

Fully-Anonymous Short Dynamic Group Signatures Without Encryption

David Derler and Daniel Slamanig

IAIK, Graz University of Technology, Austria
{[david.derler](mailto:david.derler@tugraz.at)|[daniel.slamanig](mailto:daniel.slamanig@tugraz.at)}@tugraz.at

Abstract. Group signatures are a central tool in privacy-enhancing crypto, which allow members of a group to anonymously sign on behalf of the group. Ideally, group signatures are dynamic and thus allow to dynamically and concurrently enroll new members to a group. For such schemes Bellare et al. (CT-RSA'05) proposed a strong security model (BSZ model) that preserves anonymity of a group signature even if an adversary can see arbitrary key exposures or arbitrary openings of other group signatures. All previous constructions achieving this strong anonymity notion follow the so called sign-encrypt-prove (SEP) paradigm. In contrast, all known constructions which avoid this paradigm and follow the alternative “without encryption” paradigm introduced by Bichsel et al. (SCN'10), only provide a weaker notion of anonymity (which can be problematic in practice). Until now it was not clear if constructions following this paradigm, while providing strong anonymity in the sense of BSZ even exist.

We answer this question to the affirmative by proposing a novel approach to dynamic group signature schemes following this paradigm, which is a composition of structure preserving signatures on equivalence classes (ASIACRYPT'14) and other standard primitives. Our results are interesting for various reasons: We can prove our construction following this “without encryption” paradigm secure without requiring random oracles. Moreover, when opting for an instantiation in the ROM, the so obtained scheme is extremely efficient and outperforms the fastest constructions providing anonymity in the BSZ model known to date. Regarding constructions providing a weaker anonymity notion than BSZ, we surprisingly outperform the popular short BBS group signature scheme (CRYPTO'04) and thereby even obtain shorter signatures.

1 Introduction

Group signatures, initially introduced by Chaum and van Heyst [CvH91], allow a group manager to set up a group so that every member of this group can later anonymously sign messages on behalf of the group. Thereby, a dedicated authority (called opening authority) can open a given group signature to determine the identity of the actual signer. Group signatures were first rigorously formalized

The authors have been supported by EU H2020 project PRISMACLOUD, grant agreement n°644962.

for static groups by Bellare et al. in [BMW03]. In this setting, all members are fixed at setup and also receive their honestly generated keys at setup from the group manager. This model was later extended to the dynamic case by Bellare et al. in [BSZ05] (henceforth denoted by BSZ model), where new group members can be dynamically and concurrently enrolled to the group. Further, it separates the role of the issuer and the opener so that they can operate independently. Moreover, the BSZ model requires a strong anonymity notion, where anonymity of a group signature is preserved even if the adversary can see arbitrary key exposures and arbitrary openings of other group signatures. A slightly weaker model, which is used to prove the security (and in particular anonymity) of the popular BBS group signature scheme was introduced by Boneh et al. [BBS04]. This model is a relaxation of the BSZ model, and in particular weakens anonymity so that the adversary can not request openings for signatures. As it is common, we refer to this anonymity notion as CPA-full anonymity, whereas we use CCA2-full anonymity to refer to anonymity in the sense of BSZ.

Over the years, two main construction paradigms for group signatures have been established. The first one is the widely used sign-encrypt-prove (SEP) paradigm [CS97]. Here, a signature is essentially an encrypted membership certificate together with a signature of knowledge, where the signer demonstrates knowledge of some signed value in the ciphertext [ACJT00, BBS04, NS04, BSZ05, KY05, DP06, BW07, BW06, Gro07, LPY15, LLM⁺16, LMPY16]. As an alternative to this paradigm, Bichsel et al. in [BCN⁺10] proposed an elegant design paradigm for group signatures which does not require to encrypt the membership certificate to produce signatures. Henceforth we call this paradigm sign-randomize-proof (SRP). Essentially, they use a signature scheme which supports (1) randomization of signatures so that multiple randomized versions of the same signature are unlinkable, and (2) efficiently proving knowledge of a signed value. In their construction, on joining the group, the issuer uses such a signature scheme to sign a commitment to the user’s secret key. The user can then produce a group signature for a message by randomizing the signature and computing a signature of knowledge on the message, which demonstrates knowledge of the signed secret key. To open signatures, in contrast to constructions following SEP which support constant time opening by means of decrypting the ciphertext in the signature, constructions in this paradigm require a linear scan, i.e., to check a given signature against each potential user. Bichsel et al. proposed an instantiation based on the randomizable pairing-based Camensich-Lysyanskaya (CL) signature scheme [CL04] (whose EUF-CMA security is based on the interactive LRSW assumption). Recently, Pointcheval and Sanders [PS16] proposed another randomizable signature scheme (whose EUF-CMA security is proven in the generic group model), which allows to instantiate the approach due to Bichsel et al. more efficiently. We note that while these two existing constructions do not explicitly use public key encryption, the required assumptions for the scheme imply public key encryption. Yet, it seems to be beneficial regarding performance to avoid to explicitly use public key encryption.

The main drawback of existing constructions following the SRP paradigm is that they rely on a security model that is weaker than the BSZ model [BSZ05]. In particular, anonymity only holds for users whose keys do not leak. This essentially means that once a user key leaks, all previous signatures of this user can potentially be attributed to this user. Furthermore, the model used for SRP constructions assumes that the opening authority and the issuing authority are one entity, meaning that the issuer can identify all signers when seeing group signatures. Both aforementioned weakenings can be highly problematic in practical applications of group signatures. It is thus a natural question to ask whether it is possible to prove that constructions following the SRP paradigm provide CPA- or even CCA2-full anonymity. Unfortunately, for existing constructions, we have to answer this negatively. Even when allowing to modify the existing constructions in [BCN⁺10, PS16] to allow the explicit use of encryption upon joining the group (which might solve the separability issue regarding issuer and opener), it is easy to see that knowledge of the user secret key breaks CCA2- as well as CPA-full anonymity for both constructions.¹ Since CCA2-full anonymity straight forwardly implies anonymity in the SRP model, this example confirms that CCA2-full anonymity is a strictly stronger notion. The notion of CPA-full anonymity is somewhat orthogonal to the anonymity notion used by the SRP model: it appropriately models the leakage of user secret keys, but restricts the open oracle access. Yet, in practice it seems that the risk that a user secret key leaks is extremely hard to quantify, which is why we deem CPA-full anonymity to be more desirable. This is also underpinned by the fact that—to the best of our knowledge—no attacks arising from the restriction of the open oracle access in CPA-full anonymity are known.

Motivation. Group signatures have received significant attention from the cryptographic community and also gain increasing practical relevance due to technological innovations in intelligent transportation systems (e.g., floating car data, toll systems) as well as public transportation systems (i.e., smart ticketing), where user privacy is considered to play an important role (cf. EU Directive 2010/40/EU). These developments make it important to have particularly efficient group signature candidates at hand. As an illustrative example for the importance of very fast signature generation and verification times, consider public transportation system where every user needs to sign on passing a gate.

Despite their increasing practical importance, no progress has been made with respect to computational efficiency improvements of group signature schemes providing the more desirable notions of CPA- as well as CCA2-full anonymity within the last decade. The most efficient schemes known to date are the BBS group signature scheme [BBS04] (which achieves CPA-full anonymity) and the XSGS group signature scheme [DP06] (which achieves CCA2-full anonymity).

¹ Each valid group signature contains a valid randomizable signature on the secret key of the user. While group signatures only contain a proof of knowledge of the signed secret key, being in possession of secret key candidates allows to simply test them using the verification algorithm of the randomizable signature scheme. This clearly provides a distinguisher against CCA2- as well as CPA-full anonymity.

We tackle the following open questions, which are of both theoretical and practical interest:

- *Is it possible to construct schemes providing the more desirable CPA-full and CCA2-full anonymity notions, where compelling efficiency is reached by (1) avoiding the explicit encryption of the membership certificate upon signing, yet (2) allowing to explicitly use encryption during the joining of a group?*
- *Is it possible to further push the computational efficiency limits of group signature schemes providing those more desirable anonymity notions?*

We, henceforth, refer to such schemes as “without encryption”.

Contribution. We answer both questions posed above to the affirmative by contributing a novel approach to construct group signatures “without encryption”. Our approach is a composition of structure preserving signatures on equivalence classes (SPS-EQ) [HS14, FHS14], conventional digital signatures, public key encryption, non-interactive zero-knowledge proofs, and signatures of knowledge. Although these tools may sound quite heavy, we obtain surprisingly efficient group signatures, which provably provide CCA2-full anonymity in the strongest model for dynamic group signatures, i.e., the BSZ model. In doing so, we obtain the first construction which achieves this strong security notion without an encrypted membership certificate in the signature. In addition to that, we introduce an even more efficient CPA-fully anonymous variant of our scheme.

We proceed in showing how to instantiate our constructions in the random oracle model (ROM) to obtain particularly efficient schemes. We are thereby able to further push the long standing computational efficiency limits for both CPA- and CCA2-fully anonymous schemes regarding signature generation and verification. When comparing to the popular BBS group signature scheme [BBS04] (which achieves CPA-full anonymity in the ROM), besides being more efficient we surprisingly even obtain shorter signatures. Ultimately, when comparing to instantiations in the vein of Bichsel et al. (which provide a less desirable anonymity notion), our instantiations provide comparable computational efficiency.

2 Preliminaries

In this section, we provide some preliminaries and recall the required primitives.

Notation. Let $x \xleftarrow{R} X$ denote the operation that picks an element uniformly at random from a finite set X and assigns it to x . We assume that all algorithms run in polynomial time and use $y \leftarrow A(x)$ to denote that y is assigned the output of the potentially probabilistic algorithm A on input x and fresh random coins and write $y \leftarrow A(x; r)$ to make the random coins r of A explicit. We assume that every algorithm outputs a special symbol \perp on error. We write $\Pr[\Omega : \mathcal{E}]$ to denote the probability of an event \mathcal{E} over the probability space Ω . A function $\epsilon : \mathbb{N} \rightarrow \mathbb{R}^+$ is called negligible if for all $c > 0$ there is a k_0 such that $\epsilon(k) < 1/k^c$ for all $k > k_0$. In the remainder of this paper, we use ϵ to denote such a negligible function. We use the $[\cdot]$ operator to access list entries, i.e., let $L = (a, b, \dots, z)$ then $L[1]$ refers to a .

Let $\mathbb{G}_1 = \langle P \rangle$, $\mathbb{G}_2 = \langle \hat{P} \rangle$, and \mathbb{G}_T be groups of prime order p . A bilinear map $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ is a map, where it holds for all $(P, \hat{Q}, a, b) \in \mathbb{G}_1 \times \mathbb{G}_2 \times \mathbb{Z}_p^2$

that $e(aP, b\hat{Q}) = e(P, \hat{Q})^{ab}$, and $e(P, \hat{P}) \neq 1$, and e is efficiently computable. We assume the Type-3 setting, where $\mathbb{G}_1 \neq \mathbb{G}_2$ and no efficiently computable isomorphism $\psi : \mathbb{G}_2 \rightarrow \mathbb{G}_1$ is known.

Due to the limited space, we postpone the formal definitions of the bilinear group generation, the required cryptographic hardness assumptions, and the required standard building blocks such as digital signature schemes (Σ), public key encryption schemes (Ω), non-interactive zero-knowledge proof systems (Π), and signatures of knowledge (SoK) to Appendix A.

Structure Preserving Signatures on Equivalence Classes. Subsequently, we briefly recall structure-preserving signatures on equivalence classes (SPS-EQ) as presented in [HS14, FHS14]. Therefore, let p be a prime and $\ell > 1$; then \mathbb{Z}_p^ℓ is a vector space and one can define a projective equivalence relation on it, which propagates to \mathbb{G}_i^ℓ and partitions \mathbb{G}_i^ℓ into equivalence classes. Let $\sim_{\mathcal{R}}$ be this relation, i.e., for $M, N \in \mathbb{G}_i^\ell : M \sim_{\mathcal{R}} N \Leftrightarrow \exists s \in \mathbb{Z}_p^* : M = sN$. An SPS-EQ scheme now signs an equivalence class $[M]_{\mathcal{R}}$ for $M \in (\mathbb{G}_i^*)^\ell$ by signing a representative M of $[M]_{\mathcal{R}}$. One of the design goals of SPS-EQ is to guarantee that two message-signature pairs from the same equivalence class cannot be linked. A formal definition is given in Appendix A.

3 Dynamic Group Signatures

Subsequently, we recall the established model for dynamic group signatures. We follow Bellare et al. [BSZ05] (BSZ model), with the slight difference that we relax the perfect correctness to only require computational correctness. Furthermore, we also present the weaker anonymity notion of CPA-full anonymity from [BBS04] and the notion of opening soundness [SSE⁺12], which addresses issues regarding hijacking of signatures by malicious group members. In particular, we use the notion of weak opening soundness, where the opening authority is required to be honest, since we believe that this notion provides a good tradeoff between computational efficiency of potential instantiations and expected security guarantees (even the authors of [SSE⁺12] say that already weak opening soundness addresses the attacks they had in mind).

GKeyGen(1^κ) : This algorithm takes a security parameter κ as input and outputs a triple $(\text{gpk}, \text{ik}, \text{ok})$ containing the group public key gpk , the issuing key ik as well as the opening key ok .

UKeyGen(1^κ) : This algorithm takes a security parameter κ as input and outputs a user key pair $(\text{usk}_i, \text{upk}_i)$.

Join($\text{gpk}, \text{usk}_i, \text{upk}_i$) : This algorithm takes the group public key gpk and the user's key pair $(\text{usk}_i, \text{upk}_i)$ as input. It interacts with the Issue algorithm and outputs the group signing key gsk_i of user i on success.

Issue($\text{gpk}, \text{ik}, i, \text{upk}_i, \text{reg}$) : This algorithm takes the group public key gpk , the issuing key ik , the index i of a user, user i 's public key upk_i , and the registration table reg as input. It interacts with the Join algorithm and adds an entry for user i in reg on success. In the end, it returns reg .

