

Policy-Based Sanitizable Signatures[‡]

Kai Samelin¹ and Daniel Slamanig²

¹ TÜV Rheinland i-sec GmbH, Hallbergmoos, Germany

kaispapers@gmail.com

² AIT Austrian Institute of Technology, Vienna, Austria

daniel.slamanig@ait.ac.at

Abstract. Sanitizable signatures are a variant of signatures which allow a single, and signer-defined, sanitizer to modify signed messages in a controlled way without invalidating the respective signature. They turned out to be a versatile primitive, proven by different variants and extensions, e.g., allowing multiple sanitizers or adding new sanitizers one-by-one. However, existing constructions are very restricted regarding their flexibility in specifying potential sanitizers. We propose a different and more powerful approach: Instead of using sanitizers’ public keys directly, we assign attributes to them. Sanitizing is then based on policies, i.e., access structures defined over attributes. A sanitizer can sanitize, if, and only if, it holds a secret key to attributes satisfying the policy associated to a signature, while offering full-scale accountability.

1 Introduction

Unforgeability of a digital signature scheme prevents deriving signatures for a message not explicitly endorsed by the signer. This is a desired property in many use cases of signatures. However, it turned out that certain *controlled* modifications of signed messages are beneficial in many scenarios [ABC⁺15, BPS17, DDH⁺15, GGOT16]. Over the years, different types of signature schemes supporting such modifications have been proposed, including homomorphic signatures [ABC⁺15, BFKW09], redactable signatures [DPSS15, JMSW02, SBZ01], and sanitizable signatures [ACdMT05, BFF⁺09, BFLS10]. In this paper, we focus on sanitizable signatures (3S henceforth). In a nutshell, a *standard* 3S [ACdMT05] allows for altering signer-chosen (so called admissible) blocks of signed messages by a *single* semi-trusted entity, called the sanitizer, which is specified by the signer when generating the signature. The sanitizer holds its own key pair. By using the secret key, the sanitizer can derive modified messages

[‡] The project leading to this work has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 783119 SECREDAS.

with modifiable parts (called admissible blocks) arbitrarily updated, along with corresponding valid signatures. Moreover, given a sanitizable signature, there is a (virtual) entity, dubbed the judge, who can determine whether a signature comes from the original signer or has been sanitized, providing accountability. Even though allowing arbitrary modification of signer-specified blocks seems to give too much power to the sanitizer, 3Ss have proven to be useful in numerous use-cases, as exhaustively discussed by Bilzhaue et al. [BPS17].

After 3Ss were introduced by Ateniese et al. [ACdMT05], they received a lot of attention in the recent past. The first thorough security model was given by Brzuska et al. [BFF⁺09] (later slightly modified by Gong et al. [GQZ10]). Their work was later extended for multiple signers/sanitizers [BFLS09, C JL12], unlinkability (meaning derived signatures cannot be linked to its origin) [BFLS10, BPS13, BL17, BLL⁺19, FKM⁺16], non-interactive public-accountability (every party can determine which party is accountable for a given valid message/signature pair) [BPS12], limiting the sanitizer to signer-chosen values [CJ10, DS15], invisibility (meaning that an outsider cannot determine which blocks of a message are sanitizable) [BCD⁺17, BLL⁺19, CDK⁺17, FH18], the case of strongly unforgeable signatures [KSS15], and generalizations such as merging the functionality from sanitizable and redactable signatures [KPSS18b, KPSS19]. All these extensions make 3Ss suitable for an even broader field of use-cases of (cf. [BPS17] for a discussion), and are directly applicable to our contribution.

In all of the aforementioned work on sanitizable signatures, the sanitizer(s) need(s) to be known *in advance* at signature generation, and there is no possibility to control sanitizing capabilities in a fine-grained way. We note that there is the concept of trapdoor 3Ss [CLM08, LDW13, YSL10]. Although here the signer can grant the possibility to sanitize to different entities even after generating the initial signature, existing constructions do either not provide accountability, a central feature of 3S, or require obtaining the trapdoor from the original signer before sanitizing [LDW13]. This drastically restricts the applicability of 3Ss, their flexibility, and may lead to severe problems when the specified sanitizer is not available.

Motivation and Applications. To illustrate the problem, let us consider an enterprise scenario where policies are associated to different types of documents and documents of some type can be sanitized if the person performing the sanitization fulfills the respective policy. For simplicity, assume that sanitizing should be possible if the sanitizer satisfies the policy

$P = (\text{IT department} \wedge \text{admin}) \vee (\text{team leader})$. Now, let's say that the head of IT department has previously signed a document, e.g., an order, which urgently needs to be sent to reseller but some information needs to be sanitized before, e.g., fixing the number of new PCs ordered. Unfortunately, the original signer is not available, e.g., due to vacation. Now, everyone satisfying P should be able to sanitize. Since this covers a potentially large set of persons, there is no availability issue, and the document can be sent in time. Still, the department head (the "group manager") can control via P who is trusted to sanitize the document if required, and there must be means to determine who performed the sanitization in case of a dispute. Realizing this scenario with the state-of-the-art 3S, such as using a sanitizer key per policy and giving the key to everyone satisfying it clearly destroys accountability, i.e., there is no means identifying the accountable party later on, and thus no satisfying solution can be achieved. To tackle this situation, we introduce a primitive denoted policy-based sanitizable signatures (P3S), that allows to sanitize if, and only if, the attributes associated to a sanitizer satisfy the policy associated to the signature, while at the same time providing accountability. We also want to discuss one application of P3S extending the scope of the one discussed in [DSSS19]. In particular, [DSSS19] discusses an application to updating/rewriting transactions (or more generally speaking objects) in blockchains by selectively replacing the hash function used to aggregate transactions (e.g., within a Merkle-tree) by a novel chameleon hash. This adds flexibility to the initial proposal of a redactable blockchain (where entire blocks can be rewritten) due to Ateniese et al. in [AMVA17]. In [DSSS19], everyone who wants a transaction that can be updated/rewritten can distribute attribute-keys to users who can potentially update the transactions of this entity. Using P3S instead of this novel chameleon hash allows to not only hash transactions/objects but combine it with a signature (as usual for transactions and typically also for other objects in blockchains), we can thus achieve stronger guarantees than in [DSSS19]. In addition to transparency, meaning that no outsider sees whether updates happened (as also achieved in [DSSS19]), using P3S provides accountability, i.e., it can be determined who conducted the update.

