



Group signatures and more from isogenies and lattices: generic, simple, and efficient

Ward Beullens^{1,2} · Samuel Dobson³ · Shuichi Katsumata^{4,5} · Yi-Fu Lai³ · Federico Pintore⁶

Received: 26 June 2022 / Revised: 12 January 2023 / Accepted: 15 January 2023 /
Published online: 27 February 2023
© The Author(s) 2023

Abstract

We construct an efficient dynamic group signature (or more generally an accountable ring signature) from isogeny and lattice assumptions. Our group signature is based on a simple generic construction that can be instantiated by cryptographically hard group actions such as the CSIDH group action or an MLWE-based group action. The signature is of size $O(\log N)$, where N is the number of users in the group. Our idea builds on the recent efficient OR-proof by Beullens, Katsumata, and Pintore (Asiacrypt'20), where we efficiently add a proof of valid ciphertext to their OR-proof and further show that the resulting non-interactive zero-knowledge proof system is *online extractable*. Our group signatures satisfy more ideal security properties compared to previously known constructions, while simultaneously having an attractive signature size. The signature size of our isogeny-based construction is an order of magnitude smaller than all previously known post-quantum group signatures (e.g., 6.6 KB for 64 members). In comparison, our lattice-based construction has a larger signature

Communicated by A. Enge.

✉ Federico Pintore
federico.pintore@uniba.it

Ward Beullens
wbe@zurich.ibm.com

Samuel Dobson
samuel.dobson.nz@gmail.com

Shuichi Katsumata
shuichi.katsumata@pqshield.com

Yi-Fu Lai
ylai276@aucklanduni.ac.nz

¹ imec-COSIC, KU Leuven, Leuven, Belgium

² IBM Research, Zurich, Switzerland

³ University of Auckland, Auckland, New Zealand

⁴ National Institute of Advanced Industrial Science and Technology (AIST), Tokyo, Japan

⁵ PQShield Ltd., Oxford, UK

⁶ Department of Mathematics, University of Bari, Bari, Italy

size (e.g., either 126 KB or 89 KB for 64 members depending on the satisfied security property). However, since the $O(\cdot)$ -notation hides a very small constant factor, it remains small even for very large group sizes, say 2^{20} .

Keywords Isogeny-based cryptography · Lattice-based cryptography · Linkable ring signature · Post-quantum cryptography

Mathematics Subject Classification 11T71 · 11Y40 · 13P25 · 14H52 · 14K02 · 94A60 · 94A62

1 About

An extended abstract of this work was published in EUROCRYPT 2022 [13]. This is a full version of the paper. In more details, we provide a comparison of this work with two recent concurrent and independent papers which propose isogeny-based group signatures; we prove correctness and security of the proposed generic accountable ring signature Π_{ARS} in Fig. 1; we describe in details the tight variant $\Pi_{\text{ARS}}^{\text{Tight}}$ in Fig. 2 and we prove its security; we prove correctness and security of the traceable OR sigma protocol $\Pi_{\Sigma}^{\text{base}}$ in Fig. 3; we detail how we incorporate three optimisations to the modification of $\Pi_{\Sigma}^{\text{base}}$ which enjoys a negligibly-small soundness error, and we prove its security; we give a full description of the base traceable OR sigma protocol $\Pi_{\Sigma}^{\text{baseTi}}$ for the “tight” relation $R_{\text{sig}}^{\text{Tight}}$ in Fig. 6; we formally show that the NIZK with labels $\Pi_{\text{NIZK, lbl}}$ in Fig. 7 is zero-knowledge. In Sect. 8, the building blocks required for our generic construction are instantiated via isogenies-by using the CSIDH group action—and lattices. Finally, we recall how accountable ring signatures can be naturally viewed as group signatures in the Appendix.

2 Introduction

Group signature schemes, introduced by Chaum and van Heyst [31], allow authorized members of a group to individually sign on behalf of the group while the specific identity of the signer remains anonymous. However, should the need arise, a special entity called the group manager (or sometimes the tracing authority) can trace the signature to the signer, thus holding the group members accountable for their signatures. Group signatures have been an active area of academic research for the past three decades, and have also been gathering practical attention due to the recent real-world deployment of variants of group signatures such as directed anonymous attestation (DAA) [23] and enhanced privacy ID (EPID) [24].

Currently, there are versatile constructions of *efficient* group signatures from *classical* assumptions, e.g., [6, 16, 17, 34, 39, 40, 55, 58, 71, 72]. In this work, when we argue the efficiency of a group signature, we focus on one of the quintessential metrics: the signature size. We require it to be smaller than $c \cdot \log N$ bits, where N is the group size and c is some explicit small polynomial in the security parameter. In their seminal work, Bellare, Micciancio, and Warinschi [8] provided a generic construction of a group signature with signature size $O(1)$ from any signature scheme, public-key encryption scheme, and general non-interactive zero-knowledge (NIZK) proof system. Unfortunately, this only provides an asymptotic feasibility result, and thus one of the main focuses of subsequent works, including ours, has been to construct a concretely efficient group signature.

In contrast to the classical setting, constructing efficient group signatures from any *post-quantum* assumptions has been elusive. Since the first lattice-based construction by Gordon, Katz, and Vaikuntanathan [57], there has been a rich line of subsequent works on lattice-based (and one code-based) group signatures, including but not limited to [49, 59, 65, 70, 74]. However, these results remained purely asymptotic. It was not until recently that efficient lattice-based group signatures appeared [22, 38, 47, 48]. In [47], Esgin et al. report a signature size of 12KB and 19KB for a group size of $N = 2^6$ and 2^{10} , respectively—several orders of magnitude better than prior constructions.¹ These rapid improvements in efficiency for lattices originate in the recent progress of lattice-based NIZK proof systems for useful languages [5, 21, 45, 46, 75, 76, 85], most of which rely heavily on the properties of special structured lattices. Thus, it seems impossible to import similar techniques to other post-quantum assumptions or to standard non-structured lattices. For instance, constructing efficient group signatures from isogenies—one of the promising alternative post-quantum tools to lattices—still seems out of reach using current techniques. This brings us to the main question of this work:

Can we construct an efficient group signature secure from isogenies? Moreover, can we have a generic construction that can be instantiated from versatile assumptions, including those based on less structured lattices?

In addition, as we discuss in more detail later, we notice that all works regarding efficient post-quantum group signatures [22, 38, 47, 48, 60] do not satisfy the ideal security properties (which are by now considered standard) formalized by Bootle et al. [20]. Thus, we are also interested in the following question:

Can we construct efficient post-quantum group signatures satisfying the ideal security properties formalized by Bootle et al. [20]?

To address these questions, in this work we focus on *accountable ring signatures* [84]. An accountable ring signature offers the flexibility of choosing the group of users when creating a signature (like a ring signature [82]), while also enforcing accountability by including one of the openers in the group (like a group signature). Although research on accountable ring signatures is still limited [19, 48, 62, 68, 84], we advocate that they are as relevant and interesting as group and ring signatures. As shown by Bootle et al. [19], accountable ring signatures imply group and ring signatures by naturally limiting or downgrading their functionality. Thus, an efficient post-quantum solution to an accountable ring signature implies solutions for *both* secure (dynamic) group signatures [9] and ring signatures, making it an attractive target to focus on.

Finally, as an independent interest, we are also concerned with *tightly*-secure constructions. To the best of our knowledge, all prior efficient post-quantum secure group and ring signatures are in the random oracle model and have a very loose reduction loss. In typical security proofs, given an adversary with advantage ϵ that breaks some security property of the group signature, we can only construct an adversary with advantage at most $(N^2Q)^{-1} \cdot \epsilon^2$ against the underlying hard problem, where Q is the number of random oracle queries and N is the number of users in the system. If we aim for 128-bit security (i.e., $\epsilon = 2^{-128}$), and set for example $(N, Q) = (2^{10}, 2^{50})$, then we need at least 326-bits of security for the hard problem. When aiming for a provably-secure construction, the parameters must be set much larger to compensate for this significant reduction loss, which then leads to a less efficient

¹ We note that their signature size grows by $\log^t N$ for a small constant $t > 1$ rather than simply by $\log N$.

scheme. This is especially unattractive in the isogeny setting since only the smallest among the CSIDH parameters [28] enjoys properties suitable to achieve concrete efficiency [15].

2.1 Our contribution

In this work, we construct an efficient accountable ring signature based on isogenies and lattices. This in particular implies the first efficient isogeny-based group signature. Our generic construction departs from known general feasibility results such as [8] and builds on primitives that can be efficiently instantiated. Unlike previous efficient post-quantum group signatures, our scheme satisfies all the desired properties provided by Bootle et al. [20] including *dynamicity* and *fully (CCA) anonymity*: the former states that the group members can be added and revoked dynamically and are not fixed on setup; the latter states that anonymity holds even in the presence of an adversary that sees the signing keys of all honest users, who is additionally granted access to an opening oracle. We also satisfy the ideal variant of *non-frameability* and *traceability* [20], where the former is captured by *unforgeability* in the context of accountable ring signature. Roughly, this ensures that arbitrary collusion among members, even with the help of a corrupted group manager, cannot falsely open a signature to an honest user.

Our accountable ring signature schemes are realized in three steps. We first provide a generic construction of an accountable ring signature from simple cryptographic primitives such as a public-key encryption (PKE) scheme and an accompanying NIZK for a specific language. We then show an efficient instantiation of these primitives based on a group action that satisfies certain cryptographic properties. Finally, we instantiate the group action by either the CSIDH group action or the MLWE-based group action. Our generic construction builds on the recent efficient OR-proofs for isogeny and lattice-based hard languages by Beullens, Katsumata, and Pintore [14], which were used to construct ring signatures. The most technical part of this work is to efficiently add a proof of valid ciphertext to their OR-proof and proving full anonymity, which done naively would incur an exponential security loss. At the core of our construction is an efficient *online-extractable* OR-proof that allows to also prove validity of a ciphertext.

Moreover, thanks to the online extractability, our construction achieves a much tighter reduction loss compared to prior accountable ring signatures (and also group and ring signatures). It suffices to assume that the underlying post-quantum hard problem cannot be solved with advantage more than $N^{-1} \cdot \epsilon$ rather than $(N^2 Q)^{-1} \cdot \epsilon^2$ as in prior works whose proofs rely on the forking lemma [50, 81]. Working with the above example, we only lose 10-bits rather than 198-bits of security. We further show how to remove N^{-1} using the Katz-Wang technique [61] along with some techniques unique to our NIZK. As a side product, we obtain a tightly-secure and efficient isogeny and lattice-based ring signatures, improving upon those by Beullens et al. [14] which have a loose security reduction.

Comparison to prior work. To the best of our knowledge, Esgin et al. [47, 48] are the only other work that (implicitly) provide an efficient post-quantum accountable ring signature.² Since the efficiency of an accountable ring signature is equivalent to those of the group signature obtained through limiting the functionality of the accountable ring signature, we chose to compare the efficiency of our scheme with other state-of-the-art post-quantum group signatures. Table 1 includes a comparison of the signature size and the different notions of security it satisfies. The first two schemes satisfy all the desired security properties of a

² To be precise, they consider a weaker variant of standard accountable ring signature where no Judge algorithm is considered.

Table 1 Comparison of the signature size (KB) of some concretely efficient post-quantum group signature schemes. The first three rows are our scheme

	2	2 ⁵	$\frac{N}{2^6}$	2 ¹⁰	2 ²¹	Hardness Assumption	Security Level	Anonymity	Manager Account
Isogeny	3.6	6.0	6.6	9.0	15.5	CSIDH-512	*	CCA	Yes
Lattice	124	126	126	129	134	MSIS/MLWE	NIST 2	CCA	Yes
Lattice	86	88	89	91	96	MSIS/MLWE	NIST 2	CCA	No
[47]	/	12	/	19	/	MSIS/MLWE	NIST 2	CPA	No
[60]	/	/	280	418	/	LowMC	NIST 5	selfless-CCA	No

* 128 bits of classical security and 60 bits of quantum security [80]. Manager accountability states whether the (possibly malicious) group manager is accountable when opening a signature to some user. Namely, it is “Yes” when even a malicious group manager cannot falsely accuse an honest user of signing a signature that it hasn’t signed

dynamic group signature formalized by Bootle et al. [20]. Our scheme is the only one to achieve full CCA anonymity. Esgin et al. [47] achieves full CPA anonymity, where anonymity is broken once an adversary is given access to an opening oracle; in practice, this means that if a specific signature is once opened to some user, then any signature ever signed by that particular user will lose anonymity. Here, “full” means that the signing key of all the users may be exposed to the adversary. In contrast, Katz, Kolesnikov, and Wang [60] satisfies *selfless* CCA anonymity. While their scheme supports opening oracles, anonymity no longer holds if the signing key used to sign the signature is exposed to the adversary. Moreover, our scheme is the only one that also achieves the ideal variant of non-frameability and traceability [9, 20] (illustrated in the “Manager Accountability” column). The two schemes [47, 60] assume the group manager honestly executes the opening algorithm and that everyone trusts the output. Put differently, a malicious group manager can frame any honest members in the group by simply replacing the output of the opening algorithm. In contrast, our scheme remains secure even against malicious group managers since the validity of the output of the opening algorithm is verifiable. That is, even the group manager is held *accountable* in our group signature.

Not only our group signatures satisfy more ideal security properties compared to previous constructions, Table 1 shows that our signature size remains competitive. Our isogeny-based group signature based on CSIDH provides the smallest signature size among all post-quantum group signatures, which is $0.6 \log_2(N) + 3$ KB. In contrast, our lattice signature is larger; the scheme in the second (resp. third) row has signature size $0.5 \log_2(N) + 123.5$ KB (resp. $0.5 \log_2(N) + 85.9$ KB). It is smaller compared to [60], while larger compared to [47]. Compared to the two constructions, our signature size grows much slower with the group size N (see also Footnote 1) and also satisfies stronger security. We thus leave it as an interesting open problem to lower the constants in our construction.

2.2 Technical overview

An accountable ring signature is like a standard ring signature where the ring R also includes an arbitrary opener public key opk of the signer’s choice when creating a signature σ . The signature σ remains anonymous for anybody who does not know the corresponding opener secret key osk , while the designated opener can use osk to trace the user who created σ . A ring signature can be thought of as an accountable ring signature where $\text{opk} = \perp$, while a

group signature can be thought as an accountable ring signature where there is only a single opener.

General approach. Our generic construction of an accountable ring signature follows the well-known template of the encrypt-then-prove approach to construct a group signature [27]. The high-level idea is simple. The signer encrypts its verification key vk (or another unique identifier) using the opener's public key opk for a PKE scheme and provides a NIZK proof for the following three facts: the ciphertext ct encrypts vk via opk ; vk is included in the ring R ; and that it knows a secret key sk corresponding to vk . To trace the signer, the opener simply decrypts ct to recover vk . Notice that the NIZK proof implicitly defines a verifiable encryption scheme [25, 26] since it is proving that ct is a valid encryption for some message vk in R . Below, although our construction can be based on any cryptographically-hard group action, we mainly focus on isogenies for simplicity.

One of the difficulties in instantiating this template using isogeny-based cryptography is that we do not have an efficient verifiable encryption scheme for an appropriate PKE scheme. To achieve full anonymity, most of the efficient group signatures, e.g., [38, 39, 55, 58, 71, 72], use an IND-CCA secure PKE as a building block and construct an efficient NIZK that proves validity of the ciphertext. Full anonymity stipulates that an adversary cannot de-anonymize a signature even if it is provided with an opening oracle, which traces the signatures submitted by the adversary. Roughly, by using an IND-CCA secure PKE as a building block, the reduction can simulate the opening oracle by using the decapsulation oracle provided by the IND-CCA game, rather than the opener's secret key. In the classical setting, constructing such an efficient IND-CCA secure verifiable encryption scheme is possible using the Cramer-Shoup PKE [35] that offers a rich algebraic structure. Unfortunately, in the isogeny setting, although we know how to construct an IND-CCA secure PKE based on the Fujisaki-Okamoto transform [53], it seems quite difficult to provide an accompanying verifiable encryption scheme as the construction internally uses a hash function modeled as a random oracle. Another approach is to rely on the weaker IND-CPA secure PKE but to use a stronger NIZK satisfying *online-extractability* [52]. At a high level, the reduction can use the online-extractor to extract the witness in the ciphertext ct instead of relying on the decapsulation oracle.³ However, it turns out that even this approach is still non-trivial since we do not have any efficient verifiable encryption scheme for existing isogeny-based PKEs, let alone an accompanying online-extractable NIZK. For instance, most isogeny-based IND-CPA secure PKEs are based on the *hashed* version of ElGamal, and to the best of our knowledge, there are no efficient verifiable encryption schemes for hashed ElGamal.

Verifiable encryption scheme for a limited class of PKE. In this work, we observe that in the context of accountable ring signatures and group signatures, we do not require the full decryption capability of a standard PKE. Observe that decryption is only used by the opener and that it *knows* the ciphertext ct must be an encryption of one of the verification keys included in the ring (or group) R . Therefore, given a ciphertext ct , we only require a mechanism to check if ct encrypts a particular message M , rather than being able to decrypt an arbitrary unknown message. Specifically, the opener can simply run through all the verification keys $vk \in R$ to figure out which vk was encrypted in ct . This allows us to use a simple IND-CPA secure PKE with limited decryption capability based on the CSIDH group action: Let $E_0 \in \mathcal{Ell}_p(\mathcal{O}, \pi)$ be a fixed and public elliptic curve. The public key is $pk = (E_0, E := s \star E_0)$, where $sk = s$ is sampled uniformly at random from the class group $\mathcal{Cl}(\mathcal{O})$. To encrypt a message $M \in \mathcal{Cl}(\mathcal{O})$, we sample $r \leftarrow \mathcal{Cl}(\mathcal{O})$ and set $ct = (ct_0 := r \star E_0, ct_1 := M \star (r \star E))$.

³ Note that extractability via rewinding is insufficient for full anonymity as it will cause an exponential reduction loss when trying to extract the witness from adaptively chosen signatures [11].

To check if ct decrypts to M' , we check whether ct_1 is equal to $M' \star (sk \star ct_0)$. Note that in general we cannot decrypt when M is unknown since we cannot cancel out $sk \star ct_0$ from ct_1 . Now, observe that proving ct encrypts $M \in \mathcal{C}\ell(\mathcal{O})$ is easy since there is a simple sigma protocol for the Diffie-Hellman-like statement $(ct_0, (-M) \star ct_1) = (r \star E_0, r \star E)$, where r is the witness, e.g., [43]. Although this comes closer to what we want, this simple sigma protocol is not yet sufficient since the prover must reveal the message M to run it. Specifically, it proves that ct is an encryption of M , while what we want to prove is that ct is an encryption of *some* $M \in R$. In the context of accountable ring signature and group signature, this amounts to the signer being able to hide its verification key $vk \in R$.

Constructing NIZK for accountable ring signature Let us move forward to the intermediate goal of constructing a (non-online-extractable) NIZK proof system for the following three facts: the ciphertext ct encrypts vk via pk ; vk is included in the ring R ; and that the prover knows a secret key sk corresponding to vk . Recently, Beullens, Katsumata, and Pintore [14] proposed an efficient sigma protocol (and a non-online-extractable NIZK via the Fiat-Shamir transform) for proving the last two facts, which in particular constitutes an efficient OR-proof. We show how to glue the above “weak” verifiable encryption scheme with their OR-proof.

We first review a variant of the OR-sigma protocol in [14] with proof size $O(N)$, where N is the size of the ring. Assume each user $i \in [N]$ in the ring holds $vk_i = (E_0, E_i := s_i \star E_0) \in \mathcal{E}\ell\ell p(\mathcal{O}, \pi)^2$ and $sk_i = s_i \in \mathcal{C}\ell(\mathcal{O})$. To prove $vk_I \in R$ and that it knows sk_I , the prover first sample $s' \leftarrow \mathcal{C}\ell(\mathcal{O})$ and sets $R_i = s' \star E_i$ for $i \in [N]$. It also samples randomness $rand_i$ and creates commitments $(C_i = \text{Com}(R_i, rand_i))_{i \in [N]}$, where this commitment is simply instantiated by a random oracle. It finally samples a random permutation ϕ over $[N]$ and sends a permuted tuple $(C_{\phi(i)} = \text{Com}(R_i, rand_i))_{i \in [N]}$. The verifier samples a random bit $b \in \{0, 1\}$. If $b = 0$, the prover returns all the randomness $(s', (rand_i)_{i \in [N]}, \phi)$ used to create the first message. The verifier then checks if the first message sent by the prover is consistent with this randomness. Otherwise, if $b = 1$, the prover returns $(I'', rand'', s'') := (\phi(I), rand_I, s' + s_I)$. The verifier then checks if $C_{I''} = \text{Com}(s'' \star E_0, rand'')$ holds. Notice that if the prover is honest, then $s'' \star E_0 = s' \star E_I$ as desired. It is easy to check it is honest-verifier zero-knowledge. The transcript when $b = 0$ is independent of the witness, while the transcript when $b = 1$ can be simulated if the commitment scheme is hiding. Moreover, special soundness can be checked by noticing that given s'' and s' , we can extract some (i^*, s^*) such that $(E_0, E_{i^*} = s^* \star E_0) \in R$. A full-fledged OR-sigma protocol with proof size $O(N)$ is then obtained by running this protocol λ -times in parallel, where λ denotes the security parameter. [14] showed several simple optimization techniques to compress the proof size from $O(N)$ to $O(\log N)$, but we first explain our main idea below.

We add our “weakly decryptable” PKE to this OR-sigma protocol. Since our PKE only handles messages in $\mathcal{C}\ell(\mathcal{O})$, the prover with $vk_I \in R$ encrypts the index $I \in [N]$ rather than vk_I , where we assume the verification keys in the ring R are ordered lexicographically.⁴ The statement now consists of the ring R and the ciphertext $ct = (ct_0 := r \star E_0, ct_1 = I \star (r \star E))$, where (E_0, E) is the opener’s public key opk . Recall the opener can decrypt ct with knowledge of the ring R by brute-force searching for an $i \in [N]$ such that $ct_1 = i \star (osk \star ct_0)$. Now, to prove vk_I is an entry in R and that it knows sk_I , the prover samples $s' \leftarrow \mathcal{C}\ell(\mathcal{O})$ and sets $R_i = s' \star E_i$ for $i \in [N]$ as before. It then further samples $r' \leftarrow \mathcal{C}\ell(\mathcal{O})$ and prepares $ct'_i = (r' \star ct_0, (-i) \star (r' \star ct_1))$ for all $i \in [N]$. Observe that ct'_i is an encryption of the message $(I - i)$ using randomness $(r' + r)$. Specifically, ct'_I is of the form $((r' + r) \star E_0, (r' + r) \star E)$, which admits a natural sigma protocol as explained above. Finally,

⁴ The choice of what to encrypt is rather arbitrary. The same idea works if for instance we hash vk into $\mathcal{C}\ell(\mathcal{O})$ and view the digest as the message.

the prover samples randomness rand_i and a random permutation ϕ over $[N]$, and sends the randomly permuted commitments $(C_{\phi(i)} = \text{Com}(R_i \parallel \text{ct}'_i, \text{rand}_i))_{i \in [N]}$. The verifier samples a random bit $b \in \{0, 1\}$. If $b = 0$, then similarly to the above OR-sigma protocol, the prover simply returns all the randomness and the verifier checks the consistency of the first message. Otherwise, if $b = 1$, the prover returns $(I'', \text{rand}'', s'', r'') := (\phi(I), \text{rand}_I, s' + s_I, r' + r)$. The verifier checks if $C_{I''} = \text{Com}(s'' \star E_0 \parallel (r'' \star E_0, r'' \star E), \text{rand}'')$ holds. Correctness and honest-verifier zero-knowledge holds essentially for the same reason as the above OR-sigma protocol. More importantly, special soundness holds as well. Intuitively, since the opening to $b = 0$ forces the cheating prover to commit to the proper (vk_i, i) -pair, a cheating prover cannot encrypt an index I' and prove that it has sk_I corresponding to vk_I for a different $I \neq I'$.

To compile our sigma protocol into an NIZK, we apply the Fiat-Shamir transform. Moreover, we apply similar optimization techniques used in [14] to compress the proof size from $O(N)$ to $O(\log N)$. Roughly, the prover additionally uses a pseudorandom generator to generate the randomness (i.e., $s', r', \phi, (\text{rand}_i)_{i \in [N]}$). Then, in case $b = 0$, the prover needs to reply only with the seed of size $O(1)$. The prover also uses a Merkle tree to accumulate $(C_{\phi(i)})_{i \in [N]}$ and sends the root value in the first message. It then only opens to the path necessary for verification when $b = 1$. This has a positive side-effect that we no longer require a permutation ϕ since the path hides the index if we use a slightly tweaked variant of the standard Merkle tree. Finally, we take advantage of the asymmetry in the prover's response size for $b = 0$ and $b = 1$, which are $O(1)$ and $O(\log N)$, respectively. Namely, we imbalance the challenge space so that the prover opens to more 0 than 1, while still maintaining negligible soundness error.

Adding online-extractability. To build an accountable ring signature or group signature, we require the above NIZK to be (multi-proof) *online-extractable*. This is a strengthening of standard proof of knowledge (PoK) that roughly states that the knowledge extractor, who can see what the adversary queries to the random oracle, is able to directly extract witnesses from the proofs output by the adversary. The OR-proof by [14], which our NIZK builds on, was only shown to satisfy the standard PoK, which bases on a *rewinding* extractor.

One simple way to add online-extractability to our NIZK is to apply the Unruh transform [83]. Namely, we can modify the prover to add two more commitments $h_0 = \text{Com}(s' \parallel r', \text{rand}_0)$ and $h_1 = \text{Com}(s'' \parallel r'', \text{rand}_1)$ in the first message, where Com is instantiated by the random oracle. Then, if $b = 0$ (resp. $b = 1$), the prover further opens to h_0 (resp. h_1). Recall that if the reduction obtains both (s', r') and (s'', r'') , then it can invoke the extractor provided by the underlying sigma protocol to extract some (i^*, s^*) such that $(E_0, E_{i^*} = s^* \star E_0) \in R$. Therefore, for the cheating adversary to fool the reduction, it must guess the bit b and create h_b correctly while creating h_{1-b} arbitrary. Intuitively, if we have λ -repetition of the sigma protocol, then the cheating prover cannot possibly guess all the challenge bits correctly. Therefore, there must be some challenge where it created h_0 and h_1 honestly. For that challenge bit, the reduction algorithm can then retrieve the corresponding inputs $(s' \parallel r', \text{rand}_0)$ and $(s'' \parallel r'', \text{rand}_1)$ from simply observing the random oracle, and then, run the extractor to obtain the witness.

This idea works but it comes with an extra two hashes per one execution of the binary-challenge sigma protocol. Although it may sound insignificant in an asymptotic sense, these hashes add up when we execute the sigma protocol many times, and it makes it difficult to apply some of the optimization tricks. Concretely, when we apply this change to the isogeny-based ring signature by Beullens et al. [14], the signature grows by roughly a factor of 2 to 3.

In this work, we show that we can in fact prove online-extractability *without* making any modification to the aforementioned NIZK. Our main observations are the following: if the prover uses a seed to generate the randomness used in the first message via a random oracle, then the online extractor can observe $(s', r', \phi, (\text{rand}_i)_{i \in [N]})$; and the prover must respond to some execution of the binary-challenge sigma protocol where the challenge bit is 1. The first implies that the seed implicitly acts as a type of commitment to (s', r') . The second implies the prover returns a response that includes (s'', r'') . Specifically, our online extractor only looks at all the responses for the rounds where the challenge bit was 1, and checks the random oracle for any seed that leads to the commitment provided in the first message of the sigma protocol. If such seed is found, then it succeeds in extracting a witness. The intuition is simple but it turns out that the formal proof is technically more complicated due to the several optimizations performed on the basic sigma protocol to achieve proof size $O(\log N)$.

Generalizing with group actions. Although we have been explaining our generic construction using the CSIDH group action, it is not unique to them. It works equally well for any group action that naturally induces a PKE. Specifically, we instantiate the above idea also by the MLWE group action defined roughly as $\star : R_q^{n+m} \times R_q^m : (s, \mathbf{e}) \star \mathbf{t} \rightarrow \mathbf{A} \star s + \mathbf{e} + \mathbf{t}$, where $R_q = \mathbb{Z}_q[X]/(X^d + 1)$. Since CSIDH and MLWE induce a PKE with slightly different algebraic structures, we introduce a *group-action-based* PKE defined by two group actions to formally capture both instances. This abstraction may be of an independent interest since at first glance, isogeny-based and lattice-based PKEs seem to rely on different algebraic structures. Finally, one interesting feature unique to our generic construction is that since our sigma protocol is rather combinatorial in nature, we can for instance use CSIDH for the user's public key vk and mix it with an MLWE-based PKE for the opener's public key opk . The practical impact of such mixture is that we can achieve stronger bit-security for anonymity (due to MLWE) while keeping the user's public key and signature small (due to CSIDH).

Achieving tight reduction. Since the proofs do not rely on the forking lemma [50, 81] to extract witnesses from the forged proofs, our construction achieves a tighter reduction compared to prior works on efficient group signatures. However, we still lose a factor $1/N$ in the proof of unforgeability, which may vary from $1/2$ to $1/2^{20}$.⁵ Recall N is the size of the group in group signatures but it is the size of all the users enrolled in the system for accountable ring signatures, which may be far larger than the size of the ring. The main reason for this loss was because the reduction needs to guess one user's verification key used by the adversary to create its forgery and to embed the hard problem into it.

A well known technique to obtain a tight proof is to rely on the Katz-Wang technique [61] along with the generic OR-composition of sigma protocols, and rely on a multi-instance version of the hard problem (which are believed to be as difficult as the single-instance version for specific hard problems). Namely, we modify the scheme to assign two verification keys $(\text{vk}^{(1)}, \text{vk}^{(2)})$ to each user. The users will only hold one signing key $\text{sk}^{(b)}$ for $b \in \{1, 2\}$ corresponding to the verification key $\text{vk}^{(b)}$. The user can honestly run the aforementioned sigma protocol where the statement includes $\text{vk}^{(b)}$, and a simulated sigma protocol using the ZK-simulator where the statement includes $\text{vk}^{(3-b)}$. We can then use the sequential OR-proof technique as presented in [1, 51] to bridge these two sigma protocols so that it hides the b .⁶

While this generic transform works, it unfortunately doubles the signature size, which may outweigh the motivation for having a tight reduction. In this work, we present a novel

⁵ We note that we also have some independent looseness in the anonymity proof since we rely on the "multi-challenge" IND-CPA security from our PKE. This is handled in a standard way, and this is also why we only achieve a truly tight group signature from lattices and not from isogenies.

⁶ We note that it seems difficult to use the parallel OR-proof for our sigma protocol since the challenge space is structured.

Table 2 Comparison of this work with concurrent works [33, 66] in terms of signature size, anonymity and manager accountability. The integer N represents the size of the group

Schemes	Signature size	Anonymity	Manager Accountable
[66]	$\mathcal{O}(N \log(N))$	CPA	No
[33]	$\mathcal{O}(N^2)$	CPA	Partial
This Work	$\mathcal{O}(\log(N))$	CCA	Yes

and far cheaper technique tailored to our sigma protocol. The signature size overhead is a mere 512B for our concrete lattice-based instantiation. The key observation is that we can view the set of all users' verification key $(vk^{(1)}, vk^{(2)})$ as a ring of size $2N$, rather than a ring of size N where each ring element consists of two verification keys. This observation itself is not yet sufficient since recall that we typically must encrypt some information bound to the signer for traceability, e.g., encrypt the position/index of vk in R , and it is no longer clear what to encrypt when we have two verification keys in the ring. Luckily, it turns out that our sigma protocol can be easily modified with no loss in efficiency to overcome this apparent issue. Details are provided in Sect. 6.3.

Concurrent works. There are two concurrent and independent works published on the Cryptology ePrint Archive [33, 66]. Both of these works obtain isogeny-based group signatures from variants of ring signatures. Lai et al. [66] base their group signature on *revocable* ring signatures, while Chung et al. [33] base theirs on accountable ring signatures. Chug et al. follow the security properties formalized in [84], in contrast with our accountable ring signature which follows those formalized in [20]. Since the security properties defined in [20] are stronger, our group signature satisfies more desirable security properties. Concretely, even though the opening algorithm of Chung et al. can output an opening proof, it does not enjoy the security guarantees defined in [20]. In fact, their construction does not have *tracing soundness*. In other words, their signature can be opened to two distinct parties with distinct valid opening proofs.

Table 2 compares this work with [33, 66] in terms of signature size, anonymity and manager accountability, where the integer N represents the size of the group. The construction in [33] is marked with partial manager accountability for the aforementioned reason. Ours is the only scheme that achieves an $\mathcal{O}(\log N)$ signature size and CCA anonymity. Additionally, our scheme provides a much tighter security since both [33, 66] rely on the forking lemma in their security proofs.

Structure of this paper. We begin in Sect. 3 with some preliminary background on sigma protocols, accountable ring signatures, and other mathematical content which this paper relies on. We then introduce our new, generic constructions of accountable ring signature and dynamic group signature schemes in Sect. 4. These generic constructions are built from various components put forward in the preceding sections: Sect. 5 defines group-action-based hard instance generators and public-key encryption schemes; Sect. 6 introduces our new "traceable" sigma protocol and proves its security; and Sect. 7 then constructs a NIZK proof system from said sigma protocol through the Fiat-Shamir transform. Finally, Sect. 8 details the instantiation of our schemes from isogenies and lattices.