Sign($\text{gpk}, \text{gsk}_i, m$) : This algorithm takes the group public key gpk , a group signing key gsk_i , and a message m as input and outputs a group signature σ .

$\text{Verify}(\text{gpk}, m, \sigma)$: This algorithm takes the group public key gpk , a message m and a signature σ as input and outputs a bit $b \in \{0, 1\}$.

$\text{Open}(\text{gpk}, \text{ok}, \text{reg}, m, \sigma)$: This algorithm takes the group public key gpk , the opening key ok , the registration table reg , a message m , and a valid signature σ on m under gpk as input. It extracts the identity of the signer and returns a pair (i, τ) , where τ is a proof.

$\text{Judge}(\text{gpk}, m, \sigma, i, \text{upk}_i, \tau)$: This algorithm takes the group public key gpk , a message m , a valid signature σ on m under gpk , an index i , user i 's public key upk_i , and a proof τ . It returns a bit $b \in \{0, 1\}$.

Oracles. In the following we recall the definitions of the oracles required by the security model. We assume that the keys $(\text{gpk}, \text{ik}, \text{ok})$ created in the experiments are implicitly available to the oracles. Furthermore, the environment maintains the sets HU, CU of honest and corrupted users, the set GS of message-signature tuples returned by the challenge oracle, the lists $\text{upk}, \text{usk}, \text{gsk}$ of user public keys, user private keys, and group signing keys. The list upk is publicly readable and the environment also maintains the registration table reg . Finally, SI represents a list that ensures the consistency of subsequent calls to CrptU and SndTol . All sets are initially empty and all list entries are initially set to \perp . In the context of lists, we use $\text{upk}_i, \text{usk}_i$, etc. as shorthand for $\text{upk}[i], \text{usk}[i]$, etc.

$\text{AddU}(i)$: This oracle takes an index i as input. If $i \in \text{CU} \cup \text{HU}$ it returns \perp . Otherwise it runs $(\text{usk}_i, \text{upk}_i) \leftarrow \text{UKeyGen}(1^\kappa)$ and

$$(\text{reg}, \text{gsk}_i) \leftarrow \langle \text{Issue}(\text{gpk}, \text{ik}, i, \text{upk}_i, \text{reg}) \leftrightarrow \text{Join}(\text{gpk}, \text{usk}_i, \text{upk}_i) \rangle.$$

Finally, it sets $\text{HU} \leftarrow \text{HU} \cup \{i\}$ and returns upk_i .

$\text{CrptU}(i, \text{upk}_j)$: This oracle takes an index i and user public key upk_j as input. If $i \in \text{CU} \cup \text{HU}$ it returns \perp . Otherwise it sets $\text{CU} \leftarrow \text{CU} \cup \{i\}$, $\text{SI}[i] \leftarrow \top$ and $\text{upk}_i \leftarrow \text{upk}_j$.

$\text{SndTol}(i)$: This oracle takes an index i as input. If $\text{SI}[i] \neq \top$ it returns \perp . Otherwise, it plays the role of an honest issuer when interacting with the corrupted user i . More precisely, it runs

$$\text{reg} \leftarrow \langle \text{Issue}(\text{gpk}, \text{ik}, i, \text{upk}_i, \text{reg}) \leftrightarrow \mathcal{A} \rangle.$$

In the end it sets $\text{SI}[i] \leftarrow \perp$.

$\text{SndToU}(i)$: This oracle takes an index i as input. If $i \notin \text{HU}$ it sets $\text{HU} \leftarrow \text{HU} \cup \{i\}$, runs $(\text{usk}_i, \text{upk}_i) \leftarrow \text{UKeyGen}(1^\kappa)$. Then it plays the role of the honest user i when interacting with a corrupted issuer. More precisely, it runs

$$\text{gsk}_i \leftarrow \langle \mathcal{A} \leftrightarrow \text{Join}(\text{gpk}, \text{usk}_i, \text{upk}_i) \rangle.$$

$\text{USK}(i)$: This oracle takes an index i as input and returns $(\text{gsk}_i, \text{usk}_i)$.

$\text{RReg}(i)$: This oracle takes an index i as input and returns reg_i .

$\text{WReg}(i, \rho)$: This oracle takes an index i and a registration table entry ρ as input and sets $\text{reg}_i \leftarrow \rho$.

$\text{GSig}(i, m)$: This oracle takes an index i and a message m as input. If $i \notin \text{HU}$ or $\text{gsk}_i = \perp$ it returns \perp and $\sigma \leftarrow \text{Sign}(\text{gpk}, \text{gsk}_i, m)$ otherwise.

$\text{Ch}(b, i_0, i_1, m)$: This algorithm takes a bit b , two indexes i_0 and i_1 , and a message m as input. If $\{i_0, i_1\} \not\subseteq \text{HU} \vee \text{gsk}_{i_0} = \perp \vee \text{gsk}_{i_1} = \perp$ it returns \perp . Otherwise, it computes $\sigma \leftarrow \text{Sign}(\text{gpk}, \text{gsk}_{i_b}, m)$, sets $\text{GS} \leftarrow \text{GS} \cup \{(m, \sigma)\}$ and returns σ .

$\text{Open}(m, \sigma)$: This oracle takes a message m and a signature σ as input. If $(m, \sigma) \in \text{GS}$ or $\text{Verify}(\text{gpk}, m, \sigma) = 0$ it returns \perp . Otherwise, it returns $(i, \tau) \leftarrow \text{Open}(\text{gpk}, \text{ok}, \text{reg}, m, \sigma)$.

Security Notions. We require dynamic group signatures to be correct, anonymous, traceable, non-frameable, and weakly opening sound. We recall the formal definitions below.

Correctness, requires that everything works correctly if everyone behaves honestly. Note that we relax perfect correctness to computational correctness.

Definition 1 (Correctness). *A GSS is correct, if for all PPT adversaries \mathcal{A} there is a negligible function $\epsilon(\cdot)$ such that*

$$\Pr \left[\begin{array}{l} (\text{gpk}, \text{ik}, \text{ok}) \leftarrow \text{GKeyGen}(1^\kappa), \\ \mathcal{O} \leftarrow \{\text{AddU}(\cdot), \text{RReg}(\cdot)\}, \\ (i, m) \leftarrow \mathcal{A}^{\mathcal{O}}(\text{gpk}), \\ \sigma \leftarrow \text{Sign}(\text{gpk}, \text{gsk}_i, m), \\ (j, \tau) \leftarrow \text{Open}(\text{gpk}, \text{ok}, \text{reg}, m, \sigma) \end{array} : \begin{array}{l} i \in \text{HU} \wedge \text{gsk}_i \neq \perp \wedge i = j \wedge \\ \text{Verify}(\text{gpk}, m, \sigma) = 1 \wedge \\ \text{Judge}(\text{gpk}, m, \sigma, i, \text{upk}_i, \tau) = 1 \end{array} \right] \geq 1 - \epsilon(\kappa).$$

Anonymity captures the intuition that group signers remain anonymous for everyone except the opening authority. Thereby, the adversary can see arbitrary key exposures. Furthermore, in the CCA2 case the adversary can even request arbitrary openings of other group signatures.

Definition 2 (T-Full Anonymity). *Let $\text{T} \in \{\text{CPA}, \text{CCA2}\}$. A GSS is T -fully anonymous, if for all PPT adversaries \mathcal{A} there is a negligible function $\epsilon(\cdot)$ such that*

$$\Pr \left[(\text{gpk}, \text{ik}, \text{ok}) \leftarrow \text{GKeyGen}(1^\kappa), b \xleftarrow{R} \{0, 1\}, : b = b^* \right] \leq 1/2 + \epsilon(\kappa),$$

$$b^* \leftarrow \mathcal{A}^{\mathcal{O}_\text{T}}(\text{gpk}, \text{ik})$$

where

$$\mathcal{O}_\text{T} \leftarrow \begin{cases} \left\{ \begin{array}{l} \text{Ch}(b, \cdot, \cdot, \cdot), \text{SndToU}(\cdot), \text{WReg}(\cdot, \cdot), \\ \text{USK}(\cdot), \text{CrptU}(\cdot, \cdot) \end{array} \right\} & \text{if } \text{T} = \text{CPA}, \text{ and} \\ \left\{ \begin{array}{l} \text{Ch}(b, \cdot, \cdot, \cdot), \text{Open}(\cdot, \cdot), \text{SndToU}(\cdot), \\ \text{WReg}(\cdot, \cdot), \text{USK}(\cdot), \text{CrptU}(\cdot, \cdot) \end{array} \right\} & \text{if } \text{T} = \text{CCA2}. \end{cases}$$

Traceability models the requirement that, as long as the issuer behaves honestly and its secret key remains secret, every valid signature can be traced back to a user. This must even hold if the opening authority colludes with malicious users.

Definition 3 (Traceability). A GSS is traceable, if for all PPT adversaries \mathcal{A} there is a negligible function $\epsilon(\cdot)$ such that

$$\Pr \left[\begin{array}{l} (\text{gpk}, \text{ik}, \text{ok}) \leftarrow \text{GKeyGen}(1^\kappa), \\ \mathcal{O} \leftarrow \{\text{SndTol}(\cdot), \text{AddU}(\cdot), \\ \text{RReg}(\cdot), \text{USK}(\cdot), \text{CrptU}(\cdot)\}, \\ (m, \sigma) \leftarrow \mathcal{A}^\mathcal{O}(\text{gpk}, \text{ok}), \\ (i, \tau) \leftarrow \text{Open}(\text{gpk}, \text{ok}, \text{reg}, m, \sigma) \end{array} : \begin{array}{l} \text{Verify}(\text{gpk}, m, \sigma) = 1 \wedge \\ (i = \perp \vee \\ \text{Judge}(\text{gpk}, m, \sigma, i, \text{upk}_i, \tau) = 0) \end{array} \right] \leq \epsilon(\kappa).$$

Non-frameability requires that no one can forge signatures for honest users. This must even hold if the issuing authority, the opening authority, and, other malicious users collude.

Definition 4 (Non-Frameability). A GSS is non-frameable, if for all PPT adversaries \mathcal{A} there is a negligible function $\epsilon(\cdot)$ such that

$$\Pr \left[\begin{array}{l} (\text{gpk}, \text{ik}, \text{ok}) \leftarrow \text{GKeyGen}(1^\kappa), \\ \mathcal{O} \leftarrow \{\text{SndToU}(\cdot), \text{WReg}(\cdot, \cdot), \\ \text{GSig}(\cdot, \cdot), \text{USK}(\cdot), \text{CrptU}(\cdot)\}, \\ (m, \sigma, i, \tau) \leftarrow \mathcal{A}^\mathcal{O}(\text{gpk}, \text{ok}, \text{ik}) \end{array} : \begin{array}{l} \text{Verify}(\text{gpk}, m, \sigma) = 1 \wedge \\ i \in \text{HU} \wedge \text{gsk}_i \neq \perp \wedge \\ i \notin \text{USK} \wedge (i, m) \notin \text{SIG} \wedge \\ \text{Judge}(\text{gpk}, m, \sigma, i, \text{upk}_i, \tau) = 1 \end{array} \right] \leq \epsilon(\kappa),$$

where **USK** and **SIG** denote the queries to the oracles **USK** and **Sign**, respectively.

Weak opening soundness [SSE⁺12] essentially requires that no malicious user can claim ownership of a signature issued by an honest user, as long as the opening authority behaves honestly.

Definition 5 (Weak Opening Soundness). A GSS is weakly opening sound, if for all PPT adversaries \mathcal{A} there is a negligible function $\epsilon(\cdot)$ such that

$$\Pr \left[\begin{array}{l} (\text{gpk}, \text{ik}, \text{ok}) \leftarrow \text{GKeyGen}(1^\kappa), \\ \mathcal{O} \leftarrow \{\text{AddU}(\cdot)\}, \\ (m, i, j, \text{st}) \leftarrow \mathcal{A}^\mathcal{O}(\text{gpk}), \\ \sigma \leftarrow \text{Sign}(\text{gpk}, \text{gsk}_i, m), \\ \tau \leftarrow \mathcal{A}^\mathcal{O}(\text{st}, \sigma, \text{gsk}_j) \end{array} : \begin{array}{l} i \neq j \wedge \{i, j\} \subseteq \text{HU} \wedge \\ \text{Judge}(\text{gpk}, m, \sigma, j, \text{upk}_j, \tau) = 1 \end{array} \right] \leq \epsilon(\kappa).$$

4 Construction

Our construction idea is inspired by [HS14], who use the “unlinkability” feature of SPS-EQ signatures to construct anonymous credentials. Essentially, a credential in their approach represents a signature for an equivalence class and to show a credential they always present a newly re-randomized signature to a random representative of this class. While, due to the intuitive relation of anonymous credentials and group signatures, it might seem straightforward to map this idea to group signatures, it turns out that there are various subtle, yet challenging issues which we need to solve.

First, the anonymity notion is much stronger than the one of anonymous credentials in that it does not put many restrictions on the **Ch** and the **USK** oracles. In particular, **Ch** can be called an arbitrary number of times and **USK**

can be called for all users. Thus, the user secret keys must be of a form so that it is possible to embed decision problem instances into them upon simulation, while not influencing their distribution (as the adversary sees those keys and would be able to detect the simulation otherwise). More precisely, anonymity in our paradigm seems to require that the user keys contain no \mathbb{Z}_p elements, which, in turn, renders the non-frameability proof more difficult. Second, if CCA2-full anonymity is required, the simulatability of the open oracle needs to be ensured, while the reduction must not be aware of the opening information (as otherwise the reduction could trivially break anonymity on its own and would be meaningless). This seems to crucially require a proof system providing rather strong extractability properties. To maintain efficiency, it is important to find the mildest possible requirement which still allows the security proofs to work out. Third, the non-frameability adversary is given the issuing key as well as the opening key. Thus, the reduction must be able to simulate the whole join process without knowledge of a user secret key in a way that the distribution change is not even detectable with the knowledge of these keys.