Contribution and Our Techniques. We introduce the notion of policy-based sanitizable signatures (P3S). The main idea is the following: At signing, the signer assigns some access-policy P with each generated signature. A sanitizer can sanitize such signatures, if, and only if, that sanitizer has a secret key satisfying the associated policy P . Sanitizers

can obtain new secret keys for some attributes in a dynamic fashion by a special entity named the “group manager”, essentially playing the same role as the “*issuer*” in dynamic group signatures [BSZ05].¹ The reason for this design choice stems from practical considerations: Generated sanitizing keys must only be valid for a single group; In our example mentioned above, the sanitization rights must not work for signatures for another company. However, we also allow that signers and sanitizers can re-use their keys across different groups, e.g., in an enterprise every employee can hold a single key-pair and can participate in multiple groups without generating fresh keys for every group. In our running example, this also means that, e.g., a supplier for our company could sanitize certain signatures using its long-term key (if it received the corresponding secret keys).

We provide a natural formal framework for such P3S by extending the one for 3S. We note that in the case of P3S, with a potentially large sets of sanitizers and different sanitization keys (depending on attributes), make the formal definition much trickier and somewhat involved. Still, we believe that our proposed definitions are clean and easy to comprehend. We also consider a notion analogous to opening-soundness [SSE⁺12]. Moreover, we propose very strict privacy notions, where even (most of) the keys are generated by the adversary, further strengthening already existing definitions [dMPPS14, FF15, KSS15].

Finally, we provide a construction of P3S which we rigorously analyze in the proposed framework. Technically, the heart of our construction is a recent primitive called policy-based chameleon hash (PCH) [DSSS19], which is a trapdoor collision-resistant hash-function, where the hash computation in addition to the message takes a description of a policy as input. Loosely speaking, there are many different trapdoors and collisions can be found if, and only if, a trapdoor satisfying the policy used for the computation of the hash is known. Looking ahead, the PCH proposed in [DSSS19] combines chameleon-hashes with ephemeral trapdoors (CHET) [CDK⁺17] and CCA2-secure ciphertext-policy attribute-based encryption (CP-ABE) scheme. In contrast to the original PCH definition in [DSSS19], however, we have to make some minor, yet important, alterations and show that a modified construction from [DSSS19] satisfies our stronger notions. In this regard, we also strengthen the CH and CHET definitions by Camenisch et al. [CDK⁺17] to also cover keys generated by the adversary. We believe that this strengthened definitions are also useful in many other scenarios.

¹ If wanted, a signer can also be a group manager *simultaneously, without* sacrificing accountability.

The concrete PCH construction then requires some additional tools and tricks; In order to achieve accountability, we use an “OR-trick”, and attach a non-interactive zero-knowledge proof of knowledge, demonstrating that either the signer or a sanitizer performed the signing, or the sanitization, respectively. The expressiveness of the policies supported by the P3S are determined by that of the PCH and in particular by that of the underlying CP-ABE scheme. We chose to build upon the existing PCH framework which covers (monotone) access structures as policies as this seems to be the most interesting setting for practical applications.² For a detailed intuition on the construction, see Sect. 4.

2 Preliminaries

Notation. With $\kappa \in \mathbb{N}$ we denote our security parameter. All algorithms implicitly take 1^κ as an additional input. We write $a \leftarrow A(x)$ if a is assigned to the output of algorithm A with input x . An algorithm is efficient, if it runs in probabilistic polynomial time (PPT) in the length of its input. All algorithms are PPT, if not explicitly mentioned otherwise. If we make the random coins r explicit, we use the notation $a \leftarrow A(x; r)$. Otherwise, we assume that the random coins are drawn internally. For $m = (m^1, m^2, \dots, m^l)$, we call $m^i \in \mathcal{M}$, where $\mathcal{M} = \{0, 1\}^*$, a block. Most algorithms may return a special error symbol $\perp \notin \{0, 1\}^*$, denoting an exception. Returning output ends execution of an algorithm or an oracle. If S is a set, we write $a \leftarrow_r S$ to denote that a is chosen uniformly at random from S . For a list we require that there is an injective, and efficiently reversible, encoding, mapping the list to $\{0, 1\}^*$. A function $\nu : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ is negligible, if it vanishes faster than every inverse polynomial, i.e., $\forall k \in \mathbb{N}, \exists n_0 \in \mathbb{N}$ such that $\nu(n) \leq n^{-k}, \forall n > n_0$.

Assumptions and Primitives. For our construction to work, we need a one-way function (OWF) f , an unforgeable digital signature scheme $\Sigma = \{\text{PPGen}_\Sigma, \text{KGen}_\Sigma, \text{Sign}_\Sigma, \text{Verf}_\Sigma\}$, and an IND-CCA2-secure encryption-scheme $\Pi = \{\text{PPGen}_\Pi, \text{KGen}_\Pi, \text{Enc}_\Pi, \text{Dec}_\Pi, \text{KVrf}_\Pi\}$. Key-verifiability means that for a given public key, exactly one secret key can be found (e.g., Cramer-Shoup (CS) encryption [CS98] in a setting with common group parameters suffices), while KVrf_Π checks whether a given secret key sk belongs to a pk . Moreover, we require a (labeled) simulation-sound extractable non-interactive zero-knowledge proof system $\Omega = \{\text{PPGen}_\Omega,$

² PCHs and P3S could be defined for richer policies, e.g., polynomial sized circuits.