3 Preliminaries

Notation. We begin by introducing some notation that will be used throughout the paper. For $N \in \mathbb{N}$, we denote by $[N]$ the set $\{1, \dots, N\}$. We use \parallel to represent concatenation of

two strings. We also use $\{X_i\}_{i \in S}$ to denote the set of elements X_i iterating over all values $i \in S$. For any randomized algorithm A taking as input x , we will write $A(x; r)$ to denote the execution of A on x using the randomness r . With an overload in notation, we write $A(x)$ to denote the set of all possible outputs of A on input x , and $y \in A(x)$ to indicate that there exists a randomness r such that $y = A(x; r)$. Finally, we let $\text{negl}(\lambda)$ be a *negligible* function, i.e. one dominated by $O(\lambda^{-n})$ for all $n > 0$.

A note on random oracles. Throughout the paper, we instantiate several standard cryptographic primitives, such as pseudorandom number generators (i.e., Expand) and commitment schemes, by hash functions modeled as a random oracle \mathcal{O} . We always assume the input domain of the random oracle is appropriately separated when instantiating several cryptographic primitives by one random oracle. With abuse of notation, we may occasionally write for example $\mathcal{O}(\text{Expand} \parallel \cdot)$ instead of $\text{Expand}(\cdot)$ to make the usage of the random oracle explicit. Here, we identify Expand with a unique string when inputting it to \mathcal{O} . Finally, we denote by $\mathcal{A}^{\mathcal{O}}$ an algorithm \mathcal{A} that has black-box access to \mathcal{O} , and we may occasionally omit the superscript \mathcal{O} for simplicity when the meaning is clear from context.

3.1 Sigma protocols

A sigma protocol Π_{Σ} for a NP relation $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$ is a public-coin three-move interactive protocol between a prover and a verifier that satisfies a specific flavor of soundness and zero-knowledge. The language \mathcal{L}_R is defined as $\{X \mid (X, W) \in R\}$. As standard with many sigma protocols for a language defined over post-quantum algebraic structures, we relax the soundness notion to only hold for a slightly wider relation \tilde{R} (i.e., $R \subseteq \tilde{R}$), e.g., [4, 10, 14, 36, 44, 54]. That is, a cheating prover may not be using a witness in R but is guaranteed to be using some witness in the wider relation \tilde{R} . Below, we consider a sigma protocol in the random oracle model, where the prover and verifier have access to a random oracle similarly to [14].⁷

Definition 3.1 (Sigma Protocol) A sigma protocol Π_{Σ} for the relations R and \tilde{R} such that $R \subseteq \tilde{R}$ (which are implicitly parameterized by the security parameter λ) consists of oracle-calling PPT algorithms $(P = (P_1, P_2), V = (V_1, V_2))$, where V_2 is deterministic and we assume P_1 and P_2 share states. Let ChSet denote the challenge space. Then, Π_{Σ} has the following three-move flow:

- The prover, on input $(X, W) \in R$, runs $\text{com} \leftarrow P_1^{\mathcal{O}}(X, W)$ and sends a *commitment* com to the verifier.
- The verifier runs $\text{chall} \leftarrow V_1^{\mathcal{O}}(1^{\lambda})$ to obtain a random *challenge* chall from ChSet , and sends it to the prover.
- The prover, given chall , runs $\text{resp} \leftarrow P_2^{\mathcal{O}}(X, W, \text{chall})$ and returns a *response* resp to the verifier. Here, we allow P_2 to abort with some probability. In such cases we assign resp with a special symbol \perp denoting abort.
- The verifier runs $V_2^{\mathcal{O}}(X, \text{com}, \text{chall}, \text{resp})$ and outputs \top (accept) or \perp (reject).

Here, \mathcal{O} is modeled as a random oracle and we often drop \mathcal{O} from the superscript for simplicity when the meaning is clear from context. We assume X is always given as input to P_2 and V_2 , and omit it in the following. The protocol *transcript* $(\text{com}, \text{chall}, \text{resp})$ is said to be valid in case $V_2(\text{com}, \text{chall}, \text{resp})$ outputs \top .

⁷ This should not be confused with the random oracle used to compile a sigma protocol into an NIZK proof system.

We require a sigma protocol Π_Σ in the random oracle model to satisfy the following standard properties: correctness, high min-entropy, special zero-knowledge and (relaxed) special soundness.

We require the sigma protocol to be correct conditioned on the prover not aborting the protocol. Below, if $\delta = 0$, then it corresponds to the case when the prover never aborts.

Definition 3.2 ($(1 - \delta)$ -Correctness) A sigma protocol Π_Σ is $(1 - \delta)$ -correct for $\delta \in [0, 1]$ if for all $\lambda \in \mathbb{N}$ and $(X, W) \in R$, the probability of the prover outputting \perp is at most δ , and we have

$$\Pr \left[V_2^\mathcal{O}(\text{com}, \text{chall}, \text{resp}) = \top \mid \begin{array}{l} \text{com} \leftarrow P_1^\mathcal{O}(X, W), \\ \text{chall} \leftarrow V_1^\mathcal{O}(1^\lambda), \\ \text{resp} \leftarrow P_2^\mathcal{O}(W, \text{chall}) \text{ s.t. } \text{resp} \neq \perp. \end{array} \right] = 1,$$

where the probability is taken over the randomness used by (P, V) and by the random oracle.

Definition 3.3 (High Min-Entropy) We say a sigma protocol Π_Σ has $\alpha(\lambda)$ min-entropy if for any $\lambda \in \mathbb{N}$, $(X, W) \in R$, and a possibly computationally-unbounded adversary \mathcal{A} , we have

$$\Pr [\text{com} = \text{com}' \mid \text{com} \leftarrow P_1^\mathcal{O}(X, W), \text{com}' \leftarrow \mathcal{A}^\mathcal{O}(X, W)] \leq 2^{-\alpha},$$

where the probability is taken over the randomness used by P_1 and by the random oracle. We say Π_Σ has *high min-entropy* if $2^{-\alpha}$ is negligible in λ .

Definition 3.4 (Non-Abort Special Zero-Knowledge) We say Π_Σ is (non-abort) special *zero-knowledge* if there exists a PPT simulator $\text{Sim}^\mathcal{O}$ with access to a random oracle \mathcal{O} such that for any $\lambda \in \mathbb{N}$, statement-witness pair $(X, W) \in R$, $\text{chall} \in \text{ChSet}$ and any computationally-unbounded adversary \mathcal{A} that makes at most a polynomial number of queries to \mathcal{O} , we have

$$\left| \Pr[\mathcal{A}^\mathcal{O}(1^\lambda, \tilde{P}^\mathcal{O}(X, W, \text{chall})) = 1] - \Pr[\mathcal{A}^\mathcal{O}(1^\lambda, \text{Sim}^\mathcal{O}(X, \text{chall})) = 1] \right| = \text{negl}(\lambda),$$

where \tilde{P} is a *non-aborting* prover $P = (P_1, P_2)$ run on (X, W) with a challenge fixed to chall and the probability is taken over the randomness used by (P, V) and by the random oracle.

Below, for the special soundness property, the extraction algorithm is only required to recover a “weaker” witness in \tilde{R} rather than in R used in the real protocol. In many applications, the capability of extracting from this wider relation suffices.

Definition 3.5 (Special Soundness) We say a sigma protocol Π_Σ has (relaxed) special soundness if there exists a PT extraction algorithm Extract such that, given a statement X and any two valid transcripts $(\text{com}, \text{chall}, \text{resp})$ and $(\text{com}, \text{chall}', \text{resp}')$ relative to X and such that $\text{chall} \neq \text{chall}'$, outputs a witness W satisfying $(X, W) \in \tilde{R}$.

3.2 Non-interactive Zero-knowledge proofs of knowledge in the ROM

We consider non-interactive zero-knowledge proof of knowledge protocols (or simply NIZK (proof system)) in the ROM. Below, we define a variant where the proof is generated with respect to a label. Although syntactically different, such NIZK is analogous to the notion of signature of knowledge [30]

Definition 3.6 (NIZK Proof System) Let L denote a label space, where checking membership can be done efficiently. A non-interactive zero-knowledge (NIZK) proof system Π_{NIZK} for the relations R and \tilde{R} such that $R \subseteq \tilde{R}$ (which are implicitly parameterized by λ) consists of oracle-calling PPT algorithms (Prove, Verify) defined as follows:

- $\text{Prove}^{\mathcal{O}}(\text{lbl}, X, W) \rightarrow \pi/\perp$: On input a label $\text{lbl} \in L$, a statement and witness pair $(X, W) \in R$, it outputs a proof π or a special symbol \perp denoting abort.
- $\text{Verify}^{\mathcal{O}}(\text{lbl}, X, \pi) \rightarrow \top/\perp$: On input a label $\text{lbl} \in L$, a statement X , and a proof π , it outputs either \top (accept) or \perp (reject).

We require a NIZK proof system in the random oracle model to satisfy the following standard properties: correctness, zero-knowledge, (relaxed) statistical soundness, and online extractability. We assume for simplicity that Verify always outputs \perp in case $\text{lbl} \notin L$.

Definition 3.7 ($(1 - \delta)$ -Correctness) A NIZK proof system Π_{NIZK} is $(1 - \delta)$ -correct for $\delta \in [0, 1]$ if for all $\lambda \in \mathbb{N}$, $\text{lbl} \in L$, $(X, W) \in R$, the probability of $\text{Prove}^{\mathcal{O}}(\text{lbl}, X, W)$ outputting \perp is at most δ , and we have

$$\Pr \left[\text{Verify}^{\mathcal{O}}(\text{lbl}, X, \pi) = \top \mid \begin{array}{l} \pi \leftarrow \text{Prove}^{\mathcal{O}}(\text{lbl}, X, W), \\ \pi \neq \perp. \end{array} \right] = 1,$$

where the probability is taken over the randomness used by (Prove, Verify) and by the random oracle.

Definition 3.8 (Zero-Knowledge) Let \mathcal{O} be a random oracle, Π_{NIZK} a NIZK proof system, and $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ a zero-knowledge simulator for Π_{NIZK} , consisting of two algorithms Sim_0 and Sim_1 with a shared state. We say the advantage of an adversary \mathcal{A} against Sim is

$$\text{Adv}_{\Pi_{\text{NIZK}}}^{\text{ZK}}(\mathcal{A}) = \left| \Pr \left[\mathcal{A}^{\mathcal{O}, \text{Prove}}(1^\lambda) = 1 \right] - \Pr \left[\mathcal{A}^{\text{Sim}_0, \mathcal{S}}(1^\lambda) = 1 \right] \right|,$$

where Prove and \mathcal{S} are prover oracles that on input (lbl, X, W) return \perp if $\text{lbl} \notin L \vee (X, W) \notin R$ and otherwise return $\text{Prove}^{\mathcal{O}}(\text{lbl}, X, W)$ or $\text{Sim}_1(\text{lbl}, X)$, respectively. Moreover, the probability is taken also over the randomness of sampling \mathcal{O} .

We say Π_{NIZK} for R and \tilde{R} is zero-knowledge if there exists a PPT simulator Sim such that for all (possibly computationally-unbounded) adversary \mathcal{A} making at most polynomially many queries to the random oracle and the prover oracle, we have $\text{Adv}_{\Pi_{\text{NIZK}}}^{\text{ZK}}(\mathcal{A}) \leq \text{negl}(\lambda)$.

Statistical soundness guarantees that any adversary cannot generate a proof for an invalid statement except with a negligible probability.

Definition 3.9 (Statistical Soundness) Let \mathcal{O} be a random oracle and Π_{NIZK} a NIZK proof system. We say the advantage of an adversary \mathcal{A} against soundness is

$$\text{Adv}_{\Pi_{\text{NIZK}}}^{\text{soundness}}(\mathcal{A}) = \Pr \left[\begin{array}{l} \exists W : (X, W) \in \tilde{R} \wedge \\ \text{Verify}^{\mathcal{O}}(\text{lbl}, X, \pi) = \top \end{array} \mid (\text{lbl}, X, \pi) \leftarrow \mathcal{A}^{\mathcal{O}}(1^\lambda) \right],$$

where the probability is taken also over the randomness of sampling \mathcal{O} .

We say the NIZK proof system Π_{NIZK} for R and \tilde{R} has (relaxed) statistical soundness if for all (possibly computationally-unbounded) adversary \mathcal{A} making at most polynomially many queries to the random oracle, we have $\text{Adv}_{\Pi_{\text{NIZK}}}^{\text{soundness}}(\mathcal{A}) \leq \text{negl}(\lambda)$.

Online extractability requires the existence of an extraction algorithm which, on input a valid proof π and the list or random-oracle queries made by an adversary, always extract a (relaxed) witness except with a negligible probability.

Definition 3.10 (Multi-Proof Online Extractability) A NIZK proof system Π_{NIZK} is (multi-proof) online extractable if there exists a PPT extractor `OnlineExtract` such that for any (possibly computationally-unbounded) adversary \mathcal{A} making at most polynomially-many queries has at most a negligible advantage in the following game played against a challenger (with access to a random oracle \mathcal{O}).

- (i) The challenger prepares empty lists $L_{\mathcal{O}}$ and L_P , and sets `flag` to 0.
- (ii) \mathcal{A} can make random-oracle, prove, and extract queries an arbitrary polynomial number of times:
 - (`hash`, x): The challenger updates $L_{\mathcal{O}} \leftarrow L_{\mathcal{O}} \cup \{x, \mathcal{O}(x)\}$ and returns $\mathcal{O}(x)$. We assume below that \mathcal{A} runs the verification algorithm after receiving a proof from the prover oracle and before submitting a proof to the extract oracle.⁸
 - (`prove`, lbl, X, W): The challenger returns \perp if $\text{lbl} \notin L$ or $(X, W) \notin R$. Otherwise, it returns $\pi \leftarrow \text{Prove}^{\mathcal{O}}(\text{lbl}, X, W)$ and updates $L_P \leftarrow L_P \cup \{\text{lbl}, X, \pi\}$.
 - (`extract`, lbl, X, π): The challenger checks if $\text{Verify}^{\mathcal{O}}(\text{lbl}, X, \pi) = \top$ and $(\text{lbl}, X, \pi) \notin L_P$, and returns \perp if not. Otherwise, it runs $W \leftarrow \text{OnlineExtract}^{\mathcal{O}}(\text{lbl}, X, \pi, L_{\mathcal{O}})$ and checks if $(X, W) \notin \tilde{R}$, and returns \perp if yes and sets `flag` = 1. Otherwise, if all checks pass, it returns W .
- (iii) At some point \mathcal{A} outputs 1 to indicate that it is finished with the game. We say \mathcal{A} wins if `flag` = 1. The advantage of \mathcal{A} is defined as $\text{Adv}_{\Pi_{\text{NIZK}}}^{\text{OE}}(\mathcal{A}) = \Pr[\mathcal{A} \text{ wins}]$ where the probability is also taken over the randomness used by the random oracle.

Note, importantly, that `OnlineExtract` is not given access to the queries `Prove` ^{\mathcal{O}} makes directly to \mathcal{O} . Thus, `OnlineExtract` is not guaranteed to return a valid witness W when called with any output of the `Prove` oracle. The requirement that $(\text{lbl}, X, \pi) \notin L_P$ ensures that this does not allow the adversary to trivially win the game, and in particular by extension ensures that modifying the label lbl should invalidate any proof obtained from the `Prove` oracle.

Remark 3.11 If a NIZK proof system Π_{NIZK} is (multi-proof) online extractable, it is statistically sound—that is, online extractability implies statistical soundness. This is clear, because if an adversary is able to generate an accepting tuple (lbl, X, π) for which $\nexists W : (X, W) \in \tilde{R}$ in the soundness game, then clearly `(extract, lbl, X, π)` will allow the adversary to win the online extractability game.

Remark 3.12 (NIZKs with Labels) If the label space of the NIZK is $L = \{\perp\}$, we say the NIZK is without labels (or a plain/unlabelled NIZK). In this case, we omit the `lbl` argument from the `Prove` and `Verify` functions for clarity.

3.3 Public-key encryption

We recall the standard multi-challenge IND-CPA security of a public-key encryption (PKE) scheme.

Definition 3.13 (Public-Key Encryption) A public-key encryption Π_{PKE} over a message space \mathcal{M} consists of four algorithms $\Pi_{\text{PKE}} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$:

- `Setup`(1^λ) \rightarrow `pp` : On input the security parameter 1^λ , it outputs a public parameter `pp`.

⁸ This is w.l.o.g., and guarantees that the list $L_{\mathcal{O}}$ is updated with the input/output required to verify the proof \mathcal{A} receives or sends.

- $\text{KeyGen}(\text{pp}) \rightarrow (\text{pk}, \text{sk})$: On input a public parameter pp , it outputs a pair of public key and secret key (pk, sk) .
- $\text{Enc}(\text{pk}, M) \rightarrow \text{ct}$: On input a public key pk_i and a message $M \in \mathcal{M}$, it outputs a ciphertext ct .
- $\text{Dec}(\text{sk}, \text{ct}) \rightarrow M \text{ or } \perp$: On input a secret key sk and a ciphertext ct , it outputs either $M \in \mathcal{M}$ or a special symbol $\perp \notin \mathcal{M}$.

We will denote by \mathcal{R} the set containing the randomness used by the encryption algorithm Enc .

We omit the standard definition of correctness as we provide a more generalized version in Sect. 4.1, Def. 4.1. Below, we define the standard IND-CPA security extended to the multi-challenge setting. Using a textbook hybrid argument, it is clear that the multi-challenge definition is polynomially related to the standard single-challenge definition. The motivation for introducing the multi-challenge variant is because in some cases, we can show that the two definitions are equally difficult without incurring any reduction loss.

Definition 3.14 (Multi-Challenge IND-CPA Security) A PKE scheme $\Pi_{\text{PKE}} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ is *multi-challenge IND-CPA secure* against Q challenges if, for any $\lambda \in \mathbb{N}$, any PPT adversary \mathcal{A} has at most a negligible advantage in the following game played against a challenger.

- The challenger runs $\text{pp} \leftarrow \text{Setup}(1^\lambda)$, $(\text{pk}, \text{sk}) \leftarrow \text{KeyGen}(\text{pp})$ and samples a bit $b \in \{0, 1\}$. The challenger provides (pp, pk) to \mathcal{A} .
- \mathcal{A} can adaptively query the challenge oracle at most Q times. In each query, \mathcal{A} sends a pair of messages $(M_0, M_1) \in \mathcal{M}^2$, and the challenger returns $\text{ct}_b \leftarrow \text{Enc}(\text{pk}, M_b)$ to \mathcal{A} .
- \mathcal{A} outputs a bit $b^* \in \{0, 1\}$. We say \mathcal{A} wins if $b^* = b$.

The advantage of \mathcal{A} is defined as $\text{Adv}_{\Pi_{\text{PKE}}, Q}^{\text{Multi-CPA}}(\mathcal{A}) = |\Pr[\mathcal{A} \text{ wins}] - 1/2|$.

3.4 Accountable ring signatures

We provide the definition of accountable ring signatures (ARs), following the formalization introduced by Bootle et al. [19].

Definition 3.15 (Accountable Ring Signature) An accountable ring signature Π_{ARS} consists of PPT algorithms $(\text{Setup}, \text{OKGen}, \text{UKGen}, \text{Sign}, \text{Verify}, \text{Open}, \text{Judge})$ defined as follows:

- $\text{Setup}(1^\lambda) \rightarrow \text{pp}$: On input a security parameter 1^λ , it returns a public parameter pp (sometimes implicitly) used by the scheme. We assume pp defines openers' public-key space \mathcal{K}_{opk} and users' verification-key space \mathcal{K}_{vk} , with efficient algorithms to decide membership.
- $\text{OKGen}(\text{pp}) \rightarrow (\text{opk}, \text{osk})$: On input a public parameter pp , it outputs a pair of public and secret keys (opk, osk) for an *opener*.
- $\text{UKGen}(\text{pp}) \rightarrow (\text{vk}, \text{sk})$: On input a public parameter pp , it outputs a pair of verification and signing keys (vk, sk) for a *user*.
- $\text{Sign}(\text{opk}, \text{sk}, R, M) \rightarrow \sigma$: On input an opener's public key opk , a signing key sk , a list of verification keys, i.e., a *ring*, $R = \{\text{vk}_1, \dots, \text{vk}_N\}$, and a message M , it outputs a signature σ .

- Verify(opk, R, M, σ) \rightarrow \top/\perp : On input an opener’s public key opk, a ring $R = \{vk_1, \dots, vk_N\}$, a message M, and a signature σ , it (deterministically) outputs either \top (accept) or \perp (reject).
- Open(osk, R, M, σ) \rightarrow $(vk, \pi)/\perp$: On input an opener’s secret key osk, a ring $R = \{vk_1, \dots, vk_N\}$, a message M, a signature σ , it (deterministically) outputs either a pair of verification key vk and a proof π that the owner of vk produced the signature, or \perp .
- Judge(opk, R, vk, M, σ, π) \rightarrow \top/\perp : On input an opener’s public key opk, a ring $R = \{vk_1, \dots, vk_N\}$, a verification key vk, a message M, a signature σ , and a proof π , it (deterministically) outputs either \top (accept) or \perp (reject). We assume without loss of generality that Judge(opk, R, vk, M, σ, π) outputs \perp if Verify(opk, R, M, σ) outputs \perp .

An accountable ring signature is required to satisfy the following properties: correctness, anonymity, traceability, unforgeability, and tracing soundness.

First, we require correctness to hold even if the ring contains maliciously-generated user keys or the signature has been produced for a maliciously-generated opener key. Note that the correctness guarantee for the open and judge algorithms are defined implicitly in the subsequent security definitions.

Definition 3.16 (Correctness) An accountable ring signature Π_{ARS} is *correct* if, for all $\lambda \in \mathbb{N}$, any PPT adversary \mathcal{A} has at most a negligible advantage in λ in the following game played against a challenger.

- (i) The challenger runs $pp \leftarrow \text{Setup}(1^\lambda)$ and generates a user key $(vk, sk) \leftarrow \text{UKGen}(pp)$. It then provides (pp, vk, sk) to \mathcal{A} .
- (ii) \mathcal{A} outputs an opener’s public key, a ring, and a message tuple (opk, R, M) to the challenger.
- (iii) The challenger runs $\sigma \leftarrow \text{Sign}(opk, sk, R, M)$. We say \mathcal{A} wins if
 - $opk \in \mathcal{K}_{opk}, R \subseteq \mathcal{K}_{vk},$ and $vk \in R,$
 - $\text{Verify}(opk, R, M, \sigma) = \perp.$

The advantage of \mathcal{A} is defined as $\text{Adv}_{\Pi_{ARS}}^{\text{Correct}}(\mathcal{A}) = \Pr[\mathcal{A} \text{ wins}].$

Anonymity requires that a signature does not leak any information on who signed it. We consider the standard type of anonymity notion where the adversary gets to choose the signing key used to generate the signature. Moreover, we allow the adversary to make (non-trivial) opening queries that reveal who signed the messages. This notion is often called *full (CCA) anonymity* [8, 20] to differentiate between weaker notions of anonymity such as *selfless* anonymity that restricts the adversary from exposing the signing key used to sign the signature or *CPA* anonymity where the adversary is restricted from querying the open oracle.

Definition 3.17 (Anonymity) An accountable ring signature Π_{ARS} is (CCA) *anonymous (against full key exposure)* if, for all $\lambda \in \mathbb{N}$, any PPT adversary \mathcal{A} has at most a negligible advantage in the following game played against a challenger.

- (i) The challenger runs $pp \leftarrow \text{Setup}(1^\lambda)$ and generates an opener key $(opk, osk) \leftarrow \text{OKGen}(pp)$. It also prepares an empty list Q_{sign} and samples a random bit $b \leftarrow \{0, 1\}.$
- (ii) The challenger provides (pp, opk) to $\mathcal{A}.$

(iii) \mathcal{A} can make signing and opening queries an arbitrary polynomial number of times:

- ($\text{sign}, R, M, \text{sk}_0, \text{sk}_1$): The challenger runs $\sigma_i \leftarrow \text{Sign}(\text{opk}, \text{sk}_i, R, M)$ for $i \in \{0, 1\}$ and returns \perp if $\text{Verify}(\text{opk}, R, M, \sigma_i) = \perp$ for either of $i \in \{0, 1\}$. Otherwise, it updates $Q_{\text{sign}} \leftarrow Q_{\text{sign}} \cup \{(R, M, \sigma_b)\}$ and returns σ_b .
- ($\text{open}, R, M, \sigma$): The challenger returns \perp if $(R, M, \sigma) \in Q_{\text{sign}}$. Otherwise, it returns $\text{Open}(\text{osk}, R, M, \sigma)$.

(iv) \mathcal{A} outputs a guess b^* . We say \mathcal{A} wins if $b^* = b$.

The advantage of \mathcal{A} is defined as $\text{Adv}_{\Pi_{\text{ARS}}}^{\text{Anon}}(\mathcal{A}) = |\Pr[\mathcal{A} \text{ wins}] - 1/2|$.

Unforgeability considers two types of forgeries. The first captures the natural notion of unforgeability where an adversary cannot forge a signature for a ring of honest users, i.e., a ring of users for which it does not know any of the corresponding secret keys. The second captures the fact that an adversary cannot accuse an honest user of producing a signature even if the ring contains malicious users and the opener is malicious.

Definition 3.18 (Unforgeability) An accountable ring signature scheme Π_{ARS} is *unforgeable (with respect to insider corruption)* if, for all $\lambda \in \mathbb{N}$, any PPT adversary \mathcal{A} has at most negligible advantage in the following game played against a challenger.

(i) The challenger runs $\text{pp} \leftarrow \text{Setup}(\lambda)$ and initializes an empty keyed dictionary $D_{\text{UKey}}[\cdot]$ and three empty sets $Q_{\text{UKey}}, Q_{\text{sign}}$ and Q_{cor} . It provides pp to \mathcal{A} .
 (ii) \mathcal{A} can make user key generation, signing, and corruption queries an arbitrary polynomial number of times:

- (ukeygen): The challenger runs $(\text{vk}, \text{sk}) \leftarrow \text{UKGen}(\text{pp})$. If $D_{\text{UKey}}[\text{vk}] \neq \perp$, then it returns \perp . Otherwise, it updates $D_{\text{UKey}}[\text{vk}] = \text{sk}$ and $Q_{\text{UKey}} \leftarrow Q_{\text{UKey}} \cup \{\text{vk}\}$, and returns vk .
- ($\text{sign}, \text{opk}, \text{vk}, R, M$): The challenger returns \perp if $\text{vk} \notin Q_{\text{UKey}} \cap R$. Otherwise, it runs $\sigma \leftarrow \text{Sign}(\text{opk}, D_{\text{UKey}}[\text{vk}], R, M)$. The challenger updates $Q_{\text{sign}} \leftarrow Q_{\text{sign}} \cup \{(\text{opk}, \text{vk}, R, M, \sigma)\}$ and returns σ .
- ($\text{corrupt}, \text{vk}$): The challenger returns \perp if $\text{vk} \notin Q_{\text{UKey}}$. Otherwise, it updates $Q_{\text{cor}} \leftarrow Q_{\text{cor}} \cup \{\text{vk}\}$ and returns $D_{\text{UKey}}[\text{vk}]$.

(iv) \mathcal{A} outputs $(\text{opk}, \text{vk}, R, M, \sigma, \pi)$. We say \mathcal{A} wins if

- $(\text{opk}, *, R, M, \sigma) \notin Q_{\text{sign}}, R \subseteq Q_{\text{UKey}} \setminus Q_{\text{cor}},$
- $\text{Verify}(\text{opk}, R, M, \sigma) = \top,$

or

- $(\text{opk}, \text{vk}, R, M, \sigma) \notin Q_{\text{sign}}, \text{vk} \in Q_{\text{UKey}} \setminus Q_{\text{cor}},$
- $\text{Judge}(\text{opk}, R, \text{vk}, M, \sigma, \pi) = \top.$

The advantage of \mathcal{A} is defined as $\text{Adv}_{\Pi_{\text{ARS}}}^{\text{Unf}}(\mathcal{A}) = \Pr[\mathcal{A} \text{ wins}]$.

Traceability requires that any opener key pair (opk, osk) in the range of the opener key-generation algorithm can open a valid signature σ to some user vk along with a proof valid π . This ensures that any opener can trace the user and produce a proof for its decision. Below, rather than assuming an efficient algorithm that checks set membership $(\text{opk}, \text{osk}) \in \text{OKGen}(\text{pp})$, we simply ask the adversary to output the randomness used to generate (opk, osk) . Note that this definition contains the prior definitions where opk was assumed to be uniquely defined and efficiently computable from osk [19].

Definition 3.19 (Traceability) An accountable ring signature scheme Π_{ARS} is *traceable* if, for all $\lambda \in \mathbb{N}$, any PPT adversary \mathcal{A} has at most negligible advantage in the following game played against a challenger.

- (i) The challenger runs $\text{pp} \leftarrow \text{Setup}(1^\lambda)$ and provides pp to \mathcal{A} .
- (ii) \mathcal{A} returns a randomness, a ring, a message, and a signature tuple $(\text{rr}, R, M, \sigma)$. We say \mathcal{A} wins if
 - $\text{Verify}(\text{opk}, R, M, \sigma) = \top$, where $(\text{opk}, \text{osk}) \leftarrow \text{OKGen}(\text{pp}; \text{rr})$, and
 - $\text{Judge}(\text{opk}, R, \text{vk}, M, \sigma, \pi) = \perp$, where $(\text{vk}, \pi) \leftarrow \text{Open}(\text{osk}, R, M, \sigma)$.

The advantage of \mathcal{A} is defined as $\text{Adv}_{\Pi_{\text{ARS}}}^{\text{Tra}}(\mathcal{A}) = \Pr[\mathcal{A} \text{ wins}]$.

Finally, tracing soundness requires that a signature cannot trace to two different users in the ring. This must hold even if all the users in the ring and the opener are corrupt.

Definition 3.20 (Tracing Soundness) An accountable ring signature scheme Π_{ARS} is *traceable sound* if, for all $\lambda \in \mathbb{N}$, any PPT adversary \mathcal{A} has at most negligible advantage in the following game played against a challenger.

- (i) The challenger runs $\text{pp} \leftarrow \text{Setup}(1^\lambda)$ and provides pp to \mathcal{A} .
- (ii) \mathcal{A} returns an opener’s public key, a ring, a message, a signature, and two verification keys and proofs $(\text{opk}, R, M, \sigma, \{(\text{vk}_b, \pi_b)\}_{b \in \{0,1\}})$. We say \mathcal{A} wins if
 - $\text{vk}_0 \neq \text{vk}_1$,
 - $\text{Judge}(\text{opk}, R, \text{vk}_0, M, \sigma, \pi_0) = \top$,
 - $\text{Judge}(\text{opk}, R, \text{vk}_1, M, \sigma, \pi_1) = \top$.

The advantage of \mathcal{A} is defined as $\text{Adv}_{\Pi_{\text{ARS}}}^{\text{TraS}}(\mathcal{A}) = \Pr[\mathcal{A} \text{ wins}]$.

3.5 Isogenies and ideal class group actions

Let \mathbb{F}_p be a prime field, with $p \geq 5$. In the following E and E' denote elliptic curves defined over \mathbb{F}_p . An isogeny $\varphi : E \rightarrow E'$ is a non-constant morphism mapping 0_E to $0_{E'}$. Each coordinate of $\varphi(x, y)$ is then the fraction of two polynomials in $\overline{\mathbb{F}}_p[x, y]$, where $\overline{\mathbb{F}}_p$ denotes the algebraic closure of \mathbb{F}_p . If the coefficients of the polynomials lie in \mathbb{F}_p , then φ is said to be defined over \mathbb{F}_p . We restrict our attention to *separable* isogenies (which induce separable extensions of function fields) between *supersingular* elliptic curves defined over \mathbb{F}_p , i.e., curves whose set of rational points $E(\mathbb{F}_p)$ has cardinality $p + 1$.

An isogeny $\varphi : E \rightarrow E'$ is an *isomorphism* if its kernel is equal to $\{0_E\}$, and an *endomorphism* of E if $E = E'$. The set $\text{End}_p(E)$ of all endomorphisms of E that are defined over \mathbb{F}_p , together with the zero map, form a commutative ring under pointwise addition and composition. $\text{End}_p(E)$ is isomorphic to an order \mathcal{O} of the quadratic field $\mathbb{K} = \mathbb{Q}(\sqrt{-p})$ [28]. We recall that an order is a subring of \mathbb{K} , which is also a finitely-generated \mathbb{Z} -module containing a basis of \mathbb{K} as a \mathbb{Q} -vector space. A fractional ideal \mathfrak{a} of \mathcal{O} is a finitely generated \mathcal{O} -submodule of \mathbb{K} . We say that \mathfrak{a} is invertible if there exists another fractional ideal \mathfrak{b} of \mathcal{O} such that $\mathfrak{a}\mathfrak{b} = \mathcal{O}$, and that it is principal if $\mathfrak{a} = \alpha\mathcal{O}$ for some $\alpha \in \mathbb{K}$. The invertible fractional ideals of \mathcal{O} form an Abelian group whose quotient by the subgroup of principal fractional ideals is finite. This quotient group is called the *ideal class group* of \mathcal{O} , and denoted by $\mathcal{Cl}(\mathcal{O})$.