Now, before we present our full construction, we briefly revisit our basic idea. In our scheme, each group member chooses a secret vector $(R, P) \in (\mathbb{G}_1^*)^2$ representing an equivalence class where the second component P is identical for all users. When joining the group, a blinded version $q \cdot (R, P)$ with $q \xleftarrow{R} \mathbb{Z}_p^*$ of this vector, i.e., another representative of the class, is signed by the issuer using an SPS-EQ, and, by the re-randomization property of SPS-EQ and the feature to publicly change representatives of classes, the user thus obtains a signature on the unblinded key (R, P) using $\text{ChgRep}_{\mathcal{R}}$ with q^{-1} . To provide a means to open signatures, a user additionally has to provide an encryption of a value $\hat{R} \in \mathbb{G}_2$ such that $e(R, \hat{P}) = e(P, \hat{R})$ on joining (and has to sign the ciphertext as an identity proof). The group signing key of the user is then the pair consisting of the vector (R, P) and the SPS-EQ signature on this vector. A group member can sign a message m on behalf of the group by randomizing its group signing key and computing a signature of knowledge (SoK) to the message m proving knowledge of the used randomizer.² The group signature is then the randomized group signing key and the SoK.

Very roughly, a signer remains anonymous since it is infeasible to distinguish two randomized user secret keys under DDH in \mathbb{G}_1 . The unforgeability of SPS-EQ ensures that each valid signature can be opened. Furthermore, it is hard to forge signatures of honest group members since it is hard to unblind a user secret key under co-CDHI and the signature of knowledge essentially ensures that we can extract such an unblinded user secret key from a successful adversary.

Detailed Construction. In our scheme, we require zero-knowledge proofs upon Join and Open. The NP relation R_J corresponding to the proof carried out in Join is defined as

$$((U_i, Q, \hat{C}_{J_i}, \text{pk}_O), (r, \omega)) \in R_J \iff \hat{C}_{J_i} = \Omega.\text{Enc}(\text{pk}_O, r\hat{P}; \omega) \wedge U_i = r \cdot Q.$$

² For technical reasons and in particular for extractability, we actually require a signature of knowledge for message $m' = \sigma_1 || m$, where σ_1 contains the re-randomized user secret key and SPS-EQ signature.

The **NP** relation R_O corresponding to the proof carried out upon **Open** is

$$\begin{aligned} ((\hat{C}_J, \text{pk}_O, \sigma), (\text{sk}_O, \hat{R})) \in R_O &\iff \hat{R} = \Omega.\text{Dec}(\text{sk}_O, \hat{C}_J) \wedge \\ &\text{pk}_O \equiv \text{sk}_O \wedge e(\sigma_1[1][1], \hat{P}) = e(\sigma_1[1][2], \hat{R}). \end{aligned}$$

Thereby, $\text{pk} \equiv \text{sk}$ denotes the consistency of pk and sk . Note that σ_1 represents the randomized user secret key, i.e., is of the form $((\rho R, \rho P), \sigma')$ and consists of a randomized message vector and a corresponding randomized SPS-EQ signature. We use $\sigma_1[1][j]$ to refer to the j th element in the (randomized) message vector.

Furthermore, upon **Sign** we require a signature of knowledge which is with respect to the following **NP** relation R_S .

$$((P, Q), \rho) \in R_S \iff Q = \rho \cdot P.$$

For the sake of compactness, we assume that the languages defined by R_J, R_O, R_S are implicit in the CRSs $\text{crs}_J, \text{crs}_O$, and crs_S , respectively. The full construction is presented as Scheme 1. Note that if multiple users collude and use the same value r upon **Join**⁽¹⁾, we always return the first user who registered with this particular value r in **Open**. Then, **Open** always returns the signer who initiated the collusion by sharing the r value, which, we think, is the most reasonable choice. Note that this is in line with the BSZ model: traceability only requires that every valid signature can be opened, while not requiring that it opens to one particular user out of the set of colluding users; correctness and non-frameability are defined with respect to honest users and are therefore clearly not influenced.

4.1 Security

First, note that our **Join** \leftrightarrow **Issue** protocol is inherently and trivially concurrently secure: we only have two moves which means that interleaving different **Join** \leftrightarrow **Issue** instances is impossible. Thus concurrency issues are implicitly covered by our further analysis.

Theorem 1. *If SPS-EQ is correct, SoK is correct, and Π is sound, then Scheme 1 is correct.*

Proof (Sketch). Correctness is straight forward to verify by inspection. We only have to take care of one detail: There is the possibility that two honest executions of **AddU** yield the same value r (which is chosen uniformly at random upon **Join**⁽¹⁾). Thus, the probability of two colliding r is negligible. \square

In Appendix B, we formally prove Theorem 2-5.

Theorem 2. *If Π is adaptively zero-knowledge, SoK is simulatable, Ω is IND-CPA secure, SPS-EQ perfectly adapts signatures, and the DDH assumption holds in \mathbb{G}_1 , then Scheme 1 is CPA-full anonymous.*

Theorem 3. *If Π is adaptively zero-knowledge, SoK is simulatable and straight-line f -extractable, where $f : \mathbb{Z}_p \rightarrow \mathbb{G}_2$ is defined as $r \mapsto r \cdot \hat{P}$, Ω is IND-CCA2 secure, SPS-EQ perfectly adapts signatures, and the DDH assumption holds in \mathbb{G}_1 , then Scheme 1 is CCA2-full anonymous.*

GKeyGen(1^κ) : Run $\text{BG} \leftarrow \text{BGGen}_{\mathcal{R}}(1^\kappa)$, $(\text{sk}_{\mathcal{R}}, \text{pk}_{\mathcal{R}}) \leftarrow \text{KeyGen}_{\mathcal{R}}(\text{BG}, 2)$, $(\text{sk}_0, \text{pk}_0) \leftarrow \Omega.\text{KeyGen}(1^\kappa)$, $\text{crs}_J \leftarrow \Pi.\text{Setup}(1^\kappa)$, $\text{crs}_O \leftarrow \Pi.\text{Setup}(1^\kappa)$, $\text{crs}_S \leftarrow \text{SoK}.\text{Setup}(1^\kappa)$, set $\text{gpk} \leftarrow (\text{pk}_{\mathcal{R}}, \text{pk}_0, \text{crs}_J, \text{crs}_O, \text{crs}_S)$, $\text{ik} \leftarrow \text{sk}_{\mathcal{R}}$, $\text{ok} \leftarrow \text{sk}_0$ and return $(\text{gpk}, \text{ik}, \text{ok})$.

UKeyGen(1^κ) : Return $(\text{usk}_i, \text{upk}_i) \leftarrow \Sigma.\text{KeyGen}(1^\kappa)$.

Join⁽¹⁾($\text{gpk}, \text{usk}_i, \text{upk}_i$) : Choose $q, r \xleftarrow{R} \mathbb{Z}_p^*$, set $(U_i, Q) \leftarrow (r \cdot qP, qP)$, and output $M_J \leftarrow ((U_i, Q), \hat{C}_{J_i}, \sigma_{J_i}, \pi_{J_i})$ and $\text{st} \leftarrow (\text{gpk}, q, U_i, Q)$, where

$$\hat{C}_{J_i} \leftarrow \Omega.\text{Enc}(\text{pk}_0, r\hat{P}; \omega), \sigma_{J_i} \leftarrow \Sigma.\text{Sign}(\text{usk}_i, \hat{C}_{J_i}),$$

$$\pi_{J_i} \leftarrow \Pi.\text{Proof}(\text{crs}_J, (U_i, Q, \hat{C}_{J_i}, \text{pk}_0), (r, \omega)).$$

Issue($\text{gpk}, \text{ik}, i, \text{upk}_i, \text{reg}$) : Receive $M_J = ((U_i, Q), \hat{C}_{J_i}, \sigma_{J_i}, \pi_{J_i})$, return reg and send σ' to user i , where

$$\text{reg}_i \leftarrow (\hat{C}_{J_i}, \sigma_{J_i}), \sigma' \leftarrow \text{Sign}_{\mathcal{R}}((U_i, Q), \text{sk}_{\mathcal{R}}),$$

if $\Pi.\text{Verify}(\text{crs}_J, (U_i, Q, \hat{C}_{J_i}, \text{pk}_0), \pi_{J_i}) = 1 \wedge \Sigma.\text{Verify}(\text{upk}_i, \hat{C}_{J_i}, \sigma_{J_i}) = 1$, and return \perp otherwise.

Join⁽²⁾(st, σ') : Parse st as (gpk, q, U_i, Q) and return gsk_i , where

$$\text{gsk}_i = ((rP, P), \sigma) \leftarrow \text{ChgRep}_{\mathcal{R}}((U_i, Q), \sigma', q^{-1}, \text{pk}_{\mathcal{R}}),$$

if $\text{Verify}_{\mathcal{R}}((U_i, Q), \sigma', \text{pk}_{\mathcal{R}}) = 1$, and return \perp otherwise.

Sign($\text{gpk}, \text{gsk}_i, m$) : Choose $\rho \xleftarrow{R} \mathbb{Z}_p^*$, and return $\sigma \leftarrow (\sigma_1, \sigma_2)$, where

$$\sigma_1 \leftarrow \text{ChgRep}_{\mathcal{R}}(\text{gsk}_i, \rho, \text{pk}_{\mathcal{R}}), \sigma_2 \leftarrow \text{SoK}.\text{Sign}(\text{crs}_S, (P, \sigma_1[1][2]), \rho, \sigma_1 || m).$$

Verify(gpk, m, σ) : Return 1 if the following holds, and 0 otherwise:

$$\text{Verify}_{\mathcal{R}}(\sigma_1, \text{pk}_{\mathcal{R}}) = 1 \wedge \text{SoK}.\text{Verify}(\text{crs}_S, (P, \sigma_1[1][2]), \sigma_1 || m, \sigma_2) = 1.$$

Open($\text{gpk}, \text{ok}, \text{reg}, m, \sigma$) : Parse σ as (σ_1, σ_2) , and ok as sk_0 . Obtain the lowest index i, b so that it holds for $(\hat{C}_{J_i}, \sigma_{J_i}) \leftarrow \text{reg}_i$ that $\hat{R} \leftarrow \Omega.\text{Dec}(\text{sk}_0, \hat{C}_{J_i})$ and $e(\sigma_1[1][1], \hat{P}) = e(\sigma_1[1][2], \hat{R})$. Return (i, τ) and \perp if no such entry exists, where

$$\tau \leftarrow (\pi_0, \hat{C}_{J_i}, \sigma_{J_i}), \text{ and } \pi_0 \leftarrow \Pi.\text{Proof}(\text{crs}_O, (\hat{C}_{J_i}, \text{pk}_0, \sigma), (\text{sk}_0, \hat{R})).$$

Judge($\text{gpk}, m, \sigma, i, \text{upk}_i, \tau$) : Parse τ as $(\pi_0, \hat{C}_{J_i}, \sigma_{J_i})$, and return 1 if the following holds and 0 otherwise:

$$\Sigma.\text{Verify}(\text{upk}_i, \hat{C}_{J_i}, \sigma_{J_i}) = 1 \wedge \Pi.\text{Verify}(\text{crs}_O, (\hat{C}_{J_i}, \text{pk}_0, \sigma), \pi_0) = 1.$$

^a Note that gsk_i is of the form $((R, P), \sigma)$ and σ_1 is a randomization of gsk_i . We slightly abuse the notation of $\text{Verify}_{\mathcal{R}}$ and $\text{ChgRep}_{\mathcal{R}}$ and input message-signature tuples instead of separately inputting messages and signatures.

^b We assume that the indexes are in ascending order w.r.t. the time of registration.

Scheme 1: Fully-Anonymous Dynamic Group Signature Scheme

Theorem 4. *If SPS-EQ is EUF-CMA secure, and Π is sound, then Scheme 1 is traceable.*

Theorem 5. *If Π is sound and adaptively zero-knowledge, SoK is simulatable and extractable, Σ is EUF-CMA secure, Ω is perfectly correct, and the co-CDHI assumption holds, then Scheme 1 is non-frameable.*

Theorem 6. *If Ω is perfectly correct, and Σ is EUF-CMA secure, then Scheme 1 is weakly opening sound.*

Proof (Sketch). Upon honestly executing Join for users i and j , the probability that their r (resp. \hat{R}) values collide is negligible. The perfect correctness of Ω and the EUF-CMA security of Σ thus uniquely determine user i as the signer of σ with overwhelming probability. Then, it is easy to see that an adversary against weak opening soundness is an adversary against soundness of Π . \square

5 Instantiation in the ROM

To compare our approach to existing schemes regarding signature size and computational effort upon signature generation and verification, we present the sign and verification algorithms for an instantiation of our scheme with the SPS-EQ from [FHS14, FHS15], whose security is shown to hold in the generic group model. We instantiate SoKs in the ROM by applying the transformation from [FKMV12] to Fiat-Shamir (FS) transformed Σ -protocols.

Before we introduce the approaches to obtain CPA-fully (resp. CCA2-fully) anonymous instantiations, we recall that the group signing key gsk_i consists of a vector of two group elements $(R, P) \in (\mathbb{G}_1^*)^2$ and an SPS-EQ signature $\sigma \in \mathbb{G}_1 \times \mathbb{G}_1^* \times \mathbb{G}_2^*$ on this vector. Randomization of a gsk_i with a random value $\rho \in \mathbb{Z}_p^*$, i.e., $\text{ChgRep}_{\mathcal{R}}$, requires 4 multiplications in \mathbb{G}_1 and 1 multiplication in \mathbb{G}_2 . Verification of an SPS-EQ signature on gsk_i requires 5 pairings.

We note that the proofs performed within Join and Open can straight forwardly be instantiated using standard techniques. Therefore, and since they are neither required within Sign nor Verify, we do not discuss instantiations here.