$\text{Prove}_\Omega, \text{Verify}_\Omega$ }, and a recent primitive dubbed policy-based chameleon-hash (PCH), recently introduced by Derler et al. [DSSS19].

For the sake of readability, a somewhat informal Camenisch and Stadler notation [CS97] is used. For example, the notation

$$\pi \leftarrow_r \text{Prove}_\Omega\{(g_1) : C = \text{Enc}_H(g_1)\}(\ell)$$

denotes the computation of a simulation-sound extractable non-interactive zero-knowledge proof (NIZK for short) of the plaintext g_1 contained in C (which is assumed to be public), with a non-malleable attached label $\ell \in \{0, 1\}^*$. Sometimes only “verify π ” is used for verification of a proof π . It is assumed that the public parameters, and the statement to be proven, are also input to the proof system as the label, and are public (all those values are assumed to be part of π as well). This is not made explicit to increase readability.

All primitives, but PCHs, are well-known; We give the full formal definitions of the standard building blocks in App. A, and only fully restate PCHs here. In a nutshell, a $\text{PCH} = (\text{PPGen}_{\text{PCH}}, \text{MKeyGen}_{\text{PCH}}, \text{KGen}_{\text{PCH}}, \text{Hash}_{\text{PCH}}, \text{Verify}_{\text{PCH}}, \text{Adapt}_{\text{PCH}})$ is a trapdoor collision-resistant hash-function, where the hash computation in addition to the message takes a description of a policy as input. Loosely speaking there can be many different trapdoors and collisions can be found if, and only if, a trapdoor satisfying the policy used for the computation of the hash is known.

Before we recall PCHs, we need to define what an access structure is.

Definition 1 (Access Structure). *Let \mathbb{U} denote the universe of attributes. A collection $\mathbb{A} \in 2^{\mathbb{U}} \setminus \{\emptyset\}$ of non-empty sets is an access structure on \mathbb{U} . The sets in \mathbb{A} are called the authorized sets, and the sets not in \mathbb{A} are called the unauthorized sets. A collection $\mathbb{A} \in 2^{\mathbb{U}} \setminus \{\emptyset\}$ is called monotone if $\forall B, C \in \mathbb{A} : \text{if } B \in \mathbb{A} \text{ and } B \subseteq C, \text{ then } C \in \mathbb{A}$.*

Definition 2 (Policy-Based Chameleon-Hashes [DSSS19]). *A policy-based chameleon-hash PCH consists of the following six algorithms $(\text{PPGen}_{\text{PCH}}, \text{MKeyGen}_{\text{PCH}}, \text{KGen}_{\text{PCH}}, \text{Hash}_{\text{PCH}}, \text{Verify}_{\text{PCH}}, \text{Adapt}_{\text{PCH}})$, which are defined as follows.*

$\text{PPGen}_{\text{PCH}}$. *On input a security parameter κ , $\text{PPGen}_{\text{PCH}}$ outputs the public parameters:*

$$\text{pp}_{\text{PCH}} \leftarrow_r \text{PPGen}_{\text{PCH}}(1^\kappa)$$

We assume that pp_{PCH} contains 1^κ and is implicit input to all other algorithms.

$\text{MKeyGen}_{\text{PCH}}$. On input of some global parameters pp_{PCH} , $\text{MKeyGen}_{\text{PCH}}$ outputs the master private and public key $(\text{sk}_{\text{PCH}}, \text{pk}_{\text{PCH}})$ of the scheme:

$$(\text{sk}_{\text{PCH}}, \text{pk}_{\text{PCH}}) \leftarrow_r \text{MKeyGen}_{\text{PCH}}(\text{pp}_{\text{PCH}})$$

KGen_{PCH} . On input a secret key sk_{PCH} and a set of attributes $\mathbb{S} \subseteq \mathbb{U}$ (\mathbb{U} is the universe), the key generation algorithm outputs a secret key $\text{sk}_{\mathbb{S}}$:

$$\text{sk}_{\mathbb{S}} \leftarrow_r \text{KGen}_{\text{PCH}}(\text{sk}_{\text{PCH}}, \mathbb{S})$$

Hash_{PCH} . On input a public key pk_{PCH} , access structure $\mathbb{A} \subseteq 2^{\mathbb{U}}$ and a message m , this algorithm outputs a hash h and some randomness (sometimes referred to as “check value”) r :

$$(h, r) \leftarrow_r \text{Hash}_{\text{PCH}}(\text{pk}_{\text{PCH}}, m, \mathbb{A})$$

$\text{Verify}_{\text{PCH}}$. On input a public key pk , a message m , a hash h , and a randomness r , it outputs a bit $b \in \{1, 0\}$.

$$b \leftarrow \text{Verify}_{\text{PCH}}(\text{pk}_{\text{PCH}}, m, h, r)$$

$\text{Adapt}_{\text{PCH}}$. On input a secret key $\text{sk}_{\mathbb{S}}$, messages m and m' , a hash h , and randomness value r , the adaptation algorithm outputs a new randomness r' :

$$r' \leftarrow_r \text{Adapt}_{\text{PCH}}(\text{pk}_{\text{PCH}}, \text{sk}_{\mathbb{S}}, m, m', h, r)$$

We assume that the KGen_{PCH} outputs \perp , if \mathbb{S} is not contained in \mathbb{U} .