The ideal class group $\mathcal{Cl}(\mathcal{O})$ acts freely and transitively on the set $\mathcal{Ell}_p(\mathcal{O}, \pi)$, which contains all supersingular elliptic curves E over \mathbb{F}_p - modulo isomorphisms defined over \mathbb{F}_p

- such that there exists an isomorphism between \mathcal{O} and $\text{End}_p(E)$ mapping $\sqrt{-p} \in \mathcal{O}$ into the Frobenius endomorphism $(x, y) \mapsto (x^p, y^p)$. We denote this action by $*$. Recently, it has been used to design several cryptographic primitives [15, 28, 37, 67], whose security proofs rely on (variations of) the Group Action Inverse Problem (GAIP), defined as follows.

Definition 3.21 (Group Action Inverse Problem (GAIP)) Let $[E_0]$ be an element in $\mathcal{E}llp(\mathcal{O}, \pi)$, where p is an odd prime and \mathcal{O} an order in $\mathbb{Q}(\sqrt{-p})$. Given $[E]$ sampled from the uniform distribution over $\mathcal{E}llp(\mathcal{O}, \pi)$, the GAIP_p problem consists in finding an element $[\alpha] \in \mathcal{C}l(\mathcal{O})$ such that $[\alpha] * [E_0] = [E]$.

The best known classical algorithm to solve the GAIP problem has time complexity $O(\sqrt{N})$, where $N = |\mathcal{C}l(\mathcal{O})|$. The best known quantum algorithm, on the other hand, is Kuperberg’s algorithm for the hidden shift problem [63, 64]. It has a subexponential complexity, for which the concrete security estimates are still an active area of research [12, 18, 32, 80].

For the security of the isogeny-based instantiations, we will also rely on a multi-instance variant the GAIP problem which is trivially equivalent to the GAIP problem.

Definition 3.22 (Multi-Instance GAIP (MI-GAIP) Problem) Let $[E_0]$ be an element in $\mathcal{E}llp(\mathcal{O}, \pi)$, where p is an odd prime and \mathcal{O} an order in $\mathbb{Q}(\sqrt{-p})$. Given $[E_1], \dots, [E_N]$ sampled uniformly at random from $\mathcal{E}llp(\mathcal{O}, \pi)$, where $N \in \mathbb{N}$, the $\text{MI-GAIP}_{p,N}$ problem consists in finding an element $[\alpha] \in \mathcal{C}l(\mathcal{O})$ such that $[\alpha] * [E_0] = [E_i]$ for some $i \in [N]$.

To see the equivalence (informally), given an instance of the GAIP problem ($[E_0], [E]$), sample $[\tau_1], \dots, [\tau_N] \in \mathcal{C}l(\mathcal{O})$, and compute $[E_i] = [\tau_i] * [E]$ for each i . Then a solution for the MT-GAIP on $([E_0], [E_1], \dots, [E_N])$, say $[\alpha] * [E_0] = [E_j]$, results in a solution to the GAIP by computing $[\alpha][\tau_j]^{-1}$.

We also need the following assumption, the decisional CSIDH Problem. Looking ahead, the distinguishing problems will ensure (multi-instance) IND-CPA for our PKE in Sect. 8.1 and therefore anonymity for our ring/group signature schemes. Note that we will require the class group to be of odd order to avoid the attack presented in [29]. Equivalently, we require $p \equiv 3 \pmod 4$.

Definition 3.23 (Decisional CSIDH (dCSIDH) Problem) Let $[E_0]$ be an element in $\mathcal{E}llp(\mathcal{O}, \pi)$, where p is an odd prime. The decisional CSIDH problem is that given a tuple $([\alpha_1] * [E_0], [\alpha_2] * [E_0], E)$ where $[\alpha_1], [\alpha_2]$ are sampled uniformly from $\mathcal{C}l(\mathcal{O})$ and $[E]$ is either sampled uniformly from $\mathcal{E}llp(\mathcal{O}, \pi)$ or $[E] = [\alpha_1 \alpha_2] * [E_0]$, and decide which distribution $[E]$ is drawn from.

3.6 Lattices

Let R and R_q denote the rings $\mathbb{Z}[X]/(X^n + 1)$ and $\mathbb{Z}[X]/(q, X^n + 1)$ for integers n and q , respectively. Norms over R are defined through the coefficient vectors of the polynomials, which lie over \mathbb{Z}^n . Norms over R_q are defined in the conventional way by uniquely representing coefficients of elements over R_q by elements in the range $(-q/2, q/2)$ when q is even and $[-(q - 1)/2, (q - 1)/2]$ when q is odd (see for example [42] for more details).

The hard problems we will rely on are the *module short integer solution* (MSIS) problem and *module learning with errors* (MLWE) problem, first introduced in [69].

Definition 3.24 (Module Short Integer Solution) Let n, q, k, ℓ, γ be integers. The advantage for the (Hermite normal form) *module short integer solution problem* $\text{MSIS}_{n,q,k,\ell,\gamma}$ for an algorithm \mathcal{A} is defined as

$$\text{Adv}_{n,q,k,\ell,\gamma}^{\text{MSIS}}(\mathcal{A}) = \Pr \left[\begin{array}{l} 0 < \|\mathbf{u}\|_\infty \leq \gamma \wedge \\ [\mathbf{A} \mid \mathbf{I}] \cdot \mathbf{u} = \mathbf{0} \end{array} \mid \mathbf{A} \leftarrow R_q^{k \times \ell}, \mathbf{u} \leftarrow \mathcal{A}(1^\lambda, \mathbf{A}) \right].$$

Definition 3.25 (Module Learning with Errors) Let n, q, k, ℓ be integers and D a probability distribution over R_q . For any $\mathbf{A} \in R_q^{k \times \ell}$, define two oracles as follows:

- \mathcal{O}_A : Sample $(\mathbf{s}, \mathbf{e}) \leftarrow D^k \times D^\ell$ and output $\mathbf{A}\mathbf{s} + \mathbf{e} \in R_q^k$,
- \mathcal{O}_S : Output a random $\mathbf{b} \leftarrow R_q^k$.

The advantage for the *decision module learning with errors problem* $\text{sMLWE}_{n,q,k,\ell,D}$ for an algorithm \mathcal{A} is defined as

$$\text{Adv}_{n,q,k,\ell,D}^{\text{dMLWE}}(\mathcal{A}) = \left| \Pr[\mathcal{A}^{\mathcal{O}_A}(1^\lambda, \mathbf{A}) \rightarrow 1] - \Pr[\mathcal{A}^{\mathcal{O}_S}(1^\lambda, \mathbf{A}) \rightarrow 1] \right|,$$

where the probability is taken also over the random choice of $\mathbf{A} \leftarrow R_q^{k \times \ell}$.

The advantage for the *search learning with errors problem* $\text{sMLWE}_{n,q,k,\ell,D}$ is defined as

$$\text{Adv}_{n,q,k,\ell,D}^{\text{sMLWE}}(\mathcal{A}) = \Pr \left[\begin{array}{l} \mathbf{v} = \mathbf{A}\mathbf{s} + \mathbf{e} \wedge \\ (\mathbf{s}, \mathbf{e}) \in \text{Supp}(D^\ell) \times \text{Supp}(D^k) \end{array} \mid (\mathbf{s}, \mathbf{e}) \leftarrow \mathcal{A}^{\mathcal{O}_A}(1^\lambda, \mathbf{A}) \right],$$

where \mathbf{v} is one of the vectors returned by \mathcal{O}_A .

In this work, we consider the MLWE problem where an adversary is given oracle access to a MLWE sample generator. For any PPT adversary \mathcal{A} , this is polynomially related to the conventional single-instance MLWE problem via a standard hybrid argument. There is also a simple *tight* reduction from the single-instance to the multi-instance MLWE problem *à la* “noise-flooding,” where (roughly) the support of the distribution D considered by the multi-instance problem is required to be super-polynomially larger than those considered by the single-instance problem. However, practically speaking, to the best of our knowledge, we are not aware of any attacks that exploit the multiplicity of the MLWE sample. Therefore, throughout this work, we assume the multi-instance MLWE problem to be as difficult as the single-instance MLWE problem.

The assumption on the hardness of (multi-instance) MLWE is believed to hold even when D is the uniform distribution over ring elements with infinity norm at most a fixed value B , say $B \approx 5$, for appropriate choices of n, q, k, ℓ [2]. We write $\text{MLWE}_{n,q,k,\ell,B}$ when we consider such distribution D . For example, the round-2 NIST candidate signature scheme Dilithium [42] uses such parameters for the (single-instance) MLWE problem, and in particular, our scheme borrows the same parameter sets.

4 Generic construction of accountable ring signature and dynamic group signature

In this section, we present novel generic frameworks for accountable ring signature, dynamic group signature, and their tightly secure variants. Firstly, we introduce a generic construction of an accountable ring signature in Sect. 4.1. Constructing a dynamic group signature immediately follows by limiting the functionality of accountable ring signature. Our construction achieves a tighter reduction compared to prior works on efficient group signatures

as it does not rely on the forking lemma [50, 81]. However, since we still lose a factor of $1/N$ in the reduction, we finally show how to modify our construction to be truly tight using the Katz-Wang technique [61] in Sect. 4.3.

4.1 Generic construction of accountable ring signature

In this subsection, we present our generic construction of an accountable ring signature scheme. Before diving in the details we give a brief overview of our generic construction. The setup is as follows. The opening authorities generate a PKE key-pair, denoted as (opk, osk) to indicate that they are the opener’s keys, and publish the opening public key opk . The users generate an element (x, w) in a hard relation R , and publish the statement x as verification key, and keep the witness w as secret signing key. A signature for our ARS scheme for a ring $R = \{x_1, \dots, x_N\}$ consists of a ciphertext ct , and a NIZK proof that: 1) The ciphertext is an encryption of an index $I \in [N]$ under an opener public key opk , and 2) that the signer knows a witness w corresponding to the I -th statement x_I in the ring R . The second property ensures that the signature is unforgeable, and the first property ensures that the opener (who has the secret key opk) can decrypt the ciphertext to find out who the real signer is. To convince others that a signature was produced by the I -th member of the ring, the opener uses a second NIZK proof to prove that he knows an opener secret key osk that is consistent with opk , and such that $\text{Dec}(\text{osk}, \text{ct}) = I$. If the opener could find a second secret key osk' , consistent with opk and such that ct decrypts to $I' \neq I$ under osk' , then the opener could frame I' for signing a signature, which breaks the tracing soundness of the signature scheme. To prevent this we require the PKE to satisfy a strong correctness property, which says that an encryption of I will always decrypt to I , even if the encryption randomness and decryption key are invalid (in some specific, controlled way). More formally we define the following special correctness notion for a PKE scheme.

Definition 4.1 ($(\mathcal{R}', \mathcal{KR}')$ -correctness) Consider a public-key encryption scheme $\Pi_{\text{PKE}} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$, with \mathcal{R} the set containing all possible randomness used by Enc and \mathcal{KR} the binary relation that contains all the key pairs (pk, sk) that can be generated by running KeyGen . Let \mathcal{R}' be a set containing \mathcal{R} , and \mathcal{KR}' a relation containing \mathcal{KR} . Then we say that Π_{PKE} is $(\mathcal{R}', \mathcal{KR}')$ -correct if, for all $\lambda \in \mathbb{N}$, and for all but a negligible fraction of $\text{pp} \in \text{Setup}(1^\lambda)$, we have for all $(\text{pk}, \text{sk}) \in \mathcal{KR}'$, for all messages m in the plaintext space \mathcal{M} , and all $r \in \mathcal{R}'$ that

$$\text{Dec}(\text{sk}, \text{Enc}(\text{pk}, m; r)) = m.$$

Remark 4.2 Note that pp is also implicitly used in the relations $\mathcal{KR}, \mathcal{KR}'$. If $\mathcal{R}' = \mathcal{R}$ and $\mathcal{KR}' = \mathcal{KR}$, then the $(\mathcal{R}', \mathcal{KR}')$ -correctness is exactly the standard correctness property for PKEs. If \mathcal{R}' or \mathcal{KR}' is larger than \mathcal{R} or \mathcal{KR} , respectively, then the definition becomes a stronger property, because the decryption algorithm is required to decrypt correctly even when the encryption algorithm used some invalid randomness, and/or when the keypair is invalid. (\mathcal{R}' and \mathcal{KR}' control how “invalid” randomness and secret key are allowed to be.)

Our generic construction of an accountable ring signature scheme $\Pi_{\text{ARS}} = (\text{ARS.Setup}, \text{ARS.OKGen}, \text{ARS.UKGen}, \text{ARS.Sign}, \text{ARS.Verify}, \text{ARS.Open}, \text{ARS.Judge})$, provide in Fig. 1, is based on the following building blocks:

- A hard-instance generator contains a setup algorithm RelSetup that, on input a security parameter λ , outputs a description pp of a pair of binary relations $R_{\text{pp}} \subseteq \tilde{R}_{\text{pp}}$, and an

ARS.Setup(1^λ)

1: $pp_1 \leftarrow \text{RelSetup}(1^\lambda)$
 2: $pp_2 \leftarrow \text{PKE.Setup}(1^\lambda)$
 3: **return** $pp = (pp_1, pp_2)$

ARS.UKGen(pp)

1: $(x, w) \leftarrow \text{IGen}(pp_1)$
 2: **return** $(vk := x, sk := w)$

ARS.Verify(opk, R, M, σ)

1: $(ct, \pi_{\text{sign}}) \leftarrow \sigma$
 2: $\text{tmp} \leftarrow (M, (R, \text{opk}, ct), \pi_{\text{sign}})$
 3: **return** $\text{NIZK.Verify}_{\text{lbl}}(\text{tmp})$

ARS.Judge(opk, R, vk, M, $\sigma, \pi_{\text{open}}$)

1: $(ct, \pi_{\text{sign}}) \leftarrow \sigma$
 2: **if** $\nexists I : vk = R_I$ **then**
 3: **return** \perp .
 4: **end if**
 5: $b_0 \leftarrow \text{ARS.Verify}(opk, R, M, \sigma)$
 6: $b_1 \leftarrow \text{NIZK.Verify}((opk, ct, I), \pi_{\text{open}})$
 7: **return** $b_0 \wedge b_1$

ARS.OKGen(pp)

1: $(pk, sk) \leftarrow \text{KeyGen}(pp_2)$
 2: **return** $(opk := pk, osk := sk)$

ARS.Sign(opk, sk, R, M)

1: $\{x_i\}_{i \in [N]} \leftarrow R$
 2: **if** $\nexists I : (x_I, sk) \in R_{pp_1}$ **then**
 3: **return** \perp .
 4: **end if**
 5: $r \xleftarrow{\$} \mathcal{R}$
 6: $ct = \text{Enc}(opk, I; r)$
 7: $\text{tmp} \leftarrow (M, (R, \text{opk}, ct), (I, sk, r))$
 8: $\pi_{\text{sign}} \leftarrow \text{NIZK.Prove}_{\text{lbl}}(\text{tmp})$
 9: **return** $\sigma := (ct, \pi_{\text{sign}})$

ARS.Open(osk, R, M, σ)

1: **if** $\text{ARS.Verify}(opk, R, M, \sigma) = \perp$
 then
 2: **return** \perp
 3: **end if**
 4: $(ct, \pi_{\text{sign}}) \leftarrow \sigma$
 5: $I \leftarrow \text{Dec}(osk, ct)$
 6: $\text{tmp} \leftarrow ((opk, ct, I), osk)$
 7: $\pi_{\text{open}} \leftarrow \text{NIZK.Prove}(\text{tmp})$
 8: **return** $\pi := (R_I, \pi_{\text{open}})$

Fig. 1 Generic construction of an accountable ring signature Π_{ARS} obtained from a hard instance generator (RelSetup, IGen), a public-key encryption algorithm (PKE.Setup, KeyGen, Enc, Dec) satisfying some suitable security and correctness properties, a NIZK with labels $\Pi_{\text{NIZK, lbl}}$ for R_{sig} , and a NIZK without labels Π_{NIZK} for R_{open} . The public parameter pp is provided to all algorithms where we may omit them for readability

instance generator IGen for those pairs of relations. That is, RelSetup and IGen are PPT algorithms such that $\Pr[(x, w) \in R_{pp} \mid pp \leftarrow \text{RelSetup}(1^\lambda); (x, w) \leftarrow \text{IGen}(pp)] = 1$, and such that if we define the advantage of an adversary \mathcal{A} against (RelSetup, IGen) as

$$\text{Adv}_{\text{RelSetup, IGen}}^{\text{Hard}}(\mathcal{A}) = \Pr \left[(x, w') \in \tilde{R}_{pp} \mid \begin{array}{l} pp \leftarrow \text{RelSetup}(1^\lambda) \\ (x, w) \leftarrow \text{IGen}(pp) \\ w' \leftarrow \mathcal{A}(pp, x) \end{array} \right],$$

then $\text{Adv}_{\text{RelSetup, IGen}}^{\text{Hard}}(\mathcal{A})$ is a negligible function of λ for every PPT adversary \mathcal{A} .

- A public-key encryption scheme $\Pi_{\text{PKE}} = (\text{PKE.Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ with multi-challenge IND-CPA security, and with $(\mathcal{R}', \mathcal{KR}')$ -correctness for some relaxed randomness set \mathcal{R}' and some relaxed key relation \mathcal{KR}' . The message space of the encryption scheme contains a set of indices $[N]$ for any polynomially large $N \in \mathbb{N}$.
- A multi-proof online extractable NIZK proof system with labels $\Pi_{\text{NIZK, lbl}} = (\text{NIZK.Setup}_{\text{lbl}}, \text{NIZK.Prove}_{\text{lbl}}, \text{NIZK.Verify}_{\text{lbl}})$ for the relations

$$R_{\text{sig}} = \{((\{x_i\}_{i \in [N]}, \text{pk}, ct), (I, w, r)) \mid (x_I, w) \in R_{pp} \wedge ct = \text{Enc}(\text{pk}, I; r)\}$$

$$\tilde{R}_{\text{sig}} = \{((\{x_i\}_{i \in [N]}, \text{pk}, ct), (I, w, r)) \mid (x_I, w) \in \tilde{R}_{pp} \wedge ct = \text{Enc}(\text{pk}, I; r)\}.$$

To be precise, we need to also include the public parameters output by RelSetup and PKE.Setup in the statement. We omit them for better readability.

- A statistically sound NIZK proof system (without labels) $\Pi_{\text{NIZK}} = (\text{NIZK.Setup}, \text{NIZK.Prove}, \text{NIZK.Verify})$ for the relations

$$R_{\text{open}} = \{((\text{pk}, \text{ct}, I), \text{sk})(\text{pk}, \text{sk}) \in \mathcal{KR} \wedge \text{Dec}(\text{sk}, \text{ct}) = I\}$$

$$\tilde{R}_{\text{open}} = \{((\text{pk}, \text{ct}, I), \text{sk})(\text{pk}, \text{sk}) \in \mathcal{KR}' \wedge \text{Dec}(\text{sk}, \text{ct}) = I\}.$$

Similarly to above, we omit the public parameter output by PKE.Setup in the statement. We emphasize that Π_{NIZK} does not need to be online extractable.

Correctness and security of the proposed accountable ring signature scheme Π_{ARS} are shown in the following theorems.

Theorem 4.3 *The accountable ring signature scheme Π_{ARS} in Fig. 1 is correct.*

Proof Due to the correctness of the underlying NIZK proof system, $\Pi_{\text{NIZK}, \text{lbl}}$, any signature output by ARS.Sign will be accepted by ARS.Verify with probability 1. \square

Theorem 4.4 *The accountable ring signature scheme Π_{ARS} in Fig. 1 is (CCA) anonymous (against full key exposure) in the random oracle model, assuming Π_{PKE} is multi-challenge IND-CPA secure and $(\mathcal{R}', \mathcal{KR}')$ -correct, $\Pi_{\text{NIZK}, \text{lbl}}$ is zero-knowledge, multi-challenge online-extractable, and Π_{NIZK} is zero-knowledge. Precisely, for an adversary \mathcal{A} , running in time T , there exist PPT adversaries $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3, \mathcal{B}_4$, with running times $O(T)$ such that*

$$\text{Adv}_{\Pi_{\text{ARS}}}^{\text{Anon}}(\mathcal{A}) \leq \text{Adv}_{\Pi_{\text{NIZK}}}^{\text{ZK}}(\mathcal{B}_1) + \text{Adv}_{\Pi_{\text{NIZK}, \text{lbl}}}^{\text{OE}}(\mathcal{B}_2) + \text{Adv}_{\Pi_{\text{NIZK}, \text{lbl}}}^{\text{ZK}}(\mathcal{B}_3) + \text{Adv}_{\Pi_{\text{PKE}}}^{\text{Multi-CPA}}(\mathcal{B}_4).$$

Proof We prove anonymity using a hybrid argument with the following series of games. Let the advantage of the adversary \mathcal{A} in Game_{*i*} be denoted by $\text{Adv}_i(\mathcal{A})$.

Game₁ : This is the original anonymity game defined in Def. 3.17. The adversary’s advantage in this game is $\text{Adv}_1(\mathcal{A}) = \text{Adv}_{\Pi_{\text{ARS}}}^{\text{Anon}}(\mathcal{A})$ by definition.

Game₂ : This is the same as Game₁, except that it uses the simulator $\text{NIZK.Sim} = (\text{NIZK.Sim}_0, \text{NIZK.Sim}_1)$ for Π_{NIZK} to answer random-oracle and opening queries from the adversary. When \mathcal{A} makes a random oracle query, the challenger forwards the query to NIZK.Sim_0 , records the query and answers, and forwards the answer to \mathcal{A} . When \mathcal{A} makes an opening query, rather than computing π_{open} using NIZK.Prove and osk , the challenger instead uses the output of NIZK.Sim_1 . We consider an adversary \mathcal{B}_1 against the zero-knowledge property of Π_{NIZK} which simulates Game₂ for \mathcal{A} . Let Prove and \mathcal{S} be as in the definition of zero-knowledge for the NIZK proof system. Then, if \mathcal{B}_1 ’s oracle queries are answered by $(\mathcal{O}, \text{Prove})$ the game is identical to Game₁, and if queries are answered by $(\text{NIZK.Sim}_0, \mathcal{S})$, then the game is identical to Game₂. Therefore, assuming \mathcal{B}_1 outputs 1 when \mathcal{A} wins, we have $\text{Adv}_1(\mathcal{A}) \leq \text{Adv}_2(\mathcal{A}) + \text{Adv}_{\Pi_{\text{NIZK}}}^{\text{ZK}}(\mathcal{B}_1)$.

Game₃ : This is the same as Game₂, except that the way the challenger answers opening queries is further modified. Rather than using the secret key osk to decrypt the ciphertext ct and identify the index I of the real signing key (as ARS.Open does in the honest protocol), the challenger instead runs the online extractor OnlineExtract for $\Pi_{\text{NIZK}, \text{lbl}}$ to extract the witness (I, sk, r) from $(\text{ct}, \pi_{\text{sign}})$, and then returns the user R_I . We consider an adversary \mathcal{B}_2 against the online extractability of $\Pi_{\text{NIZK}, \text{lbl}}$ that simulates Game₃ for \mathcal{A} such that

- random-oracle queries from \mathcal{A} are replied by querying (hash, \cdot) (see Def. 3.10);
- instead of computing π_{sign} when answering a signing query, \mathcal{B}_2 makes a query (prove, M, x, w) , where $(x, w) = ((R, \text{opk}, \text{ct}), (I, \text{sk}, r))$, and
- instead of running OnlineExtract , \mathcal{B}_2 makes a query $(\text{extract}, M, x, \pi_{\text{sign}})$.

Note that extract for proofs originating from prove queries are answered with \perp , which is compatible with the fact that the challenger outputs \perp for opening queries that correspond to signatures originating from the signing oracle in Game_3 . If \mathcal{B}_2 loses the multi-proof online extractability game (i.e., \mathcal{B}_2 did not cause the extractor to fail), then it follows from the $(\mathcal{R}', \mathcal{KR}')$ -correctness of Π_{PKE} that for each extraction $W = (I, \text{sk}, r)$ we have $\text{Dec}(\text{osk}, \text{ct}) = \text{Dec}(\text{osk}, \text{Enc}(\text{opk}, I; r)) = I$, so the view of \mathcal{A} is not affected by whether I was obtained from OnlineExtract or by decrypting ct with osk . Therefore, we have $\text{Adv}_2(\mathcal{A}) \leq \text{Adv}_3(\mathcal{A}) + \text{Adv}_{\Pi_{\text{NIZK}, \text{lbl}}}^{\text{OE}}(\mathcal{B}_2)$.

Game₄ : This is the same as Game_3 , except that we change how the challenger answers signing queries from the adversary: The challenger generates ct as in Game_3 , but uses the zero-knowledge simulator Sim for $\Pi_{\text{NIZK}, \text{lbl}}$ to create the proof π_{sign} rather than using $\text{NIZK.Prove}_{\text{lbl}}$. It then outputs $(\text{ct}, \pi_{\text{sign}})$ as the signature. Similarly to the transition from Game_1 to Game_2 , we can define an adversary \mathcal{B}_3 against the zero-knowledge property of $\Pi_{\text{NIZK}, \text{lbl}}$ such that $\text{Adv}_3(\mathcal{A}) \leq \text{Adv}_4(\mathcal{A}) + \text{Adv}_{\Pi_{\text{NIZK}, \text{lbl}}}^{\text{ZK}}(\mathcal{B}_3)$.

Game₅ : This is the same as Game_4 , except we further change how the challenger answers signing queries: Instead of encrypting the correct index I to obtain ct , the challenger encrypts a random index I' . We define a multi-challenge IND-CPA adversary \mathcal{B}_4 for Π_{PKE} that simulates Game_5 for \mathcal{A} , but instead of generating (opk, osk) , the adversary \mathcal{B}_4 receives opk from the multi-challenge IND-CPA challenger, and instead of producing the ciphertexts ct the adversary \mathcal{B}_4 makes encryption queries (I, I') , where I is the correct index, and I' is a random index. Note that, say on input $(\text{sign}, \mathbf{R}, M, \text{sk}_0, \text{sk}_1)$, the I -th key in \mathbf{R} is the verification key corresponding to sk_0 . We can make this replacement because in Game_5 , the challenger does not use osk . (The purpose of Game_2 and Game_3 were to remove the use of osk for this reason.) If the hidden bit b in the IND-CPA game is 0, then the IND-CPA experiment is identical to Game_4 , and if the bit is 1, then the experiment is equal to Game_5 . Therefore, we have that $\text{Adv}_4(\mathcal{A}) \leq \text{Adv}_5(\mathcal{A}) + \text{Adv}_{\Pi_{\text{PKE}}}^{\text{Multi-CPA}}(\mathcal{B}_4)$.

Finally, observe that in Game_5 the challenger leaks no information about the secret bit b because b is not used. Hence, $\text{Adv}_5(\mathcal{A}) = 0$. □

Remark 4.5 In the previous proof we really relied on the online extractability property (without rewinding). This is because, even if we allow for a non-tight reduction, we cannot resort to rewinding (i.e., the forking lemma) since there can be polynomially many open queries and the reduction loss will be exponential if we try to extract from all of them. Here, keep in mind that the online extractor must succeed with (roughly) $1 - \text{negl}(\lambda)$ rather than any non-negligible function $1/\text{poly}(\lambda)$ since there can be polynomially many open queries. Namely, even a success probability of $1/2$ will not be good enough. Most, if not all, prior works circumvent this issue by using an IND-CCA PKE as building block rather than a (possibly inefficient) online extractable NIZK to simulate the decryption of ct .

Theorem 4.6 *The accountable ring signature scheme Π_{ARS} in Fig. 1 is unforgeable in the random oracle model. More precisely, for any adversary \mathcal{A} that runs in time T and makes Q_u queries to the ukeygen oracle, there exist adversaries $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$, running in time $O(T)$,*

such that

$$Adv_{\Pi_{ARS}}^{Unf}(\mathcal{A}) \leq Adv_{\Pi_{NIZK,|bl}}^{OE}(\mathcal{B}_1) + Adv_{\Pi_{NIZK,|bl}}^{ZK}(\mathcal{B}_2) + Q_u Adv_{RelSetup, IGen}^{Hard}(\mathcal{B}_3)$$

Proof We prove unforgeability using a hybrid argument with the following series of games. Let the advantage of the adversary \mathcal{A} in Game_{*i*} be denoted by $Adv_i(\mathcal{A})$.

- Game₁ : This is the original unforgeability game defined in Def. 3.18. The adversary’s advantage in this game is $Adv_1(\mathcal{A}) = Adv_{\Pi_{ARS}}^{Unf}(\mathcal{A})$ by definition.
- Game₂ : This is the same as Game₁, but the winning condition is changed. We let the challenger maintain a list $L_{\mathcal{O}}$ of all the random oracle queries that \mathcal{A} makes. When \mathcal{A} finishes the game by outputting $(\text{opk}, \text{vk}, R, M, \sigma = (\text{ct}, \pi_{\text{sign}}), \pi)$, the challenger runs $(I, \text{sk}, r) \leftarrow \text{OnlineExtract}(M, (R, \text{opk}, \text{ct}), \pi_{\text{sign}}, L_{\mathcal{O}})$. The game results in a loss if $((R, \text{opk}, \text{ct}), (I, \text{sk}, r)) \notin \tilde{R}_{\text{sig}}$, otherwise, the winning condition is not changed. We construct an online-extractability adversary \mathcal{B}_1 for $\Pi_{NIZK,|bl}$ that simulates Game₂ for \mathcal{A} . He replies random-oracle queries from \mathcal{A} by querying (hash, \cdot) (see Def. 3.10), signing queries by making an oracle call $(\text{prove}, M, (R, \text{opk}, \text{ct}), (I, \text{sk}, r))$ instead of computing π_{sign} himself, and makes the oracle call $(\text{extract}, M, (R, \text{opk}, \text{ct}), \pi_{\text{sign}})$ instead of running OnlineExtract . The view of \mathcal{A} during the game simulated by \mathcal{B}_1 is identical to its view during Game₁ and Game₂. Suppose that the output received by \mathcal{A} is a win for the winning condition of Game₁, but a loss for the winning condition of Game₂. This means that $\text{NIZK.Verify}_{|bl}^{\mathcal{O}}(M, (R, \text{opk}, \text{ct}), \pi_{\text{sign}}) = \top$ and $(\text{ct}, \pi_{\text{sign}})$ was not the output of a query $(\text{sign}, \text{opk}, \text{vk}', R, M)$ for any vk' , otherwise the winning condition of Game₁ would not be met. Moreover, we would have $((R, \text{opk}, \text{ct}), (I, \text{sk}, r)) \notin \tilde{R}_{\text{sig}}$, otherwise the winning condition of Game₂ would be met. This is precisely the situation \mathcal{B}_1 needs in order to win the online extractability game. Therefore, we have $Adv_1(\mathcal{A}) \leq Adv_2(\mathcal{A}) + Adv_{\Pi_{NIZK,|bl}}^{OE}(\mathcal{B}_1)$
- Game₃ : This is the same as Game₂ except that we change the way the challenger answers signing queries from \mathcal{A} . Specifically, the challenger generates ct as in Game₂ but uses the zero-knowledge simulator $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ for $\Pi_{NIZK,|bl}$ to create the proof π_{sign} . That is, it forwards the random-oracle queries to Sim_0 , and runs Sim_1 to get π_{sign} . It then outputs $(\text{ct}, \pi_{\text{sign}})$ as the signature. Let \mathcal{B}_2 be an adversary against the zero-knowledge property of $\Pi_{NIZK,|bl}$, which simulates Game₃ for \mathcal{A} by forwarding random-oracle queries and proving queries to the oracles Sim_0 and Sim_1 , respectively. If \mathcal{B}_2 is given access to oracles \mathcal{O} and Prove (see Def. 3.8), then \mathcal{A} ’s view is identical to Game₂, and if \mathcal{B}_2 is run with access to $\text{Sim}_0, \text{Sim}_1$, then \mathcal{A} ’s view is identical to Game₃. Therefore, we have $Adv_2(\mathcal{A}) \leq Adv_3(\mathcal{A}) + Adv_{\Pi_{NIZK,|bl}}^{ZK}(\mathcal{B}_2)$.
- Game₄ : This is the same as Game₃ except that we change the winning condition again: the challenger guesses a random index $\tilde{I} \in \{1, \dots, Q_u\}$ at the outset of the game. If \mathcal{A} makes a corruption query to corrupt the verification key returned in the \tilde{I} -th user key generation query, then Game₄ aborts. The game results in a win if the winning condition of Game₃ is met and if $\tilde{I} = I$. Since I is information-theoretically hidden during the execution of the game, we have $\tilde{I} = I$ with probability $1/Q_u$. Therefore, we have $Adv_3(\mathcal{A}) = Q_u Adv_4(\mathcal{A})$.