5.1 CPA-Full Anonymity

Subsequently, we show how Sign and Verify are instantiated in the CPA-full anonymity setting. Therefore, let $H : \{0, 1\}^* \rightarrow \mathbb{Z}_p$ be a random oracle and let x be the proven statement (which is implicitly defined by the scheme):

$\text{Sign}(\text{gpk}, \text{gsk}_i, m) : \text{Parse } \text{gsk}_i \text{ as } ((R, P), \sigma)$, choose $\rho \xleftarrow{R} \mathbb{Z}_p$, compute $\sigma_1 = ((R', P'), \sigma') \leftarrow \text{ChgRep}_{\mathcal{R}}(\text{gsk}_i, \rho, \text{pk}_{\mathcal{R}})$. Choose $\nu \xleftarrow{R} \mathbb{Z}_p$, compute $N \leftarrow \nu P$, $c \leftarrow H(N || \sigma_1 || m || x)$, $z \leftarrow \nu + c \cdot \rho$, set $\sigma_2 \leftarrow (c, z)$, and return $\sigma \leftarrow (\sigma_1, \sigma_2)$.

$\text{Verify}(\text{gpk}, m, \sigma) : \text{Parse } \sigma \text{ as } (\sigma_1, \sigma_2) = (((R', P'), \sigma'), (c, z))$, return 0 if $\text{Verify}_{\mathcal{R}}(\sigma_1, \text{pk}_{\mathcal{R}}) = 0$. Otherwise compute $N \leftarrow zP - cP'$ and check whether $c = H(N || \sigma_1 || m || x)$ holds. If so return 1 and 0 otherwise.

Since the used Σ -protocol is a standard proof of knowledge of the discrete logarithm $\log_P P'$, it is easy to see that applying the transformations from [FKMV12] yields a SoK in the ROM with the properties we require. All in all, group signatures contain 4 elements in \mathbb{G}_1 , 1 element in \mathbb{G}_2 and 2 elements in \mathbb{Z}_p . Counting only the expensive operations, signing costs 5 multiplications in \mathbb{G}_1 and 1 multiplication in \mathbb{G}_2 , and verification costs 2 multiplications in \mathbb{G}_1 and 5 pairings.

5.2 CCA2-Full Anonymity

CCA2-full anonymity requires straight-line extractable SoKs, as standard rewinding would lead to an exponential blowup in the reduction (cf. [BFW15]). One possibility would be to rely on the rather inefficient approach to straight-line extraction due to Fischlin [Fis05]. However, as we do not need to straight-line extract the full witness w , but it is sufficient to straight-line extract an image of w under a one-way function $f : \rho \mapsto \rho \cdot \hat{P}$, we can use the notion of straight-line f -extractable SoKs as recently proposed by Cerulli et al. [BCC⁺15]. This allows us to still use the FS paradigm with good efficiency. The construction uses the generic conversion in [FKMV12, BPW12]. The generic trick in [BCC⁺15] to obtain straight-line f -extractability is by computing an extractable commitment to the image of the witness w under a function f with respect to an extraction key in the CRS and proving consistency with the witness.³

For straight-line extractability, we let \hat{Y} be a public key for the ElGamal variant in \mathbb{G}_2 from [BCC⁺15], which is generated upon SoK.Setup and represents the CRS of SoK. SoK.SimSetup additionally returns τ such that $\hat{Y} = \tau \cdot \hat{P}$. Furthermore, let x be the proven statement (implicitly defined by the scheme and the generic compiler). Subsequently, we show how Sign and Verify are instantiated in this setting, where $H : \{0, 1\}^* \rightarrow \mathbb{Z}_p$ is modelled as a random oracle:

$\text{Sign}(\text{gpk}, \text{gsk}_i, m)$: Parse gsk_i as $((R, P), \sigma)$, choose $\rho \xleftarrow{R} \mathbb{Z}_p$, compute $\sigma_1 = ((R', P'), \sigma') \leftarrow \text{ChgRep}_{\mathcal{R}}(\text{gsk}_i, \rho, \text{pk}_{\mathcal{R}})$. Choose $u, \nu, \eta \xleftarrow{R} \mathbb{Z}_p$, compute $(\hat{C}_1, \hat{C}_2) = (u\hat{Y}, \rho\hat{P} + u\hat{P})$, $N \leftarrow \nu P$, $\hat{M}_1 \leftarrow \eta\hat{Y}$, $\hat{M}_2 \leftarrow (\nu + \eta)\hat{P}$, $c \leftarrow H(N || \hat{M}_1 || \hat{M}_2 || \sigma_1 || m || x)$, $z_1 \leftarrow \nu + c \cdot \rho$, $z_2 \leftarrow \eta + c \cdot u$, set $\sigma_2 \leftarrow (\hat{C}_1, \hat{C}_2, c, z_1, z_2)$, and return $\sigma \leftarrow (\sigma_1, \sigma_2)$.

$\text{Verify}(\text{gpk}, m, \sigma)$: Parse σ as $(\sigma_1, \sigma_2) = (((R', P'), \sigma), (c, z_1, z_2))$, return 0 if $\text{Verify}_{\mathcal{R}}(\sigma_1, \text{pk}_{\mathcal{R}}) = 0$. Otherwise compute $N \leftarrow z_1 P - c P'$, $\hat{M}_1 \leftarrow z_2 \cdot \hat{Y} - c \cdot \hat{C}_1$, $\hat{M}_2 \leftarrow (z_1 + z_2) \cdot \hat{P} - c \cdot \hat{C}_2$, and check whether $c = H(N || \hat{M}_1 || \hat{M}_2 || \sigma_1 || m || x)$ holds. If so return 1 and 0 otherwise.

In Appendix B.4 and Appendix B.5 we prove the following lemmas. Note that we additionally require the Σ -protocol to provide quasi-unique responses [Fis05], i.e., given an accepting proof it should be computationally infeasible to find a new valid response for that proof, in order for the compiler in [BCC⁺15] to apply.

Lemma 1. *The above Σ -protocol is perfectly complete, SHVZK, special-sound and has quasi-unique responses.*

Lemma 2. *Applying the generic conversions from [FKMV12] to the Fiat-Shamir transformed version of the above Σ -protocol with the setup SoK.Setup as described in Section 5.2 produces a signature of knowledge in the random oracle model, that is extractable and straight-line f -extractable.*

Switching Groups. The protocol presented above requires more operations in the more expensive group \mathbb{G}_2 than in \mathbb{G}_1 . As we work in the SXDH setting, we

³ Note that one can still obtain the full witness w using a rewinding extractor.

can simply switch the roles of \mathbb{G}_1 and \mathbb{G}_2 and thus all elements in \mathbb{G}_1 to \mathbb{G}_2 and vice versa. This allows us to trade computational efficiency for signature size.

6 Evaluation and Discussion

Subsequently, we discuss our work in the light of some recent concurrent and independent work and provide a performance evaluation.

The [BCC⁺16] Model. In independent and concurrent work, a new model for fully-dynamic group signatures was proposed by Bootle et al. in [BCC⁺16]. Bootle et al. address maliciously generated issuer and opener keys, include the notion of opening soundness from [SSE⁺12] and formally model revocation by means of epochs. Although we target security in a different model, we want to briefly put our construction in context of their recent model.

In our scheme, one can straight forwardly incorporate the requirement to support maliciously generated keys in the fashion of [BCC⁺16] by extending the actual public keys of issuer and opener by a (straight-line extractable) zero-knowledge proof of knowledge of the respective secret issuer and opener key.

For a practical revocation approach, it seems to be reasonable to choose a re-issuing based approach, i.e., to set up a new group after every epoch, as also used in [BCC⁺16]. Their group signature construction being secure in their model builds upon accountable ring signatures [BCC⁺15]. It comes at the cost of a group public key size linear in the number of group members as well as a signature size logarithmic in the number of group members, and the revocation related re-issuing requires every group member to obtain the new group public key. Applying the same revocation approach to our scheme yields public keys as well as signatures of constant size, and re-issuing requires each group member which is still active to re-join the new group.

While our scheme provides weak opening soundness, achieving the stronger notion for our scheme (where the opening authority may be malicious) would require the opening authority to additionally prove that the opened index i corresponds to the lowest index in \mathbf{reg} so that the respective entry together with the signature in question satisfies the relation $R_{\mathcal{O}}$. Such a proof could efficiently be instantiated using non-interactive plaintext in-equality proofs [BDSS16]. Nevertheless, we opted to stick with weak opening soundness because: (1) The only benefit of strong opening soundness would be to also cover dishonest opening authorities, while we believe that assuming the opening authority’s honesty—given its power to deanonymize every user—is a crucial and very reasonable assumption. (2) Even [SSE⁺12], who introduced opening soundness emphasize that already weak opening soundness addresses all the attacks that motivated opening soundness in the first place. (3) Strong opening soundness would unnecessarily degrade the simplicity of our scheme.

6.1 Performance Evaluation and Comparison

To underline the practical efficiency of our approach, we provide a comparison of our ROM instantiation with other schemes in the ROM. In particular we use two schemes who follow the approach of Bichsel et al., i.e., [BCN⁺10, PS16], which provide less desirable anonymity guarantees (denoted CCA^-), and the

well known BBS scheme [BBS04] (with and without precomputations) providing CPA-full anonymity. We note that we use the plain BBS scheme for comparison, which does not even provide non-frameability and the non-frameable version would be even more expensive. Moreover, we use the group signature scheme with the shortest known signatures [DP06] (with and without precomputations) being secure in the strong BSZ model and thus providing CCA2-full anonymity. Finally, we also compare our scheme to the recent CCA2-fully anonymous scheme by Libert et al. [LMPY16] which is secure in the ROM under SXDH.

In Table 1 we provide a comparison of the estimated efficiency in a 254bit BN-pairing setting, where we highlight the values where our scheme is currently the best known scheme among other existing schemes providing the same security guarantees. Our estimations are based on performance values on an ARM-Cortex-M0+ with drop-in hardware accelerator [UW14]. This processor is small enough to be suited for smart cards or wireless sensor nodes [UW14]. Table 2 in the Appendix provides an abstract comparison regarding signature size, computational costs, and the type of the underlying hardness assumption.

Scheme	Anon.	Signature Size	Signature Cost	Verification Cost
[BCN ⁺ 10]	CCA ⁻	1273bit	351ms	1105ms
[PS16]	CCA ⁻	1018bit	318ms	777ms
[BBS04]	CPA	2289bit	1545ms	2092ms
[BBS04] (prec.)	CPA	2289bit	1053ms	1600ms
This paper	CPA	2037bit	266ms	886ms
This paper	CCA2	3309bit	771ms	1290ms
This paper (switch)	CCA2	3563bit	703ms	1154ms
[DP06]	CCA2	2290bit	1380ms	2059ms
[DP06] (prec.)	CCA2	2290bit	1020ms	1353ms
[LMPY16]	CCA2	2547bit	1688ms	2299ms

Table 1. Estimations based on a BN-pairing implementation on an ARM-Cortex-M0+ with drop-in hardware accelerator, operating at 48MHz [UW14]. The performance figures using 254-bit curves are 33ms-101ms-252ms-164ms (\mathbb{G}_1 - \mathbb{G}_2 - \mathbb{G}_T -pairing). For the estimation of signature sizes, we use 255bit for elements in \mathbb{G}_1 , 509bit for elements in \mathbb{G}_2 and 254bit for elements in \mathbb{Z}_p . We note that [BBS04] is defined for a Type-2 pairing setting, which means that our performance estimation for this scheme is rather optimistic and likely to be worse in practice. The bold values highlight where our schemes are currently the fastest and have the shortest signatures.

Computational Efficiency. When comparing our CPA-fully anonymous scheme as well as our CCA2-fully anonymous scheme to other schemes providing the same anonymity guarantees, ours are the *by now fastest ones regarding signature generation and verification costs*. While some of the schemes used for comparison use slightly less progressive assumptions, it seems that very good performance requires more progressive assumptions. When looking for instance at the most compact CCA2-fully anonymous group signatures in the standard model under standard assumptions (SXDH and XDLIN) by Libert et al. [LPY15], signature sizes in the best case will have 30 \mathbb{G}_1 and 14 \mathbb{G}_2 elements (≈ 15000 bit when

taking the setting in Table 1), large public keys and computation times that are far from being feasible for resource constrained devices.

Regarding signature generation, we want to emphasize that our CPA-fully anonymous instantiation is the fastest among all schemes used for comparison (even among the ones providing CCA^- anonymity), and, to the best of our knowledge, *the fastest among all existing schemes*. This is of particular importance since signature generation is most likely to be executed on a constrained device. Regarding signature verification our CPA-fully anonymous instantiation is only outperformed by the CCA^- anonymous instantiation in [PS16].

Signature Size. Comparing schemes providing the same anonymity guarantees, our CPA-fully anonymous instantiation even provides shorter signature sizes than the popular BBS scheme [BBS04] and, to the best of our knowledge, *the shortest signature sizes among all CPA-fully anonymous schemes*. Regarding CCA^2 -fully anonymous schemes, it seems that gained efficiency in the “without encryption” paradigm comes at the cost of larger signatures compared to instantiations following the SEP paradigm. It is interesting to note that the schemes in vein of Bichsel et al. providing only CCA^- anonymity have the smallest signatures among all schemes.

7 Conclusion

In this paper we further pushed the efficiency limits of CPA- as well as CCA^2 -fully anonymous group signature schemes with respect to signature generation and verification in favor of accepting slightly more progressive, yet very plausible, assumptions. We also pushed the limits with respect to signature size for CPA-fully anonymous group signature schemes. Besides that, our results answer the theoretical question whether CPA- as well as CCA^2 -fully anonymous schemes following the “without encryption” paradigm are possible at all.