Note, we have added an additional algorithm $\text{PPGen}_{\text{PCH}}$ which outputs some additional global parameters, which was not part of the original description in [DSSS19], as we work in a slightly different setting. For correctness, we require that for all $\kappa \in \mathbb{N}$, for all $\text{pp}_{\text{PCH}} \leftarrow_r \text{PPGen}_{\text{PCH}}(1^\kappa)$, for all $(\text{sk}_{\text{PCH}}, \text{pk}_{\text{PCH}}) \leftarrow_r \text{MKeyGen}_{\text{PCH}}(\text{pp}_{\text{PCH}})$, for all $\mathbb{A} \subseteq 2^{\mathbb{U}}$, for all $\mathbb{S} \in \mathbb{A}$, for all $\text{sk}_{\mathbb{S}} \leftarrow_r \text{KGen}_{\text{PCH}}(\text{sk}_{\text{PCH}}, \mathbb{S})$, for all $m \in \mathcal{M}$, for all $(h, r) \leftarrow_r \text{Hash}_{\text{PCH}}(\text{pk}_{\text{PCH}}, m, \mathbb{A})$, for all $m' \in \mathcal{M}$, for all $r' \leftarrow_r \text{Adapt}_{\text{PCH}}(\text{pk}_{\text{PCH}}, \text{sk}_{\mathbb{S}}, m, m', h, r)$, we have that $1 = \text{Verify}_{\text{PCH}}(\text{pk}_{\text{PCH}}, m, h, r) = \text{Verify}_{\text{PCH}}(\text{pk}_{\text{PCH}}, m', h, r')$.

Furthermore, we require the following security properties, where our notion of indistinguishability below is stronger than the one introduced in [DSSS19]. We also restate the black-box construction from [DSSS19] (with some minor rephrasing and slightly stronger primitives) in App. C. The security proof in our stronger model is given in App. B.

Full Indistinguishability. Informally, indistinguishability requires that it be intractable to decide whether for a chameleon-hash its randomness is fresh or was created using the adaption algorithm. Full indistinguishability even lets the adversary choose the secret key used in the `HashOrAdapt` oracle. The security experiment grants the adversary access to a left-or-right style `HashOrAdapt` oracle and requires that the randomnesses r does not reveal whether it was obtained through `HashPCH` or `AdaptPCH`. The messages and secret keys are adaptively chosen by the adversary.

```

Exp $\mathcal{A}, \text{PCH}$ Flndistinguishability( $\kappa$ )
   $\text{pp}_{\text{PCH}} \leftarrow_r \text{PPGen}_{\text{PCH}}(1^\kappa)$ 
   $b \leftarrow_r \{0, 1\}$ 
   $b^* \leftarrow_r \mathcal{A}^{\text{HashOrAdapt}(\cdot, \cdot, \cdot, \cdot, b)}(\text{pp}_{\text{PCH}})$ 
  where HashOrAdapt on input  $\text{pk}_{\text{PCH}}, m, m', \text{sk}_S, \mathbb{A}, b$ :
     $(h_0, r_0) \leftarrow_r \text{Hash}_{\text{PCH}}(\text{pk}_{\text{PCH}}, m', \mathbb{A})$ 
     $(h_1, r_1) \leftarrow_r \text{Hash}_{\text{PCH}}(\text{pk}_{\text{PCH}}, m, \mathbb{A})$ 
     $r_1 \leftarrow_r \text{Adapt}_{\text{PCH}}(\text{pk}_{\text{PCH}}, \text{sk}_S, m, m', h_1, r_1)$ 
    return  $\perp$ , if  $r_0 = \perp \vee r_1 = \perp$ 
    return  $(h_b, r_b)$ 
  return 1, if  $b = b^*$ 
  return 0

```

Fig. 1: PCH Full Indistinguishability

Definition 3 (PCH Full Indistinguishability). *We say a PCH scheme is fully indistinguishable, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:*

$$\left| \Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{PCH}}^{\text{Flndistinguishability}}(\kappa) = 1 \right] - \frac{1}{2} \right| \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 1.

Insider Collision-Resistance. Insider collision-resistance addresses the requirement that not even insiders who possess secret keys with respect to some attributes can find collisions for hashes which were computed with respect to policies which are not satisfied by their keys (oracle $\text{KGen}'_{\text{PCH}}$). Intuitively, this notion enforces the attribute-based access-control policies, even if the adversary sees collisions for arbitrary attributes (oracles $\text{KGen}''_{\text{PCH}}$ and $\text{Adapt}'_{\text{PCH}}$).

$\mathbf{Exp}_{\mathcal{A}, \text{PCH}}^{\text{CRIns}}(\kappa)$
 $\text{pp}_{\text{PCH}} \leftarrow_r \text{PPGen}_{\text{PCH}}(1^\kappa)$
 $(\text{sk}_{\text{PCH}}, \text{pk}_{\text{PCH}}) \leftarrow_r \text{MKeyGen}_{\text{PCH}}(\text{pp}_{\text{PCH}})$
 $\mathcal{S} = \mathcal{H} = \mathcal{Q} \leftarrow \emptyset$
 $i \leftarrow 0$
 $(m^*, r^*, m'^*, r'^*, h^*) \leftarrow_r \mathcal{A}_{\text{Hash}'_{\text{PCH}}(\text{pk}_{\text{PCH}}, \cdot), \text{KGen}'_{\text{PCH}}(\text{sk}_{\text{PCH}}, \cdot), \text{KGen}''_{\text{PCH}}(\text{sk}_{\text{PCH}}, \cdot), \text{Adapt}'_{\text{PCH}}(\text{pk}_{\text{PCH}}, \cdot, \cdot, \cdot)}(\text{pk}_{\text{PCH}})$
 where $\text{KGen}'_{\text{PCH}}$ on input $\text{sk}_{\text{PCH}}, \mathbb{S}$:
 $\text{sk}_{\mathbb{S}} \leftarrow_r \text{KGen}_{\text{PCH}}(\text{sk}, \mathbb{S})$
 $\mathcal{S} \leftarrow \mathcal{S} \cup \{\mathbb{S}\}$
 return $\text{sk}_{\mathbb{S}}$
 and $\text{KGen}''_{\text{PCH}}$ on input $\text{sk}_{\text{PCH}}, \mathbb{S}$:
 $\text{sk}_{\mathbb{S}} \leftarrow_r \text{KGen}_{\text{PCH}}(\text{sk}, \mathbb{S})$
 $\mathcal{Q} \cup \{(i, \text{sk}_{\mathbb{S}})\}$
 $i \leftarrow i + 1$
 return \perp
 and $\text{Hash}'_{\text{PCH}}$ on input $\text{pk}_{\text{PCH}}, m, \mathbb{A}$:
 $(h, r) \leftarrow_r \text{Hash}_{\text{PCH}}(\text{pk}_{\text{PCH}}, m, \mathbb{A})$
 if $r \neq \perp$, $\mathcal{H} \leftarrow \mathcal{H} \cup \{(h, \mathbb{A}, m)\}$
 return (h, r)
 and $\text{Adapt}'_{\text{PCH}}$ on input $\text{pk}_{\text{PCH}}, m, m', h, r, j$:
 return \perp , if $(j, \text{sk}_{\mathbb{S}}) \notin \mathcal{Q}$ for some $\text{sk}_{\mathbb{S}}$
 $r' \leftarrow_r \text{Adapt}_{\text{PCH}}(\text{pk}_{\text{PCH}}, \text{sk}_{\mathbb{S}}, m, m', h, r)$
 if $r' \neq \perp \wedge (h, \mathbb{A}, m) \in \mathcal{H}$ for some \mathbb{A} , $\mathcal{H} \leftarrow \mathcal{H} \cup \{(h, \mathbb{A}, m')\}$
 return r'
 return 1, if
 $\text{Verify}_{\text{PCH}}(\text{pk}, m^*, h^*, r^*) = \text{Verify}_{\text{PCH}}(\text{pk}, m'^*, h^*, r'^*) = 1 \wedge$
 $(h^*, \mathbb{A}, \cdot) \in \mathcal{H}$, for some $\mathbb{A} \wedge m^* \neq m'^* \wedge \mathbb{A} \cap \mathcal{S} = \emptyset \wedge (h^*, \cdot, m^*) \notin \mathcal{H}$
 return 0