Finally, let \mathcal{B}_3 be an adversary against $(\text{RelSetup}, \text{IGen})$ which simulates Game₄ for \mathcal{A} . At the beginning of the game, \mathcal{B}_3 is given an instance (pp_1, x) . The adversary \mathcal{B}_3 simulates an execution of Game₄ by using the public parameter pp_1 that is given to him, rather than generating a new pp_1 himself using RelSetup , and by answering the \tilde{I} -th ukeygen query assigning $\text{vk}_{\tilde{I}} = x$ instead of running $(x, w) \leftarrow \text{IGen}(\text{pp}_1)$. Note that \mathcal{B}_3 does not need

w because if \mathcal{A} makes a query to corrupt $vk_{\tilde{\gamma}}$ then the game aborts. The view of \mathcal{A} during \mathcal{B}_3 's simulation is the same as its view during a real execution of Game_4 , so OnlineExtract outputs a valid witness (\tilde{I}, sk, r) with probability at least $\text{Adv}_4(\mathcal{A})$. If this is the case, then \mathcal{B}_3 wins his game against the hardness of $(\text{RelSetup}, \text{lGen})$ by outputting sk . Therefore, we have $\text{Adv}_4(\mathcal{A}) \leq \text{Adv}_{\text{RelSetup}, \text{lGen}}^{\text{Hard}}(\mathcal{B}_3)$. \square

Theorem 4.7 *The accountable ring signature scheme Π_{ARS} in Fig. 1 is traceable and tracing sound in the random oracle model. More precisely, for any adversary \mathcal{A} that runs in time T , we have adversaries $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ that run in time $O(T)$, such that*

$$\text{Adv}_{\Pi_{\text{ARS}}}^{\text{Tra}}(\mathcal{A}) \leq \text{Adv}_{\Pi_{\text{NIZK}, \text{lbl}}}^{\text{soundness}}(\mathcal{B}_1)$$

and

$$\text{Adv}_{\Pi_{\text{ARS}}}^{\text{TraS}}(\mathcal{A}) \leq \text{Adv}_{\Pi_{\text{NIZK}, \text{lbl}}}^{\text{soundness}}(\mathcal{B}_2) + 2\text{Adv}_{\Pi_{\text{NIZK}}}^{\text{soundness}}(\mathcal{B}_3)$$

Proof We prove the two properties separately as follows:

Traceability Traceability follows from the statistical soundness of $\Pi_{\text{NIZK}, \text{lbl}}$, the $(\mathcal{R}', \mathcal{KR}')$ -correctness of Π_{PKE} , and the correctness of Π_{NIZK} . Observe that if \mathcal{A} wins an execution of the traceability game, then $\text{NIZK.Verify}_{\text{lbl}}(M, X = (R, \text{opk}, \text{ct}), \pi_{\text{sign}}) = \top$, but still there cannot be a witness $W = (I, sk, r)$ such that $(X, W) \in \tilde{R}_{\text{sig}}$. Towards a contradiction, suppose that such a witness does exist, then the $(\mathcal{R}', \mathcal{KR}')$ -correctness of the PKE implies that $\text{Dec}(\text{osk}, \text{ct} = \text{Enc}(\text{opk}, I; r)) = I$, which implies that $((\text{opk}, \text{ct}, I), \text{osk}) \in R_{\text{open}}$, so the correctness of Π_{NIZK} implies that $\text{NIZK.Verify}((\text{opk}, \text{ct}, I), \pi_{\text{open}}) = \top$. This means that \mathcal{A} did not win the traceability game. Therefore, \mathcal{A} produces valid proofs for statements not in \tilde{R}_{sig} with probability at least $\text{Adv}_{\Pi_{\text{ARS}}}^{\text{Tra}}(\mathcal{A})$. We can use this to construct an adversary \mathcal{B}_1 against the statistical soundness of $\Pi_{\text{NIZK}, \text{lbl}}$ that generates $\text{pp} \leftarrow \text{ARS.Setup}(1^\lambda)$ for a security parameter λ , runs $(rr, R, M, \sigma) \leftarrow \mathcal{A}(\text{pp})$ where $\sigma = (\text{ct}, \pi_{\text{sign}})$, and $(\text{osk}, \text{opk}) \leftarrow \text{ARS.OKGen}(\text{pp}; rr)$, and outputs $(M, x := (R, \text{opk}, \text{ct}), \pi_{\text{sign}})$, which makes \mathcal{B}_1 win. \mathcal{B}_1 's advantage is therefore $\text{Adv}_{\Pi_{\text{ARS}}}^{\text{Tra}}(\mathcal{A}) \leq \text{Adv}_{\Pi_{\text{NIZK}, \text{lbl}}}^{\text{soundness}}(\mathcal{B}_1)$.

Tracing soundness Similarly, tracing soundness follows from the statistical soundness of Π_{NIZK} and $\Pi_{\text{NIZK}, \text{lbl}}$, and the $(\mathcal{R}', \mathcal{KR}')$ -correctness of the Π_{PKE} . In order for \mathcal{A} to win the tracing soundness game, it needs to output valid proofs $\pi_{\text{sign}}, \pi_0, \pi_1$ (the former is part of the produced signature $\sigma = (\text{ct}, \pi_{\text{sign}})$) such that there exist witnesses $(I, sk, r), \text{osk}_0$ and osk_1 where

$$\begin{aligned} ((R, \text{opk}, \text{ct}), (I, sk, r)) &\in \tilde{R}_{\text{sig}} \\ ((\text{opk}, \text{ct}, I_0), \text{osk}_0) &\in \tilde{R}_{\text{open}} \\ ((\text{opk}, \text{ct}, I_1), \text{osk}_1) &\in \tilde{R}_{\text{open}}, \end{aligned}$$

with $I_0 \neq I_1$. However, it follows from the $(\mathcal{R}', \mathcal{KR}')$ -correctness of Π_{PKE} that no three such witnesses can exist. Suppose, towards a contradiction, that those witnesses exist. Then we have $I_0 = \text{Dec}(\text{osk}_0, \text{ct} = \text{Enc}(\text{opk}, I; r))$, so the $(\mathcal{R}', \mathcal{KR}')$ -correctness implies that $I_0 = I$, and similarly it follows from $I_1 = \text{Dec}(\text{osk}_1, \text{ct} = \text{Enc}(\text{opk}, I; r))$ that $I_1 = I$, which contradicts $I_0 \neq I_1$. Therefore, at least one of $\pi_{\text{sign}}, \pi_0, \pi_1$ is a valid proof of an invalid statement, i.e. a X for which does not exist W such that $(X, W) \in \tilde{R}_{\text{sig}}$ (or $(v) \in \tilde{R}_{\text{open}}$), with probability at least $\text{Adv}_{\Pi_{\text{ARS}}}^{\text{TraS}}(\mathcal{A})$. Let \mathcal{B}_2 and \mathcal{B}_3 be statistical-soundness adversaries for $\Pi_{\text{NIZK}, \text{lbl}}$ and Π_{NIZK} , respectively, that simulate the tracing soundness game and output π_{sign} or π_b , respectively, where b is a random bit. Then we have $\text{Adv}_{\Pi_{\text{ARS}}}^{\text{TraS}}(\mathcal{A}) \leq \text{Adv}_{\Pi_{\text{NIZK}, \text{lbl}}}^{\text{soundness}}(\mathcal{B}_2) + 2\text{Adv}_{\Pi_{\text{NIZK}}}^{\text{soundness}}(\mathcal{B}_3)$. \square

4.2 Accountable ring signature to dynamic group signature

Accountable ring signatures are known to trivially imply dynamic group signatures [19, 20]. A formal treatment is provided by Bootle et al. [20]. We remark that the transformation provided in [20] retains the same level of security provided by the underlying accountable ring signature. That is, all reductions between unforgeability, full-anonymity and traceability are tight. For completeness, we provide more details on group signatures and the transformation in App. B.

4.3 Tightly secure variant

Observe the only source of loose reduction in the previous section was in the unforgeability proof (see Theorem 4.6), where we assume each building blocks, i.e., NIZK and PKE, are tightly reduced to concrete hardness assumptions. In this subsection, we apply the Katz-Wang technique [61] to modify our construction in Fig. 1 to obtain a tight reduction.

We firstly give an intuition of the method. Recall that in the proof of Theorem 4.6, the reduction is given a challenge instance x , guesses which user's signature the adversary will forge, and assigns x to the verification key vk of the selected user. If the adversary queries the corruption oracle on the key vk , the reduction fails and aborts since it will not be able to produce the corresponding secret key for vk . If the guess is correct and the adversary successfully forges the signature, then the reduction can recover a witness w' such that (x, w') is in the relation \tilde{R}_{pp_1} . Therefore, if the adversary makes Q_u user key generation queries and its advantage is ϵ , then the reduction can extract a witness with probability roughly ϵ/Q_u . A high-level viewpoint of the Katz-Wang method is that each user is given a pair of statements $(x^{(1)}, x^{(2)})$ as the verification key vk , with only one witness w as the secret signing key, such that either $(x^{(1)}, w)$ or $(x^{(2)}, w)$ is in the relation \tilde{R}_{pp_1} . Also, we assume that now the reduction is given Q_u challenge instances $\{x_i\}_{i \in [Q_u]}$ and it is required to solve any one of them. The reduction in this case needs no guessing steps as above. Specifically, the reduction can use IGen to generate pairs $(\tilde{x}_i, \tilde{w}_i)$ for $i \in [Q_u]$, randomly permutes x_i, \tilde{x}_i and assigns the obtained ordered pair to vk_i . Therefore, the reduction can always answer any corruption query with \tilde{w}_i . As long as the adversary wins the unforgeability game by forging a signature, the reduction can return a witness for one of the $\{x_i\}_{i \in [Q_u]}$ with probability $1/2$. Roughly speaking, if the success rate of the adversary is ϵ , then the reduction can extract the answer for the challenge $(\star, X_0, \{x_i\}_{i \in [Q_u]})$ with probability around $\epsilon/2$. Here, it is important that the information on which verification key the user knows the corresponding signing key remains hidden from the adversary. Otherwise, the adversary may always create a forgery with respect to the signing key the reduction already knows.

To turn the above idea into a formal proof, we require two new ingredients: an instance generator that outputs multiple challenges and a NIZK that additionally hides the information on which signing key is used. More formally, we build a tightly secure accountable ring signature scheme $\Pi_{ARS}^{\text{Tight}} = (\text{ARS.Setup}, \text{ARS.OKGen}, \text{ARS.UKGen}, \text{ARS.Sign}, \text{ARS.Verify}, \text{ARS.Open}, \text{ARS.Judge})$ based on the following tools. The only difference between the tools used in Sect. 4.1 are the hard multi-instance generator and the NIZK for the relation $R_{\text{sig}}^{\text{Tight}}$.

- A hard *multi-instance* generator (RelSetup, IGen) contains a setup algorithm RelSetup that outputs a description pp of a pair of relations $R_{pp} \subseteq \tilde{R}_{pp}$, and an instance generator IGen for these pairs of relations. That is, RelSetup and IGen are PPT algorithms such that $\Pr[(x_i, w_i) \in R_{pp} \mid pp \leftarrow \text{RelSetup}(1^\lambda); \{(x_i, w_i)\}_{i \in [N]} \leftarrow \text{IGen}(pp, N)] = 1$.

Moreover, if we define the advantage of an adversary \mathcal{A} against $(\text{RelSetup}, \text{IGen})$ as

$$\text{Adv}_{\text{RelSetup}, \text{IGen}, N}^{\text{Multi-Hard}}(\mathcal{A}) = \Pr \left[(x_i, w') \in \tilde{R}_{\text{pp}} \mid \begin{array}{l} \text{pp} \leftarrow \text{RelSetup}(1^\lambda) \\ \{(x_i, w_i)\}_{i \in [N]} \leftarrow \text{IGen}(\text{pp}, N) \\ (i, w') \leftarrow \mathcal{A}(\text{pp}, \{x_i\}_{i \in [N]}) \end{array} \right]$$

then $\text{Adv}_{\text{RelSetup}, \text{IGen}, N}^{\text{Multi-Hard}}(\mathcal{A})$ is a negligible function in λ for every PPT adversary \mathcal{A} .

- A public-key encryption scheme $\Pi_{\text{PKE}} = (\text{PKE.Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ with multi-challenge IND-CPA security, and with $(\mathcal{R}', \mathcal{KR}')$ -correctness for some relaxed randomness set \mathcal{R}' and some relaxed key relation \mathcal{KR}' . The message space of the encryption scheme contains a set of indices $[N]$ for any polynomially large $N \in \mathbb{N}$.
- A multi-proof online extractable NIZK proof system with labels $\Pi_{\text{NIZK}, \text{lbl}} = (\text{NIZK.Setup}_{\text{lbl}}, \text{NIZK.Prove}_{\text{lbl}}, \text{NIZK.Verify}_{\text{lbl}})$ for the family of relations

$$R_{\text{sig}}^{\text{Tight}} = \left\{ \left((\text{pp}, \{x_i^{(j)}\}_{(i,j) \in [N] \times [2]}, \text{pk}, \text{ct}), (I, b, w, r) \right) \mid \begin{array}{l} (I, r) \in [N] \times \mathcal{R} \wedge (x_I^{(b)}, w) \in R_{\text{pp}} \\ \wedge \text{ct} = \text{Enc}(\text{pk}, I; r) \end{array} \right\}$$

$$\tilde{R}_{\text{sig}}^{\text{Tight}} = \left\{ \left((\text{pp}, \{x_i^{(j)}\}_{(i,j) \in [N] \times [2]}, \text{pk}, \text{ct}), (I, b, w, r) \right) \mid \begin{array}{l} (I, r) \in [N] \times \mathcal{R}' \wedge (x_I^{(b)}, w) \in \tilde{R}_{\text{pp}} \\ \wedge \text{ct} = \text{Enc}(\text{pk}, I; r) \end{array} \right\}.$$

- A second NIZK proof system (without labels) $\Pi_{\text{NIZK}} = (\text{NIZK.Setup}, \text{NIZK.Prove}, \text{NIZK.Verify})$ for the family of relations

$$R_{\text{open}} = \{((\text{pk}, \text{ct}, I), \text{sk})(\text{pk}, \text{sk}) \in \mathcal{KR} \wedge \text{Dec}(\text{sk}, \text{ct}) = I\}$$

$$\tilde{R}_{\text{open}} = \{((\text{pk}, \text{ct}, I), \text{sk})(\text{pk}, \text{sk}) \in \mathcal{KR}' \wedge \text{Dec}(\text{sk}, \text{ct}) = I\},$$

with statistical soundness (Def. 3.9).

The building blocks listed above are combined similarly to Fig. 1. For the sake of completeness, we detail the resulting protocol in Fig. 2. For the security properties, we only focus on unforgeability. The others are a direct consequence of the proofs given for the non-tight construction in Fig. 1.

Theorem 4.8 *The accountable ring signature scheme $\Pi_{\text{ARS}}^{\text{Tight}}$ in Fig. 2 is unforgeable in the random oracle model. More precisely, for any adversary \mathcal{A} that runs in time T and makes Q_u queries to the `ukeygen` oracle, there exist adversaries $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$, running in time $O(T)$, such that*

$$\text{Adv}_{\Pi_{\text{ARS}}}^{\text{Unf}}(\mathcal{A}) \leq \text{Adv}_{\Pi_{\text{NIZK}, \text{lbl}}}^{\text{OE}}(\mathcal{B}_1) + \text{Adv}_{\Pi_{\text{NIZK}, \text{lbl}}}^{\text{ZK}}(\mathcal{B}_2) + 2\text{Adv}_{\text{RelSetup}, \text{IGen}, Q_u}^{\text{Multi-Hard}}(\mathcal{B}_3).$$

Proof We prove unforgeability using a hybrid argument with the following series of games. Let the advantage of an adversary \mathcal{A} in Game_i be denoted by $\text{Adv}_i(\mathcal{A})$.

- The first game, Game_1 , is the original unforgeability game defined in Def. 3.18. The adversary's advantage in this game is $\text{Adv}_1(\mathcal{A}) = \text{Adv}_{\Pi_{\text{ARS}}}^{\text{Unf}}(\mathcal{A})$ by definition.
- Game_2 is the same as Game_1 , but with a modified winning condition. We let the challenger maintain a list $L_{\mathcal{O}}$ of all the random-oracle queries that \mathcal{A} makes. When \mathcal{A} finishes the game by outputting $(\text{opk}, \text{vk}, R, M, \sigma = (\text{ct}, \pi_{\text{sign}}), \pi)$, the challenger runs $(I, b, \text{sk}, r) \leftarrow \text{OnlineExtract}(M, (\text{pp}_1, R, \text{opk}, \text{ct}), \pi_{\text{sign}}, L_{\mathcal{O}})$. The game results in a loss if $((\text{pp}_1, R, \text{opk}, \text{ct}), (I, b, \text{sk}, r)) \notin \tilde{R}_{\text{sig}}^{\text{Tight}}$, otherwise, the winning condition is not changed. As we have shown in the proof of Theorem 4.6, there exists

ARS.Setup(1^λ)

1: $pp_1 \leftarrow \text{RelSetup}(1^\lambda)$
 2: $pp_2 \leftarrow \text{PKE.Setup}(1^\lambda)$
 3: **return** $pp = (pp_1, pp_2)$

ARS.UKGen(pp)

1: $(pp_1, pp_2) \leftarrow pp$
 2: $b \xleftarrow{\$} \{1, 2\}$
 3: $(x^{(1)}, w^{(1)}), (x^{(2)}, w^{(2)}) \leftarrow \text{IGen}(pp_1)$
 4: **return** $(vk := (X^{(1)}, X^{(2)}),$
 $sk := (b, w^{(b)}))$

ARS.Verify(opk, R, M, σ)

1: $(pp_1, pp_2) \leftarrow pp$
 2: $(ct, \pi_{\text{sign}}) \leftarrow \sigma$
 3: $tmp \leftarrow (M, (pp_1, R, opk, ct), \pi_{\text{sign}})$
 4: **return** $\text{NIZK.Verify}_{\text{lbl}}(tmp)$

ARS.Judge($opk, R, vk, M, \sigma, \pi$)

1: $(ct, \pi_{\text{sign}}) \leftarrow \sigma$
 2: **if** $\nexists I : vk = R_I$ **then**
 3: **return** \perp .
 4: **end if**
 5: $b_0 \leftarrow \text{ARS.Verify}(opk, R, M, \sigma)$
 6: $b_1 \leftarrow \text{NIZK.Verify}((opk, ct, I), \pi_{\text{open}})$
 7: **return** $b_0 \wedge b_1$

ARS.OKGen(pp)

1: $(pp_1, pp_2) \leftarrow pp$
 2: $(pk, sk) \leftarrow \text{KeyGen}(pp_2)$
 3: **return** $(opk := pk, osk := sk)$

ARS.Sign(opk, sk, R, M)

1: $(pp_1, pp_2) \leftarrow pp$
 2: $\{(x_i^{(1)}, x_i^{(2)})\}_{i \in [N]} \leftarrow R$
 3: **if** $\nexists (I, b) : (x_I^{(b)}, sk) \in R_{pp_1}$
 then
 4: **return** \perp .
 5: **end if**
 6: $r \xleftarrow{\$} \mathcal{R}$
 7: $ct \leftarrow \text{Enc}(opk, I; r)$
 8: $tmp \leftarrow (M, (pp_1, R, opk, ct),$
 $(I, b, sk, r))$
 9: $\pi_{\text{sign}} \leftarrow \text{NIZK.Prove}_{\text{lbl}}(tmp)$
 10: **return** $\sigma := (ct, \pi_{\text{sign}})$

ARS.Open(osk, R, M, σ)

1: **if** $\text{ARS.Verify}(opk, R, M, \sigma) = \perp$
 then
 2: **return** \perp
 3: **end if**
 4: $(ct, \pi_{\text{sign}}) \leftarrow \sigma$
 5: $I \leftarrow \text{Dec}(osk, ct)$
 6: $\pi_{\text{open}} \leftarrow \text{NIZK.Prove}((opk, ct, I), osk)$
 7: **return** $\pi := (R_I, \pi_{\text{open}})$

Fig. 2 Modified tightly-secure construction of an accountable ring signature $\Pi_{\text{ARS}}^{\text{Tight}}$ obtained from a hard multi-instance generator (RelSetup, IGen), a public-key encryption algorithm (PKE.Setup, KeyGen, Enc, Dec) satisfying some suitable correctness and security properties, a NIZK proof system with labels $\Pi_{\text{NIZK, lbl}}$ for $R_{\text{sig}}^{\text{Tight}}$, and a NIZK proof system without labels Π_{NIZK} for R_{open}

an online-extractability adversary \mathcal{B}_1 for $\Pi_{\text{NIZK, lbl}}$ running in time $O(T)$ such that $\text{Adv}_1(\mathcal{A}) \leq \text{Adv}_2(\mathcal{A}) + \text{Adv}_{\Pi_{\text{NIZK, lbl}}}^{\text{OE}}(\mathcal{B}_1)$.

- The third game, Game_3 , is the same as Game_2 except that we change the way the challenger answers signing queries from \mathcal{A} . Specifically, the challenger generates ct as in Game_2 but uses the $\Pi_{\text{NIZK, lbl}}$ zero-knowledge simulator $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ to create the proof π_{sign} . As we have shown in the proof of Theorem 4.6, there exists a zero-knowledge adversary \mathcal{B}_2 for $\Pi_{\text{NIZK, lbl}}$ running in time $O(T)$ and such that $\text{Adv}_2(\mathcal{A}) \leq \text{Adv}_3(\mathcal{A}) + \text{Adv}_{\Pi_{\text{NIZK, lbl}}}^{\text{ZK}}(\mathcal{B}_2)$.
- Finally, we consider an adversary \mathcal{B}_3 against the hardness of (RelSetup, IGen) which simulates Game_3 for \mathcal{A} . At the beginning of the game, the adversary \mathcal{B}_3 is given the instances $(pp_1, \{x\}_{i \in [Q_u]})$. \mathcal{B}_3 uses the public parameter pp_1 that is given to him, rather than generating new pp_1 himself using RelSetup. Moreover, when answering the i -th `keygen` query, \mathcal{B}_3 uniformly draws b_i from $\{1, 2\}$, generates $(\tilde{x}_i, \tilde{w}_i) \leftarrow \text{IGen}(pp_1)$, and assigns $vk_i = (x_i^{(1)}, x_i^{(2)})$ where $(x_i^{(b_i)}, x_i^{(3-b_i)}) = (\tilde{x}_i, \tilde{x}_i)$. Note that now \mathcal{B}_3 is able to respond to any valid corruption query `corrupt`. In fact, for any $i \in [Q_u]$, if \mathcal{A} makes a

corruption query to corrupt vk_i , then \mathcal{B}_3 responds by $sk = (b_i, \tilde{w}_i)$. The view of \mathcal{A} during \mathcal{B}_3 's simulation is the same as its view during a real execution of Game_3 , so OnlineExtract outputs a valid witness $(\tilde{I}, sk = (b', w'), r)$ with probability at least $\text{Adv}_3(\mathcal{A})$. Since the sampling of the statements and witnesses follows the same distribution determined by $\text{lGen}(\text{pp}_1)$ in the real execution, there is an $1/2$ chance that $b' = (3 - b\tilde{\gamma})$. That is, $(x\tilde{\gamma}, w') \in \tilde{R}_{\text{pp}_1}$. Therefore, we have $\text{Adv}_3(\mathcal{A})/2 \leq \text{Adv}_{\text{RelSetup, lGen, } \mathcal{Q}_u}^{\text{Multi-Hard}}(\mathcal{B}_3)$.

□

5 Group-action-based hard instance generators and PKEs

In this section, we introduce group-action-based hard instance generators (HIGs) and group-action-based PKEs. These are classes of HIGs and PKEs, that derive their security from cryptographic group actions, and which have some specific internal structure. We define these concepts because, as we will see in Sects. 6 and 7, if we instantiate our generic accountable ring signature construction with a group-action-based HIG and a group-action-based PKE, then we can construct a very efficient multi-proof online extractable NIZK for the R_{sig} relation. We provide concrete instantiations of group-action-based HIGs and PKEs from lattices and isogenies in Sect. 8.

5.1 Group-action-based hard instance generator

We consider a special class of hard instance generators naturally induced by cryptographic hard actions.

Definition 5.1 (Group-Action-based Hard Instance Generator) *A group-action-based hard instance generator, GA-HIG in short, is a pair of efficient algorithms (RelSetup, lGen) with the following properties:*

- On input a security parameter λ , RelSetup outputs $\text{pp} = (G, S_1, S_2, \delta, X_0, \mathcal{X}, \star)$ such that: G is an additive group whose elements can be represented uniquely, $S_1 \subseteq S_2$ are symmetric subsets of G , such that membership in S_1 and S_2 can be decided efficiently, and such that the group law can be computed efficiently for elements in $S_1 \cup S_2$. Moreover, the intersection $S_3 = \bigcap_{g \in S_1} g + S_2$ has cardinality $\delta|S_2|$ and membership of S_3 can be decided efficiently. \star is an action $\star : G \times \mathcal{X} \rightarrow \mathcal{X}$ of G on a set \mathcal{X} that contains the element X_0 . \star can be evaluated efficiently on elements of $S_1 \cup S_2$. These parameters describe an NP-relation

$$R_{\text{pp}} = \{(X, s) \mid s \in S_1 : s \star X_0 = X\},$$

and a relaxed NP-relation

$$\tilde{R}_{\text{pp}} = \{(X, s) \mid s \in S_2 + S_3 : s \star X_0 = X\}.$$

- On input pp , lGen samples an element s from S_1 and outputs $(s \star X_0, s) \in R_{\text{pp}}$.
- (RelSetup, lGen) is a hard instance generator as defined in Sect. 4.

5.2 Group-action-based PKE

We also consider group actions provided with a corresponding public-key encryption scheme, as specified in the following definition.

Definition 5.2 (Group-action-based PKE) A *group-action-based public-key encryption scheme*, GA-PKE in short, is a public-key encryption scheme $\Pi_{\text{GA-PKE}} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ with the following properties:

$\text{Setup}(1^\lambda) \rightarrow \text{pp}$: On input a security parameter 1^λ , it returns the public parameter $\text{pp} = (G, G_M, \mathcal{X}, S_1, S_2, \delta, D_{\mathcal{X}}, \star_M, \mathcal{M})$ (sometimes implicitly) used by the scheme. Here, G, G_M are additive groups, S_1, S_2 two symmetric subsets of G , \mathcal{X} a finite set, δ a real number in $[0, 1]$, $D_{\mathcal{X}}$ a distribution over a set of group actions $\star_{\text{pk}} : G \times \mathcal{X} \rightarrow \mathcal{X}$ and elements in \mathcal{X} , $\star_M : G_M \times \mathcal{X} \rightarrow \mathcal{X}$ a group action, $\mathcal{M} \subseteq G_M$ a message space. For any polynomially large $N \in \mathbb{N}$, we assume that there exists a feasible and invertible embedding τ from the set of index $[N]$ into the message space \mathcal{M} . For simplicity, we will write $\tau(i)\star_{\mathcal{M}}X, \text{Enc}(\text{pk}, \tau(i))$ as $i\star_{\mathcal{M}}X, \text{Enc}(\text{pk}, i)$ respectively without causing confusion.

$\text{KeyGen}(\text{pp}) \rightarrow (\text{pk}, \text{sk})$: On input a public parameter pp , it returns a public key pk and a secret key sk . We assume $\text{pk} = (\star_{\text{pk}}, X_{\text{pk}})$ to be drawn from $D_{\mathcal{X}}$, where $\star_{\text{pk}} : G \times \mathcal{X} \rightarrow \mathcal{X}$ is a group action and $X_{\text{pk}} \in \mathcal{X}$, and $\text{sk} \in G$. We also assume pk includes pp w.l.o.g.

$\text{Enc}(\text{pk}, M; r) \rightarrow \text{ct}$: On input a public key $\text{pk} = (\star_{\text{pk}}, X_{\text{pk}})$ and a message $M \in \mathcal{M}$, it returns a ciphertext ct . We assume ct is generated as $M\star_M(r\star_{\text{pk}}X_{\text{pk}}) \in \mathcal{X}$, where the encryption randomness is sampled as $r \xleftarrow{\$} S_1$.

$\text{Dec}(\text{sk}, \text{ct}) \rightarrow M$: On input a secret key sk and a ciphertext ct , it (deterministically) returns a message $M \in \mathcal{M}$.

In addition, we assume the following properties hold for the group actions defined by pp .

1. There exists a positive-valued polynomial T such that for all $\lambda \in \mathbb{N}$, $\text{pp} \in \text{Setup}(1^\lambda)$, and $(\text{pk}, \text{sk}) \in \text{KeyGen}(\text{pp})$, one can efficiently compute $g\star_{\text{pk}}X$ for all $g \in S_1 \cup S_2$ and all $X \in \mathcal{X}$ in time at most $T(\lambda)$, sample uniformly from S_1 and S_2 , and represent elements of G and \mathcal{X} uniquely. It is also efficient to compute the action \star_M for every possible input.
2. The intersection S_3 of the sets $S_2 + g$, with g varying in S_1 , is such that its cardinality is equal to $\delta|S_2|$. Furthermore, it is efficient to check whether an element $g \in G$ belongs to S_3 .

We further require a group-action-based PKE to satisfy standard correctness and decryption efficiency.

Definition 5.3 (Correctness and Decryption Efficiency) We say a group-action-based PKE $\Pi_{\text{GA-PKE}}$ is *correct* if for all $\lambda \in \mathbb{N}$, and for all but a negligible fraction of $\text{pp} \in \text{Setup}(1^\lambda)$, we have $\text{Dec}(\text{sk}, \text{Enc}(\text{pk}, M)) = M$ for all $(\text{pk}, \text{sk}) \in \text{KeyGen}(\text{pp})$ and $M \in \mathcal{M}$.

Moreover, we require Dec to run in $\text{poly}(\lambda)$ for a fixed polynomial function poly and for all possible inputs.

As we show in Sect. 4.1, in order to construct an accountable ring signature, a group-action-based PKE is also required to be (multi-challenge) IND-CPA secure and $(\mathcal{R}', \mathcal{KR}')$ -correct for some relaxed randomness set \mathcal{R}' and some relaxed key relation \mathcal{KR}' (Def. 4.1).

The concrete choice of $(\mathcal{R}', \mathcal{KR}')$ may depend on the instantiation. For instance, while we define $(\mathcal{R}', \mathcal{KR}') = (\mathcal{R}, \mathcal{KR})$ for our isogeny-based instantiation in Sect. 8.1, we must rely on a strictly wider relation for our lattice-based instantiation to compensate for the *relaxed* soundness. In slightly more detail, in our lattice-based NIZK, we are only able to argue that an adversary created a ciphertext ct using message M and randomness $r \in \mathcal{R}'$, and/or that a ct can be decrypted to M using secret key sk such that $(\text{pk}, \text{sk}) \in \mathcal{KR}'$. Roughly, $(\mathcal{R}', \mathcal{KR}')$ -correctness guarantees that such an argument suffices to prove that ct can only be decrypted to a unique M .

6 Sigma protocol for a “Traceable” OR relation

In this section, we present an efficient sigma protocol for the relation R_{sig} introduced in Sect. 4.1, using group-action-based HIG and a group-action-based PKE from the previous section. Recall this relation was used to define the multi-proof online extractable NIZK with labels Π_{NIZK} , which allowed an OR proof along with a proof of opening to a ciphertext. Looking ahead, in Sect. 7, we show that our sigma protocol can be turned into a multi-proof online extractable NIZK using the Fiat-Shamir transform. This is in contrast to the common application of the Fiat-Shamir transform that only provides a proof of knowledge via the rewinding argument [50, 81]. We note that we do not focus on the other NIZK for the relation R_{open} in Sect. 4.1 since they can be obtained easily from prior works.

We call the sigma protocol we present in this section as a *traceable* OR sigma protocol since it allows to trace the prover. This section is structured as follows. Firstly, we introduce a *base* traceable OR sigma protocol $\Pi_{\Sigma}^{\text{base}}$ for the relation R_{sig} with proof size $O(\log N)$ but with a binary challenge space. Secondly, we amplify the soundness of the sigma protocol by performing parallel repetitions. Here, instead of applying λ -parallel repetitions naively, we optimize it using three approaches developed in [14] to obtain our *main* traceable OR sigma protocol Π_{Σ}^{OR} . Finally, we show a sigma protocol for the “tight” relation $R_{\text{sig}}^{\text{Tight}}$ introduced in Sect. 4.3.

6.1 From a group-action-based HIG and PKE to base traceable OR sigma protocol

In this section, we present a *base* OR sigma protocol for the relation R_{sig} with a binary challenge space from which the main OR sigma protocol will be deduced.

Parameters and Binary Relation The sigma protocol is based on a group-action-based HIG and PKE. Let $\text{pp}_1 = (G, \mathcal{X}, S_1, S_2, \delta_x, \star, X_0)$ and $\text{pp}_2 = (\overline{G}, \overline{G}_T, \mathcal{Y}, \overline{S}_1, \overline{S}_2, \delta_y, D_{\mathcal{Y}}, \star_M, \mathcal{M})$ be public parameters in the image of RelSetup and PKE.Setup , respectively. Moreover, let $(\text{pk}, \text{sk}) \in \text{KeyGen}(\text{pp}_2)$. The relation R_{sig} in Sect. 4.1 can be equivalently rewritten as follows:

$$R_{\text{sig}} = \left\{ \left((\{X_i\}_{i \in [N]}, \text{pk}, \text{ct}), (I, s, r) \right) \mid \begin{array}{l} (I, s, r) \in [N] \times S_1 \times \overline{S}_1 \wedge \\ X_I = s \star X_0 \wedge \text{ct} = \text{Enc}(\text{pk}, I; r) \end{array} \right\}.$$

Recall that by definition of GA-PKE (Def. 5.2), the ciphertext ct is restricted to the simple form $I \star_M(r \star_{\text{pk}} Y_{\text{pk}}) \in \mathcal{Y}$, where $r \in \overline{S}_1 \subseteq \overline{G}$.