Finally, we point out a few interesting observations. First, due to our modular construction, coming up with a suitable SPS-EQ under weaker assumptions would directly imply an instantiation of our scheme under weaker assumptions. Besides that, our construction paradigm may as well be inspiring for efficient standard model instantiations. Yet, it seems that moving to weaker assumptions would almost certainly mean to give up (some of) our efficiency. Additionally, we observe that our construction neither requires any pairing computations nor computations in the target group \mathbb{G}_T upon signature creation, which makes it especially suitable for constrained devices. Finally, it seems that one could exploit the \hat{R} values in a relatively straight forward manner to obtain traceable signatures [KTY04, Cho09].

Acknowledgements. We thank the anonymous referees from ASIACRYPT’16 for their valuable comments.

References

- [ACJT00] Giuseppe Ateniese, Jan Camenisch, Marc Joye, and Gene Tsudik. A Practical and Provably Secure Coalition-Resistant Group Signature Scheme. In *CRYPTO*, 2000.

- [BBS04] Dan Boneh, Xavier Boyen, and Hovav Shacham. Short Group Signatures. In *CRYPTO*, 2004.
- [BCC⁺15] Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Essam Ghadafi, Jens Groth, and Christophe Petit. Short Accountable Ring Signatures Based on DDH. In *ESORICS*, 2015.
- [BCC⁺16] Jonathan Bootle, Andrea Cerulli, Pyrros Chaidos, Essam Ghadafi, and Jens Groth. Foundations of fully dynamic group signatures. In *ACNS*, 2016. Full Version: IACR ePrint Report 2016/368.
- [BCN⁺10] Patrik Bichsel, Jan Camenisch, Gregory Neven, Nigel P. Smart, and Bogdan Warinschi. Get Shorty via Group Signatures without Encryption. In *SCN*, 2010.
- [BDSS16] Olivier Blazy, David Derler, Daniel Slamanig, and Raphael Spreitzer. Non-Interactive Plaintext (In-)Equality Proofs and Group Signatures with Verifiable Controllable Linkability. In *CT-RSA*, 2016.
- [BFW15] David Bernhard, Marc Fischlin, and Bogdan Warinschi. Adaptive proofs of knowledge in the random oracle model. In *PKC*, 2015.
- [BGI14] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional Signatures and Pseudorandom Functions. In *PKC*, 2014.
- [BMW03] Mihir Bellare, Daniele Micciancio, and Bogdan Warinschi. Foundations of Group Signatures: Formal Definitions, Simplified Requirements, and a Construction Based on General Assumptions. In *EUROCRYPT*, 2003.
- [BPW12] David Bernhard, Olivier Pereira, and Bogdan Warinschi. How not to prove yourself: Pitfalls of the fiat-shamir heuristic and applications to helios. In *ASIACRYPT*, 2012.
- [BSZ05] Mihir Bellare, Haixia Shi, and Chong Zhang. Foundations of Group Signatures: The Case of Dynamic Groups. In *CT-RSA*, 2005.
- [BW06] Xavier Boyen and Brent Waters. Compact Group Signatures Without Random Oracles. In *EUROCRYPT*, 2006.
- [BW07] Xavier Boyen and Brent Waters. Full-Domain Subgroup Hiding and Constant-Size Group Signatures. In *PKC*, 2007.
- [Cho09] Sherman S. M. Chow. Real Traceable Signatures. In *SAC*, 2009.
- [CL04] Jan Camenisch and Anna Lysyanskaya. Signature Schemes and Anonymous Credentials from Bilinear Maps. In *CRYPTO*, 2004.
- [CL06] Melissa Chase and Anna Lysyanskaya. On Signatures of Knowledge. In *CRYPTO*, 2006.
- [CS97] Jan Camenisch and Markus Stadler. Efficient Group Signature Schemes for Large Groups (Extended Abstract). In *CRYPTO*, 1997.
- [CvH91] David Chaum and Eugène van Heyst. Group Signatures. In *EUROCRYPT*, 1991.
- [DP06] Cécile Delerablée and David Pointcheval. Dynamic Fully Anonymous Short Group Signatures. In *VIETCRYPT*, 2006.
- [FHS14] Georg Fuchsbauer, Christian Hanser, and Daniel Slamanig. Structure-Preserving Signatures on Equivalence Classes and Constant-Size Anonymous Credentials. IACR Cryptology ePrint Archive, Report 2014/944, 2014. <http://eprint.iacr.org/>.
- [FHS15] Georg Fuchsbauer, Christian Hanser, and Daniel Slamanig. Practical Round-Optimal Blind Signatures in the Standard Model. In *CRYPTO*, 2015.
- [Fis05] Marc Fischlin. Communication-Efficient Non-interactive Proofs of Knowledge with Online Extractors. In *CRYPTO*, 2005.

- [FKMV12] Sebastian Faust, Markulf Kohlweiss, Giorgia Azzurra Marson, and Daniele Venturi. On the non-malleability of the fiat-shamir transform. In *INDOCRYPT*, 2012.
- [GMR88] S. Goldwasser, S. Micali, and R. L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM JoC*, 17(2), 1988.
- [Gro07] Jens Groth. Fully Anonymous Group Signatures Without Random Oracles. In *ASIACRYPT*, 2007.
- [HS14] Christian Hanser and Daniel Slamanig. Structure-Preserving Signatures on Equivalence Classes and Their Application to Anonymous Credentials. In *ASIACRYPT*, 2014.
- [KTY04] Aggelos Kiayias, Yiannis Tsiounis, and Moti Yung. Traceable Signatures. In *EUROCRYPT*, 2004.
- [KY05] Aggelos Kiayias and Moti Yung. Group signatures with efficient concurrent join. In *EUROCRYPT*, 2005.
- [LLM⁺16] Benoît Libert, San Ling, Fabrice Mouhartem, Khoa Nguyen, and Huaxiong Wang. Signature Schemes with Efficient Protocols and Dynamic Group Signatures from Lattice Assumptions. Cryptology ePrint Archive, Report 2016/101, 2016.
- [LMPY16] Benoît Libert, Fabrice Mouhartem, Thomas Peters, and Moti Yung. Practical “Signatures with Efficient Protocols” from Simple Assumptions. In *Asia CCS*, 2016.
- [LPY15] Benoît Libert, Thomas Peters, and Moti Yung. Short Group Signatures via Structure-Preserving Signatures: Standard Model Security from Simple Assumptions. In *CRYPTO*, 2015.
- [NS04] Lan Nguyen and Reihaneh Safavi-Naini. Efficient and provably secure trapdoor-free group signature schemes from bilinear pairings. In *ASIACRYPT*, 2004.
- [PS16] David Pointcheval and Olivier Sanders. Short Randomizable Signatures. In *CT-RSA*, 2016.
- [SSE⁺12] Yusuke Sakai, Jacob C. N. Schuldt, Keita Emura, Goichiro Hanaoka, and Kazuo Ohta. On the Security of Dynamic Group Signatures: Preventing Signature Hijacking. In *PKC*, 2012.
- [UW14] Thomas Unterluggauer and Erich Wenger. Efficient Pairings and ECC for Embedded Systems. In *CHES*, 2014.

A Hardness Assumptions and Building Blocks

Definition 6 (Bilinear Group Generator). *Let BGGen be an algorithm which takes a security parameter κ and generates a bilinear group $\text{BG} = (p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, P, \hat{P})$ in the Type-3 setting, where the common group order p of the groups $\mathbb{G}_1, \mathbb{G}_2$ and \mathbb{G}_T is a prime of bitlength κ , e is a pairing and P and \hat{P} are generators of \mathbb{G}_1 and \mathbb{G}_2 , respectively.*

Definition 7 (Decisional Diffie-Hellman Assumption (DDH)). *Let $\mathbb{G} = \langle P \rangle$ be a group of prime order p , such that $\log_2 p = \kappa$. Then, for all PPT adversaries \mathcal{A} there exists a negligible function $\epsilon(\cdot)$ such that:*

$$\Pr \left[b \xleftarrow{R} \{0, 1\}, r, s, t \xleftarrow{R} \mathbb{Z}_p, \right. \\ \left. b^* \leftarrow \mathcal{A}(P, rP, sP, (b \cdot (rs) + (1 - b) \cdot t)P) : b = b^* \right] \leq 1/2 + \epsilon(\kappa).$$

Definition 8 (Symmetric External Diffie-Hellman Assumption (SXDH)).

Let BG be a bilinear group generated by BGGen . Then, the SXDH assumption states that the DDH assumption holds in \mathbb{G}_1 and \mathbb{G}_2 .

Additionally, we introduce a natural assumption in the Type-3 bilinear group setting.

Definition 9 (Computational co-Diffie-Hellman Inversion Assumption (co-CDHI)). Let $\text{BG} \leftarrow \text{BGGen}(1^\kappa)$. The co-CDHI assumption states that for all PPT adversaries \mathcal{A} there exists a negligible function $\epsilon(\cdot)$ such that:

$$\Pr \left[a \leftarrow^R \mathbb{Z}_p, C \leftarrow \mathcal{A}(\text{BG}, aP, 1/a\hat{P}) : C = 1/aP \right] \leq \epsilon(\kappa).$$

Structure Preserving Signatures on Equivalence Classes. Below, we formally recall the definition of an SPS-EQ scheme.

Definition 10. An SPS-EQ on \mathbb{G}_i^* (for $i \in \{1, 2\}$) consists of the following PPT algorithms:

$\text{BGGen}_{\mathcal{R}}(1^\kappa)$: This algorithm on input of a security parameter κ outputs a bilinear group BG .

$\text{KeyGen}_{\mathcal{R}}(\text{BG}, \ell)$: This algorithm on input of a bilinear group BG and a vector length $\ell > 1$ outputs a key pair (sk, pk) .

$\text{Sign}_{\mathcal{R}}(M, \text{sk})$: This algorithm on input a representative $M \in (\mathbb{G}_i^*)^\ell$ and a secret key sk outputs a signature σ for the equivalence class $[M]_{\mathcal{R}}$.

$\text{ChgRep}_{\mathcal{R}}(M, \sigma, \rho, \text{pk})$: This algorithm on input of a representative $M \in (\mathbb{G}_i^*)^\ell$ of class $[M]_{\mathcal{R}}$, a signature σ for M , a scalar ρ and a public key pk returns an updated message-signature pair (M', σ') , where $M' = \rho \cdot M$ is the new representative and σ' its updated signature.

$\text{Verify}_{\mathcal{R}}(M, \sigma, \text{pk})$: This algorithm on input of a representative $M \in (\mathbb{G}_i^*)^\ell$, a signature σ and a public key pk outputs a bit $b \in \{0, 1\}$.

$\text{VKey}_{\mathcal{R}}(\text{sk}, \text{pk})$ This algorithm on input a secret key sk and a public key pk outputs a bit $b \in \{0, 1\}$.

For security, one requires the following properties.

Definition 11 (Correctness). An SPS-EQ scheme on $(\mathbb{G}_i^*)^\ell$ is called correct if for all security parameters $\kappa \in \mathbb{N}$, $\ell > 1$, $\text{BG} \leftarrow \text{BGGen}_{\mathcal{R}}(1^\kappa)$, $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}_{\mathcal{R}}(\text{BG}, \ell)$, $M \in (\mathbb{G}_i^*)^\ell$ and $\rho \in \mathbb{Z}_p^*$:

$$\begin{aligned} \text{VKey}_{\mathcal{R}}(\text{sk}, \text{pk}) = 1 \quad \wedge \quad \Pr \left[\text{Verify}_{\mathcal{R}}(M, \text{Sign}_{\mathcal{R}}(M, \text{sk}), \text{pk}) = 1 \right] = 1 \quad \wedge \\ \Pr \left[\text{Verify}_{\mathcal{R}}(\text{ChgRep}_{\mathcal{R}}(M, \text{Sign}_{\mathcal{R}}(M, \text{sk}), \rho, \text{pk}), \text{pk}) = 1 \right] = 1. \end{aligned}$$

For EUF-CMA security, outputting a valid message-signature pair, corresponding to an unqueried equivalence class, is considered to be a forgery:

Definition 12 (EUF-CMA). An SPS-EQ over $(\mathbb{G}_i^*)^\ell$ is existentially unforgeable under adaptively chosen-message attacks, if for all PPT adversaries \mathcal{A} with access to a signing oracle $\mathcal{O}^{\text{Sign}\mathcal{R}}$, there is a negligible function $\epsilon(\cdot)$ such that:

$$\Pr \left[\begin{array}{l} \text{BG} \leftarrow \text{BGGen}_{\mathcal{R}}(1^\kappa), \\ (\text{sk}, \text{pk}) \leftarrow \text{KeyGen}_{\mathcal{R}}(\text{BG}, \ell), : [M^*]_{\mathcal{R}} \neq [M]_{\mathcal{R}} \quad \forall M \in Q^{\text{Sign}\mathcal{R}} \quad \wedge \\ (M^*, \sigma^*) \leftarrow \mathcal{A}^{\mathcal{O}^{\text{Sign}\mathcal{R}}(\text{sk}, \cdot)}(\text{pk}) \quad \text{Verify}_{\mathcal{R}}(M^*, \sigma^*, \text{pk}) = 1 \end{array} \right] \leq \epsilon(\kappa),$$

where $Q^{\text{Sign}\mathcal{R}}$ is the set of queries that \mathcal{A} has issued to the signing oracle $\mathcal{O}^{\text{Sign}\mathcal{R}}$.

Besides EUF-CMA security, an additional security property for SPS-EQ was introduced in [FHS15].

Definition 13 (Perfect Adaption of Signatures). An SPS-EQ scheme on $(\mathbb{G}_i^*)^\ell$ perfectly adapts signatures if for all tuples $(\text{sk}, \text{pk}, M, \sigma, \rho)$ where it holds that $\text{VKey}_{\mathcal{R}}(\text{sk}, \text{pk}) = 1$, $\text{Verify}_{\mathcal{R}}(M, \sigma, \text{pk}) = 1$, $M \in (\mathbb{G}_i^*)^\ell$, and $\rho \in \mathbb{Z}_p^*$, the distributions $(\rho M, \text{Sign}_{\mathcal{R}}(\rho M, \text{sk}))$ and $\text{ChgRep}_{\mathcal{R}}(M, \sigma, \rho, \text{pk})$ are identical.