Fig. 2: PCH Insider Collision-Resistance

Definition 4 (PCH Insider Collision-Resistance). *We say a PCH scheme is insider collision-resistant, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:*

$$\Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{PCH}}^{\text{CRIns}}(\kappa) = 1 \right] \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 2.

Uniqueness. We also introduce the new notion of uniqueness for PCHs, which basically requires that it is hard to find different randomness yielding the same hash for an adversarial chosen message and public key.

```

Exp $\mathcal{A}, \text{PCH}$ Uniqueness( $\kappa$ )
  ppPCH  $\leftarrow_r$  PPGenPCH( $1^\kappa$ )
  ( $\text{pk}^*, m^*, r^*, r'^*, h^*$ )  $\leftarrow_r$   $\mathcal{A}(\text{pp}_{\text{PCH}})$ 
  return 1, if VerifyPCH( $\text{pk}^*, m^*, h^*, r^*$ ) = VerifyPCH( $\text{pk}^*, m^*, h^*, r'^*$ ) = 1  $\wedge$   $r^* \neq r'^*$ 
  return 0

```

Fig. 3: PCH Uniqueness

Definition 5 (PCH Uniqueness). *We say a PCH scheme is unique, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:*

$$\Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{PCH}}^{\text{Uniqueness}}(\kappa) = 1 \right] \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 3.

Note, we do not require the outsider collision-resistance notion from [DSSS19].

3 Our Framework for P3Ss

Additional Notation. We need to introduce some additional notation, to make our representation more compact. Our notation is taken from existing work, making reading more accessible [BCD⁺17, BFF⁺09, CDK⁺17]. The variable \mathbf{A} contains the set of indices of the modifiable blocks, as well as l denoting the total number of blocks in the message m . We write $\mathbf{A}(m) = 1$, if \mathbf{A} is valid w.r.t. m , i.e., \mathbf{A} contains a fitting l , i.e., the correct length of m , and the indices of the admissible blocks are actually part of m . For example, let $\mathbf{A} = (\{1, 2, 3, 5\}, 5)$. Then, m must contain five blocks, and all but the fourth can be modified. If we write $m^i \in \mathbf{A}$, we mean that m^i is admissible. We also use $m_{\mathbf{A}}$ for the list of blocks in m which are admissible w.r.t. \mathbf{A} . Likewise, we use $m_{\neg \mathbf{A}}$ for the list of blocks of m which are not admissible w.r.t. to \mathbf{A} . Moreover, \mathbf{M} is a set containing pairs (i, m^i) for those blocks that are modified, meaning that m^i is replaced with m'^i . We write $\mathbf{M}(\mathbf{A}) = 1$, if \mathbf{M} is valid w.r.t. \mathbf{A} , meaning that the indices to be modified are contained in \mathbf{A} , i.e., admissible.

Definitional Framework. We now introduce our definitional framework. It is based on existing work [BCD⁺17, BFF⁺09, CDK⁺17]. The main idea is following the line of reasoning of group signatures. Namely, a designated entity, which we name “the group manager” generates a key

pair for its group. The group manager can use its secret key to assign secret keys to sanitizers which are identified by their own key pair. In contrast, signers can create signatures for a signer-chosen group, identified by a public key. Moreover, signers do not require any prior interaction, i.e., knowledge of the group public-key is sufficient, which is a major difference to group signatures, and any sanitizer “authorized” by the manager of that group can then sanitize the generated signatures. Moreover, in contrast to group signatures, *only the signer* can decide which party has generated a signature, essentially it is also the “opener” in group signatures, but the group manager has no opening capabilities. These proofs, however, can be verified by anyone. We keep the wording of the algorithms mostly consistent with existing work to ease readability [BFF⁺09].