Sigma Protocol for R_{sig} . We now sketch the base traceable OR sigma protocol $\Pi_{\Sigma}^{\text{base}}$.

A prover with witness $(I, s, r) \in [N] \times S_1 \times \overline{S}_1$ first samples $(s', r') \xleftarrow{\$} S_2 \times \overline{S}_2$, and $(\{\text{bits}_i\}_{i \in [N]}) \leftarrow \{0, 1\}^{\lambda N}$. Then, it computes commitments

$$C_i = \mathcal{O}(\text{Com} \parallel s' \star X_i \parallel r' \star_{\text{pk}}(-i \star_M \text{ct}) \parallel \text{bits}_i) \quad \forall i \in [N],$$

and builds a Merkle tree with (C_1, \dots, C_N) as its leaves, obtaining root . Here, notice $r' \star_{\text{pk}}(-i \star_M \text{ct}) = r' \star_{\text{pk}}(-i + I) \star_M(r \star_{\text{pk}} Y_{\text{pk}})$ is simply $(r' + r) \star_{\text{pk}} Y_{\text{pk}}$ when $i = I$. Then, the prover sends $\text{com} = \text{root}$ to the verifier as the commitment of the sigma protocol. The verifier, in turn, responds with a uniform challenge $\text{chall} \in \{0, 1\}$.

If the challenge bit chall is 0, then the prover sends (s', r') and the commitment randomness $\{\text{bits}_i\}_{i \in [N]}$. That is, all the randomness it generated in the first round. The verifier then can reconstruct the Merkle tree and verify that the root of the obtained tree is equal to root .

If the challenge bit chall is equal to 1, then the prover computes $s'' = s' + s$, $r'' = r' + r$. The prover aborts the protocol if $s'' \notin S_3$ or $r'' \notin \overline{S}_3$. The first event will occur with probability

$(1 - \delta_x)$ and, similarly, the second event will occur with probability $(1 - \delta_y)$. Otherwise, the prover sends (r'', s'') together with the path connecting C_I to root in the Merkle tree, and the corresponding commitment randomness bits $_I$ to the verifier. The verifier computes $\tilde{C}_I = \mathcal{O}(\text{Com} \parallel s'' \star X_0 \parallel r'' \star_{\text{p}} k Y_{\text{p}} k \parallel \text{bits}_I)$ and uses the received path to reconstruct $\widetilde{\text{root}}$ of the Merkle tree. The verifier checks whether $\widetilde{\text{root}} = \text{root}$.

To reduce the communication cost, a pseudorandom number generator (PRG) Expand can be run over a uniform seed $\text{seed} \in \{0, 1\}^\lambda$ to produce the group elements s', r' and all commitment randomness values $\text{bits}_1, \dots, \text{bits}_N$ (part of the response for $\text{chall} = 0$). As a consequence, if the challenge bit is 0, the prover responds with seed so that the verifier can generate $(s', r', \text{bits}_1, \dots, \text{bits}_N)$ with the PRG Expand. The response corresponding to the challenge bit $\text{chall} = 1$ remains unchanged. We instantiate the PRG by a random oracle $\mathcal{O}(\text{Expand} \parallel \cdot)$. Looking ahead, using a PRG not only provides efficiency, but it proves to be essential when proving multi-proof online extractability when compiled into a NIZK. Roughly, the seed binds the cheating prover from using arbitrary $(s', r', \text{bits}_1, \dots, \text{bits}_N)$ and the random oracle allows for efficient extraction. Finally, we instantiate the collision-resistant hash function $\mathcal{H}_{\text{Coll}}(\cdot)$ used in our Merkle tree by a random oracle $\mathcal{O}(\text{Coll} \parallel \cdot)$.

A formal description of $\Pi_{\Sigma}^{\text{base}}$ is provided in Fig. 3

Security of Sigma Protocol $\Pi_{\Sigma}^{\text{base}}$. The following Theorems 6.1 and 6.2 summarize the security of our sigma protocol. We point out that in Theorem 6.1, we show our sigma protocol satisfies special soundness for the relations R_{sig} and \tilde{R}'_{sig} such that $R_{\text{sig}} \subset \tilde{R}'_{\text{sig}}$, rather than for the relations R_{sig} and \tilde{R}_{sig} such that $R_{\text{sig}} \subseteq \tilde{R}_{\text{sig}}$, where \tilde{R}_{sig} is the relaxed relation introduced in Sect. 4.1. The subtle difference is that \tilde{R}'_{sig} captures the scenario where the extractor may extract a witness that forms a collision in the random oracle. This has no concrete impact as we are able to turn such a sigma protocol into a multi-proof online extractable NIZK for the relations R_{sig} and \tilde{R}_{sig} .

Theorem 6.1 *The sigma protocol $\Pi_{\Sigma}^{\text{base}}$ has correctness with abort rate $(1 - \delta_x \delta_y)/2$ and relaxed special soundness for the relations R_{sig} and \tilde{R}'_{sig} , where*

$$\tilde{R}'_{\text{sig}} = \left\{ \left((\{X_i\}_{i \in [N]}, \text{pk}, \text{ct}), W \right) \left| \begin{array}{l} W = (I, s, r) \in [N] \times (S_2 + S_3) \times (\overline{S_2} + \overline{S_3}) \\ \wedge X_I = s \star X_0 \wedge \text{ct} = \text{Enc}(\text{pk}, I; r) \text{ or} \\ W = (x_1, x_2) \in \{0, 1\}^* \wedge \\ x_1 \neq x_2 \wedge \mathcal{O}(x_1) = \mathcal{O}(x_2) \end{array} \right. \right\}.$$

Here, \tilde{R}'_{sig} is identical to the one defined in Sect. 4.1 if we ignore the hash collision $W = (x_1, x_2)$ and set $\mathcal{R}' = \overline{S_2} + \overline{S_3}$ in the $(\mathcal{R}', \mathcal{KR}')$ -correctness of GA-PKE.

Proof Correctness. Say the prover honestly runs $\Pi_{\Sigma}^{\text{base}}$ on an input (I, s, r) satisfying $X_I = s \star X_0$ and $\text{ct} = \text{Enc}(\text{pk}, I; r)$, and does not abort. If $\text{chall} = 0$, then the verifier repeats the computation in the commitment phase (see Round 1 in Fig. 3) and therefore obtains the same output. If $\text{chall} = 1$, then the verifier computes $\tilde{T} = s'' \star X_0$ and $\tilde{\text{ct}} = r'' \star_{\text{p}} k Y_{\text{p}} k$ where $s'' = s' + s$ and $r'' = r' + r$. Besides, since \tilde{T} is equal to $T_I = s' \star X_I$, $\tilde{\text{ct}}$ is equal to $\text{ct}_I = r' \star_{\text{p}} k (-I \star_{\text{M}} \text{ct})$ and $\tilde{C} = \mathcal{O}(\text{Com} \parallel \tilde{T} \parallel \tilde{\text{ct}} \parallel \text{bits})$ is equal to the leaf $\tilde{C} = C_I \in \{C_1, \dots, C_N\}$, the verifier reconstructs the root $\widetilde{\text{root}}$ which is equal to root. Hence, the protocol has (non-abort) correctness.

Abort rate. The prover will not abort in the case $\text{chall} = 0$. When $\text{chall} = 1$ (which occurs with probability 1/2) the prover aborts if $s'' = s' + s \notin S_3$ or $r'' = r' + r \notin \overline{S_3}$. We note that s is in S_1 and s' is drawn uniformly at random from S_2 (in the random oracle model). We can therefore say s'' is drawn uniformly at random from $S_2 + s$, which contains S_3 as a subset.

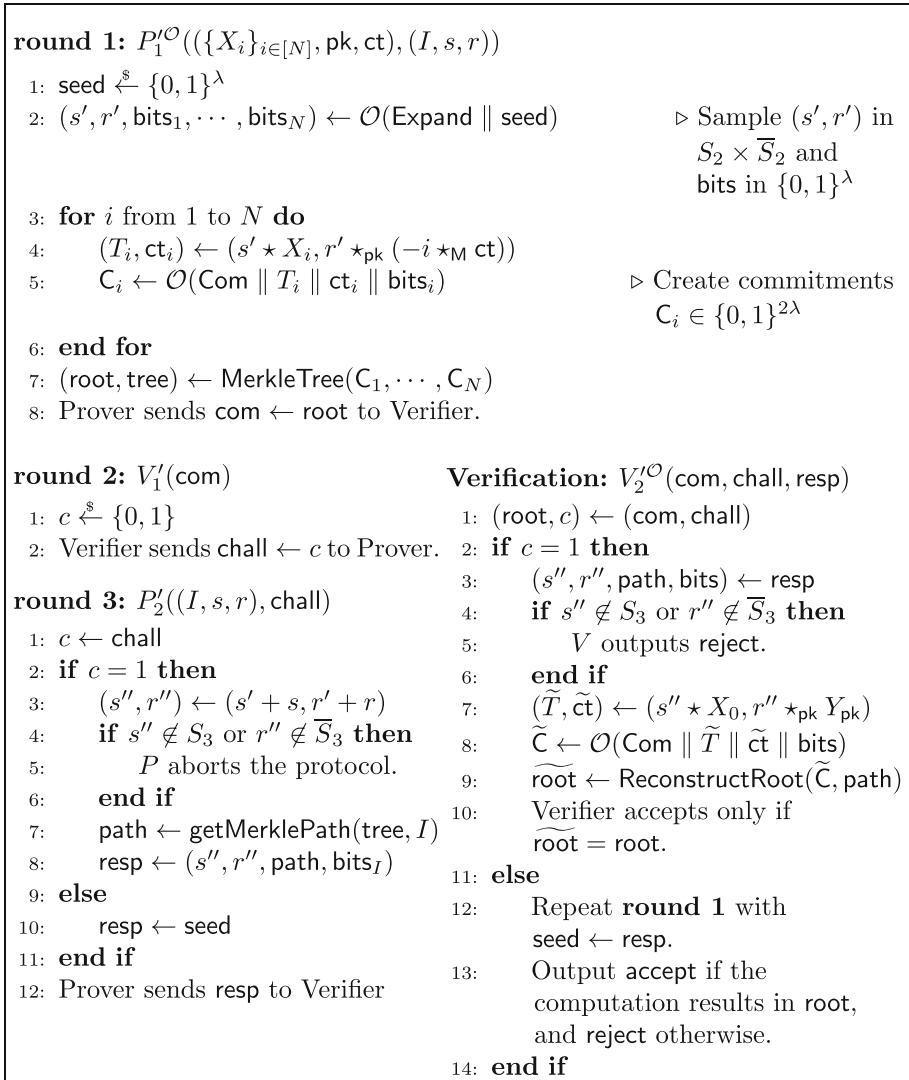


Fig. 3 Construction of the base traceable OR sigma protocol $\Pi_{\Sigma}^{\text{base}} = (P' = (P'_1, P'_2), V' = (V'_1, V'_2))$ for the relation R_{sig} . Informally, $\mathcal{O}(\text{Expand} \parallel \cdot)$ and $\mathcal{O}(\text{Com} \parallel \cdot)$ are a PRG and a commitment scheme instantiated by the random oracle, respectively

So the probability that $s'' = s' + s \in S_3$ is $|S_3|/|S_2| = \delta_x$. The same reasoning applies to r'' , so the probability of both s'', r'' lying in S_3, \bar{S}_3 respectively is $\delta_x \delta_y$ and the total abort rate is $(1 - \delta_x \delta_y)/2$

Relaxed special soundness. Given two valid transcripts for the same statement and on the same commitment, $(\text{com}, 0, \text{seed})$ and $(\text{com}, 1, (s'', r'', \text{path}, \text{bits}))$ where $\text{com} = \text{root}$, an extraction algorithm Extract for a witness for the relation \tilde{R}'_{pk} proceeds as follows. Extract firstly generates $(s', r', \text{bits}_1, \dots, \text{bits}_N) \leftarrow \mathcal{O}(\text{Expand} \parallel \text{seed})$ and constructs C_1, \dots, C_N such that the Merkle Tree with leaves (C_1, \dots, C_N) has root equal to root . Extract outputs $W = (\text{Coll} \parallel x_1, \text{Coll} \parallel x_2)$ as the witness if it there exists $x_1 \neq x_2$ such that $\mathcal{O}(\text{Coll} \parallel x_1) =$

$\mathcal{O}(\text{Coll} \parallel x_2)$. Otherwise, by Lem. A.1, we can assume $\tilde{C} = \mathcal{O}(\text{Com} \parallel s'' \star X_0 \parallel r'' \star_{\text{pk}} Y_{\text{pk}} \parallel \text{bits})$ is equal to $C_{\tilde{\gamma}} \in \{C_1, \dots, C_N\}$ for some $\tilde{\gamma} \in [N]$. Then, Extract outputs $W = (\text{Com} \parallel x_1, \text{Com} \parallel x_2)$ as the witness if it there exists $x_1 \neq x_2$ such that $\mathcal{O}(\text{Com} \parallel x_1) = \mathcal{O}(\text{Com} \parallel x_2)$. Otherwise, from $\tilde{C} = C_{\tilde{\gamma}}$, we can assume $s' \star X_{\tilde{\gamma}} = s'' \star X_0, r' \star_{\text{pk}} k(\tilde{\gamma} \star_{\text{Mct}}) = r'' \star_{\text{pk}} Y_{\text{pk}}$, and $\text{bits} = \text{bits}_{\tilde{\gamma}}$. Let $\tilde{s} = -s' + s'' \in S_2 + S_3$ and $\tilde{r} = -r' + r'' \in \bar{S}_2 + \bar{S}_3$. Finally, Extract outputs $W = (\tilde{I}, \tilde{s}, \tilde{r})$. Here, the equalities $\tilde{s} \star X_0 = X_{\tilde{\gamma}}$ and $\tilde{I} \star_{\text{M}}(\tilde{r} \star_{\text{pk}} Y_{\text{pk}}) = \text{ct}$ follow directly from the relations $s' \star X_{\tilde{\gamma}} = s'' \star X_0$ and $r' \star_{\text{pk}} k(\tilde{\gamma} \star_{\text{Mct}}) = r'' \star_{\text{pk}} Y_{\text{pk}}$, respectively. Therefore, $W = (\tilde{I}, \tilde{s}, \tilde{r})$ is a witness for the “relaxed” relation \tilde{R}'_{pk} . Hence, the protocol $\Pi_{\Sigma}^{\text{base}}$ has relaxed special soundness. \square

Theorem 6.2 *The sigma protocol $\Pi_{\Sigma}^{\text{base}}$ has non-abort special zero-knowledge. Precisely, there exists a PPT simulator $\text{Sim}^{\mathcal{O}}$ with access to a random oracle \mathcal{O} such that, for any statement-witness pair $(X, W) \in R_{\text{sig}}$, $\text{chall} \in \{0, 1\}$, and any computationally-unbounded adversary \mathcal{A} that makes at most Q queries to the random oracle \mathcal{O} , we have*

$$\left| \Pr[\mathcal{A}^{\mathcal{O}}(1^\lambda, \tilde{P}^{\mathcal{O}}(X, W, \text{chall})) = 1] - \Pr[\mathcal{A}^{\mathcal{O}}(1^\lambda, \text{Sim}^{\mathcal{O}}(X, \text{chall})) = 1] \right| \leq \frac{Q}{2^\lambda},$$

where \tilde{P} is a non-aborting prover $P' = (P'_1, P'_2)$ run on (X, W) with a challenge fixed to chall .

Proof Assume the adversary makes Q_{Expand} and Q_{Com} queries to the random oracles of the form $\mathcal{O}(\text{Expand} \parallel \cdot)$ and $\mathcal{O}(\text{Com} \parallel \cdot)$, respectively. We have $Q_{\text{Expand}} + Q_{\text{Com}} \leq Q$. The PPT simulator $\text{Sim}^{\mathcal{O}}$, on input (X, chall) , proceeds as follows.

- If $\text{chall} = 0$, the simulator executes as $P'^{\mathcal{O}}(X, \perp, \text{chall})$, where notice P' does not require the witness when $\text{chall} = 0$. Concretely, the simulator outputs $(\text{com} = \text{root}, \text{chall} = 0, \text{resp} = \text{seed})$ where root, seed are honestly generated as in the execution of $P'_1{}^{\mathcal{O}}$.
- If $\text{chall} = 1$, the simulator uniformly samples (s'', r'') from $S_3 \times \bar{S}_3$, and bits from $\{0, 1\}^\lambda$. It computes $C_1 = \mathcal{O}(\text{Com} \parallel s'' \star X_0 \parallel r'' \star_{\text{pk}} Y_{\text{pk}} \parallel \text{bits})$. It then uniformly samples dummy commitments C_i for $i \in \{2, \dots, N\}$ from $\{0, 1\}^{2\lambda}$, and computes the (index-hiding) Merkle tree $(\text{root}, \text{tree}) \leftarrow \text{MerkleTree}(C_1, \dots, C_N)$. After that, it extracts the path $\text{path} \leftarrow \text{getMerklePath}(\text{tree}, 1)$ in the tree and sets $\text{com} = \text{root}$, and $\text{resp} = (s'', r'', \text{path}, \text{bits})$. Finally, the simulator returns $(\text{com}, \text{chall} = 1, \text{resp})$.

In the first case, the whole transcript is generated exactly as in the protocol. Hence transcripts generated by $\tilde{P}^{\mathcal{O}}$ and $\text{Sim}^{\mathcal{O}}$ are indistinguishable to the adversary \mathcal{A} . Therefore, we have

$$\left| \Pr[\mathcal{A}^{\mathcal{O}}(1^\lambda, \tilde{P}^{\mathcal{O}}(X, W, \text{chall} = 0)) = 1] \right| = \left| \Pr[\mathcal{A}^{\mathcal{O}}(1^\lambda, \text{Sim}^{\mathcal{O}}(X, \text{chall} = 0)) = 1] \right|.$$

To conclude the proof, it suffices to show that the difference between the probabilities that the adversary \mathcal{A} outputs 1 for the other case, $\text{chall} = 1$, is also bounded by $\frac{Q}{2^\lambda}$.

We use a hybrid argument by introducing a series of simulators $\text{Sim}_0 = \tilde{P}, \dots, \text{Sim}_4 = \text{Sim}$, gradually changing from the honest prover \tilde{P} to Sim , to show that they are indistinguishable with overwhelming probability. We fix an adversary $\mathcal{A}, (X, W) \in R_{\text{sig}}$, and for each $i \in \{0, 1, \dots, 4\}$, we denote by E_i the event that $\mathcal{A}^{\mathcal{O}}(1^\lambda, \text{Sim}_i^{\mathcal{O}}(X, \text{chall} = 1)) = 1$.

- Sim_1 is identical to Sim_0 except that instead of using Expand to generate $s', r', \{\text{bits}_i\}_{i \in [N]}$, the simulator generates these by sampling uniformly at random from the corresponding domains. This does not change the view of \mathcal{A} , unless the adversary queries \mathcal{O} on input

(Expand || seed). Since seed has λ bits of min-entropy and because it is information-theoretically hidden from \mathcal{A} , the probability that \mathcal{A} queries \mathcal{O} on this input is bounded by $Q_{\text{Expand}}/2^\lambda$. That is, $|\Pr[E_1] - \Pr[E_0]| \leq \frac{Q_{\text{Expand}}}{2^\lambda}$.

- Sim₂ is identical to Sim₂ except that all the commitments C_i for $i \in [N] \setminus \{I\}$ are generated uniformly at random. This does not change the view of \mathcal{A} , unless the adversary queries \mathcal{O} on input (Com || T_i || ct_i || bits _{i}) for any $i \in [N] \setminus \{I\}$, where $T_i = s' \star X_i$ and $ct_i = r' \star_{\text{pk}}(-i \star_{\text{Mct}}$). Since for any $i \in [N] \setminus \{I\}$ the string bits _{i} has λ bits of min-entropy and because it is information-theoretically hidden from \mathcal{A} , the probability that \mathcal{A} queries \mathcal{O} on input (Com || T_i || ct_i || bits _{i}) is bounded by $Q_{\text{Com}}/2^\lambda$. That is, $|\Pr[E_2] - \Pr[E_1]| \leq \frac{Q_{\text{Com}}}{2^\lambda}$.
- Sim₃ is identical to Sim₃ except that instead of computing s'', r'' as $s' + s, r' + r$ (conditioned on them respectively lying in S_3, \bar{S}_3 , due to non-aborting transcripts), the simulator generates these two values by sampling uniformly at random from S_3, \bar{S}_3 , respectively. Both the distributions are uniform over S_3 and \bar{S}_3 . Therefore, we have $|\Pr[E_3] - \Pr[E_2]| = 0$.
- Sim₄ = Sim is identical to Sim₄ except that the simulator uses $I = 1$ instead of the value I in the witness W . These two simulators are indistinguishable because the Merkle tree is index-hiding (by Lemma A.2). Precisely, we have $|\Pr[E_4] - \Pr[E_3]| = 0$.

Collecting the bounds, we obtain the bound in the statement. □

6.2 From base to main traceable OR sigma protocol

In this section, compile Π_Σ^{base} to make the soundness error negligibly small. This is straightforward if we run the OR sigma protocol in parallel λ -times. However, we show how to do much better by incorporating the three optimizations developed in [14] explained in the technical overview. Our main traceable OR sigma protocol, denote by Π_Σ^{TOR} , is detailed in Fig. 4.

Unbalanced Challenge Space Given the construction Π_Σ^{base} , one can observe that the response produced by the prover by running P'_2 when the challenge is 1 is larger than the response produced when the challenge is 0, which is a single seed of λ bits. Concretely, the response for the challenge $\text{chall} = 1$ consists of a Merkle tree path, two elements in S_3, \bar{S}_3 respectively, and a λ bit string. We leverage this fact by preparing an unbalanced challenge space $C_{M,K}$, where each element in $C_{M,K}$ is a string containing K 1's and $M - K$ 0's. We chose $K \ll M$ to chose more 0's, while satisfying $\binom{M}{K} \geq 2^\lambda$ for negligible soundness error.

Seed Trees The seed tree described in Sect. A.2 allows the prover to generate all seeds seed by using a single seed $\text{seed}_{\text{root}}$, and reveal parts of the tree according to the challenge. A smaller signature size follows directly from this approach.

Adding Salt We prefix a salt and the session identifier, i.e. (salt || i), to the random oracle when used within the i -th parallel execution of Π_Σ^{base} . In particular, throughout such execution, $\mathcal{O}_i(\cdot) = \mathcal{O}(\text{salt} || i || \cdot)$ is used. The salt is used as a prefix also within the construction of Merkle trees and seed trees. Adding salt benefits the protocol in having a tighter reduction and resisting multi-target attacks, such as those in [41]. The approach appears to make no difference in a sigma protocol but it is quite beneficial for a group (ring) signature scheme after we apply the Fiat-Shamir transform. Roughly, in the anonymity game (Def. 3.17) each oracle \mathcal{O} query made by the adversary will only give useful information to at most one challenge signature due to distinct prefix salts. In contrast, without salts an oracle query of \mathcal{O} can give useful information to *each* challenge signature.

round 1: $P_1^{\mathcal{O}}((\{X_i\}_{i \in [N]}, \text{pk}, \text{ct}), (I, s, r))$

- 1: $\text{seed}_{\text{root}} \leftarrow \{0, 1\}^\lambda$
- 2: $\text{salt} \xleftarrow{\$} \{0, 1\}^{2\lambda}$
- 3: $\mathcal{O}'(\cdot) := \mathcal{O}(\text{salt} \parallel 0 \parallel \cdot)$
- 4: $(\text{seed}_1, \dots, \text{seed}_M) \leftarrow \text{SeedTree}^{\mathcal{O}'}(\text{seed}_{\text{root}}, M)$
- 5: **for** j from 1 to M **do**
- 6: $\mathcal{O}_j(\cdot) := \mathcal{O}(\text{salt} \parallel j \parallel \cdot)$
- 7: $\text{com}_j \leftarrow P_1^{\mathcal{O}_j}((\{X_i\}_{i \in [N]}, \text{pk}, \text{ct}), (I, s, r); \text{seed}_j)$
- 8: **end for**
- 9: Prover sends $\text{com} \leftarrow (\text{salt}, \text{com}_1, \dots, \text{com}_M)$ to Verifier.

round 2: $V_1(\text{com})$

- 1: $\mathbf{c} \xleftarrow{\$} C_{M,K}$
- 2: Verifier sends $\text{chall} \leftarrow \mathbf{c}$ to Prover.

round 3: $P_2^{\mathcal{O}}((I, s, r), \text{chall})$

- 1: $\mathbf{c} = (c_1, \dots, c_M) \leftarrow \text{chall}$
- 2: $\mathcal{O}'(\cdot) := \mathcal{O}(\text{salt} \parallel \cdot)$
- 3: **for** j s.t. $c_j = 1$ **do**
- 4: $\text{resp}_j \leftarrow P_2^{\mathcal{O}_j}((I, s, r), c_j; \text{seed}_j)$ ▷ Run P_2' on state seed_j
- 5: **end for**
- 6: $\text{seeds}_{\text{internal}} \leftarrow \text{ReleaseSeeds}^{\mathcal{O}'}(\text{seed}_{\text{root}}, 1^M \oplus \mathbf{c})$
- 7: Prover sends $\text{resp} \leftarrow (\text{seeds}_{\text{internal}}, \{\text{resp}_j\}_{j \text{ s.t. } c_j=1})$ to Verifier

Verification: $V_2^{\mathcal{O}}(\text{com}, \text{chall}, \text{resp})$

- 1: $(\text{salt}, \text{com}_1, \dots, \text{com}_M), \mathbf{c} = (c_1, \dots, c_M) \leftarrow (\text{com}, \text{chall})$
- 2: $(\text{seeds}_{\text{internal}}, \{\text{resp}_j\}_{j \text{ s.t. } c_j=1}) \leftarrow \text{resp}$
- 3: $\mathcal{O}'(\cdot) := \mathcal{O}(\text{salt} \parallel 0 \parallel \cdot)$
- 4: $\{\text{resp}_j\}_{j \text{ s.t. } c_j=0} \leftarrow \text{RecoverLeaves}^{\mathcal{O}'}(\text{seeds}_{\text{internal}}, 1^M \oplus \mathbf{c})$
- 5: **for** j from 1 to M **do**
- 6: $\mathcal{O}_j(\cdot) := \mathcal{O}(\text{salt} \parallel j \parallel \cdot)$
- 7: Verifier outputs reject if $V_2^{\mathcal{O}_j}(\text{com}_j, c_j, \text{resp}_j)$ outputs reject.
- 8: **end for**
- 9: Verifier outputs accept.

Fig. 4 Main traceable OR sigma protocol $\Pi_{\Sigma}^{\text{OR}} = (P = (P_1, P_2), V = (V_1, V_2))$ for the relation R_{sig} built on the the base traceable OR sigma protocol $\Pi_{\Sigma}^{\text{base}} = (P' = (P'_1, P'_2), V' = (V'_1, V'_2))$ in Fig. 3. The challenge space is defined as $C_{M,K} := \{\mathbf{c} \in \{0, 1\}^M \mid \|\mathbf{c}\|_1 = K\}$. Both the seed tree and $\Pi_{\Sigma}^{\text{base}}$ have access to salted random oracles derived from \mathcal{O}

Theorem 6.3 *The sigma protocol $\Pi_{\Sigma}^{\text{TOR}}$ has correctness with abort rate $(1 - \delta_x^K \delta_y^K)$, high min-entropy, and relaxed special soundness for the relations R_{sig} and \tilde{R}'_{sig} , where the relations are identical to those used in Theorem 6.1.*

Proof As a starting remark, we note that in the following lines we will use the notation of Fig. 4.

Correctness and abort rate. If the execution of $\Pi_{\Sigma}^{\text{TOR}}$ does not abort, then the verifier will accept with probability 1 due to the correctness of $\Pi_{\Sigma}^{\text{base}}$ and SeedTree. We recall that in the case of challenge equal to 1 the execution of $\Pi_{\Sigma}^{\text{base}}$ will abort with probability $(1 - \delta_x \delta_y)$. Since the challenge \mathbf{c} , sampled from $C_{M,K}$, is of Hamming weight K , the abort rate of $\Pi_{\Sigma}^{\text{TOR}}$ is $(1 - \delta_x^K \delta_y^K)$.

High min-entropy. Since a random salt of length 2λ is included in the commitment com , it has at least 2λ bits of min-entropy.

Relaxed special Soundness. The proof is similar to the one for the relaxed special soundness of $\Pi_{\Sigma}^{\text{base}}$. Let $(\text{com}, \text{chall} = \mathbf{c}, \text{resp})$ $(\text{com}, \text{chall}' = \mathbf{c}', \text{resp}')$ be two accepting transcripts for the same statement. Without loss of generality, say $c_j = 0, c'_j = 1$, i.e. the j^{th} components of \mathbf{c} and \mathbf{c}' are different. By computing $\{\text{resp}_i\}_{i \text{ s.t. } c_i=0} \leftarrow \text{RecoverLeaves}^{\mathcal{O}'}(\text{seeds}_{\text{internal}}, 1^M \oplus \mathbf{c})$, the extraction algorithm gets resp_j . In this way, two valid transcripts $(\text{com}_j, 0, \text{resp}_j)$ and $(\text{com}_j, 1, \text{resp}'_j)$ for $\Pi_{\Sigma}^{\text{base}}$ have been obtained, and the extractor of $\Pi_{\Sigma}^{\text{base}}$ in Theorem 6.1 can be invoked to extract the witness for the relation \tilde{R}_{sig} . To be concrete, in case a witness $W = (x_1, x_2)$ is extracted by the extractor of Theorem 6.1 such that it forms a collision in the random oracle $\mathcal{O}_j = \mathcal{O}(\text{salt} \parallel j \parallel \cdot)$, then the extractor appends x_1 and x_2 by either $\text{salt} \parallel j \parallel \text{Coll}$ or $\text{salt} \parallel j \parallel \text{Com}$ to produce a collision in \mathcal{O} . \square

Theorem 6.4 *The sigma protocol $\Pi_{\Sigma}^{\text{TOR}}$ has non-abort special zero-knowledge. Precisely, there exists a PPT simulator $\text{Sim}^{\mathcal{O}}$ with access to a random oracle \mathcal{O} such that, for any statement-witness pair $(X, W) \in R_{\text{sig}}$, $\text{chall} \in C_{M,K}$ and any computationally-unbounded adversary \mathcal{A} that makes at most Q queries of the form $(\text{salt} \parallel \cdot)$ to the random oracle \mathcal{O} , where salt is the salt value included in the transcript returned by \tilde{P} or Sim , we have*

$$\left| \Pr[\mathcal{A}^{\mathcal{O}}(1^\lambda, \tilde{P}^{\mathcal{O}}(X, W, \text{chall})) = 1] - \Pr[\mathcal{A}^{\mathcal{O}}(1^\lambda, \text{Sim}^{\mathcal{O}}(X, \text{chall})) = 1] \right| \leq \frac{Q}{2^\lambda},$$

where \tilde{P} is a non-aborting prover $P = (P_1, P_2)$ run on (X, W) with a challenge fixed to chall .

Proof The PPT simulator $\text{Sim}^{\mathcal{O}}(X, \text{chall})$ for the main sigma protocol $\Pi_{\Sigma}^{\text{TOR}}$ proceeds as in Fig. 5, where the simulator used for the base sigma protocol $\Pi_{\Sigma}^{\text{base}}$ in Theorem 6.2, denoted by Sim' is a subroutine. Say the adversary makes Q_i queries to the random oracle of the form $\mathcal{O}(\text{salt} \parallel i \parallel \cdot)$ for $i \in \{0\} \cup [M]$. We have $\sum_0^M Q_i \leq Q$.

We use a hybrid argument by introducing a sequence of simulators $\text{Sim}_0, \dots, \text{Sim}_2$ that gradually change from $\text{Sim}_0 = \tilde{P}$ to $\text{Sim}_2 = \text{Sim}$. We fix an adversary $\mathcal{A}, (X, W) \in R_{\text{sig}}$, and for each $i \in \{0, 1, 2\}$, we denote by E_i the event $\mathcal{A}^{\mathcal{O}}(1^\lambda, \text{Sim}_i^{\mathcal{O}}(X, \text{chall})) = 1$.

- Sim_1 is identical to Sim_0 , except that, rather than using a SeedTree with root $\text{seed}_{\text{root}}$ to generate $\text{seeds}_{\text{internal}}$ and $\{\text{seed}_i\}_{i \text{ s.t. } c_i=0}$, the simulator instead runs $\text{SimulateSeeds}(1^M \oplus \mathbf{c})$ to obtain $\text{seeds}_{\text{internal}}$, and then $\{\text{seed}_i\}_{i \text{ s.t. } c_i=0}$ via $\text{RecoverLeaves}(\text{seeds}_{\text{internal}}, 1^M \oplus \mathbf{c})$. The simulator picks the remaining seeds (for the challenge components c_i equal to 1) $\{\text{seed}_i\}_{i \text{ s.t. } c_i=1}$ uniformly at random from $\{0, 1\}^\lambda$. Lemma A.3 for the bit string $1^M \oplus \mathbf{c}$ implies that the distributions of $\text{seeds}_{\text{internal}}$ and

```

SimO(X, chall)
1: c = (c1, ..., cM) ← chall
2: salt ← {0, 1}2λ
3: O'(·) := O(salt || 0 || ·)
4: seedsinternal ← SimulateSeedsO'(1M ⊕ c)
5: {seedi}i s.t. ci=0 ← RecoverLeavesO'(seedsinternal, 1M ⊕ c)
6: for i ∈ [M] do
7:   O'i(·) := O(salt || i || ·)
8:   if ci = 1 then
9:     seedi ← {0, 1}λ
10:  end if
11:  (comi, ci, respi) ← SimO'i(X, ci; seedi)
12: end for
13: com ← (salt, com1, ..., comM)
14: resp ← (seedsinternal, {respi}i s.t. ci=1)
15: return (com, chall, resp)
    
```

Fig. 5 Zero-knowledge simulator Sim for the main sigma protocol $\Pi_{\Sigma}^{\text{TOR}}$

{seed_i}_{i s.t. c_i=1} generated in this way rather than as in the honest protocol can be distinguished with an advantage not greater than $\frac{Q_0}{2^\lambda}$. That is, $|\Pr[E_1] - \Pr[E_0]| \leq \frac{Q_0}{2^\lambda}$.

