{(\text{id}, w)}, \text{ct}_{\text{id}}, \boldsymbol{\pi}), sig) \quad (13)$$

We note that we need two instantiations of the Groth-Sahai dual-mode argument system: One to prove the satisfiability of scalar product equations such as in (8), (9) and (11), where we have to commit to scalars; and another one to prove the satisfiability of PPEs like (10) and (12), where we have to commit to group elements. Furthermore, the argument  $\boldsymbol{\pi}$  must show the consistency between the committed values in both instances, namely, with respect to  $g^{\text{id}} = \prod_{i=1}^{\ell} (g^{\text{id}[i]})^{2^{i-1}}$  and  $\text{id}[1], \dots, \text{id}[\ell]$ , and to  $Z_{c_{ev}} = Z_0^{1-c_{ev}} \cdot Z_1^{c_{ev}}$  and  $c_{ev} = P(\text{id}, M, w)$ .

We can now more precisely evaluate the size of the BiAS signature in the pairing-based instantiation of the generic construction:

- The Groth-Sahai argument  $\boldsymbol{\pi}$  contains  $\ell_w + \ell$  commitments to the bits of  $\mu = (\text{id}, w) \in \{0, 1\}^{\ell + \ell_w}$  represented by scalars. Since commitments to scalars take 3 group elements each, we already have  $3\ell_w + 3\ell$  group elements. In order to prove that these scalars are bits, we need  $\ell_w + \ell$  proofs for quadratic scalar equations, which require 6 group elements each. This amounts to  $6\ell_w + 6\ell$  group elements. The commitments are compatible with the argument for multi-exponentiation multiplication equations such as (8), (9) and (11).
- The Groth-Sahai argument  $\boldsymbol{\pi}$  also contains 31 commitments to the (vector of) group elements  $g^{\text{id}}, \text{pk}_{\text{id}}, \text{cert}_{\text{id}}, \sigma, Z_{c_{ev}}, K_1, K_2, K_3$ , which leads to 93 group elements. Let  $\text{Com}_{\text{id}} = \iota(g^{\text{id}}) \cdot \mathbf{h}_1^{\rho_1} \cdot \mathbf{h}_2^{\rho_2} \cdot \mathbf{h}_3^{\rho_3}$  and  $\text{Com}_z = \iota(Z_{c_{ev}}) \cdot \mathbf{h}_1^{\eta_1} \cdot \mathbf{h}_2^{\eta_2} \cdot \mathbf{h}_3^{\eta_3}$  denote the commitments to  $g^{\text{id}}$  and  $Z_{c_{ev}}$  over  $\mathbb{G}^3$ , where  $\iota(X) = (1, 1, X) \in \mathbb{G}^3$
- We further need  $2\ell_w + 5\ell + 8$  commitments to scalars  $(r_i, s_i)_{i=1}^{\ell_w + \ell}, r^{\text{hec}}$ ,  $(r_{\text{id}, i}, s_{\text{id}, i})_{i=1}^{\ell}$ ,  $c_{ev}$  including the local variables  $\rho_1, \rho_2, \rho_3, \eta_1, \eta_2, \eta_3$  and  $\{\gamma_i = (1 - c_{ev}) \cdot \text{id}[i]\}_{i=1}^{\ell}$ . This introduces  $6\ell_w + 15\ell + 24$  additional group elements.
- Based on the above commitments, the proof that  $c_{ev}$  is a bit and that  $\gamma_i = (1 - c_{ev}) \cdot \text{id}[i]$ , for each  $i = 1$  to  $\ell$ , costs  $6(\ell + 1)$  group elements.
- Since the remaining  $4\ell_w + 7\ell + 6$  scalar relations (namely, (8), with  $3(\ell_w + \ell)$  equations, (9), with  $\ell_w + \ell$  equations and where  $V_{i, \mu_i} = V_{i, 0}^{1-\mu[i]} \cdot V_{i, 1}^{\mu[i]}$ , and (11) with  $3\ell$  equations, as well as  $\text{Com}_{\text{id}}$  and  $\text{Com}_z$ , with 3 equations for both) are linear, the total amount of group elements for this part raises to  $8\ell_w + 14\ell + 12$ .
- Finally, we have 6 PPEs from (10) and (12) out of which 3 are linear after a partial re-randomization of the variables. Then, we get  $3 \cdot 9 + 3 \cdot 3 = 36$  group elements to show that these PPEs are satisfied.

Overall, our pairing-based BiAS system incurs signatures comprised of about  $23\ell_w + 44\ell + 171$  group elements under the DLIN assumption. We note that it is possible to more efficiently instantiate the scheme in asymmetric pairings (which are known to offer more compact group elements for the same security level) using the SXDH-based variant of Groth-Sahai proof systems [32]. This requires

to adapt the HEC system of Katsumata *et al.* [34] in asymmetric pairings, which we believe possible.

Although our concrete pairing-based construction is not truly practical, it is already much more efficient than in an instantiation obtained from generic dual-mode NIZK proofs. As mentioned by Groth [30], generically applying the Groth-Ostrovsky-Sahai techniques [31] in the context of group signatures would translate into thousands of group elements per signature.