Definition 6 (P3S). *A sanitizable signature with attribute-based sanitizing P3S consists of the algorithms $\{\text{ParGen}_{\text{P3S}}, \text{Setup}_{\text{P3S}}, \text{KGenSig}_{\text{P3S}}, \text{KGenSan}_{\text{P3S}}, \text{Sign}_{\text{P3S}}, \text{AddSan}_{\text{P3S}}, \text{Sanitize}_{\text{P3S}}, \text{Verify}_{\text{P3S}}, \text{Proof}_{\text{P3S}}, \text{Judge}_{\text{P3S}}\}$ such that:*

$\text{ParGen}_{\text{P3S}}$. *The algorithm $\text{ParGen}_{\text{P3S}}$ generates the public parameters:*

$$\text{pp}_{\text{P3S}} \leftarrow_r \text{ParGen}_{\text{P3S}}(1^\kappa)$$

We assume that pp_{P3S} contains 1^κ and is implicit input to all other algorithms.

$\text{Setup}_{\text{P3S}}$. *The algorithm $\text{Setup}_{\text{P3S}}$ outputs the global public key pk_{P3S} of a P3S, and some master secret key sk_{P3S} , i.e., it generates the group manager’s key pair:*

$$(\text{sk}_{\text{P3S}}, \text{pk}_{\text{P3S}}) \leftarrow_r \text{Setup}_{\text{P3S}}(\text{pp}_{\text{P3S}})$$

$\text{KGenSig}_{\text{P3S}}$. *The algorithm $\text{KGenSig}_{\text{P3S}}$ generates a key-pair for a signer:*

$$(\text{sk}_{\text{P3S}}^{\text{Sig}}, \text{pk}_{\text{P3S}}^{\text{Sig}}) \leftarrow_r \text{KGenSig}_{\text{P3S}}(\text{pp}_{\text{P3S}})$$

$\text{KGenSan}_{\text{P3S}}$. *The algorithm $\text{KGenSan}_{\text{P3S}}$ generates a key-pair for a sanitizer:*

$$(\text{sk}_{\text{P3S}}^{\text{San}}, \text{pk}_{\text{P3S}}^{\text{San}}) \leftarrow_r \text{KGenSan}_{\text{P3S}}(\text{pp}_{\text{P3S}})$$

Sign_{P3S} . *The algorithm Sign_{P3S} generates a signature σ , on input of a master public key pk_{P3S} , a secret key $\text{sk}_{\text{P3S}}^{\text{Sig}}$, a message m , \mathbb{A} , and some access structure \mathbb{A} :*

$$\sigma \leftarrow_r \text{Sign}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, m, \mathbb{A}, \mathbb{A})$$

$\text{AddSan}_{\text{P3S}}$. The algorithm $\text{AddSan}_{\text{P3S}}$ allows to the group manager to generate a secret sanitizing key $\text{sk}_{\mathbb{S}}$ for a particular sanitizer, on input of sk_{P3S} , a public key $\text{pk}_{\text{P3S}}^{\text{San}}$, and some set of attributes $\mathbb{S} \subseteq \mathbb{U}$:

$$\text{sk}_{\mathbb{S}} \leftarrow_r \text{AddSan}_{\text{P3S}}(\text{sk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{San}}, \mathbb{S})$$

$\text{Verify}_{\text{P3S}}$. The deterministic algorithm $\text{Verify}_{\text{P3S}}$ allows to verify a signature σ on input of a master public key pk_{P3S} , a signer public key $\text{pk}_{\text{P3S}}^{\text{Sig}}$, and a message m . It outputs a decision $b \in \{0, 1\}$:

$$b \leftarrow \text{Verify}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \sigma, m)$$

$\text{Sanitize}_{\text{P3S}}$. The algorithm $\text{Sanitize}_{\text{P3S}}$ allows to derive a new signature on input of a master public key pk_{P3S} , a signer's public key $\text{pk}_{\text{P3S}}^{\text{Sig}}$, a sanitizer's secret key $\text{sk}_{\text{P3S}}^{\text{San}}$, a token $\text{sk}_{\mathbb{S}}$, some modification instruction M , a message m , and a signature σ :

$$(\sigma', m') \leftarrow_r \text{Sanitize}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \text{sk}_{\text{P3S}}^{\text{San}}, \text{sk}_{\mathbb{S}}, m, \sigma, M)$$

$\text{Proof}_{\text{P3S}}$. The algorithm $\text{Proof}_{\text{P3S}}$ allows to generate a proof π_{P3S} and some public pk , used by the next algorithm, to find the accountable party, on input of a master public key pk_{P3S} , a signer's secret key $\text{sk}_{\text{P3S}}^{\text{Sig}}$, a signature σ , and a message m :

$$(\pi_{\text{P3S}}, \text{pk}) \leftarrow_r \text{Proof}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, \sigma, m)$$

$\text{Judge}_{\text{P3S}}$. The algorithm $\text{Judge}_{\text{P3S}}$ allows to verify whether a proof π_{P3S} is valid. The inputs are a master public key pk_{P3S} , a signer's public key $\text{pk}_{\text{P3S}}^{\text{Sig}}$, some other public key pk , a proof π_{P3S} , a signature σ , and a message m . It outputs a decision $b \in \{0, 1\}$, stating whether π_{P3S} is a valid proof that the holder of pk is accountable for σ :

$$b \leftarrow_r \text{Judge}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \text{pk}, \pi_{\text{P3S}}, \sigma, m)$$