- Sim₂ is identical to Sim₁ except that the simulator uses the base simulator subroutine Sim' to compute, for each $i \in [M]$ such that $c_i = 1$, com_i and resp_i on randomness bits_i by seed_i $\xleftarrow{\$}$ {0, 1}^λ. By Theorem 6.2, the distinguishing advantage of the adversary is bounded by $\frac{Q_i}{2^\lambda}$ for each $i \in [M]$ such that $c_i = 1$. That is, $|\Pr[E_3] - \Pr[E_2]| \leq \frac{\sum_1^M Q_i}{2^\lambda}$.

Collecting the bounds, we obtain the bound in the statement. □

6.3 Base sigma protocol for the “Tight” relation $R_{\text{sig}}^{\text{Tight}}$

In this section, we show how to slightly tweak our base sigma protocol for the relation R_{sig} to obtain a sigma protocol for the “tight” relation $R_{\text{sig}}^{\text{Tight}}$ (see Sect. 4.3). This can then be used to construct the desired NIZK for $R_{\text{sig}}^{\text{Tight}}$ required for our tightly secure accountable ring signature construction (see the full version of this paper).

As explained in the technical overview, we can use the sigma protocol for R_{sig} along with the sequential OR-proof [51] to construct a sigma protocol for the “tight” relation $R_{\text{sig}}^{\text{Tight}}$. Unfortunately, this approach requires to double the proof size. Instead, we present a small tweak to our sigma protocol for R_{sig} to directly support statements in $R_{\text{sig}}^{\text{Tight}}$. Concretely, we use the same Merkle tree to commit to the $2N$ instances $\{X_i^{(j)}\}_{(i,j) \in [N] \times [2]}$ and for each $X_i^{(1)}$ and $X_i^{(2)}$, we encrypt the *same* index i . The main observation is that when the prover opens to the challenge bit 1 (which is the only case that depends on the witness), the path does not leak which $X_i^{(1)}$ and $X_i^{(2)}$ it opened to, and hence hides $b \in [2]$.

Notice the only increase in the size of the response is due to the path. Since the accumulated commitment only grows from N to $2N$, the overhead in the size of the path is merely 2λ bits. By using the unbalanced challenge space $C_{M,K}$ for the optimized parallel repetition, which consists of M -bit strings of Hamming weight K , the additional cost is only $2K\lambda$ where we typically set K to be a small constant (e.g., $K \leq 20$ for our concrete instantiation). This is much more efficient than the generic approach that doubles the proof size.

Formally, the sigma protocol for the “tight” relation $R_{\text{sig}}^{\text{Tight}}$, denoted as $\Pi_{\Sigma}^{\text{baseTi}}$, is provided in Fig. 6. We can turn it into a full-fledged sigma protocol with negligible soundness error by applying exactly the same argument in Sect. 6.1. We omit the proof of correctness and security for $\Pi_{\Sigma}^{\text{baseTi}}$ as they are almost identical to those of our sigma protocol $\Pi_{\Sigma}^{\text{base}}$ for R_{sig} .

7 Multi-proof online extractable NIZK from sigma protocol $\Pi_{\Sigma}^{\text{TOR}}$

In this section, we show that applying the Fiat-Shamir transform to our traceable OR sigma protocol $\Pi_{\Sigma}^{\text{TOR}}$ from the previous section results in a multi-proof online extractable NIZK with labels $\Pi_{\text{NIZK,lbl}}$. The construction of our $\Pi_{\text{NIZK,lbl}}$ for the relation R_{sig} is provide in Fig. 7.⁹ We assume the output of $\mathcal{O}(\text{FS} \parallel \cdot)$ is an M -bit string of Hamming weight K , i.e., the image is the challenge set $C_{M,K}$.

Correctness of $\Pi_{\text{NIZK,lbl}}$ for the relation R_{sig} follows directly from the correctness of the underlying traceable OR sigma protocol $\Pi_{\Sigma}^{\text{TOR}}$. We show in Theorems 7.1 and 7.4 that $\Pi_{\text{NIZK,lbl}}$ is multi-proof online extractable and zero-knowledge. We highlight that while we show special soundness for $\Pi_{\Sigma}^{\text{TOR}}$ with respect to the relaxed relation \tilde{R}'_{sig} (see Theorem 6.1), $\Pi_{\text{NIZK,lbl}}$ is multi-proof online extractable with respect to the relaxed relation \tilde{R}_{sig} originally considered in Sect. 4.1 for the generic construction of accountable ring signature. At a high level, we upper bound the probability that a cheating prover finds a collision in the random oracle, which was the only difference between \tilde{R}_{sig} and \tilde{R}'_{sig} . This subtle difference makes the resulting NIZK more handy to use as a building block, since we can ignore the edge case where the extractor accidentally extracts a collision in the random oracle. Below, we provide the proof of the multi-proof online extractability.

Theorem 7.1 *The NIZK with labels $\Pi_{\text{NIZK,lbl}}$ in Fig. 7 is multi-proof online extractable for the family of relations R_{sig} and \tilde{R}_{sig} considered in Sect. 4.1, where R_{sig} was formally redefined using notations related to group actions in Sect. 6.1 and \tilde{R}_{sig} is formally redefined as follows:*

$$\tilde{R}_{\text{sig}} = \left\{ \left((\{X_i\}_{i \in [N]}, \text{pk}, \text{ct}), W \right) \mid \begin{array}{l} W = (I, s, r) \in [N] \times (S_2 + S_3) \times (\bar{S}_2 + \bar{S}_3) \\ \wedge X_I = s \star X_0 \wedge \text{ct} = \text{Enc}(\text{pk}, I; r) \end{array} \right\}.$$

More precisely, for any (possibly computationally-unbounded) adversary \mathcal{A} making at most Q queries to the random oracle and T queries to the extract oracle, we have

$$\text{Adv}_{\Pi_{\text{NIZK,lbl}}}^{\text{OE}}(\mathcal{A}) \leq T \cdot (Q^2/2^{2\lambda-2} + (M \cdot Q)/2^{\lambda} + 1/|C_{M,K}|),$$

where $C_{M,K}$ is the challenge space (or equivalently the output space of $\mathcal{O}(\text{FS} \parallel \text{cdot})$).

Proof We begin the proof by providing the description of the online extractor `OnlineExtract`. Below, it is given as input $(\text{lbl}, X, \pi, L_{\mathcal{O}})$, where π is guaranteed to be valid by definition.

1. It parses $(\{X_i\}_{i \in [N]}, \text{pk}, \text{ct}) \leftarrow X, (\overline{\text{com}}, \overline{\text{chall}}, \overline{\text{resp}}) \leftarrow \pi, ((\text{salt}, \text{com}_1, \dots, \text{com}_M), \mathbf{c} = (c_1, \dots, c_M)) \leftarrow (\overline{\text{com}}, \overline{\text{chall}}), (\text{seeds}_{\text{internal}}, \{\text{resp}_j\}_{j \text{ s.t. } c_j=1}) \leftarrow \overline{\text{resp}}$, and $\text{root}_j \leftarrow \text{com}_j$ for $j \in [M]$.¹⁰

⁹ An astute reader may notice that the prover is only *expected* polynomial time. We can always assign an upper bound on the runtime of the prover, but did not do so for better readability. In practice, for concrete choices of the parameter, the number of repetition never exceeds, say 10.

¹⁰ Throughout the proof, we use overlines for $(\overline{\text{com}}, \overline{\text{chall}}, \overline{\text{resp}})$ to indicate that it is a transcript of $\Pi_{\Sigma}^{\text{TOR}}$. We use resp_j without overlines to indicate elements of $\overline{\text{resp}}$.

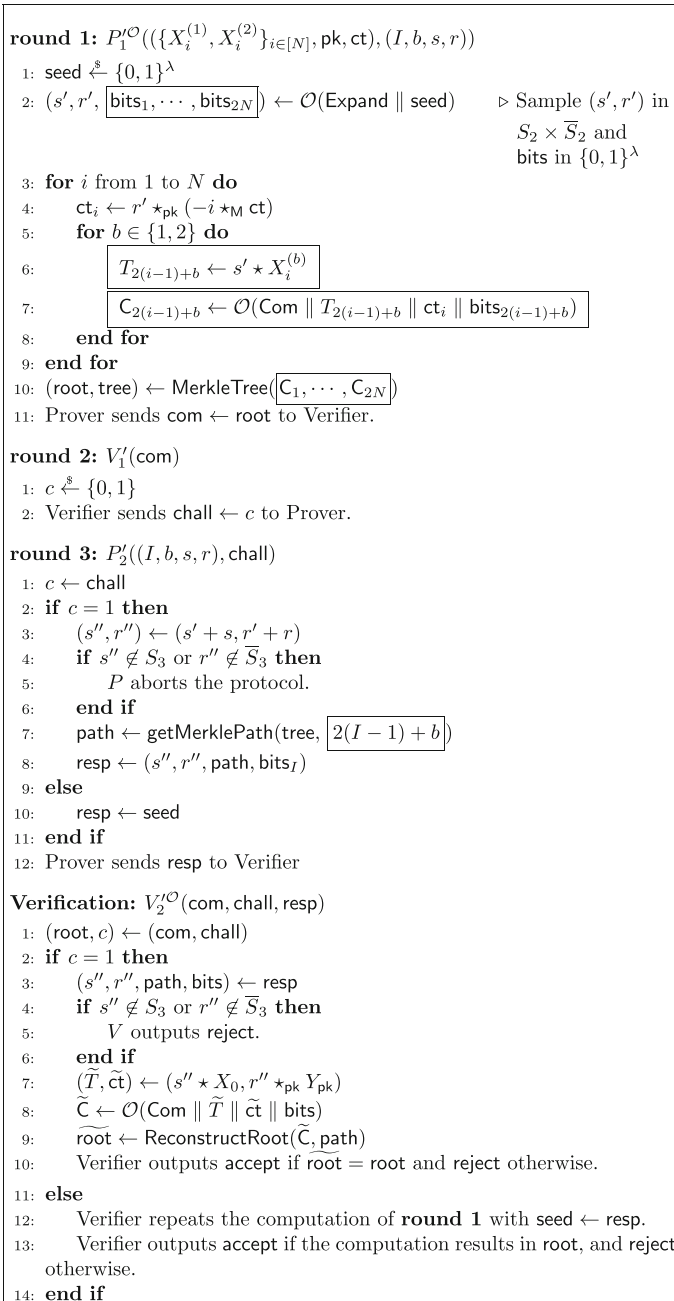


Fig. 6 Construction of the base traceable OR sigma protocol $\Pi_{\Sigma}^{\text{baseTi}} = (P' = (P_1', P_2'), V' = (V_1', V_2'))$ for the “tight” relation $R_{\text{sig}}^{\text{tight}}$. The box indicates the difference between the “non-tight” relation R_{sig} . Informally, $\mathcal{O}(\text{Expand} \parallel \cdot)$ and $\mathcal{O}(\text{Com} \parallel \cdot)$ are a PRG and a commitment scheme instantiated by the random oracle, respectively

```

ProveO(lbl, ({Xi}i∈[N], pk, ct), (I, sI, r))
1: resp := ⊥
2: while resp = ⊥ do
3:   com ← P1O(({Xi}i∈[N], pk, ct), (I, sI, r))
4:   chall ← O(FS || lbl || ({Xi}i∈[N], pk, ct) || com)
5:   resp ← P2O((I, sI, r), chall)
6: end while
7: return π ← (com, chall, resp)

VerifyO(lbl, ({Xi}i∈[N], pk, ct), π)
1: (com, chall, resp) ← π
2: if accept ← V2O(com, chall, resp) ∧ chall = O(FS || lbl ||
   ({Xi}i∈[N], pk, ct) || com) then
3:   return ⊤
4: else
5:   return ⊥
6: end if
    
```

Fig. 7 A multi-proof online extractable NIZK with labels $\Pi_{\text{NIZK, lbl}}$ for the relation R_{sig} obtained by applying the Fiat-Shamir transform to the traceable OR sigma protocol $\Pi_{\Sigma}^{\text{TOR}} = (P = (P_1, P_2), V = (V_1, V_2))$ in Fig. 4

2. For $j \in [M]$ such that $c_j = 1$, it proceeds as follows:

- (a) It parses $(s''_j, r''_j, \text{path}_j) \leftarrow \text{resp}_j$.
- (b) For every $((\text{salt} \parallel j \parallel \text{Expand} \parallel \text{seed}), (s', r', \text{bits}_1, \dots, \text{bits}_N)) \in L_{\mathcal{O}}$, where $\text{salt} \parallel j \parallel \text{Expand}$ is fixed, it proceeds as follows:
 - (i) It sets $(s, r) = (s''_j - s', r''_j - r')$ and checks if $(s, r) \in (S_2 + S_3) \times (\bar{S}_2 + \bar{S}_3)$.
 - (ii) It then checks if there exists $I \in [N]$ such that $X_I = s \star X_0$ and $\text{ct} = \text{Enc}(\text{pk}, I; r)$.
 - (iii) If all the check above passes, it returns $W = (I, s, r)$.

3. If it finds no witness W of the above form, then it returns $W = \perp$.

We analyze the probability of \mathcal{A} winning the multi-proof online extractability game with the above online extractor OnlineExtract . Below, P' and V' are the prover and verifier of the base traceable OR sigma protocol $\Pi_{\Sigma}^{\text{base}}$ in Fig. 3.

- We say a tuple $\text{input}_{\text{base}} = (X, \text{salt}, j, \text{com}, \text{chall}, \text{resp})$ is valid if the following properties hold:
 - $\text{chall} = 1$;
 - $V_2'^{\mathcal{O}(\text{salt} \parallel j \parallel \cdot)}(\text{com}, \text{chall}, \text{resp})$ outputs accept (i.e., it is a valid transcript for $\Pi_{\Sigma}^{\text{base}}$ with challenge 1);
 - there exists $(\text{seed}, s', r', \text{bits}_1, \dots, \text{bits}_N)$ such that $((\text{salt} \parallel j \parallel \text{Expand} \parallel \text{seed}), (s', r', \text{bits}_1, \dots, \text{bits}_N)) \in L_{\mathcal{O}}$, and if we execute $P_1'^{\mathcal{O}(\text{salt} \parallel j \parallel \cdot)}$ with randomness seed , it produces com . Here, we use the fact that $P_1'^{\mathcal{O}(\text{salt} \parallel j \parallel \cdot)}$ can be executed without the witness. By correctness of $\Pi_{\Sigma}^{\text{base}}$, this implies that $(\text{com}, 0, \text{seed})$ is a valid transcript.
- We say a tuple $\text{input}_{\text{base}} = (X, \text{salt}, j, \text{com}, \text{chall}, \text{resp})$ is invalid if $\text{chall} = 1$, $V_2'^{\mathcal{O}(\text{salt} \parallel j \parallel \cdot)}(\text{com}, \text{chall}, \text{resp})$ outputs accept, but it is not valid.

Observe that if $\text{input}_{\text{base}}$ is valid, then the online extractor can recover a valid transcript $(\text{com}, 0, \text{seed})$ from $\text{input}_{\text{base}}$. Then, it can (informally) extract a witness by combining

it with $(\text{com}, 1, \text{resp})$ and using the extractor from $\Pi_{\Sigma}^{\text{base}}$ constructed in Theorem 6.1. In contrast, if $\text{input}_{\text{base}}$ is invalid, then intuitively, no adversary would be able to prepare a valid response $\text{resp} = \text{seed}$ for the challenge $\text{chall} = 0$ since $L_{\mathcal{O}}$ (i.e., the random oracle query the adversary makes) does not contain a valid response. However, to make this claim formal, we need to also take into account the fact that the adversary may learn non-trivial information about $\text{resp} = \text{seed}$ via the proof output by the prove query. That is, when the challenger runs $P^{\mathcal{O}}$, the adversary may learn non-trivial input/output pairs without directly querying the random oracle itself. In this case, even though no useful information is stored in $L_{\mathcal{O}}$, the adversary may still be able to forge a proof.

We formally show in Lem. 7.2 below that if an adversary \mathcal{A} submits an extract query on a valid input (lbl, X, π) , then a valid $\text{input}_{\text{base}}$ must be included in π (i.e., it cannot consist of $\text{input}_{\text{base}}$ that are all invalid). This allows us to argue that the online extractor will be able to recover two valid transcripts with overwhelming probability, which then further allows the online extractor to extract the witness by running the extractor for the special soundness of the base traceable OR sigma protocol $\Pi_{\Sigma}^{\text{base}}$.

Lemma 7.2 *Assume an adversary \mathcal{A} submits a total of T extract queries of the form $\{(\text{lbl}_k, X_k, \pi_k)\}_{k \in [T]}$, where every π_k is a valid proof including the same salt and satisfies $(\text{lbl}_k, X_k, \pi_k) \notin L_P$. Let $\{(\text{com}_{k,j}, \text{chall}_{k,j}, \text{resp}_{k,j})\}_{j \in [M]}$ be the transcript of the base traceable OR sigma protocol $\Pi_{\Sigma}^{\text{base}}$ that the verification algorithm reconstructs when verifying π_k (see Line 7 of **Verification** $V_2^{\mathcal{O}}$ in Fig. 4). Then, with probability at least $1 - T \cdot (Q_{\text{salt}}/2^{2\lambda-1} + (M \cdot Q_{\text{salt}})/2^{\lambda} + 1/|C_{M,K}|)$, for all $k \in T$ there exists at least one $j \in [M]$ such that $\text{input}_{\text{base}} = (X_k, \text{salt}, j, \text{com}_{k,j}, \text{chall}_{k,j} = 1, \text{resp}_{k,j})$ is valid.*

Proof For any $k \in [T]$, let us redefine $\pi_k = (\overline{\text{com}}, \overline{\text{chall}}, \overline{\text{resp}})$, $(\overline{\text{com}}, \overline{\text{chall}}) = ((\text{salt}, \text{com}_1, \dots, \text{com}_M), \mathbf{c} = (c_1, \dots, c_M))$ where $\mathbf{c} = \mathcal{O}(\text{FS} \parallel \text{lbl} \parallel X \parallel \overline{\text{com}})$, $\overline{\text{resp}} = (\text{seeds}_{\text{internal}}, \{\text{resp}_j\}_{j \text{ s.t. } c_j=1})$. Namely, we omit the subscript k for better readability. We consider two cases: (1) there exists $(\text{lbl}, X, \pi') \in L_P$ such that $\pi' = (\overline{\text{com}}, \overline{\text{chall}}, \overline{\text{resp}'})$ and $\overline{\text{resp}'} \neq \overline{\text{resp}}$ and (2) no such entry in L_P exists.

We consider the first case (1). This corresponds to the case where \mathcal{A} reuses the proof π' obtained through the prove query by simply modifying the response. We claim that this cannot happen with overwhelming probability. Let $\overline{\text{resp}'} = (\text{seed}'_{\text{internal}}, \{\text{resp}'_j\}_{j \text{ s.t. } c_j=1})$. It is clear if $\text{seed}'_{\text{internal}}$ is different from $\text{seeds}_{\text{internal}}$, then \mathcal{A} finds a collision in the random oracle. Since we use a seed tree to generate the randomness used in each base sigma protocol, we can very loosely upper bound the probability of \mathcal{A} outputting such transcript for any $k \in [T]$ by $Q_{\text{salt}}/2^{2\lambda}$. Similarly, consider $\text{resp}'_j \neq \text{resp}_j$ for some j such that $c_j = 1$. Then, it either finds a collision in $\mathcal{O}(\text{Coll} \parallel \cdot)$ (used by the Merkle tree) or $\mathcal{O}(\text{Com} \parallel \cdot)$. We can again very loosely upper bound the probability of \mathcal{A} outputting such transcript for any $k \in [T]$ by $Q_{\text{salt}}/2^{2\lambda}$. Thus, case (1) occurs with probability at most $Q_{\text{salt}}/2^{2\lambda-1}$.

We next consider the second case (2). If $\overline{\text{com}}$ included in π is the same as π' , then $\overline{\text{chall}}$ is the same challenge included in π since the challenge is generated as $\mathcal{O}(\text{FS} \parallel \text{lbl} \parallel X \parallel \overline{\text{com}})$. However, this results in a tuple that falls in the first case (1). Therefore, there exists no π' in L_P that contains the same $\overline{\text{com}}$ as π . This, in particular, implies that the output $\text{chall} \leftarrow \mathcal{O}(\text{FS} \parallel \text{lbl} \parallel X \parallel \overline{\text{com}})$ is distributed uniform random from the view of \mathcal{A} before it makes the hash query.

Now, for the sake of contradiction, we assume $\text{input}_{\text{base},j} = (X, \text{salt}, j, \text{com}_j, c_j, \text{resp}_j)$ is invalid for all $j \in [M]$ such that $c_j = 1$. Let $L_{\mathcal{O}_p}$ be a list that contains all the inputs/outputs of the random oracle queries $\text{Prove}^{\mathcal{O}}$ makes when the challenger answers the prove query made by \mathcal{A} . We prove the following corollary.

Corollary 7.3 For any $j^* \in [M]$, if $\text{input}_{\text{base},j^*}$ is invalid, then either of the following holds:

- there exists no tuple $(s', r', \text{bits}_1, \dots, \text{bits}_N, \text{seed})$ and $j' \in [M]$ such that $((\text{salt} \parallel j' \parallel \text{Expand} \parallel \text{seed}), (s', r', \text{bits}_1, \dots, \text{bits}_N)) \in L_{\mathcal{O}_p}$, but if we execute $P_1^{\mathcal{O}(\text{salt} \parallel j' \parallel \cdot)}$ with randomness seed , it produces com_{j^*} ;
- there exists such a tuple but seed retains λ -bits of min-entropy from the view of \mathcal{A} except with probability at most $(M Q_{\text{salt}})/2^\lambda$.

Proof Assume such an entry is found in $L_{\mathcal{O}_p}$. This corresponds to the case \mathcal{A} is reusing com_{j^*} that was included in a proof π obtained through the prove query. Let $\{(\text{com}'_{j'}, c'_{j'}, \text{resp}'_{j'})\}_{j' \in [M]}$ be the transcript of the base traceable OR sigma protocol $\Pi_{\Sigma}^{\text{base}}$ that the verification algorithm reconstructs from such π (see Line 7 of **Verification** $V_2^{\mathcal{O}}$ in Fig. 4), where $\text{com}'_{j'} = \text{com}_{j^*}$. Our current goal is to prove that $c'_{j'} = 1$ (i.e., seed was not used as a response). Since $\text{com}'_{j'}$ and com_{j^*} are roots of a Merkle tree and the indices j' and j^* are used as prefix to the hash when constructing the roots, respectively, the probability of \mathcal{A} outputting com_{j^*} such that $j' \neq j^*$ is upper bounded by $((M - 1)Q_{\text{salt}})/2^\lambda$. Below, we assume $j' = j^*$. Recall by definition of the online extractability game (see Def. 3.10), \mathcal{A} runs the verification algorithm to check if π is valid. Therefore, if $\text{input}_{\text{base},j^*}$ is invalid, then we have $c'_{j'} = 1$. Otherwise, there must exist an entry $((\text{salt} \parallel j^* \parallel \text{Expand} \parallel \text{seed}), (s', r', \text{bits}_1, \dots, \text{bits}_N)) \in L_{\mathcal{O}}$, which contradicts that $\text{input}_{\text{base},j^*}$ is invalid. This further implies that $\text{resp}'_{j'}$ does not include seed . Then, by Lem. A.3 regarding the seed tree, seed that was used to construct $\text{com}_{j'} = \text{com}_{j^*}$ is statistically hidden to the adversary with all but probability $Q_{\text{salt}}/2^\lambda$. The proof is completed by collecting all the bounds. \square

By Lem. 7.3, if $\text{input}_{\text{base},j}$ is invalid, then \mathcal{A} cannot prepare a valid response for the challenge $c_j = 0$ with all but probability at most $(M Q_{\text{salt}})/2^\lambda$. This is because such response is either not recorded in both $L_{\mathcal{O}}$ and $L_{\mathcal{O}_p}$, or it is recorded in $L_{\mathcal{O}_p}$ but the seed retains λ -bits of min-entropy from the view of \mathcal{A} except with probability $(M Q_{\text{salt}})/2^\lambda$. Moreover, since chall is statistically hidden to \mathcal{A} before it queries the random oracle, the probability that chall coincides with challenges for which \mathcal{A} can open to is at most $1 - 1/|C_{M,K}|$, where recall $C_{M,K}$ is the challenge space (or equivalently the output space of $\mathcal{O}(\text{FS} \parallel \cdot)$).

Taking the union bound and collecting all the bounds together, at least one of the $\text{input}_{\text{base}}$ must be valid with the probability stated in the statement. This completes the proof of the lemma. \square

We are now prepared to analyze the probability that \mathcal{A} wins the multi-proof online extractability game with the aforementioned online extractor OnlineExtract . By Lem. 7.2, if \mathcal{A} makes at most T extract queries, then by a simple union bound and using the inequality $\sum_i Q_{\text{salt}_i} \leq Q$, with probability at least $1 - T \cdot ((2Q)/2^{2\lambda} + (M \cdot Q)/2^\lambda + 1/|C_{M,K}|)$, all the $\text{input}_{\text{base}}$ included in the queried proof are valid. Then, by the definition of valid and the description of OnlineExtract , OnlineExtract is able to extract two valid transcripts for all T proofs queried by \mathcal{A} . Recalling Thms. 6.1 and 6.3, OnlineExtract either succeeds in extracting a witness $W = (I, s, r) \in [N] \times (S_2 + S_3) \times (\bar{S}_2 + \bar{S}_3)$ or a witness that consists of a collision in $\mathcal{O}(\text{salt} \parallel j \parallel \text{Coll} \parallel \cdot)$ or $\mathcal{O}(\text{salt} \parallel j \parallel \text{Com} \parallel \cdot)$ for some $j \in [M]$. Hence, with all but probability $Q^2/2^{2\lambda}$, OnlineExtract succeeds in extracting a witness $W = (I, s, r)$ as desired, conditioned on all the $\text{input}_{\text{base}}$ included in the queried proof are valid. Collecting the bounds, we arrive at our statement. \square

Theorem 7.4 The NIZK with labels $\Pi_{\text{NIZK},|b|}$ in Fig. 7 is zero-knowledge. Precisely, there exists a PPT simulator $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ such that, for any statement-witness pair $(X, W) \in R_{\text{sig}}$

$\text{Sim}_0(x)$ 1: if $x \in L$ then 2: return $L[x]$ 3: end if 4: $y \leftarrow D_x$ 5: $L[x] := y$ 6: return y	$\text{Sim}_1(\text{lbl}, X)$ 1: $\text{chall} \leftarrow C_{M,K}$ 2: $(\text{com}, \text{chall}, \text{resp}) \leftarrow \text{Sim}_\Sigma^{\text{Sim}_0}(X, \text{chall})$ 3: if $(\text{FS} \parallel \text{lbl} \parallel X \parallel \text{com}) \in L$ then 4: return \perp 5: end if 6: $L[(\text{FS} \parallel \text{lbl} \parallel X \parallel \text{com})] := \text{chall}$ 7: return $\pi \leftarrow (\text{com}, \text{chall}, \text{resp})$
---	--

Fig. 8 Zero-knowledge simulator $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ for $\Pi_{\text{NIZK}, \text{lbl}}$

Fig. 9 Intermediate simulator Sim_{int} , where $P = (P_1, P_2)$ is the prover of the traceable OR sigma protocol Π_Σ^{TOR} in Fig. 4

$$\text{Sim}_{\text{int}}(\text{lbl}, X, W)$$

```

1: com ← P1Sim0(X, W)
2: chall ← CM,K
3: if (FS || lbl || X || com) ∈ L then
4:   return ⊥
5: end if
6: L[(FS || lbl || X || com)] := chall
7: resp ← P2Sim0(X, chall)
8: return π ← (com, chall, resp)
    
```

and any computationally-unbounded adversary \mathcal{A} that makes at most Q_1 queries to \mathcal{O} or Sim_0 , and Q_2 queries to Prove or \mathcal{S} , we have

$$\begin{aligned} \text{Adv}_{\Pi_{\text{NIZK}, \text{lbl}}}^{\text{ZK}}(\mathcal{A}) &= \left| \Pr \left[\mathcal{A}^{\mathcal{O}, \text{Prove}}(1^\lambda) = 1 \right] - \Pr \left[\mathcal{A}^{\text{Sim}_0, \mathcal{S}}(1^\lambda) = 1 \right] \right| \\ &\leq \frac{Q_2 \cdot (Q_1 + Q_2)}{2^{2\lambda}} + \frac{Q_1}{2^\lambda}. \end{aligned}$$

Proof To prove the zero-knowledge property of $\Pi_{\text{NIZK}, \text{lbl}} = (\text{Prove}^\mathcal{O}, \text{Verify}^\mathcal{O})$, we define a zero-knowledge simulator $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ in Fig. 8, where Sim_0 and Sim_1 share states, including a list L which is initially empty. At a high level, Sim_0 simulates the random oracle \mathcal{O} in an on-the-fly manner but replaces certain queries for consistency with Sim_1 . On the other hand, Sim_1 simulates the prover oracle using the simulator from the underlying sigma protocol, which we denote here by Sim_Σ (see Theorem 6.4), as a subroutine. Specifically, Sim_1 is given a valid statement $X = (\{X_i\}_{i \in [N]}, \text{pk}, \text{ct})$, and samples a random challenge chall from the challenge space $C_{M,K}$, which is also the output space of $\mathcal{O}(\text{FS} \parallel \cdot)$. It then runs Sim_Σ on challenge chall by providing it oracle access to Sim_0 , and updates the list L accordingly. In Fig. 8, we denote by D_x the distribution of $\mathcal{O}(x)$, where the probability is taken over the random choice of the random oracle \mathcal{O} . Without loss of generality, we assume D_x to be efficiently sampleable.

To show the indistinguishability of $(\mathcal{O}, \text{Prove})$ and $(\text{Sim}_0, \mathcal{S})$, we use a hybrid argument by introducing an intermediate pair of simulators $(\text{Sim}_0, \text{Sim}_{\text{int}})$, where Sim_{int} is defined in Fig. 9. Let \mathcal{S}_{int} , analog to Prove and \mathcal{S} , be an oracle that on input (lbl, X, W) returns \perp if $\text{lbl} \notin L \vee (X, W) \notin R_{\text{sig}}$ and otherwise returns $\text{Sim}_{\text{int}}(\text{lbl}, X, W)$.

Suppose \mathcal{A} makes Q_1 queries to the oracles \mathcal{O} or Sim_0 , and Q_2 queries to the oracles Prove , \mathcal{S}_{int} , or \mathcal{S} . For each $i \in \{1, 2, 3\}$, we denote by E_i the event that \mathcal{A} returns 1 respectively. We analyze the differences by defining three games as follows:

- Game₁ : This is the real zero-knowledge game where \mathcal{A} is given access to \mathcal{O} and Prove .
- Game₂ : The game is modified to provide \mathcal{A} access to Sim_0 and \mathcal{S}_{int} instead. The view of \mathcal{A} is identical to the previous game unless Sim_{int} outputs \perp in Line 4. Roughly,

this occurs when the reprogramming of the random oracle fails due to the input being already defined. By Theorem 6.3, com has 2λ bits of min-entropy. Since at most $Q_1 + Q_2$ queries of the form $(F5 \parallel \text{lbl} \parallel X \parallel \text{com})$ are made in this game, we have $|\Pr[E_1] - \Pr[E_2]| \leq \frac{Q_2 \cdot (Q_1 + Q_2)}{2^{2\lambda}}$.

Game₃ : The game is modified to provide \mathcal{A} access to Sim_0 and \mathcal{S} instead. The only difference is that rather than computing honestly via (P_1, P_2) from the traceable OR sigma protocol $\Pi_{\Sigma}^{\text{TOR}}$, the simulator Sim_1 simulates these using the simulator Sim_{Σ} provided by $\Pi_{\Sigma}^{\text{TOR}}$.

Let salt_i represent the salt that Sim_{int} or Sim_1 samples on its i -th invocation. For $i \in [Q_2]$, let Q'_i be the number of queries the adversary makes to oracle Sim_0 of the form $(\text{salt}_i \parallel \cdot)$. By Theorem 6.4, the advantage of the adversary in distinguishing Sim_{int} or Sim_1 is bounded by $\frac{Q'_i}{2^{\lambda}}$ for each $i \in [Q_2]$. Therefore,

$$|\Pr[E_2] - \Pr[E_3]| \leq \frac{\sum_1^{Q_2} Q'_i}{2^{\lambda}} \leq \frac{Q_1}{2^{\lambda}}$$

Collecting the bounds, we obtain the bound in the statement. □

8 Instantiations

We instantiate the building blocks required for our generic construction of an accountable ring signature scheme presented in Sect. 4 via isogenies based on CSIDH group action and lattices.