An instantiation providing all above security properties is provided in [FHS14, FHS15]. Here, assuming the DDH assumption to hold on the message space yields that different message-signature pairs from the same equivalence class cannot be linked.

Digital Signature Schemes. Subsequently, we recall a definition of digital signature schemes.

Definition 14 (Digital Signatures). A digital signature scheme Σ is a triple $(\text{KeyGen}, \text{Sign}, \text{Verify})$ of PPT algorithms, which are defined as follows:

$\text{KeyGen}(1^\kappa)$: This algorithm takes a security parameter κ as input and outputs a secret (signing) key sk and a public (verification) key pk with associated message space \mathcal{M} (we may omit to mention the message space \mathcal{M}).

$\text{Sign}(\text{sk}, m)$: This algorithm takes a secret key sk and a message $m \in \mathcal{M}$ as input and outputs a signature σ .

$\text{Verify}(\text{pk}, m, \sigma)$: This algorithm takes a public key pk , a message $m \in \mathcal{M}$ and a signature σ as input and outputs a bit $b \in \{0, 1\}$.

Besides correctness we require existential unforgeability under adaptively chosen message attacks (EUF-CMA) [GMR88]. Subsequently, we recall formal definitions of these properties.

Definition 15 (Correctness). A digital signature scheme Σ is correct, if for all κ , all $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(1^\kappa)$ and all $m \in \mathcal{M}$ it holds that

$$\Pr[\text{Verify}(\text{pk}, m, \text{Sign}(\text{sk}, m)) = 1] = 1.$$

Definition 16 (EUF-CMA). A digital signature scheme Σ is EUF-CMA secure, if for all PPT adversaries \mathcal{A} there is a negligible function $\epsilon(\cdot)$ such that

$$\left[\begin{array}{l} (\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(1^\kappa), \\ (m^*, \sigma^*) \leftarrow \mathcal{A}^{\mathcal{O}^{\text{Sign}(\text{sk}, \cdot)}}(\text{pk}) \end{array} : \text{Verify}(\text{pk}, m^*, \sigma^*) = 1 \quad \wedge \quad m^* \notin Q^{\text{Sign}} \right] \leq \epsilon(\kappa),$$

where \mathcal{A} has access to an oracle $\mathcal{O}^{\text{Sign}}$ that allows to execute the **Sign** algorithm and the environment keeps track of all message queried to $\mathcal{O}^{\text{Sign}}$ via Q^{Sign} .

Public Key Encryption. We also require public key encryption, which we recall below.

Definition 17. A public key encryption scheme Ω is a triple $(\text{KeyGen}, \text{Enc}, \text{Dec})$ of PPT algorithms, which are defined as follows:

$\text{KeyGen}(1^\kappa)$: This algorithm takes a security parameter κ as input and outputs a secret decryption key sk and a public encryption key pk (and we assume that the message space \mathcal{M} is implicitly defined by pk).

$\text{Enc}(\text{pk}, m)$: This algorithm takes a public key pk and a message $m \in \mathcal{M}$ as input and outputs a ciphertext c .

$\text{Dec}(\text{sk}, c)$: This algorithm takes a secret key sk and a ciphertext c as input and outputs a message $m \in \mathcal{M}$ or \perp .

We require a public key encryption scheme to be correct and IND-T secure and recall the formal definitions below.

Definition 18 (Correctness). A public key encryption scheme Ω is correct if it holds for all κ , for all $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(1^\kappa)$, and for all messages $m \in \mathcal{M}$ that

$$\Pr[\text{Dec}(\text{sk}, \text{Enc}(\text{pk}, m)) = m] = 1.$$

Definition 19 (IND-T Security). Let $\text{T} \in \{\text{CPA}, \text{CCA2}\}$. A public key encryption scheme Ω is IND-T secure, if for all PPT adversaries \mathcal{A} there exists a negligible function $\epsilon(\cdot)$ such that

$$\Pr \left[\begin{array}{l} (\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(1^\kappa), \\ (m_0, m_1, \text{st}) \leftarrow \mathcal{A}^{\mathcal{O}_\text{T}}(\text{pk}), \\ b \xleftarrow{R} \{0, 1\}, c \leftarrow \text{Enc}(\text{pk}, m_b), \\ b^* \leftarrow \mathcal{A}^{\mathcal{O}_\text{T}}(c, \text{st}) \end{array} : c \notin Q^{\text{Dec}} \wedge |m_0| = |m_1| \wedge b = b^* \right] \leq 1/2 + \epsilon(\kappa),$$

where the adversary runs in two stages,

$$\mathcal{O}_\text{T} \leftarrow \begin{cases} \emptyset & \text{if } \text{T} = \text{CPA}, \text{ and} \\ \{\mathcal{O}^{\text{Dec}}(\text{sk}, \cdot)\} & \text{if } \text{T} = \text{CCA2}, \end{cases}$$

and Q^{Dec} denotes the list of queries to \mathcal{O}^{Dec} and we set $Q^{\text{Dec}} \leftarrow \emptyset$ if $\text{T} = \text{CPA}$.

Non-Interactive Zero-Knowledge Proof Systems. Now, we recall a standard definition of non-interactive zero-knowledge proof systems. Therefore, let L_R be an NP-language with witness relation $R : L_R = \{x \mid \exists w : R(x, w) = 1\}$.

Definition 20 (Non-Interactive Zero-Knowledge Proof System). A non-interactive proof system Π is a tuple of algorithms $(\text{Setup}, \text{Proof}, \text{Verify})$, which are defined as follows:

$\text{Setup}(1^\kappa)$: This algorithm takes a security parameter κ as input, and outputs a common reference string crs .

$\text{Proof}(\text{crs}, x, w)$: This algorithm takes a common reference string crs , a statement x , and a witness w as input, and outputs a proof π .

$\text{Verify}(\text{crs}, x, \pi)$: This algorithm takes a common reference string crs , a statement x , and a proof π as input, and outputs a bit $b \in \{0, 1\}$.

Subsequently, we recall formal definition of those properties (adapted from [BGI14]).

Definition 21 (Completeness). A non-interactive proof system Π is complete, if for every adversary \mathcal{A} it holds that

$$\Pr \left[\begin{array}{l} \text{crs} \leftarrow \text{Setup}(1^\kappa), (x, w) \leftarrow \mathcal{A}(\text{crs}), \\ \pi \leftarrow \text{Proof}(\text{crs}, x, w) \end{array} : \begin{array}{l} \text{Verify}(\text{crs}, x, \pi) = 1 \\ \wedge (x, w) \in R \end{array} \right] \approx 1.$$

Definition 22 (Soundness). A non-interactive proof system Π is sound, if for every PPT adversary \mathcal{A} there is a negligible function $\epsilon(\cdot)$ such that

$$\Pr \left[\text{crs} \leftarrow \text{Setup}(1^\kappa), (x, \pi) \leftarrow \mathcal{A}(\text{crs}) : \text{Verify}(\text{crs}, x, \pi) = 1 \wedge x \notin L_R \right] \leq \epsilon(\kappa).$$

If we quantify over all adversaries \mathcal{A} and require $\epsilon = 0$, we have perfect soundness, but we present the definition for computationally sound proofs (arguments).

Definition 23 (Adaptive Zero-Knowledge). A non-interactive proof system Π is adaptively zero-knowledge, if there exists a PPT simulator $S = (S_1, S_2)$ such that for every PPT adversary \mathcal{A} there is a negligible function $\epsilon(\cdot)$ such that

$$\left| \begin{array}{l} \Pr \left[\text{crs} \leftarrow \text{Setup}(1^\kappa) : \mathcal{A}^{\mathcal{P}(\text{crs}, \cdot, \cdot)}(\text{crs}) = 1 \right] \\ \Pr \left[(\text{crs}, \tau) \leftarrow S_1(1^\kappa) : \mathcal{A}^{S(\text{crs}, \tau, \cdot, \cdot)}(\text{crs}) = 1 \right] \end{array} \right| \leq \epsilon(\kappa),$$

where, τ denotes a simulation trapdoor. Thereby, \mathcal{P} and S return \perp if $(x, w) \notin R$ or $\pi \leftarrow \text{Proof}(\text{crs}, x, w)$ and $\pi \leftarrow S_2(\text{crs}, \tau, x)$, respectively, otherwise.

If $\epsilon = 0$, we have perfect adaptive zero-knowledge.

Signatures of Knowledge. Below we recall signatures of knowledge (SoKs) [CL06], where L_R is as above. For the formal notions we follow [BCC⁺15] and use a stronger generalization of the original extraction property termed f -extractability. A signature of knowledge (SoK) for L_R is defined as follows.

Definition 24. A SoK is a tuple of PPT algorithms ($\text{Setup}, \text{Sign}, \text{Verify}$), which are defined as follows:

$\text{Setup}(1^\kappa)$: This algorithm takes a security parameter κ as input and outputs a common reference string crs . We assume that the message space \mathcal{M} is implicitly defined by crs .

$\text{Sign}(\text{crs}, x, w, m)$: This algorithm takes a common reference string crs , a word x , a witness w , and a message m as input and outputs a signature σ .

$\text{Verify}(\text{crs}, x, m, \sigma)$: This algorithm takes a common reference string crs , a word x , a message m , and a signature σ as input and outputs a bit $b \in \{0, 1\}$.

Definition 25 (Correctness). A SoK with respect to L_R is correct, if there exists a negligible function $\epsilon(\cdot)$ such that for all $x \in L_R$, for all w such that $(x, w) \in R$, and for all $m \in \mathcal{M}$ it holds that

$$\Pr[\text{crs} \leftarrow \text{Setup}(1^\kappa), \sigma \leftarrow \text{Sign}(\text{crs}, x, w, m) : \text{Verify}(\text{crs}, x, m, \sigma) = 1] \geq 1 - \epsilon(\kappa).$$

Definition 26 (Simulatability). A SoK with respect to L_R is simulatable, if there exists a PPT simulator $\mathcal{S} = (\text{SimSetup}, \text{SimSign})$ such that for all PPT adversaries \mathcal{A} there exists a negligible function $\epsilon(\cdot)$ such that it holds that

$$\left| \Pr[\text{crs} \leftarrow \text{Setup}(1^\kappa), b \leftarrow \mathcal{A}^{\text{Sign}(\text{crs}, \cdot, \cdot, \cdot)}(\text{crs}) : b = 1] - \Pr[(\text{crs}, \tau) \leftarrow \text{SimSetup}(1^\kappa), b \leftarrow \mathcal{A}^{\text{Sim}(\text{crs}, \tau, \cdot, \cdot)}(\text{crs}) : b = 1] \right| \leq \epsilon(\kappa),$$

where $\text{Sim}(\text{crs}, \tau, x, w, m) := \text{SimSign}(\text{crs}, \tau, x, m)$ and Sim only responds if $(x, w) \in R$.

Definition 27 (f -Extractability). A SoK with respect to L_R is f -extractable, if in addition to \mathcal{S} there exists a PPT extractor Extract , such that for all PPT adversaries \mathcal{A} there exists a negligible function $\epsilon(\cdot)$ such that it holds that

$$\Pr \left[\begin{array}{l} (\text{crs}, \tau) \leftarrow \text{SimSetup}(1^\kappa), \\ (x, m, \sigma) \leftarrow \mathcal{A}^{\text{Sim}(\text{crs}, \tau, \cdot, \cdot)}(\text{crs}), \\ y \leftarrow \text{Extract}(\text{crs}, \tau, x, m, \sigma) \end{array} : \begin{array}{l} \text{Verify}(\text{crs}, x, m, \sigma) = 0 \vee \\ (x, m, \sigma) \in Q^{\text{Sim}} \vee \\ (\exists w : (x, w) \in R \wedge \\ y = f(w)) \end{array} \right] \geq 1 - \epsilon(\kappa),$$

where Q^{Sim} denotes the queries (resp. answers) of Sim .

We note that, as illustrated in [BCC⁺15], this notion is a generalization of the original extractability notion from [CL06] which implies the original extractability notion if f is the identity. In this case, we simply call the f -extractability property *extractability*. Analogous to [BCC⁺15], we require the used SoK to be at the same time extractable and *straight-line f -extractable* with respect to some f other than the identity, where straight-line as usual says that the extractor runs without rewinding the adversary [Fis05].

B Security Proofs

In our proofs, we omit to make the negligible distribution switches which arise when sampling uniformly random from \mathbb{Z}_p instead of \mathbb{Z}_p^* explicit and instead treat them as conceptual changes for the sake of compactness.

B.1 Proof of Theorem 2 and Theorem 3

Proof (Anonymity). We prove Theorem 2 and 3 by showing that the output distributions of the Ch oracle are (computationally) independent of the bit b , where we highlight the parts of the proof which are specific to Theorem 3 and can be omitted to prove Theorem 2. Therefore, let $q_{\text{Ch}} \leq \text{poly}(\kappa)$ be the number of queries to Ch, $q_{\text{O}} \leq \text{poly}(\kappa)$ be the number of queries to Open, and $q_{\text{SndToU}} \leq \text{poly}(\kappa)$ be the number of queries to SndToU.

Game 0: The original anonymity game.

Game 1: As Game 0, but we run $(\text{crs}_J, \tau_J) \leftarrow \Pi.\mathcal{S}_1(1^\kappa)$ instead of $\text{crs}_J \leftarrow \Pi.\text{Setup}(1^\kappa)$ upon running GKeyGen and store the trapdoor τ_J . Then, we simulate all calls to $\Pi.\text{Proof}$ executed in Join using the simulator (without a witness).