For each P3S it is required that the correctness properties hold. In particular, it is required that for all $\kappa \in \mathbb{N}$, for all $\text{pp}_{\text{P3S}} \leftarrow_r \text{ParGen}_{\text{P3S}}(1^\kappa)$, for all $(\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}) \leftarrow_r \text{Setup}_{\text{P3S}}(\text{pp}_{\text{P3S}})$, for all $(\text{sk}_{\text{P3S}}^{\text{Sig}}, \text{pk}_{\text{P3S}}^{\text{Sig}}) \leftarrow_r \text{KGenSig}_{\text{P3S}}(\text{pp}_{\text{P3S}})$, for all $l \in \mathbb{N}$, for all $m \in \mathcal{M}^l$, for all $\mathbb{A} \in 2^{\mathbb{U}}$, for all $A \in \{A_i \mid A_i(m) = 1\}$, for all $\sigma \leftarrow_r \text{Sign}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, m, A, \mathbb{A})$, we have that $\text{Verify}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \sigma, m) = 1$ and for all $(\pi_{\text{P3S}}, \text{pk}) \leftarrow_r \text{Proof}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, \sigma, m)$ we have that $\text{Judge}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \text{pk}, \pi_{\text{P3S}}, \sigma, m) = 1$ and $\text{pk} = \text{pk}_{\text{P3S}}^{\text{Sig}}$. We also require that for all $(\text{sk}_{\text{P3S}}^{\text{San}}, \text{pk}_{\text{P3S}}^{\text{San}}) \leftarrow_r \text{KGenSan}_{\text{P3S}}(\text{pp}_{\text{P3S}})$,

for all $\mathbb{S} \in \mathbb{A}$, for all $\text{sk}_{\mathbb{S}} \leftarrow_r \text{AddSan}_{\text{P3S}}(\text{sk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{San}}, \mathbb{S})$, for all $M \in \{M_i \mid M_i(\mathbb{A}) = 1\}$, for all $(\sigma', m') \leftarrow_r \text{Sanitize}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \text{sk}_{\text{P3S}}^{\text{San}}, \text{sk}_{\mathbb{S}}, m, \sigma, M)$ we have that $\text{Verify}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \sigma', m') = 1$ and that for all $(\pi'_{\text{P3S}}, \text{pk}') \leftarrow_r \text{Proof}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, \sigma', m')$, we have that $\text{Judge}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \text{pk}_{\text{P3S}}^{\text{San}}, \pi'_{\text{P3S}}, \sigma', m') = 1$ and $\text{pk}' = \text{pk}_{\text{P3S}}^{\text{San}}$.

Security Definitions. We now introduce our security definitions. To increase readability, we keep the naming close to the already existing definitions for standard 3Ss [BFF⁺09]. However, due to the increased expressiveness of our new primitive, this is not always possible. Namely, we require new unforgeability and privacy definitions not considered before. This also has the effect that the implications and separations by Brzuska et al. [BFF⁺09] have to be revisited.

Overview. We first briefly introduce each security notion to ease understanding of the formal definitions given afterwards.

- **Unforgeability.** Unforgeability requires that an adversary cannot (except with negligible probability) generate any valid signature, if it does not hold enough attributes to do so. We explicitly include the case that the adversary can be group manager of other groups, but the challenge one.
- **Immutability.** Immutability requires that an adversarial group manager cannot (except with negligible probability) create signatures with altered immutable parts. This also includes appending or removing blocks.
- **Privacy.** Privacy requires that an adversary does not learn (except with negligible probability) anything about sanitized parts, even if it can generate all keys.
- **Transparency.** Transparency requires that an adversary cannot decide (except with negligible probability) whether it sees a freshly signed signature or a sanitized one, even if it can generate all keys, but the signer's one.
- **Pseudonymity.** Pseudonymity requires that an adversary does not learn (except with negligible probability) which party is accountable for a given sanitized signature, even if it can generate all keys, but the signer's one.
- **Signer-Accountability.** Signer-Accountability requires that an adversary cannot (except with negligible probability) blame an honest sanitizer for a signature it did not create, even if it can generate all keys but the sanitizer's one.

- **Sanitizer-Accountability.** Sanitizer-Accountability requires that an adversary cannot (except with negligible probability) blame an honest signer for a signature it did not create, even if it can generate all keys but the signer’s one.
- **Proof-Soundness.** Proof-Soundness requires that an adversary cannot (except with negligible probability) generate a proof for an adverserially chosen signature/message pair that points to different entities, even if it can generate all keys.
- **Traceability.** Traceability requires that an adversary cannot (except with negligible probability) generate a verifying signature such that an honest signer cannot identify the accountable party, even if it can generate all keys, but the signer’s one.

Unforgeability. The property of unforgeability prohibits that an adversary, which is not a signer, or the entity holding sk_{P3S} , i.e., the group manager, can generate any validating signature which verifies for honestly generated keys. This also includes messages for which the adversary does not hold enough attributes for, even if it sees sanitizations of such signatures. We define it in such a way that $(\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}})$, and $(\text{sk}_{\text{P3S}}^{\text{Sig}}, \text{pk}_{\text{P3S}}^{\text{Sig}})$, are generated honestly. The adversary gets access to the following oracles: (1) $\text{Sign}'_{\text{P3S}}$ (where it can even use different pk_{P3S} s, which models the case that secret signing keys can be re-used across multiple “groups”), (2) GetSan which generates a new sanitizer (tracked by \mathcal{S}), (3) $\text{AddSan}'_{\text{P3S}}$ which allows to decide which attributes a given sanitizer holds (tracked by \mathcal{R}), (4) $\text{Sanitize}'_{\text{P3S}}$ which allows sanitizing signatures for an honest sanitizer (generated by GetSan) for the challenge group, and (5) $\text{Sanitize}''_{\text{P3S}}$ which allows sanitizing for signatures from any other group (i.e., where the adversary is the group manager). The adversary wins, if it can generate a valid signature for the defined group which has never been output by either $\text{Sign}'_{\text{P3S}}$ or $\text{Sanitize}'_{\text{P3S}}$ (tracked by the set \mathcal{M} ; Note, this set may be exponential in size, but membership is trivial to decide by checking whether the element could have been derived using \mathbf{A} and \mathbf{A}), and the adversary \mathcal{A} does not hold enough attributes itself.

Definition 7 (P3S Unforgeability). *We say a P3S scheme is unforgeable, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:*

$$\Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{P3S}}^{\text{Unforgeability}}(\kappa) = 1 \right] \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 4.