8.1 Instantiation from isogenies

We instantiate a group-action-based HIG and PKE, and the corresponding NIZKs for the relations R_{sig} and R_{open} based on the CSIDH paradigm. In particular we assume that the structure of the ideal class group $\mathcal{Cl}(\mathcal{O})$ is known, and cyclic of odd order n , so that it is isomorphic to \mathbb{Z}_n . Given a generator \mathfrak{g} of $\mathcal{Cl}(\mathcal{O})$, \mathbb{Z}_n acts freely and transitively on $\mathcal{Ell}p(\mathcal{O}, \pi)$ via the group action $\star : (a, E) \mapsto \mathfrak{g}^a * E$, which we can compute efficiently. Note that in case the class group structure is not known (e.g., at higher security levels where computing the class group is currently not feasible.) we can still instantiate all the building blocks using rejection sampling *à la* SeaSign.

Group-Action-Based HIG. We instantiate the group-action-based HIG defined by the algorithms $(\text{RelSetup}, \text{IGen})$ as follows. The output of RelSetup describes a setup for a CSIDH group action $\star : \mathcal{Cl}(\mathcal{O}) \times \mathcal{Ell}p(\mathcal{O}, \pi) \rightarrow \mathcal{Ell}p(\mathcal{O}, \pi)$, sets $G = S_1 = S_2 = \mathcal{Cl}(\mathcal{O})$, $\delta = 1$, $\mathcal{X} = \mathcal{Ell}p(\mathcal{O}, \pi)$, and $X_0 = E_0$, where E_0 is the elliptic curve $E_0 : y^2 = x^3 + x$ over \mathbb{F}_p . The output of IGen is then $(E_0, \mathfrak{a} * E_0)$, where \mathfrak{a} is uniformly sampled from $\mathcal{Cl}(\mathcal{O})$. Then the properties of Def. 5.1 are easily verified. In particular, the security of the hard instance generator is equivalent to the hardness of GAIP for CSIDH. Moreover, it is not difficult to see that the group-action-based HIG is also a hard *multi-instance* generator based on the same assumption. Concretely, given one instance (E_0, E) , the reduction can rerandomize this arbitrarily many times to obtain fresh statements $(E_0, \mathfrak{b} * E)$, where \mathfrak{b} is uniformly sampled from $\mathcal{Cl}(\mathcal{O})$. If an adversary succeeds in breaking any of these instances, then the reduction can subtract \mathfrak{b} from it to solve its original instance.

Group-Action-Based PKE. We can define an ElGamal-like public-key encryption scheme $\Pi_{\text{GA-PKE}} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ based on the CSIDH group action, as follows. Note that the decryption algorithm works by enumerating the message space, so the PKE is only

efficient when the message space \mathcal{M} (which is a subset of $\mathcal{Cl}(\mathcal{O})$) is polynomially large. This relaxed notion of decryption suffices for our ARS generic construction.

- Setup(1^λ) \rightarrow pp: On input a security parameter 1^λ , it returns the setup for a CSIDH group action $\star : \mathcal{Cl}(\mathcal{O}) \times \mathcal{Ellp}(\mathcal{O}, \pi) \rightarrow \mathcal{Ellp}(\mathcal{O}, \pi)$, and sets $G = G_M = S_1 = S_2 = \mathcal{Cl}(\mathcal{O})$, $\mathcal{Y} = \mathcal{Ellp}(\mathcal{O}, \pi) \times \mathcal{Ellp}(\mathcal{O}, \pi)$, $\delta_y = 1$. The “message” group action $\star_M : G \times \mathcal{Y} \rightarrow \mathcal{Y}$ is defined as $(a, (E_1, E_2)) \mapsto (E_1, a \star E_2)$ (i.e., \star_M acts on the second component only).
- KeyGen(pp) \rightarrow (pk, sk): On input a public parameter pp, it returns a secret key sk sampled uniformly from $\mathcal{Cl}(\mathcal{O})$, and a public key pk = (\star_{pk} , Y_{pk}), where $\star_{pk} : G \times \mathcal{Y} \rightarrow \mathcal{Y}$ is defined as $(a, (E_1, E_2)) \mapsto (a \star E_1, a \star E_2)$ (i.e., \star_{pk} acts on both components), and $Y_{pk} = (E_0, sk \star E_0)$.
- Enc(pk, M; r) \rightarrow ct: On input a public key pk = (\star_{pk} , Y_{pk}) and a message $M \in \mathcal{M}$, it returns the ciphertext ct = ($M \star_M (r \star_{pk} Y_{pk})$) $\in \mathcal{Y}$, where $r \leftarrow G$.
- Dec(sk, ct) \rightarrow M: On input a secret key sk and a ciphertext ct = (ct₁, ct₂), the decryption algorithm tries all messages $M \in \mathcal{M}$ until it finds a message M such that $M \star ct_1 = -sk \star ct_2$. If such a message exists, it is, unique, and the algorithm outputs it; otherwise, \perp is output.

It is not difficult to verify that the above-defined Π_{GA} -PKE is correct (with probability 1). The decryption scheme of Π_{GA} -PKE differs from that of ElGamal since it is not possible to *divide out* $sk \star ct_1$ from ct_2 . Therefore, retrieving M from ct_1, ct_2, sk requires the resolution of an instance of GAIP with input $(sk \star ct_1, ct_2)$. Dec solves this problem by a brute force over the message space \mathcal{M} . In case \mathcal{M} is polynomially large, then we have efficient decryption as desired.

Multi-challengeIND-CPA Security. The scheme is multi-challenge IND-CPA secure based on the dCSIDH assumption. Since Π_{GA} -PKE is an ElGamal-like encryption scheme in the CSIDH setting — where each exponentiation is replaced by a group action — for the security proof it is sufficient to adapt the usual proof for the group-based ElGamal encryption scheme. Note that the the reduction loses a factor $1/Q_{ct}$, where Q_{ct} is the number of challenge ciphertext the adversary observes. This is the only reason why we do not achieve tight security for our accountable ring signature and group signature.

We point out that by ignoring the PKE, we obtain a ring signature identical to Beullens et al. [14]. Thus we obtain the first tightly secure and efficient isogeny-based ring signature in this work.

($\mathcal{R}', \mathcal{KR}'$)-correctness. In the isogeny setting, it is not needed to relax the key relation (contrary to our lattice instantiation where some relaxation is necessary in order to get an efficient opening proof). We can simply set $\mathcal{KR}' = \mathcal{KR} = \{(E, sk) \mid sk \star E_0 = E\} \subseteq \mathcal{Ellp}(\mathcal{O}, \pi) \times \mathcal{Cl}$. Similarly, since $S_2 = S_1$, there is no relaxation in the encryption randomness. Therefore ($\mathcal{R}', \mathcal{KR}'$)-correctness is equivalent to the standard correctness property (with probability 1), which is satisfied by our PKE.

Multi-Proof Online Extractable NIZK with Labels $\Pi_{NIZK,|bl}$. Using the group-action-based HIG and PKE, we can instantiate $\Pi_{NIZK,|bl}$ for the signing relation R_{sig} (see Sect. 4.1) as explained in Sects. 6 and 7.

Statistically Sound NIZK without Labels Π_{NIZK} . The last ingredient for our ARS is a NIZK for the opening relation R_{open} , which in our instantiation is

$$R_{open} = \{((pk, ct = (E_1, E_2), M), sk) \mid sk \star E_0 = pk \wedge M \star sk \star E_1 = E_2\}.$$

A sigma protocol for this relation was introduced in [43, Sect. 3.2]. We can then turn this sigma protocol into an NIZK by applying the Fiat-Shamir transform. (Note that we do not need this NIZK to be online-extractable.)

Concrete Instantiation for Tab. 1. For our isogeny-based instantiation, we chose an HIG and a PKE based on the CSIDH-512 group action. The structure of this class group has been computed [15], which allows for more efficient proofs. We chose the challenge space as string of length $M = 855$ with Hamming weight $K = 19$. Most of the signature is independent of N , and contains a fixed number of curves and class group elements as well as some overhead from the generic construction such as a hash value, the internal nodes in the seed tree, and commitment randomness to open the commitments. The only reason the signature size increases with N is that the signature contains a fixed amount of paths in a Merkle tree of depth $\log_2 N$. This makes for a very mild dependence on N .

8.2 Instantiation from lattices

We instantiate a group-action-based HIG and PKE, and the corresponding NIZKs for the relations R_{sig} and R_{open} based on lattices under the MSIS and MLWE assumptions. The choices for the integer n , modulus q , and ring R_q are provided in Sect. 3.6.

Group-Action-Based HIG. By Def. 5.1, it suffices to define the public parameter $\text{pp}_1 = (G, S_1, S_2, \delta_x, X_0, \mathcal{X}, \star)$ generated by RelSetup and to check that the output of IGen defines a hard relation. The public parameters pp are defined as follows:

- $(G, \mathcal{X}) = (R_q^\ell \times R_q^k, R_q^k)$, where X_0 is an arbitrary element in \mathcal{X} ,
- For $b \in \{0, 1\}$, $S_b = \{(\mathbf{s}, \mathbf{z}) \in G \mid \|\mathbf{s}\|_\infty, \|\mathbf{e}\|_\infty \leq B_b\}$, where B_1, B_2 are positive integers such that $B_1 < B_2 < q$,
- $\delta_x = \left(\frac{2(B_2 - B_1) + 1}{2B_2 + 1}\right)^{n(k + \ell)}$,
- The group action $\star : G \times \mathcal{X} \rightarrow \mathcal{X}$ is defined as $(\mathbf{s}, \mathbf{e}) \star \mathbf{w} = (\mathbf{A}\mathbf{s} + \mathbf{z}) + \mathbf{w}$, where $\mathbf{A} \in R_q^{k \times \ell}$ is a fixed matrix sampled uniformly by RelSetup .

We define S_3 to be a subset of G with coefficients all bounded by $B_2 - B_1$. It can be checked that pp satisfies all the conditions in Def. 5.1, where δ_x follows by simply counting the points included in S_2 and S_3 . It remains to check that the relation $\tilde{R}_{\text{pp}} = \{(\mathbf{b}, (\mathbf{s}, \mathbf{z})) \mid \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \wedge (\mathbf{s}, \mathbf{e}) \in S_2 + S_3\}$ defines a hard relation as defined in Sect. 4.1, where $S_2 + S_3$. Note that if the adversary \mathcal{A} is restricted to output a witness $(\mathbf{s}, \mathbf{e}) \in S_1$, then this follows directly from the MLWE_{n,q,B_1} assumption. For our application, we have to further consider the scenario where \mathcal{A} may output a witness (\mathbf{s}, \mathbf{e}) outside of S_1 . We need to consider this case since our online extractor for the NIZK can only extract a witness in the relaxed relation \tilde{R}_{pp} rather than R_{pp} .

The hardness of our group-action-based HIG follows naturally from the $\text{MSIS}_{n,q,k,\ell,2B_2}$ and $\text{sMLWE}_{n,q,k,\ell,B_1}$ assumptions. We only focus on an adversary \mathcal{A} that outputs a witness (\mathbf{s}, \mathbf{e}) outside of S_1 , since the other case simply follows from MLWE as we seen above. Let us construct an adversary \mathcal{B} against the $\text{MSIS}_{n,q,k,\ell,2B_2}$ problem by using \mathcal{A} as a subroutine. \mathcal{B} , given \mathbf{A} as input, samples a random $(\mathbf{s}, \mathbf{e}) \leftarrow S_1$, sets $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$ and invokes \mathcal{A} on input pp, \mathbf{b} , where pp includes \mathbf{A} . When \mathcal{A} outputs $(\mathbf{s}', \mathbf{e}')$, \mathcal{B} submits $(\mathbf{s} + \mathbf{s}', \mathbf{e} + \mathbf{e}')$ as its solution. By assumption, $\|\mathbf{s} + \mathbf{s}'\|_\infty, \|\mathbf{e} + \mathbf{e}'\|_\infty \leq B_1 + B_2 + B_3 = 2B_2$ and they are non-zero. Therefore, \mathcal{B} breaks the $\text{MSIS}_{n,q,k,\ell,2B_2}$ problem as desired.

Finally, the same proof shows that our group-action-based HIG is a hard *multi-instance* generator based on the same assumptions.

Group-Action-Based PKE. We use a PKE scheme based on the Lindner-Peikert framework [73]. We first explain the public parameters $\text{pp}_2 = (\overline{G}, \overline{G}_T, \mathcal{Y}, \overline{S}_1, \overline{S}_2, \delta_y, D_{\mathcal{Y}}, \star_M, \mathcal{M})$ generated by PKE.Setup .¹¹

- $(\overline{G}, \overline{G}_T, \mathcal{Y}) = (R_q^k \times R_q^\ell \times R_q, R_q, R_q^k \times R_q)$,
- For $b \in \{0, 1\}$, $\overline{S}_b = \{(\mathbf{r}, \mathbf{e}, e) \in \overline{G} \mid \|\mathbf{r}\|_\infty, \|\mathbf{e}\|_\infty, \|e\|_\infty \leq B_b\}$, where B_1, B_2 are positive integers such that $B_1 < B_2 < q$ and $4(nk + 1)(2B_2 - B_1) \leq q$,
- $\delta_y = \left(\frac{2(B_2 - B_1) + 1}{2B_2 + 1}\right)^{n(k+\ell+1)}$,
- $D_{\mathcal{Y}}$ is a distribution that samples a uniform random $(\mathbf{A}, \mathbf{s}, \mathbf{z}) \in R^{k \times \ell} \times R_q^\ell \times R_q^k$ and outputs a group action $\star : \overline{G} \times \mathcal{Y} \rightarrow \mathcal{Y}$ defined as $(\mathbf{r}, \mathbf{e}, e)\star(\mathbf{w}, w) = ((\mathbf{A}^\top \mathbf{r} + \mathbf{e} + \mathbf{w}, \mathbf{b}^\top \mathbf{r} + e + w)$ and an element $Y = (\mathbf{w}, w) \in \mathcal{Y}$, where $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{z}$,
- $\star_M : \overline{G}_T \times \mathcal{Y} \rightarrow \mathcal{Y}$ is a group action defined as $M\star_M(\mathbf{c}, c) = (\mathbf{c}, c + M \cdot \lfloor q/2 \rfloor)$,
- The message space \mathcal{M} is a subset of $\overline{G}_T = R_q$ with coefficients in $\{0, 1\}$.

We define S_3 to be a subset of G with coefficients all bounded by $B_2 - B_1$. It can be checked that pp satisfies the conditions in Def. 5.2, where δ_y follows by simply counting the points included in S_2 and S_3 . The remaining algorithms (KeyGen, Enc, Dec) are defined as follows, where $U(B)$ denotes elements in R_q with infinity norm at most $B \in \mathbb{N}$:

KeyGen(pp): It samples a uniform random $(\mathbf{A}, \mathbf{s}, \mathbf{z}) \in R^{k \times \ell} \times U(B_1)^\ell \times U(B_1)^k$ and outputs $(\text{pk}, \text{sk}) = ((\star_{\text{pk}}, \mathbf{0}), \mathbf{s})$, where $\mathbf{0}$ is the zero polynomial in \mathcal{Y} and \star_{pk} is a group action defined as $(\mathbf{r}, \mathbf{e}, e)\star_{\text{pk}}(\mathbf{w}, w) = (\mathbf{A}^\top \mathbf{r} + \mathbf{e} + \mathbf{w}, \mathbf{b}^\top \mathbf{r} + e + w)$, where $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{z}$. Note that pk is distributed as a sample from $D_{\mathcal{Y}}$.

Enc(pk, M): On input a public key $\text{pk} = (\star_{\text{pk}}, Y_{\text{pk}} = \mathbf{0})$ and a message $M \in \mathcal{M}$, it samples $(\mathbf{r}, \mathbf{e}, e) \leftarrow \overline{S}_1$ and returns $\text{ct} = M\star_M((\mathbf{r}, \mathbf{e}, e)\star_{\text{pk}}\mathbf{0}) = (\mathbf{A}^\top \mathbf{r} + \mathbf{e}, \mathbf{b}^\top \mathbf{r} + e + M \cdot \lfloor q/2 \rfloor) \in \mathcal{Y}$.

Dec(sk, ct) → M: It parses $(\mathbf{c}, c) \leftarrow \text{ct}$ and computes $w = c - \mathbf{c}^\top \mathbf{s}$ over R_q . It rounds each coefficient back to either 0 or $\lfloor q/2 \rfloor$ whichever is closest modulo q and outputs the polynomial.

Correctness is a consequence of $(\mathcal{R}', \mathcal{KR}')$ -correctness, which we show below, and decryption efficiency clearly holds as well. We discuss the remaining properties.

Multi-challengeIND-CPA Security. The security follows by a standard proof using dMLWE. For completeness, we provide the proof: We consider a sequence of games and prove that the adversary’s advantage only changes negligibly in each adjacent games. The first game is the original security game. In the second game, we modify the group action \star_{pk} included in the public key to be defined by a random $(\mathbf{A}, \mathbf{b}) \leftarrow R^{k \times \ell} \times R_q^k$. By the dMLWE $_{n,q,k,\ell,B_1}$ assumption, this game is indistinguishable from the previous game. In the final game, we sample each ciphertext as $\text{ct} \leftarrow R^k \times R_q$. By the dMLWE $_{n,q,\ell+1,k,B_1}$ assumption, this game is indistinguishable from the previous game. Note that we appropriately parse the matrix $\mathbf{A}' \in R_q^{(\ell+1) \times k}$ provided by the challenge as \mathbf{A} and \mathbf{b} , and query the oracle once for each ciphertext. Since the challenge bit b is statistically hidden from the adversary, no adversary has advantage in winning this game. This concludes the proof.

We note that we can prove multi-challenge IND-CPA security while only relying on the dMLWE assumption with a fixed number of instances (i.e., those that do not rely on the number of challenge ciphertexts), if we can tolerate choosing slightly less efficient parameters. Specifically, we can use the dual-Regev encryption [56], where \mathbf{A} is a tall matrix. When \mathbf{A}

¹¹ Note that although we use the same (q, B_1, B_2) as those used by the group-action-based HIG, they can be set differently. We only use the same notations for better readability.

is tall enough, $\mathbf{A}^\top \mathbf{r}$ and $\mathbf{b}^\top \mathbf{r}$ is distributed statistically close to random under appropriate choices of parameters owing to the regularity lemma [77]. Hence, we only need the dMLWE assumption to jump from the first to second game above.

$(\mathcal{R}', \mathcal{KR}')$ – correctness We define \mathcal{R}' and \mathcal{KR}' as follows, where the choice of \mathcal{R}' coincides with those considered in Theorem 6.1:

- $(\mathcal{R}', \mathcal{KR}') = (\bar{S}_2 + \bar{S}_3, U(2B_2 - B_1)^\ell \times U(2B_2 - B_1)^k)$, where recall S_3 is a subset of G with ring elements whose coefficients are all bounded by $B_2 - B_1$. Specifically, $\bar{S}_2 + \bar{S}_3 = \{(\mathbf{r}, \mathbf{e}, e) \in \bar{G} \mid \|\mathbf{r}\|_\infty, \|\mathbf{e}\|_\infty, \|e\|_\infty \leq 2B_2 - B_1\}$.

We check that correctness holds even if the ciphertext is encrypted using randomness $(\mathbf{r}, \mathbf{e}, e) \in \mathcal{R}'$ and a secret key $\text{sk} = (\mathbf{s}, \mathbf{e}) \in \mathcal{KR}'$. Let $\text{ct} = (\mathbf{A}^\top \mathbf{r} + \mathbf{e}, \mathbf{b}^\top \mathbf{r} + e + M \cdot \lfloor q/2 \rfloor)$, then $c - \mathbf{c}^\top \mathbf{s} = M \cdot \lfloor q/2 \rfloor + e + \mathbf{e}^\top \mathbf{s} - \mathbf{z}^\top \mathbf{r}$. Then, $\|e + \mathbf{e}^\top \mathbf{s} - \mathbf{z}^\top \mathbf{r}\|_\infty \leq \|e\|_\infty + \|\mathbf{e}^\top \mathbf{s}\|_\infty + \|\mathbf{z}^\top \mathbf{r}\|_\infty \leq (2B_2 - B_1) + 2nk(2B_2 - B_1)^2 \leq q/4$, where the last inequality follows from our parameter choice. Thus, M can be correctly decrypted with probability 1.

Multi-Proof Online Extractable NIZK with Labels $\Pi_{\text{NIZK}, \text{lbl}}$. Using the group-action-based HIG and PKE, we can instantiate $\Pi_{\text{NIZK}, \text{lbl}}$ for the signing relations R_{sig} and R'_{sig} (see Sect. 4.1) as explained in Sects. 6 and 7.

Statistically Sound NIZK without Labels Π_{NIZK} . It remains to show how to construct Π_{NIZK} for the opening relations R_{open} and R'_{open} . We can rewrite the relation R_{open} (see Sect. 4.1) as follows:

$$R_{\text{open}} = \left\{ ((\text{pk} = \mathbf{b}, \text{ct} = (\mathbf{c}, c), M), \text{sk} = (\mathbf{s}, \mathbf{z})) \mid \begin{array}{l} \|\mathbf{s}\|_\infty, \|\mathbf{e}\|_\infty \leq B_1 \wedge \mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{z} \\ \wedge \|c - \mathbf{c}^\top \mathbf{s} - M \cdot \lfloor q/2 \rfloor\|_\infty \leq q/4 \end{array} \right\}.$$

Notice we can rewrite the righthand side as

$$\underbrace{\begin{bmatrix} \mathbf{A} \\ \mathbf{c}^\top \end{bmatrix}}_{\tilde{\mathbf{A}}} \mathbf{s} + \underbrace{\begin{bmatrix} \mathbf{z} \\ 0 \end{bmatrix}}_{\tilde{\mathbf{z}}} = \underbrace{\begin{bmatrix} \mathbf{b} \\ c - M \cdot \lfloor q/2 \rfloor + d \end{bmatrix}}_{\tilde{\mathbf{b}}},$$

where d is some element in R_q such that $\|d\|_\infty \leq q/4$. Since d is not secret, we can think d is included in the statement $(\text{pk}, \text{ct}, M)$. Then, Π_{NIZK} can simply viewed as an NIZK for the standard MLWE-based statement $\tilde{\mathbf{A}}\mathbf{s} + \tilde{\mathbf{z}} = \tilde{\mathbf{b}}$, where $\|\mathbf{s}\|_\infty, \|\tilde{\mathbf{z}}\|_\infty \leq B_1$. Notice that such a statement is implicitly used in $\Pi_{\text{NIZK}, \text{lbl}}$ for the relation R_{sig} since this statement is essentially the group-action-based HIG. Specifically, if we remove all the components regarding the OR proof and leave the proof regarding the group-action-based HIG from Figs. 3,4 and 7, we arrive at our desired NIZK. Similarly to $\Pi_{\text{NIZK}, \text{lbl}}$ for the relation R_{sig} , we can only prove that a cheating prover was using a witness (i.e., secret key) satisfying $\|\mathbf{s}\|_\infty, \|\tilde{\mathbf{z}}\|_\infty \leq B_2 + B_3$. This is exactly the \mathcal{KR}' defined above and coincides with the relaxed relation \tilde{R}'_{open} .

One may wonder if we can construct an NIZK for this standard MLWE relation based on a sigma protocol with a non-binary challenge set. Although the proof size of Π_{NIZK} is already constant, this may further minimize the proof size of the opening proof. We claim that this may be difficult. The main reason is that when we use a non-binary challenge space, the extracted witness $(\mathbf{s}, \tilde{\mathbf{z}})$ typically comes from a *furthered* relaxed relation such that not only they have a larger norm, they are guaranteed to only satisfy $\tilde{\mathbf{A}}\mathbf{s} + \tilde{\mathbf{z}} = t \cdot \tilde{\mathbf{b}}$ for some short $t \in R_q$. This relaxation may suffice in some settings but it turns out that it won't for ours as we can no longer prove $(\mathcal{R}', \mathcal{KR}')$ -correctness. When restricted to binary challenges, we can control t to be $1 \in R_q$.

Remark 8.1 (Bai-Galbraith Optimization [7]) We can apply the Bai-Galbraith optimization [7] by exploiting the lattice structure. This is a common and simple optimization used in

various lattice-based interactive protocols based on the Fiat-Shamir with aborts paradigm [78] that allows to roughly halve the proof size, or signature size when viewing the proof as a signature, with no additional cost. Intuitively, for MLWE, proving knowledge of a short \mathbf{s} indirectly proves knowledge of a short \mathbf{e} since it is uniquely defined as $\mathbf{b} - \mathbf{A}\mathbf{s}$. Therefore, we can remove the components that are used to explicitly prove that \mathbf{e} is short. Since the size of \mathbf{s} and \mathbf{e} are about the same in our construction, this allows to almost halve the proof size. For further details, see for example [7, 14, 42].

Concrete Instantiation for Tab. 1. For the concrete instantiation in Tab. 1, we use $M = 1749$, $K = 16$. For the HIG, we chose the parameters according to the parameters used in the Security Level II variant of the (round 3) NIST submission of the Dilithium signature scheme. Concretely, we use the ring $R_q = \mathbb{Z}_q[X]/(X^n + 1)$, with $n = 256$ and $q = 2^{23} - 2^{13} + 1$, and we put $l = k = 4$, $B_1 = 2$, $B_2 = 2^{17}$. These parameters are chosen by the Dilithium team such that the relevant MLWE and MSIS problems are hard enough to reach NIST SL II.

For the PKE, we use the ring R'_q with $n = 256$ and $q' \approx 2^{49}$, and we put $k = l = 8$, $B_1 = 1$, $B_2 \approx 2^{16.3}$. The LWE estimator of Albrecht et al. estimates that this MLWE instance has 141 bits of security [3]. Moreover, the $(\mathcal{R}', \mathcal{KR}')$ -correctness holds, because we have $(2B_2 - B_1) + 2nk(2B_2 - B_1)^2 \leq q/4$. For the parameter set without manager accountability, we only require $(\mathcal{R}', \mathcal{KR})$ -correctness, so we only need $(2B_2 - B_1) + 2nk(2B_2 - B_1)B_1 \leq q/4$. Therefore, we can choose our parameters as $q' \approx 2^{30}$, $l = k = 5$, $B_1 = 1$, and $B_2 = 2^{15.9}$ for better signature sizes. The LWE estimator of Albrecht et al. estimates that this MLWE instance has also 141 bits of security. In either cases, we use an optimization due to Bai and Galbraith to reduce the size of the proofs (and therefore the size of the signature).

Similar to the isogeny instantiation, the signature size depends very mildly on N because N only affects the length of some paths in the signature. Finally, we can use Sect. 6.3 to obtain a tightly secure scheme. Since $K = 16$, the overhead compared to the non-tight scheme is a mere 512B.

Acknowledgements Yi-Fu Lai was supported by the Ministry for Business, Innovation and Employment in New Zealand. Shuichi Katsumata was supported by JST CREST Grant Number JPMJCR19F6, Japan. This work was supported by CyberSecurity Research Flanders with reference number VR20192203, and in part by the Research Council KU Leuven grant C14/18/067 on Cryptanalysis of post-quantum cryptography. Ward Beullens is funded by FWO Junior Postdoc- toral Fellowship 1S95620N.

Funding Open access funding provided by Università degli Studi di Bari Aldo Moro within the CRUI-CARE Agreement.

Declaration

Data sharing Data sharing not applicable to this article as no datasets were generated or analysed during the current study.

Conflict of interest The authors declare they have no financial and non-financial interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix

A Omitted primitives

A.1 Index-hiding Merkle trees

The definition an *index-hiding* Merkle tree is taken almost verbatim from [14]. Merkle trees [79] allow one to hash a list of elements $A = (a_0, \dots, a_N)$ into one hash value (often called the root). At a later point, one can efficiently prove to a third party that an element a_i was included at a certain position in the list A . In the following, we consider a slight modification of the standard Merkle tree construction, such that one can prove that a single element a_i was included in the tree without revealing its position in the list.

Formally, the Merkle tree technique consists of three algorithms (MerkleTree, getMerklePath, ReconstructRoot) with access to a common hash function $\mathcal{H}_{\text{Coll}} : \{0, 1\}^* \rightarrow \{0, 1\}^{2\lambda}$.

- **MerkleTree**(A) \rightarrow (root, tree): On input a list of 2^k elements $A = (a_1, \dots, a_{2^k})$, with $k \in \mathbb{N}$, it constructs a binary tree of height k with $\{l_i = \mathcal{H}_{\text{Coll}}(a_i)\}_{i \in [2^k]}$ as its leaf nodes, and where every internal node h with children h_{left} and h_{right} equals the hash digest of a concatenation of its two children. While it is standard to consider the concatenation $h_{\text{left}} \parallel h_{\text{right}}$, we consider a variation which consists in ordering the two children according to the lexicographical order (or any other total order on binary strings). We denote by $(h_{\text{left}}, h_{\text{right}})_{\text{lex}}$ this modified concatenation. The algorithm then outputs the root root of the Merkle tree, as well as a description of the entire tree tree .
- **getMerklePath**(tree, I) \rightarrow path: On input the description of a Merkle tree tree and an index $i \in [2^k]$, it outputs the list path , which contains the sibling of l_i (i.e. a node, different from l_i , that has the same parent as l_i), as well as the sibling of any ancestor of l_i , ordered by decreasing height.
- **ReconstructRoot**(a, path) \rightarrow root: On input an element a in the list of elements $A = (a_1, \dots, a_{2^k})$ and $\text{path} = (n_1, \dots, n_k)$, it outputs a reconstructed root $\text{root}' = h_k$, which is calculated by putting $h_0 = \mathcal{H}_{\text{Coll}}(a)$ and defining h_i for $i \in [k]$ recursively as $h_i = \mathcal{H}_{\text{Coll}}((h_{i-1}, n_i)_{\text{lex}})$.

If the hash function $\mathcal{H}_{\text{Coll}}$ that is used in the Merkle tree is collision-resistant, then the following easy lemma implies that the Merkle tree construction is *binding*, i.e. that one cannot construct a path that “proves” that a value $b \notin A = (a_1, \dots, a_N)$ is part of the list A that was used to construct the Merkle tree without breaking the collision-resistance of the underlying hash function $\mathcal{H}_{\text{Coll}}$.

Lemma A.1 (Binding for merkle tree) *There is an efficient extractor algorithm that, given the description tree of a Merkle tree (having root root and constructed using the list of elements A) and (b, path) such that $b \notin A$ and $\text{ReconstructRoot}(b, \text{path}) = \text{root}$, outputs a collision for the hash function $\mathcal{H}_{\text{Coll}}$.*

The use of the lexicographical order to concatenate two children nodes in the Merkle tree construction implies that the output path of the **getMerklePath** algorithm information-theoretically hides the index $i \in [N]$ given as input. Formally, we have the following.

Lemma A.2 (*Index Hiding for Merkle Tree*) *Let $N \in \mathbb{N}$ be a power of 2, D, D' be two arbitrary distributions over $\{0, 1\}^*$ and D_I , with $I \in [N]$, be the distribution defined as*

$$D_I = \left[(a_I, \text{path}, \text{root}) \left| \begin{array}{l} a_I \leftarrow D, \\ a_i \leftarrow D' \ \forall 1 \leq i \neq I \leq N, \\ (\text{tree}, \text{root}) \leftarrow \text{MerkleTree}(A), \\ \text{path} \leftarrow \text{getMerklePath}(\text{tree}, I) \end{array} \right. \right]$$

where $A = (a_1, \dots, a_N)$. Then we have $D_I = D_J$ for all $I, J \in [N]$.

A.2 Seed tree

The definition seed tree is taken almost verbatim from [14]. The purpose of a seed tree is to first generate a number of pseudorandom values and later disclose an arbitrary subset of them, without revealing information on the remaining values. The seed tree is a complete binary tree¹² of λ -bit seed values such that the left (resp. right) child of a seed seed_h is the left (resp. right) half of $\text{Expand}(\text{seed} \parallel h)$, where Expand is a pseudorandom generator (PRG). The unique identifier h of the parent seed is appended to separate the input domains of the different calls to the PRG. A sender can efficiently reveal the seed values associated with a subset of the set of leaves by revealing the appropriate set of internal seeds in the tree. We provide the full detail of the seed tree below. Let $\text{Expand} : \{0, 1\}^{\lambda + \lceil \log_2(M-1) \rceil} \rightarrow \{0, 1\}^{2\lambda}$ be a PRG for any $\lambda, M \in \mathbb{N}$, instantiated by a random oracle \mathcal{O} . Then, a seed tree consists of the following four oracle-calling algorithms.