Transition - Game 0 \rightarrow Game 1: A distinguisher $\mathcal{D}^{0 \rightarrow 1}$ is an adversary against adaptive zero-knowledge of Π , and, therefore, the probability to distinguish Game 0 and Game 1 is negligible, i.e., $|\Pr[S_1] - \Pr[S_0]| \leq \epsilon_{\text{ZK}_J}(\kappa)$.

Game 2: As Game 1, but we run $(\text{crs}_O, \tau_O) \leftarrow \Pi.\mathcal{S}_1(1^\kappa)$ instead of $\text{crs}_O \leftarrow \Pi.\text{Setup}(1^\kappa)$ upon running GKeyGen and store the trapdoor τ_O . Then, we simulate all calls to $\Pi.\text{Proof}$ in Open using the simulator (without a witness).

Transition - Game 1 \rightarrow Game 2: A distinguisher $\mathcal{D}^{1 \rightarrow 2}$ is an adversary against adaptive zero-knowledge of Π , and, therefore, the probability to distinguish Game 1 and Game 2 is negligible, i.e., $|\Pr[S_2] - \Pr[S_1]| \leq \epsilon_{\text{ZK}_O}(\kappa)$.

Game 3: As Game 2, but we run $(\text{crs}_S, \tau_S) \leftarrow \text{SoK}.\text{SimSetup}(1^\kappa)$ instead of $\text{crs}_S \leftarrow \text{SoK}.\text{Setup}(1^\kappa)$ upon running GKeyGen and store the trapdoor τ_S . Then we simulate all calls to $\text{SoK}.\text{Sign}$ using the simulator (without a witness).

Transition - Game 2 \rightarrow Game 3: A distinguisher $\mathcal{D}^{2 \rightarrow 3}$ is an adversary against simulatability of SoK. Therefore, the distinguishing probability is negligible, i.e., $|\Pr[S_3] - \Pr[S_2]| \leq \epsilon_{\text{SIM}}(\kappa)$.

Game 4: As Game 3, but instead of setting $(\text{sk}_O, \text{pk}_O) \leftarrow \Omega.\text{KeyGen}(1^\kappa)$ in GKeyGen, we obtain pk_O from an IND-CPA (resp. IND-CCA2) challenger and set $\text{sk}_O \leftarrow \perp$.

In the CCA2 case, we additionally maintain secret lists AU and OI, and upon each call to the SndToU oracle we store $\text{AU}[i] \leftarrow (\text{gsk}_i, \hat{C}_{J_i}) = (((R, P), \sigma), \hat{C}_{J_i})$. Then, we simulate the WReg oracle as follows

WReg(i, ρ): As the original oracle, but we additionally parse ρ as $(\hat{C}_{J_i}, \sigma_{J_i})$. If there exists an index j so that $\text{AU}[j][2] = \hat{C}_{J_i}$, we parse $\text{AU}[j][1]$ as $((R, P), \sigma)$ and set $\text{OI}[i] \leftarrow (R, \perp)$. If there exists no such index, we obtain \hat{R} using the decryption oracle and set $\text{OI}[i] \leftarrow (\perp, \hat{R})$.

Furthermore, we simulate the Open algorithm within the Open oracle as follows.

Open(gpk, ok, reg, m, σ): First, obtain $\hat{\Psi} = \rho \hat{P}$ using the straight-line f -extractor. Then, obtain the lowest index i where either $e(\sigma_1[1][1], \hat{P}) = e(\sigma_1[1][2], \text{OI}[i][2])$ holds, or $e(\text{OI}[i][1], \hat{\Psi}) = e(\sigma_1[1][1], \hat{P})$ holds. Compute a simulated proof τ and return (i, τ) and \perp if no such index exists.

If the extractor fails at some point, we choose $b \xleftarrow{R} \{0, 1\}$ and return b .

Transition - Game 3 \rightarrow Game 4 (CPA): In the CPA case, we do not have to simulate the open oracle, and we only obtain the opening key from an IND-CPA challenger. Thus, this change is conceptual, i.e., $\Pr[S_3] = \Pr[S_4]$.

Transition - Game 3 \rightarrow Game 4 (CCA2): By the straight-line f -extractability of the SoK, one can extract a witness ρ in every call to **Open** with overwhelming probability $1 - \epsilon_{\text{EXT}}(\kappa)$. Thus, both games proceed identically unless the extraction fails, i.e., $|\Pr[S_3] - \Pr[S_4]| \leq q_{\text{O}} \cdot \epsilon_{\text{EXT}}(\kappa)$.

Game 5: As Game 4, but we compute the ciphertext \hat{C}_{J_i} in the Join algorithm (executed within the SndToU oracle) as $\hat{C}_{J_i} \leftarrow \Omega.\text{Enc}(\text{pk}, \hat{P})$, i.e., with a constant message that is independent of the user.

Transition - Game 4 \rightarrow Game 5: A distinguisher $\mathcal{D}^{4 \rightarrow 5}$ is a distinguisher for the IND-CPA (resp. IND-CCA2) game of Ω , i.e., $|\Pr[S_5] - \Pr[S_4]| \leq q_{\text{SndToU}} \cdot \epsilon_{\text{CPA}}(\kappa)$ (resp. $|\Pr[S_5] - \Pr[S_4]| \leq q_{\text{SndToU}} \cdot \epsilon_{\text{CCA2}}(\kappa)$).⁴

Game 6: As Game 5, but we re-add sk_{O} , i.e., we again obtain $(\text{sk}_{\text{O}}, \text{pk}_{\text{O}}) \leftarrow \Omega.\text{KeyGen}(1^\kappa)$. In the CCA2 case, we again decrypt ourselves with in the WReg simulation instead of using the decryption oracle.

Transition - Game 5 \rightarrow Game 6: This change is conceptual, i.e., $\Pr[S_5] = \Pr[S_6]$.

Game 7: As Game 6, but all calls to $\text{ChgRep}_{\mathcal{R}}(M, \rho, \text{pk}_{\mathcal{R}})$ are replaced by $\text{Sign}_{\mathcal{R}}(\rho \cdot M, \text{sk}_{\mathcal{R}})$.

Transition - Game 6 \rightarrow Game 7: Under perfect adaption of signatures, the output distributions in Game 6 and Game 7 are identical, i.e., $\Pr[S_7] = \Pr[S_6]$.

Game 8: As Game 7, but we modify the Ch oracle as follows. Instead of running $\sigma_1 \leftarrow \text{Sign}_{\mathcal{R}}(\rho \cdot \text{gsk}_{i_b}[1], \text{sk}_{\mathcal{R}})$, we choose $S, T \xleftarrow{R} \mathbb{G}_1$, and compute $\sigma_1 \leftarrow \text{Sign}_{\mathcal{R}}((T, S), \text{sk}_{\mathcal{R}})$.

Transition - Game 7 \rightarrow Game 8: We claim that $|\Pr[S_7] - \Pr[S_8]| \leq q_{\text{Ch}} \cdot \epsilon_{\text{DDH}}(\kappa)$. We will below proof this claim separately.

In Game 8, the simulation is independent of the bit b , i.e., $\Pr[S_8] = 1/2$; what remains is to obtain a bound on the success probability in Game 0. In the CPA case, we have that $\Pr[S_0] \leq 1/2 + q_{\text{SndToU}} \cdot \epsilon_{\text{CPA}}(\kappa) + q_{\text{Ch}} \cdot \epsilon_{\text{DDH}}(\kappa) + \epsilon_{\text{ZK}_J}(\kappa) + \epsilon_{\text{ZK}_O}(\kappa) + \epsilon_{\text{SIM}}(\kappa)$, which proves Theorem 2. In the CCA2 case, we have that $\Pr[S_0] \leq 1/2 + q_{\text{SndToU}} \cdot \epsilon_{\text{CCA2}}(\kappa) + q_{\text{Ch}} \cdot \epsilon_{\text{DDH}}(\kappa) + \epsilon_{\text{ZK}_J}(\kappa) + \epsilon_{\text{ZK}_O}(\kappa) + \epsilon_{\text{SIM}}(\kappa) + q_{\text{O}} \cdot \epsilon_{\text{EXT}}(\kappa)$, which proves Theorem 3. \square

⁴ For compactness, we collapsed the q_{SndToU} game changes into a single game change and note that one can straight forwardly unroll this to q_{SndToU} game changes where a single ciphertext is exchanged in each game.

Proof (of Claim). Below we will show that Game 7 and Game 8 are indistinguishable by introducing further intermediate hybrid games.

Game 7₁: As Game 7, but we introduce a conceptual change which will make the subsequent distribution changes easier to follow. In particular upon each SndToU, we modify the simulation of Join so that we no longer choose $r \xleftarrow{R} \mathbb{Z}_p$ to obtain $(U_i, Q) \leftarrow (r \cdot qP, qP)$, but choose $R \xleftarrow{R} \mathbb{G}_1$ and obtain $(U_i, Q) \leftarrow (qR, qP)$.

Transition - Game 7 \rightarrow Game 7₁: This is a conceptual change, i.e., $\Pr[S_7] = \Pr[S_{7_1}]$. Observe that we do not need to know r , as the proofs upon Join are simulated without a witness. Also the user secret keys $\text{gsk}_i = ((R, P), \sigma)$ are exactly the same as honest secret keys.

Game 7_j ($2 \leq j \leq q_{\text{Ch}} + 1$): As Game 7₁, but we modify the Ch oracle as follows. For the first $j-1$ queries, instead of running $\sigma_1 \leftarrow \text{Sign}_{\mathcal{R}}(\rho \cdot \text{gsk}_{i_b}[1], \text{sk}_{\mathcal{R}})$, we choose $S, T \xleftarrow{R} \mathbb{G}_1$, and compute $\sigma_1 \leftarrow \text{Sign}_{\mathcal{R}}((T, S), \text{sk}_{\mathcal{R}})$.

Transition - 7_j \rightarrow 7_{j+1} ($1 \leq j \leq q_{\text{Ch}} + 1$): For each transition, we present a hybrid game, which uses a DDH challenger to interpolate between Game 7_j and Game 7_{j+1}. First, we obtain a DDH instance $(aP, bP, cP) \in \mathbb{G}_1^3$ relative to BG. Then we proceed as follows:

- Upon each SndToU, we modify the simulation of Join as follows. Let i be the index of the user to join. We use the random self reducibility of DDH to obtain an independent DDH instance $(R_i, S_i, T_i) \xleftarrow{R, S, R} (aP, bP, cP)$ and set $\text{CH}[i] \leftarrow (R_i, S_i, T_i)$. Then, we let $(U_i, Q) \leftarrow (qR_i, qP)$.
- Up to the $j-1$ th query to Ch (i.e., for all queries where the answers are already random in Game 7_j), we compute σ_1 by choosing $S, T \xleftarrow{R} \mathbb{G}_1$, and compute $\sigma_1 \leftarrow \text{Sign}_{\mathcal{R}}((T, S), \text{sk}_{\mathcal{R}})$.
- Upon the j th query to Ch, we obtain $(\cdot, S_{i_b}, T_{i_b}) \leftarrow \text{CH}[i_b]$ and set $\sigma_1 \leftarrow \text{Sign}_{\mathcal{R}}((T_{i_b}, S_{i_b}), \text{sk}_{\mathcal{R}})$.
- Starting from the $j+1$ th query to Ch (i.e., for all queries where the answers are still honest in Game 7_j), we obtain $(R_{i_b}, \cdot, \cdot) \leftarrow \text{CH}[i_b]$, choose $\rho \xleftarrow{R} \mathbb{Z}_p$ and set $\sigma_1 \leftarrow \text{Sign}_{\mathcal{R}}((\rho R_{i_b}, \rho P), \text{sk}_{\mathcal{R}})$.

In Game 7_j the first $j-1$ answers are already random due to the previous switches. Furthermore, the validity of the DDH instance (aP, bP, cP) provided by the challenger determines whether the answer of Ch for the j th query are for user i_b or random, i.e., if we are in Game j or in Game $j+1$. That is, $|\Pr[S_j] - \Pr[S_{j+1}]| \leq \epsilon_{\text{DDH}}(\kappa)$.

In Game 7 _{$q_{\text{Ch}}+1$} all answers of Ch are random, i.e., this Game is equal to Game 8, i.e., $\Pr[S_8] = \Pr[7_{q_{\text{Ch}}+1}]$. We can conclude the proof by summing over the distinguishing probabilities of all game changes which yields $|\Pr[S_7] - \Pr[S_8]| \leq q_{\text{Ch}} \cdot \epsilon_{\text{DDH}}(\kappa)$. \square

B.2 Proof of Theorem 4

Proof (Traceability). We show that traceability holds using a sequence of games, where we let $q \leq \text{poly}(\kappa)$ be the number of queries to the SndTol oracle.

Game 0: The original traceability game.

Game 1: As Game 0, but we obtain crs_J from a soundness challenger of Π .

Transition - Game 0 \rightarrow Game 1: This change is conceptual, i.e., $\Pr[S_0] = \Pr[S_1]$.

Game 2: As Game 1, but after every successful execution of SndTol , we obtain $\hat{R} \leftarrow \Omega.\text{Dec}(\text{sk}_O, C_{J_i})$ and abort if $e(U_i, \hat{P}) \neq e(Q, \hat{R})$.

Transition - Game 0 \rightarrow Game 1: If we abort we have a valid proof π_{J_i} attesting that $(U_i, Q, \hat{C}_{J_i}, \text{pk}_O) \in L_{R_J}$, but by the perfect correctness of Ω there exists no ω such that $C_{J_i} = \Omega.\text{Enc}(\text{pk}_O, r \cdot \hat{P}; \omega) \wedge U_i = r \cdot Q$, i.e., $(U_i, Q, \hat{C}_{J_i}, \text{pk}_O)$ is actually not in L_{R_J} . Thus, both games proceed identically unless the adversary breaks the soundness of Π in one oracle query, i.e., $|\Pr[S_1] - \Pr[S_2]| \leq q \cdot \epsilon_S(\kappa)$.

Game 3: As Game 2, but we obtain BG and a public key $\text{pk}_{\mathcal{R}}$ from an EUF-CMA challenger of the SPS-EQ. Whenever an SPS-EQ signature is required, the message to be signed is forwarded to the signing oracle provided by the EUF-CMA challenger.

Transition - Game 2 \rightarrow Game 3: This change is conceptual, i.e., $\Pr[S_2] = \Pr[S_3]$.

If the adversary eventually outputs a valid forgery (m, σ) , we know that σ contains an SPS-EQ signature σ_1 for some (rP, P) such that we have never seen a corresponding $r\hat{P}$, i.e., there is no entry i in the registration table where \hat{C}_{J_i} contains $r\hat{P}$ s.t. $e(\sigma_1[1][1], \hat{P}) = e(\sigma_1[1][2], r\hat{P})$ holds. Consequently, σ_1 is a valid SPS-EQ signature for an unqueried equivalence class and we have that $\Pr[S_3] \leq \epsilon_F(\kappa)$. This yields $\Pr[S_0] \leq \epsilon_F(\kappa) + q \cdot \epsilon_S(\kappa)$, which proves the theorem. \square

B.3 Proof of Theorem 5

Proof (Non-frameability). We prove non-frameability using a sequence of games. Thereby we let the number of users in the system be $q \leq \text{poly}(\kappa)$.

Game 0: The original non-frameability game.

Game 1: As Game 0, but we guess the index i^* that will be attacked by the adversary. If the adversary attacks another index, we abort.

Transition - Game 0 \rightarrow Game 1: The winning probability in Game 1 is the same as in Game 0, unless an abort event happens, i.e., $\Pr[S_1] = \Pr[S_0] \cdot 1/q$.

Game 2: As Game 1, but we run $(\text{crs}_J, \tau_J) \leftarrow \Pi.\mathcal{S}_1(1^\kappa)$ instead of $\text{crs}_J \leftarrow \Pi.\text{Setup}(1^\kappa)$ upon running GKeyGen and store the trapdoor τ_J . Then, we simulate all calls to $\Pi.\text{Proof}$ in Join using the simulator (without a witness).

Transition - Game 1 \rightarrow Game 2: A distinguisher $\mathcal{D}^{1 \rightarrow 2}$ is an adversary against adaptive zero-knowledge of Π , and, therefore, the probability to distinguish Game 1 and Game 2 is negligible, i.e., $|\Pr[S_2] - \Pr[S_1]| \leq \epsilon_{ZK_J}(\kappa)$.

Game 3: As Game 2, but we obtain crs_O from a soundness challenger upon running GKeyGen .

Transition - Game 2 \rightarrow Game 3: This change is conceptual, i.e., $\Pr[S_3] = \Pr[S_2]$.

Game 4: As Game 3, but we setup the SoK in simulation mode, i.e., we run $(\text{crs}_S, \tau_S) \leftarrow \text{SoK}.\text{SimSetup}(1^\kappa)$ instead of $\text{crs}_S \leftarrow \text{SoK}.\text{Setup}(1^\kappa)$ upon running GKeyGen and store the trapdoor τ_S . Then, we simulate all calls to $\text{SoK}.\text{Sign}$ using the simulator, i.e., without a witness.

Transition - Game 3 \rightarrow Game 4: A distinguisher $\mathcal{D}^{3 \rightarrow 4}$ is an adversary against simulatability of SoK. Therefore, the distinguishing probability is negligible, i.e., $|\Pr[S_4] - \Pr[S_3]| \leq \epsilon_{\text{SIM}}(\kappa)$.

Game 5: As Game 4, but we choose the values $r, q \xleftarrow{R} \mathbb{Z}_p$ used in the Join algorithm (executed within the SndToU oracle) when queried for user with index i^* beforehand and let (U_{i^*}, Q_{i^*}) denote (rqP, qP) . Then, on every Join (within SndToU) for a user $i \neq i^*$ we check whether we have incidentally chosen the same class as for user i^* . This check is implemented as follows: with r_i being the value for r chosen upon Join for user i , we check whether $U_{i^*} = r_i \cdot Q_{i^*}$ (note that this check does not require to know the discrete logarithms q and r for user i^*).

Transition - Game 4 \rightarrow Game 5: Both games proceed identically unless we have to abort. An abort happens with probability $\epsilon_{\text{guess}}(\kappa) = q/p-1$ and we have that $|\Pr[S_4] - \Pr[S_5]| \leq \epsilon_{\text{guess}}(\kappa)$.

Game 6: As Game 5, but we obtain a co-CDHI instance $(aP, 1/a\hat{P})$ relative to BG and choose $\tau \xleftarrow{R} \mathbb{Z}_p$. Then, we modify the Join algorithm (executed within the SndToU oracle) when queried for user with index i^* as follows. We set $(U_{i^*}, Q_{i^*}) \leftarrow (\tau \cdot P, aP)$, and compute $\hat{C}_{J_{i^*}} \leftarrow \Omega.\text{Enc}(\text{pk}_O, \tau \cdot 1/a\hat{P})$ and store τ . On successful execution we set $\text{gsk}_{i^*} \leftarrow ((U_{i^*}, Q_{i^*}), \sigma')$ (note that $\pi_{J_{i^*}}$ as well as the signatures in the GSig oracle are already simulated, i.e., the discrete log of no Q_i value is required to be known to the environment).

Transition - Game 5 \rightarrow Game 6: Since τ is uniformly random, we can write it as $\tau = ra$ for some $r \in \mathbb{Z}_p$. Then it is easy to see that the game change is conceptual, i.e., $\Pr[S_5] = \Pr[S_4]$.

Game 7: As Game 6, but for every forgery output by the \mathcal{A} , we extract $\rho \leftarrow \text{SoK.Extract}(\text{crs}_S, \tau_S, (P, \sigma_1[1][2]), \sigma_1 || m, \sigma_2)$ and abort if the extraction fails.

Transition - Game 6 \rightarrow Game 7: By the extractability of the SoK, one can extract a witness ρ with overwhelming probability $1 - \epsilon_{\text{EXT}}(\kappa)$. Thus, both games proceed identically unless the extractor fails $|\Pr[S_6] - \Pr[S_7]| \leq \epsilon_{\text{EXT}}(\kappa)$.

Game 8: As Game 7, but we further modify the Join algorithm when queried for user with index i^* (executed within the SndToU oracle) as follows. Instead of choosing $(\text{usk}_{i^*}, \text{upk}_{i^*}) \leftarrow \text{UKeyGen}(1^\kappa)$, we engage with an EUF-CMA challenger, obtain upk_{i^*} and set $\text{usk}_{i^*} \leftarrow \emptyset$. If any signature is required, we obtain it using the oracle provided by the EUF-CMA challenger.

Transition Game 7 \rightarrow Game 8: This change is conceptual, i.e., $\Pr[S_7] = \Pr[S_6]$.

At this point we have three possibilities if \mathcal{A} outputs a valid forgery.

1. If a signature for $\hat{C}_{J_{i^*}}$ was never requested, \mathcal{A} is an EUF-CMA forger for Σ and the forgery is $(\hat{C}_{J_{i^*}}, \sigma_{J_{i^*}})$. The probability for this to happen is upper bounded by $\epsilon_f(\kappa)$.
2. Otherwise, we know that $\hat{C}_{J_{i^*}}$ is honestly computed by the environment and—by the perfect correctness of Ω —thus contains $\tau/a\hat{P}$, which leaves us two possibilities:
 - (a) If $e(\sigma[1][1], \hat{P}) = e(\sigma[1][2], \tau/a\hat{P})$, \mathcal{A} is an adversary against co-CDHI, since we can obtain $((\tau \cdot 1/aP, P), \sigma') \leftarrow \text{ChgRep}_{\mathcal{R}}(\sigma_1, \rho^{-1}, \text{pk}_{\mathcal{R}})$ and

use τ to output $\tau^{-1} \cdot (\tau \cdot 1/aP) = 1/aP$. The probability for this to happen is upper bounded by $\epsilon_{\text{co-CDH}}(\kappa)$.

- (b) If $e(\sigma[1][1], \hat{P}) \neq e(\sigma[1][2], \tau/a\hat{P})$, \mathcal{A} has produced an opening proof for a statement which is actually not in L_{R_0} . The probability for this to happen is upper bounded by $\epsilon_S(\kappa)$.

Taking the union bound we obtain $\epsilon_{\text{nf8}}(\kappa) \leq \epsilon_f(\kappa) + \epsilon_{\text{co-CDH}}(\kappa) + \epsilon_S(\kappa)$, which yields the following bound for the success probability in Game 1: $\Pr[S_0] \leq q \cdot (\epsilon_{\text{nf8}}(\kappa) + \epsilon_{\text{ZK}_j}(\kappa) + \epsilon_{\text{SIM}}(\kappa) + \epsilon_{\text{guess}}(\kappa) + \epsilon_{\text{EXT}}(\kappa))$, which is negligible.⁵ \square

B.4 Proof of Lemma 1

Proof. We investigate all the properties below.

Perfect completeness. Is straight forward to verify and omitted.

SHVZK. We describe a simulator which outputs transcripts being indistinguishable from real transcripts. First, it chooses $P' \xleftarrow{R} \mathbb{G}_1, \hat{C}_1 \xleftarrow{R} \mathbb{G}_2, \hat{C}_2 \xleftarrow{R} \mathbb{G}_2$. While P' and \hat{C}_1 are identically distributed as in a real transcript, the random choice of \hat{C}_2 is not detectable under DDH in \mathbb{G}_2 which holds in the SXDH setting (more generally under IND-CPA of the used encryption scheme). Then, the simulator chooses $z_1, z_2, c \xleftarrow{R} \mathbb{Z}_p$ and computes $N \leftarrow z_1 \cdot P - c \cdot P', \hat{M}_1 \leftarrow z_2 \cdot \hat{Y} - c \cdot \hat{C}_1, \hat{M}_2 \leftarrow (z_1 + z_2) \cdot \hat{P} - c \cdot \hat{C}_2$. It is easy to see that the transcript $(P', \hat{C}_1, \hat{C}_2, N, \hat{M}_1, \hat{M}_2, z_1, z_2, c)$ represents a valid transcript and its distribution is computationally indistinguishable from a real transcript.

Special soundness. Let us consider that we have two accepting answers (z_1, z_2, c) and (z'_1, z'_2, c') from the prover for distinct challenges $c \neq c'$. Then we have that

$$z_1 - c \cdot \rho = z'_1 - c' \cdot \rho \text{ and } z_2 - c \cdot u = z'_2 - c' \cdot u,$$

and extract a witness as $\rho \leftarrow \frac{z_1 - z'_1}{c - c'}, u \leftarrow \frac{z_2 - z'_2}{c - c'}$.

Quasi-unique responses. The answers z_1 and z_2 are uniquely determined by the word $\hat{Y}, P', \hat{C}_1, \hat{C}_2$, the commitments N, \hat{M}_1, \hat{M}_2 as well as the challenge c (and thus the verification equation). \square

B.5 Proof of Lemma 2

The proof is analogous to [BCC+15], but we re-state it for completeness.

Proof. For simulatability, we observe that the CRS output by `SoK.SimSetup` is identical to the CRS output by `SoK.Setup` and `SoK.SimSign` programs the random oracle to simulate proofs. Simulatability then follows from SHVZK. For extractability we rely on rewinding, special soundness and quasi-unique responses, using the results from [FKMV12]. For straight-line f -extractability, we use the trapdoor τ to decrypt (\hat{C}_1, \hat{C}_2) in the proof transcript and obtain $\rho\hat{P} = f(\rho)$. \square

⁵ We note that we could also write the three cases in the final step as three additional game changes where we abort upon the respective forgeries. However, we opted for this more compact presentation, which also gives us the same bound.

Scheme	Anon.	Signature Size	Signature Cost	Verification Cost	Assumption Type
[BCN ⁺ 10]	CCA ⁻	$3G_1 + 2Z_p$	$1G_T + 3G_1$	$5P + 1G_T + 1G_1$	Interactive
[PS16]	CCA ⁻	$2G_1 + 2Z_p$	$1G_T + 2G_1$	$3P + 1G_T + 1G_1$	GGM
[BBS04]	CPA	$3G_1 + 6Z_p$	$3P + 3G_T + 9G_1$	$5P + 4G_T + 8G_1$	q-Type (non-static)
[BBS04] (prec.)	CPA	$3G_1 + 6Z_p$	$3G_T + 9G_1$	$4G_T + 8G_1$	q-Type (non-static)
This paper	CPA	$1G_2 + 4G_1 + 2Z_p$	$1G_2 + 5G_1$	$5P + 2G_1$	GGM
This paper	CCA2	$3G_2 + 4G_1 + 3Z_p$	$6G_2 + 5G_1$	$5P + 4G_2 + 2G_1$	GGM
This paper (switch)	CCA2	$4G_2 + 3G_1 + 3Z_p$	$5G_2 + 6G_1$	$5P + 2G_2 + 4G_1$	GGM
[DP06]	CCA2	$4G_1 + 5Z_p$	$3P + 3G_T + 4G_1$	$5P + 4G_T + 7G_1$	q-Type (non-static)
[DP06] (prec.)	CCA2	$4G_1 + 5Z_p$	$3G_T + 8G_1$	$1P + 3G_T + 2G_2 + 7G_1$	q-Type (non-static)
[LMPY16]	CCA2	$7G_1 + 3Z_p$	$4P + 2G_T + 16G_1$	$8P + 3G_T + 7G_1$	Standard

Table 2. Comparison of related group signature schemes in the ROM regarding signature size, signing and verification cost, and required hardness assumptions, where, in terms of computational costs, we only count the expensive operations in G_1 , G_2 , and G_T as well as the pairings. The values for [BCN⁺10] and [PS16] are taken from [PS16]. We use ‘CCA⁻’ to denote anonymity in the sense of [BCN⁺10] and note that precomputation in [BBS04, DP06] requires to store extra elements in G_T .