- $\text{SeedTree}^{\mathcal{O}}(\text{seed}_{\text{root}}, M) \rightarrow \{\text{leaf}_i\}_{i \in [M]}$: On input a root seed $\text{seed}_{\text{root}} \in \{0, 1\}^\lambda$ and an integer $M \in \mathbb{N}$, it constructs a complete binary tree with M leaves by recursively expanding each seed to obtain its children seeds. Calls are of the form $\mathcal{O}(\text{Expand} \parallel \text{seed}_h \parallel h)$, where $h \in [M - 1]$ is a unique identifier for the position of seed in the binary tree.
- $\text{ReleaseSeeds}^{\mathcal{O}}(\text{seed}_{\text{root}}, \mathbf{c}) \rightarrow \text{seeds}_{\text{internal}}$: On input a root seed $\text{seed}_{\text{root}} \in \{0, 1\}^\lambda$, and a challenge $\mathbf{c} \in \{0, 1\}^M$, it outputs the list of seeds $\text{seeds}_{\text{internal}}$ that covers all the leaves with index i such that $c_i = 1$. Here, we say that a set of nodes D covers a set of leaves S if the union of the leaves of the subtrees rooted at each node $v \in D$ is exactly the set S .
- $\text{RecoverLeaves}^{\mathcal{O}}(\text{seeds}_{\text{internal}}, \mathbf{c}) \rightarrow \{\text{leaf}_i\}_{i \text{ s.t. } c_i=1}$: On input a set $\text{seeds}_{\text{internal}}$ and a challenge $\mathbf{c} \in \{0, 1\}^M$, it computes and outputs all the leaves of subtrees rooted at seeds in $\text{seeds}_{\text{internal}}$. By construction, this is exactly the set $\{\text{leaf}_i\}_{i \text{ s.t. } c_i=1}$.
- $\text{SimulateSeeds}^{\mathcal{O}}(\mathbf{c}) \rightarrow \text{seeds}_{\text{internal}}$: On input a challenge $\mathbf{c} \in \{0, 1\}^M$, it computes the set of nodes covering the leaves with index i such that $c_i = 1$. It then randomly samples a seed from $\{0, 1\}^\lambda$ for each of these nodes, and finally outputs the set of these seeds as $\text{seeds}_{\text{internal}}$.

By construction, the leaves $\{\text{leaf}_i\}_{i \text{ s.t. } c_i=1}$ output by $\text{SeedTree}(\text{seed}_{\text{root}}, M)$ are the same as those output by $\text{RecoverLeaves}(\text{ReleaseSeeds}(\text{seed}_{\text{root}}, \mathbf{c}), \mathbf{c})$ for any $\mathbf{c} \in \{0, 1\}^M$. The last algorithm SimulateSeeds can be used to argue that the seeds associated with all the leaves with index i such that $c_i = 0$ are indistinguishable from uniformly random values for a recipient that is only given $\text{seeds}_{\text{internal}}$ and \mathbf{c} . More formally, we have the following.

¹² A complete binary tree is a binary tree in which every level, except possibly the last, is completely filled, and all nodes are as far left as possible.

Lemma A.3 Fix any $M \in \mathbb{N}$ and any $\mathbf{c} \in \{0, 1\}^M$. If we model *Expand* as a random oracle \mathcal{O} , then any (computationally unbounded) adversary $A^{\mathcal{O}}$ that makes Q queries to the random oracle \mathcal{O} can distinguish the following two distributions D_1 and D_2 with distinguishing advantage bounded by $\frac{Q}{2^\lambda}$:

$$\begin{aligned}
 D_1 : & \left\{ \text{seeds}_{\text{internal}}, \{\text{leaf}_i\}_{i \text{ s.t. } c_i=0} \middle| \begin{array}{l} \text{seed}_{\text{root}} \leftarrow \{0, 1\}^\lambda \\ \{\text{leaf}_i\}_{i \in [M]} \leftarrow \text{SeedTree}^{\mathcal{O}}(\text{seed}_{\text{root}}, M) \\ \text{seeds}_{\text{internal}} \leftarrow \text{ReleaseSeeds}^{\mathcal{O}}(\text{seed}_{\text{root}}, \mathbf{c}) \end{array} \right\} \\
 D_2 : & \left\{ \text{seeds}_{\text{internal}}, \{\text{leaf}_i\}_{i \text{ s.t. } c_i=0} \middle| \begin{array}{l} \forall i \text{ s.t. } c_i = 0 : \text{leaf}_i \leftarrow \{0, 1\}^\lambda \\ \text{seeds}_{\text{internal}} \leftarrow \text{SimulateSeeds}^{\mathcal{O}}(\mathbf{c}) \end{array} \right\}
 \end{aligned}$$

Here, the distributions take into account the randomness used by the random oracle as well.

B Dynamic group signatures from accountable ring signatures

In this section, we review briefly the definition of group signatures and explain how accountable ring signatures can be naturally viewed as group signatures. A formal treatment can be found in Bootle et al. [20]

B.1 Preliminaries on group signatures

Group signatures can be divided into two primary types: static schemes [8] and dynamic schemes [9]. Roughly, while static group signature require the group to be fixed at setup, dynamic group signatures allow members to join and leave the group at any time. This joining and leaving is administered by the group manager, who has the power to add and revoke membership — as well as the ability to revoke anonymity and reveal the specific signer of a certain signature. For a dynamic group signature scheme, the revocation mechanism can be handled by a separate entity called opening or tracing authority to offer better flexibility in the scheme and this makes only little difference regarding the security notions.

Informally, a dynamic group signature scheme consists of a setup algorithm *Setup*, key generation algorithms *MKGen* and *UKGen* for the group manager and group members (or users) respectively, and *Sign*, *Verify*, *Open*, and *Judge* algorithms which are counterparts of the ARS scheme functions of the same names. Additionally, an interactive *Join* protocol run between the group manager and a user allows users to be added to the group, while an *UpdateGroup* function allows the group manager to revoke a user’s membership in the group dynamically (this is done via some publicly-published *group info* info).

Dynamic group signature schemes should satisfy standard security properties of correctness, anonymity, traceability and non-frameability [9, 20]. Correctness ensures that a signature produced by a user running *Sign* after joining the group via *Join* is accepted by *Verify*. The inclusion of the *Join* function in this definition ensures joining works as intended, beyond just guaranteeing the signing algorithms’s correctness. Full CCA-anonymity (often refereed simply as *full* anonymity) states that even under full key exposure of all group members (other than the group manager, who can trivially revoke anonymity via *Open*), and with access to an opening oracle, the user who generated a certain signature cannot be identified. More specifically, an adversary should be unable to distinguish between signatures generated by any two members of the adversary’s choice— even if the adversary knows all secret keys involved. This notion is almost identical to its namesake in the ARS setting (Sect. 3.4). In

contrast, CPA-anonymity is a weaker notion which still allows the adversary to learn all group members' keys, but removes access to the opening oracle. Weaker variants of these two are *selfless* CCA-anonymity and *selfless* CPA-anonymity where the adversary cannot obtain any secret keys of targeted members in the anonymity game. Traceability states that an adversary who is able to corrupt any members is not able to produce a signature for which Open fails to return an active member of the group even if the group manager's secret key is leaked. Finally, non-frameability states that even if the group manager and all but one of the group members are corrupted, they cannot forge or falsely attribute a signature to an honest member who did not produce it. These properties also imply what is usually called *unforgeability*, because if an adversary could produce a signature for a group they knew no secret keys for, the signature must either fail to Open to an active user, or would frame an honest member of the group—violating either traceability or non-frameability. We also remark a difference, usually being neglected, that the group manager can be corrupted in the security model of a dynamic group while a static variant only takes into account the exposure of the opening secret key [8]. We refer the reader to [20] for more thorough definitions.

B.2 Constructing group signatures from ARS

For completeness, we now review the generic construction of a dynamic group signature scheme from an accountable ring signature scheme, by Bootle et al. [19, 20]. Let Π_{ARS} be a secure ARS scheme, then we define a group signature scheme Π_{GS} as follows:

Let the group manager be the opening authority of Π_{ARS} , and let the group manager's keypair be $(\text{gmpk} = \text{opk}, \text{gmsk} = \text{osk})$. The group public key gpk is then set to (gmpk, pp) , where pp is the output of $\text{GS.Setup} := \text{ARS.Setup}$. Define $\text{GS.UKGen} := \text{ARS.UKGen}$, so that users generate their own keypairs directly. The Join protocol proceeds by a user submitting their public key pk to the group manager, who appends it to the list of keys in $\text{info}_\tau := [\text{vk}_0, \dots, \text{vk}_i]$ (the group info at epoch τ) and publishes $\text{info}_{\tau+1}$. Membership is similarly revoked by the group manager via UpdateGroup by removing the user's public key from info_τ and publishing the updated info. Finally, define:

- $\text{GS.Sign}(\text{gpk}, \text{info}_\tau, \text{sk}_i, M) := \text{ARS.Sign}(\text{gmpk}, \text{sk}_i, \text{info}_\tau, M)$.
- $\text{GS.Verify}(\text{gpk}, \text{info}_\tau, M, \sigma) := \text{ARS.Verify}(\text{gmpk}, \text{info}_\tau, M, \sigma)$.
- $\text{GS.Open}(\text{gpk}, \text{info}_{\text{au}}, \text{gmsk}, M, \sigma)$ calls $(\text{vk}_j, \pi) \leftarrow \text{ARS.Open}(\text{gmsk}, \text{info}_\tau, M, \sigma)$ and returns (j, π) .
- $\text{GS.Judge}(\text{gpk}, \text{info}_\tau, M, \sigma, (j, \pi)) := \text{ARS.Judge}(\text{gmpk}, \text{info}_\tau, \text{vk}_j, M, \sigma, \pi)$.

Note that info_τ defines the ring of signers at epoch τ and should be publicly accessible, as too should be the index-to-public-key ($j \leftrightarrow \text{vk}_j$) correspondence table, maintained by the group manager. As shown in [20], this generic construction of a group signature from an ARS is *tightly* secure assuming the ARS is secure. Hence, our ARS construction in Sect. 4.1 implies a secure dynamic group signature scheme. The type of security notions satisfied by the resulting group signature, e.g., full or selfless, CCA or CPA anonymity, is directly inherited from the ARS.

We note that this scheme's group info grows linearly in the number of group members. This is the same as all other proposed efficient post-quantum group signature constructions such as [48]. It remains an interesting open problem to construct a efficient group signature where the group info grows at most logarithmically in the number of group members.

References

1. Abe M., Ohkubo M., Suzuki K.: 1-out-of-n signatures from a variety of keys. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 415–432. Springer (2002). https://doi.org/10.1007/3-540-36178-2_26
2. Albrecht M.R., Curtis B.R., Deo A., Davidson A., Player, R., Postlethwaite E.W., Virdia F., Wunderer T.: Estimate all the LWE, NTRU schemes! In: Catalano D., De Prisco R. (eds.) SCN 18. LNCS, vol. 11035, pp. 351–367. Springer (2018). https://doi.org/10.1007/978-3-319-98113-0_19
3. Albrecht M.R., Player R., Scott S.: On The Concrete Hardness Of Learning With Errors. Cryptology ePrint Archive, Report 2015/046. <https://eprint.iacr.org/2015/046> (2015)
4. Asharov G., Jain A., López-Alt A., Tromer E., Vaikuntanathan V., Wichs D.: Multiparty computation with low communication, computation and interaction via threshold FHE. In: Pointcheval D., Johansson T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 483–501. Springer (2012). https://doi.org/10.1007/978-3-642-29011-4_29
5. Attema T., Lyubashevsky V., Seiler G.: Practical product proofs for lattice commitments. In: Micciancio D., Ristenpart T. (eds.) CRYPTO 2020, Part II. LNCS, vol. 12171, pp. 470–499. Springer (2020). https://doi.org/10.1007/978-3-030-56880-1_17
6. Backes M., Hanzlik L., Schneider-Bensch J.: Membership privacy for fully dynamic group signatures. In: Cavallaro L., Kinder J., Wang X., Katz, J. (eds.) ACM CCS 2019, pp. 2181–2198. ACM Press (2019). <https://doi.org/10.1145/3319535.3354257>
7. Bai S., Galbraith S.D.: An improved compression technique for signatures based on learning with errors. In: Benaloh J. (ed.) CT-RSA 2014. LNCS, vol. 8366, pp. 28–47. Springer (2014). https://doi.org/10.1007/978-3-319-04852-9_2
8. Bellare M., Micciancio D., Warinschi, B.: Foundations of group signatures: Formal definitions, simplified requirements, and a construction based on general assumptions. In: Biham E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 614–629. Springer (2003). https://doi.org/10.1007/3-540-39200-9_38
9. Bellare M., Shi H., Zhang C.: Foundations of group signatures: The case of dynamic groups. In: Menezes A. (ed.) CT-RSA 2005. LNCS, vol. 3376, pp. 136–153. Springer (2005). https://doi.org/10.1007/978-3-540-30574-3_11
10. Benhamouda F., Camenisch J., Krenn S., Lyubashevsky V., Neven G.: Better zero-knowledge proofs for lattice encryption and their application to group signatures. In: Sarkar P., Iwata T. (eds.) ASIACRYPT 2014, Part I. LNCS, vol. 8873, pp. 551–572. Springer (2014). https://doi.org/10.1007/978-3-662-45611-8_29
11. Bernhard D., Fischlin M., Warinschi B.: Adaptive proofs of knowledge in the random oracle model. In: Katz J. (ed.) PKC 2015. LNCS, vol. 9020, pp. 629–649. Springer (2015). https://doi.org/10.1007/978-3-662-46447-2_28
12. Bernstein D.J., Lange T., Martindale C., Panny L.: Quantum circuits for the CSIDH: Optimizing quantum evaluation of isogenies. In: Ishai Y., Rijmen V. (eds.) EUROCRYPT 2019, Part II. LNCS, vol. 11477, pp. 409–441. Springer (2019). https://doi.org/10.1007/978-3-030-17656-3_15
13. Beullens W., Dobson S., Katsumata S., Lai Y.-F., Pintore F.: Group signatures and more from isogenies and lattices: Generic, simple, and efficient. In: EUROCRYPT 2022, Part II. LNCS, pp. 95–126. Springer, (2022). https://doi.org/10.1007/978-3-031-07085-3_4
14. Beullens W., Katsumata S., Pintore F.: Calamari and Falafel: Logarithmic (linkable) ring signatures from isogenies and lattices. In: Moriai S., Wang H. (eds.) ASIACRYPT 2020, Part II. LNCS, vol. 12492, pp. 464–492. Springer (2020). https://doi.org/10.1007/978-3-030-64834-3_16
15. Beullens W., Kleinjung T., Vercauteren F.: CSI-FiSh: Efficient isogeny based signatures through class group computations. In: Galbraith S.D., Moriai S. (eds.) ASIACRYPT 2019, Part I. LNCS, vol. 11921, pp. 227–247. Springer (2019). https://doi.org/10.1007/978-3-030-34578-5_9
16. Bichsel P., Camenisch J., Neven G., Smart, N.P., Warinschi, B.: Get shorty via group signatures without encryption. In: Garay J.A., Prisco R.D. (eds.) SCN 10. LNCS, vol. 6280, pp. 381–398. Springer (2010). https://doi.org/10.1007/978-3-642-15317-4_24
17. Boneh D., Boyen X., Shacham H.: Short group signatures. In: Franklin M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 41–55. Springer (2004). https://doi.org/10.1007/978-3-540-28628-8_3
18. Bonnetain X., Schrottenloher A.: Quantum security analysis of CSIDH. In: Canteaut A., Ishai Y. (eds.) EUROCRYPT 2020, Part II. LNCS, vol. 12106, pp. 493–522. Springer (2020). https://doi.org/10.1007/978-3-030-45724-2_17
19. Bootle J., Cerulli A., Chaidos P., Ghadafi E., Groth J., Petit C.: Short accountable ring signatures based on DDH. In: Pernul G., Ryan P.Y.A., Weippl E.R. (eds.) ESORICS 2015, Part I. LNCS, vol. 9326, pp. 243–265. Springer (2015). https://doi.org/10.1007/978-3-319-24174-6_13

20. Bootle J., Cerulli A., Chaidos P., Ghadafi E., Groth J.: Foundations of fully dynamic group signatures. In: Manulis M., Sadeghi A.-R., Schneider S. (eds.) ACNS 16. LNCS, vol. 9696, pp. 117–136. Springer (2016). https://doi.org/10.1007/978-3-319-39555-5_7
21. Bootle J., Lyubashevsky V., Seiler G.: Algebraic techniques for short(er) exact lattice-based zero-knowledge proofs. In: Boldyreva A., Micciancio D. (eds.) CRYPTO 2019, Part I. LNCS, vol. 11692, pp. 176–202. Springer (2019). https://doi.org/10.1007/978-3-030-26948-7_7
22. Boschini C., Camenisch J., Neven G.: Floppy-sized group signatures from lattices. In: Preneel B., Vercauteren F. (eds.) ACNS 18. LNCS, vol. 10892, pp. 163–182. Springer (2018). https://doi.org/10.1007/978-3-319-93387-0_9
23. Brickell E.F., Camenisch J., Chen L.: Direct anonymous attestation. In: Atluri V., Pfizmann B., McDaniel P. (eds.) ACM CCS 2004, pp. 132–145. ACM Press, (2004). <https://doi.org/10.1145/1030083.1030103>
24. Brickell E., Li J.: Enhanced privacy id: A direct anonymous attestation scheme with enhanced revocation capabilities. In: Proceedings of the 2007 ACM Workshop on Privacy in Electronic Society, pp. 21–30 (2007)
25. Camenisch J., Damgård I.: Verifiable encryption, group encryption, and their applications to separable group signatures and signature sharing schemes. In: Okamoto T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 331–345. Springer (2000). https://doi.org/10.1007/3-540-44448-3_25
26. Camenisch J., Shoup, V.: Practical verifiable encryption and decryption of discrete logarithms. In: Boneh D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 126–144. Springer (2003). https://doi.org/10.1007/978-3-540-45146-4_8
27. Camenisch J.: Efficient and generalized group signatures. In: Fumy W. (ed.) EUROCRYPT'97. LNCS, vol. 1233, pp. 465–479. Springer (1997). https://doi.org/10.1007/3-540-69053-0_32
28. Castryck W., Lange T., Martindale C., Panny L., Renes J.: CSIDH: An efficient post-quantum commutative group action. In: Peyrin T., Galbraith S. (eds.) ASIACRYPT 2018, Part III. LNCS, vol. 11274, pp. 395–427. Springer (2018). https://doi.org/10.1007/978-3-030-03332-3_15
29. Castryck W., Sotáková J., Vercauteren, F.: Breaking the decisional Diffie-Hellman problem for class group actions using genus theory. In: Micciancio D., Ristenpart T. (eds.) CRYPTO 2020, Part II. LNCS, vol. 12171, pp. 92–120. Springer (2020). https://doi.org/10.1007/978-3-030-56880-1_4
30. Chase M., Lysyanskaya A.: On signatures of knowledge. In: Dwork C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 78–96. Springer (2006). https://doi.org/10.1007/11818175_5
31. Chaum D., van Heyst E.: Group signatures. In: Davies D.W. (ed.) EUROCRYPT'91. LNCS, vol. 547, pp. 257–265. Springer, (1991). https://doi.org/10.1007/3-540-46416-6_22
32. Chávez-Saab J., Chi-Domínguez J.-J., Jaques S., Rodríguez-Henríquez F.: The SQALE of CSIDH: square-root vélu quantum-resistant isogeny action with low exponents. IACR Cryptol. ePrint Arch. **2020**, 1520 (2020).
33. Chung K.-M., Hsieh Y.-C., Huang M.-Y., Huang Y.-H., Lange T., Yang B.-Y.: Group signatures and accountable ring signatures from isogeny-based assumptions. Cryptology ePrint Archive, Report 2021/1368. <https://ia.cr/2021/1368> (2021)
34. Clarisse R., Sanders O.: Group signature without random oracles from randomizable signatures. In: Nguyen K., Wu W., Lam K.-Y., Wang H. (eds.) ProvSec 2020. LNCS, vol. 12505, pp. 3–23. Springer, (2020). https://doi.org/10.1007/978-3-030-62576-4_1
35. Cramer R., Shoup V.: A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In: Krawczyk H. (ed.) CRYPTO'98. LNCS, vol. 1462, pp. 13–25. Springer (1998). <https://doi.org/10.1007/BFb0055717>
36. Damgård I., Fujisaki E.: A statistically-hiding integer commitment scheme based on groups with hidden order. In: Zheng Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 125–142. Springer (2002). https://doi.org/10.1007/3-540-36178-2_8
37. De Feo L., Galbraith S.D.: SeaSign: Compact isogeny signatures from class group actions. In: Ishai Y., Rijmen V. (eds.) EUROCRYPT 2019, Part III. LNCS, vol. 11478, pp. 759–789. Springer (2019). https://doi.org/10.1007/978-3-030-17659-4_26
38. del Pino R., Lyubashevsky V., Seiler G.: Lattice-based group signatures and zero-knowledge proofs of automorphism stability. In: Lie D., Mannan M., Backes M., Wang X. (eds.) ACM CCS 2018, pp. 574–591. ACM Press (2018). <https://doi.org/10.1145/3243734.3243852>
39. Delerablée C., Pointcheval, D.: Dynamic fully anonymous short group signatures. In: Nguyen P.Q. (ed.) Progress in Cryptology - VIETCRYPT 06. LNCS, vol. 4341, pp. 193–210. Springer (2006)
40. Derler D., Slamanig D.: Highly-efficient fully-anonymous dynamic group signatures. In: Kim J., Ahn G.-J., Kim S., Kim Y., López J., Kim T. (eds.) ASIACCS 18, pp. 551–565. ACM Press, (2018)
41. Dinur I., Nadler N.: Multi-target attacks on the Picnic signature scheme and related protocols. In: Ishai Y., Rijmen V. (eds.) EUROCRYPT 2019, Part III. LNCS, vol. 11478, pp. 699–727. Springer (2019). https://doi.org/10.1007/978-3-030-17659-4_24

42. Ducas L., Kiltz E., Lepoint T., Lyubashevsky V., Schwabe P., Seiler G., Stehlé D.: CRYSTALS-Dilithium: a lattice-based digital signature scheme. *IACR TCHES* **1**, 238–268 (2018). <https://doi.org/10.13154/tches.v2018.i1.238-268>.
43. El Kaafarani A., Katsumata S., Pintore F.: Lossy CSI-FiSh: Efficient signature scheme with tight reduction to decisional CSIDH-512. In: Kiayias A., Kohlweiss M., Wallden P., Zikas V. (eds.) PKC 2020, Part II. LNCS, vol. 12111, pp. 157–186. Springer (2020). https://doi.org/10.1007/978-3-030-45388-6_6
44. El Kaafarani A., Katsumata S.: Attribute-based signatures for unbounded circuits in the ROM and efficient instantiations from lattices. In: Abdalla M., Dahab R. (eds.) PKC 2018, Part II. LNCS, vol. 10770, pp. 89–119. Springer (2018). https://doi.org/10.1007/978-3-319-76581-5_4
45. Esgin M.F., Nguyen N.K., Seiler G.: Practical exact proofs from lattices: New techniques to exploit fully-splitting rings. In: Moriai S., Wang H. (eds.) ASIACRYPT 2020, Part II. LNCS, vol. 12492, pp. 259–288. Springer (2020). https://doi.org/10.1007/978-3-030-64834-3_9
46. Esgin M.F., Steinfeld R., Liu J.K., Liu D.: Lattice-based zero-knowledge proofs: New techniques for shorter and faster constructions and applications. In: Boldyreva A., Micciancio D. (eds.) CRYPTO 2019, Part I. LNCS, vol. 11692, pp. 115–146. Springer (2019). https://doi.org/10.1007/978-3-030-26948-7_5
47. Esgin M.F., Steinfeld R., Zhao R.K.: MatRiCT+: More Efficient Post-Quantum Private Blockchain Payments. Cryptology ePrint Archive, Report 2021/545 (To appear in IEEE S&P 2022)
48. Esgin M.F., Zhao R.K., Steinfeld R., Liu J.K., Liu D.: MatRiCT: Efficient, scalable and post-quantum blockchain confidential transactions protocol. In: Cavallaro L., Kinder J., Wang X., Katz J. (eds.) ACM CCS 2019, pp. 567–584. ACM Press (2019). <https://doi.org/10.1145/3319535.3354200>
49. Ezerman M.F., Lee H.T., Ling S., Nguyen K., Wang, H.: A provably secure group signature scheme from code-based assumptions. In: Iwata T., Cheon J.H. (eds.) ASIACRYPT 2015, Part I. LNCS, vol. 9452, pp. 260–285. Springer (2015). https://doi.org/10.1007/978-3-662-48797-6_12
50. Fiat A., Shamir A.: How to prove yourself: Practical solutions to identification and signature problems. In: Odlyzko A.M. (ed.) CRYPTO'86. LNCS, vol. 263, pp. 186–194. Springer (1987). https://doi.org/10.1007/3-540-47721-7_12
51. Fischlin M., Harasser P., Janson C.: Signatures from sequential-OR proofs. In: Canteaut A., Ishai Y. (eds.) EUROCRYPT 2020, Part III. LNCS, vol. 12107, pp. 212–244. Springer (2020). https://doi.org/10.1007/978-3-030-45727-3_8
52. Fischlin M.: Communication-efficient non-interactive proofs of knowledge with online extractors. In: Shoup V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 152–168. Springer (2005). https://doi.org/10.1007/11535218_10
53. Fujisaki E., Okamoto T.: Secure integration of asymmetric and symmetric encryption schemes. In: Wiener M.J. (ed.) CRYPTO'99. LNCS, vol. 1666, pp. 537–554. Springer (1999). https://doi.org/10.1007/3-540-48405-1_34
54. Fujisaki E., Okamoto T.: Statistical zero knowledge protocols to prove modular polynomial relations. In: Kaliski Jr., B.S. (ed.) CRYPTO'97. LNCS, vol. 1294, pp. 16–30. Springer (1997). <https://doi.org/10.1007/BFb0052225>
55. Furukawa J., Imai H.: An efficient group signature scheme from bilinear maps. *IEICE Trans. Fund. Electron. Commun. Comput. Sci.* **89**(5), 1328–1338 (2006).
56. Gentry C., Peikert C., Vaikuntanathan V.: Trapdoors for hard lattices and new cryptographic constructions. In: Ladner R.E., Dwork C. (eds.) 40th ACM STOC, pp. 197–206. ACM Press (2008). <https://doi.org/10.1145/1374376.1374407>
57. Gordon S.D., Katz J., Vaikuntanathan V.: A group signature scheme from lattice assumptions. In: Abe M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 395–412. Springer (2010). https://doi.org/10.1007/978-3-642-17373-8_23
58. Groth J.: Fully anonymous group signatures without random oracles. In: Kurosawa K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 164–180. Springer (2007). https://doi.org/10.1007/978-3-540-76900-2_10
59. Katsumata S., Yamada S.: Group signatures without NIZK: From lattices in the standard model. In: Ishai Y., Rijmen V. (eds.) EUROCRYPT 2019, Part III. LNCS, vol. 11478, pp. 312–344. Springer (2019). https://doi.org/10.1007/978-3-030-17659-4_11
60. Katz J., Kolesnikov V., Wang X.: Improved non-interactive zero knowledge with applications to post-quantum signatures. In: Lie, D., Mannan M., Backes M., Wang X. (eds.) ACM CCS 2018, pp. 525–537. ACM Press (2018). <https://doi.org/10.1145/3243734.3243805>
61. Katz J., Wang N.: Efficiency improvements for signature schemes with tight security reductions. In: Jajodia S., Atluri V., Jaeger T. (eds.) ACM CCS 2003, pp. 155–164. ACM Press (2003). <https://doi.org/10.1145/948109.948132>
62. Kumawat S., Paul S.: A new constant-size accountable ring signature scheme without random oracles. In: International Conference on Information Security and Cryptology, pp. 157–179 (2017). Springer

63. Kuperberg G.: Another subexponential-time quantum algorithm for the dihedral hidden subgroup problem. In: 8th Conference on the Theory of Quantum Computation, vol. 22, pp. 20–34. Communication and Cryptography (TQC 2013). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl (2013).
64. Kuperberg G.: A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM J. Comput.* **35**(1), 170–188 (2005).
65. Laguillaumie F., Langlois A., Libert B., Stehlé D.: Lattice-based group signatures with logarithmic signature size. In: Sako K., Sarkar P. (eds.) ASIACRYPT 2013, Part II. LNCS, vol. 8270, pp. 41–61. Springer (2013). https://doi.org/10.1007/978-3-642-42045-0_3
66. Lai Y.-F., Dobson S.: Collusion Resistant Revocable Ring Signatures and Group Signatures from Hard Homogeneous Spaces. Cryptology ePrint Archive, Report 2021/1365. <https://ia.cr/2021/1365> (2021)
67. Lai Y.-F., Galbraith S.D., de Saint Guilhem C.: Compact, efficient and UC-secure isogeny-based oblivious transfer. In: Canteaut A., Standaert F.-X. (eds.) EUROCRYPT 2021, Part I. LNCS, vol. 12696, pp. 213–241. Springer (2021). https://doi.org/10.1007/978-3-030-77870-5_8
68. Lai R.W.F., Zhang T., Chow S.S.M., Schröder D.: Efficient sanitizable signatures without random oracles. In: Askoxylakis I.G., Ioannidis S., Katsikas S.K., Meadows C.A. (eds.) ESORICS 2016, Part I. LNCS, vol. 9878, pp. 363–380. Springer (2016). https://doi.org/10.1007/978-3-319-45744-4_18
69. Langlois A., Stehlé D.: Worst-case to average-case reductions for module lattices. *Des. Codes Cryptogr.* **75**(3), 565–599 (2015). <https://doi.org/10.1007/s10623-014-9938-4>.
70. Libert B., Ling S., Nguyen K., Wang H.: Zero-knowledge arguments for lattice-based accumulators: Logarithmic-size ring signatures and group signatures without trapdoors. In: Fischlin M., Coron J.-S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 1–31. Springer (2016). https://doi.org/10.1007/978-3-662-49896-5_1
71. Libert B., Mouhartem F., Peters T., Yung M.: Practical “signatures with efficient protocols” from simple assumptions. In: Chen X., Wang X., Huang X. (eds.) ASIACCS 16, pp. 511–522. ACM Press (2016)
72. Libert B., Peters T., Yung M.: Short group signatures via structure-preserving signatures: Standard model security from simple assumptions. In: Gennaro R., Robshaw M.J.B. (eds.) CRYPTO 2015, Part II. LNCS, vol. 9216, pp. 296–316. Springer (2015). https://doi.org/10.1007/978-3-662-48000-7_15
73. Lindner R., Peikert C.: Better key sizes (and attacks) for LWE-based encryption. In: Kiayias A. (ed.) CT-RSA 2011. LNCS, vol. 6558, pp. 319–339. Springer (2011). https://doi.org/10.1007/978-3-642-19074-2_21
74. Ling S., Nguyen K., Wang H., Xu Y.: Constant-size group signatures from lattices. In: Abdalla M., Dahab R. (eds.) PKC 2018, Part II. LNCS, vol. 10770, pp. 58–88. Springer (2018). https://doi.org/10.1007/978-3-319-76581-5_3
75. Lyubashevsky V., Nguyen N.K., Seiler G.: Practical lattice-based zero-knowledge proofs for integer relations. In: Ligatti J., Ou X., Katz J., Vigna G. (eds.) ACM CCS 2020, pp. 1051–1070. ACM Press (2020). <https://doi.org/10.1145/3372297.3417894>
76. Lyubashevsky V., Nguyen N.K., Seiler G.: SMILE: Set membership from ideal lattices with applications to ring signatures and confidential transactions. In: Malkin T., Peikert C. (eds.) CRYPTO 2021, Part II. LNCS, vol. 12826, pp. 611–640. Springer, Virtual Event (2021). https://doi.org/10.1007/978-3-030-84245-1_21
77. Lyubashevsky V., Peikert C., Regev O.: A toolkit for ring-LWE cryptography. In: Johansson T., Nguyen P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 35–54. Springer (2013). https://doi.org/10.1007/978-3-642-38348-9_3
78. Lyubashevsky V.: Lattice signatures without trapdoors. In: Pointcheval D., Johansson T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 738–755. Springer (2012). https://doi.org/10.1007/978-3-642-29011-4_43
79. Merkle R.C.: A digital signature based on a conventional encryption function. In: Pomerance C. (ed.) CRYPTO’87. LNCS, vol. 293, pp. 369–378. Springer (1988). https://doi.org/10.1007/3-540-48184-2_32
80. Peikert C.: He gives C-sieves on the CSIDH. In: Canteaut A., Ishai Y. (eds.) EUROCRYPT 2020, Part II. LNCS, vol. 12106, pp. 463–492. Springer (2020). https://doi.org/10.1007/978-3-030-45724-2_16
81. Pointcheval D., Stern J.: Security arguments for digital signatures and blind signatures. *J. Cryptol.* **13**(3), 361–396 (2000). <https://doi.org/10.1007/s001450010003>.
82. Rivest R.L., Shamir A., Tauman Y.: How to leak a secret. In: Boyd C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 552–565. Springer (2001). https://doi.org/10.1007/3-540-45682-1_32
83. Unruh D.: Non-interactive zero-knowledge proofs in the quantum random oracle model. In: Oswald E., Fischlin M. (eds.) EUROCRYPT 2015, Part II. LNCS, vol. 9057, pp. 755–784. Springer (2015). https://doi.org/10.1007/978-3-662-46803-6_25
84. Xu S., Yung M.: Accountable ring signatures: A smart card approach. In: Smart Card Research and Advanced Applications VI, pp. 271–286. Springer (2004)

85. Yang R., Au M.H., Zhang Z., Xu Q., Yu Z., Whyte W.: Efficient lattice-based zero-knowledge arguments with standard soundness: Construction and applications. In: Boldyreva A., Micciancio D. (eds.) CRYPTO 2019, Part I. LNCS, vol. 11692, pp. 147–175. Springer (2019). https://doi.org/10.1007/978-3-030-26948-7_6

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.