$\mathbf{Exp}_{\mathcal{A}, P3S}^{\text{Unforgeability}}(\kappa)$
 $\text{pp}_{P3S} \leftarrow_r \text{ParGen}_{P3S}(1^\kappa)$
 $(\text{sk}_{P3S}, \text{pk}_{P3S}) \leftarrow_r \text{Setup}_{P3S}(\text{pp}_{P3S})$
 $(\text{sk}_{P3S}^{\text{Sig}}, \text{pk}_{P3S}^{\text{Sig}}) \leftarrow_r \text{KGenSig}_{P3S}(\text{pp}_{P3S})$
 $\mathcal{Q} = \mathcal{S} = \mathcal{R} = \mathcal{M} = \mathcal{Z} \leftarrow \emptyset$
 $i \leftarrow 0$
 $(m^*, \sigma^*) \leftarrow_r \mathcal{A}_{\text{Sanitize}'_{P3S}(\cdot, \text{sk}_{P3S}^{\text{Sig}}, \cdot, \cdot), \text{GetSan}(\cdot), \text{AddSan}'_{P3S}(\text{sk}_{P3S}, \cdot, \cdot), \text{Sanitize}'_{P3S}(\text{pk}_{P3S}, \cdot, \cdot, \cdot, \cdot)}(\text{pk}_{P3S}, \text{pk}_{P3S}^{\text{Sig}})$
 where Sign'_{P3S} on input $\text{pk}'_{P3S}, \text{sk}_{P3S}^{\text{Sig}}, m, A, \mathbb{A}$:
 $\sigma \leftarrow_r \text{Sign}_{P3S}(\text{pk}'_{P3S}, \text{sk}_{P3S}^{\text{Sig}}, m, A, \mathbb{A})$
 if $\text{pk}'_{P3S} = \text{pk}_{P3S} \wedge \sigma \neq \perp$:
 $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(\sigma, m, \mathbb{A}, A)\}$
 if $\mathbb{A} \in \mathcal{R}, \mathcal{M} \leftarrow \mathcal{M} \cup \{M(m) \mid M(A) = 1\}$
 return σ
 and GetSan :
 $(\text{sk}_{P3S}^{\text{San}}, \text{pk}_{P3S}^{\text{San}}) \leftarrow_r \text{KGenSan}_{P3S}(\text{pp}_{P3S})$
 $\mathcal{S} \leftarrow \mathcal{S} \cup \{(\text{sk}_{P3S}^{\text{San}}, \text{pk}_{P3S}^{\text{San}})\}$
 return $\text{pk}_{P3S}^{\text{San}}$
 and AddSan'_{P3S} on input $\text{sk}_{P3S}, \text{pk}_{P3S}^{\text{San}}, \mathbb{S}$
 if $\neg \exists(\cdot, \text{pk}_{P3S}^{\text{San}}) \in \mathcal{S}$:
 $\text{sk}_{\mathbb{S}} \leftarrow_r \text{AddSan}_{P3S}(\text{sk}_{P3S}, \text{pk}_{P3S}^{\text{San}}, \mathbb{S})$
 return \perp , if $\text{sk}_{\mathbb{S}} = \perp$
 $\mathcal{R} \leftarrow \mathcal{R} \cup \{\mathbb{S}\}$
 for all $(\sigma_i, m_i, \mathbb{A}_i, A_i) \in \mathcal{Q}$, where $\mathbb{S} \in \mathbb{A}_i, \mathcal{M} \cup \{M(m_i) \mid M(A_i) = 1\}$
 return $\text{sk}_{\mathbb{S}}$
 $\text{sk}_{\mathbb{S}} \leftarrow_r \text{AddSan}_{P3S}(\text{sk}_{P3S}, \text{pk}_{P3S}^{\text{San}}, \mathbb{S})$
 $\mathcal{Z} \leftarrow \mathcal{Z} \cup \{(i, \text{sk}_{\mathbb{S}})\}$
 $i \leftarrow i + 1$
 return $(i - 1, \text{pk}_{P3S}^{\text{San}})$
 and $\text{Sanitize}'_{P3S}$ on input $\text{pk}_{P3S}, \text{pk}_{P3S}^{\text{Sig}}, \text{pk}_{P3S}^{\text{San}}, j, m, \sigma, M$:
 return \perp , if $\neg \exists(\text{sk}_{P3S}^{\text{San}}, \text{pk}_{P3S}^{\text{San}}) \in \mathcal{S}$ for some $\text{sk}_{P3S}^{\text{San}}$
 $(\sigma', m') \leftarrow_r \text{Sanitize}_{P3S}(\text{pk}_{P3S}, \text{pk}_{P3S}^{\text{Sig}}, \text{sk}_{P3S}^{\text{San}}, \text{sk}_{\mathbb{S}}, m, \sigma, M)$
 where $\text{sk}_{\mathbb{S}}$ is taken from $(j, \text{sk}_{\mathbb{S}}) \in \mathcal{Z}$
 if $\sigma' \neq \perp$:
 $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(\sigma', m', \perp, \perp)\}$
 return σ'
 and $\text{Sanitize}''_{P3S}$ on input $\text{pk}'_{P3S}, \text{pk}_{P3S}^{\text{Sig}}, \text{pk}_{P3S}^{\text{San}}, \text{sk}_{\mathbb{S}}, m, \sigma, M$:
 return \perp , if $\neg \exists(\text{sk}_{P3S}^{\text{San}}, \text{pk}_{P3S}^{\text{San}}) \in \mathcal{S} \vee \text{pk}'_{P3S} = \text{pk}_{P3S}$
 $(\sigma', m') \leftarrow_r \text{Sanitize}_{P3S}(\text{pk}'_{P3S}, \text{pk}_{P3S}^{\text{Sig}}, \text{sk}_{P3S}^{\text{San}}, \text{sk}_{\mathbb{S}}, m, \sigma, M)$
 return σ'
 return 0, if $\text{Verify}_{P3S}(\text{pk}_{P3S}, \text{pk}_{P3S}^{\text{Sig}}, \sigma^*, m^*) = 0 \vee m^* \in \mathcal{M}$
 return 1, if $(\sigma^*, m^*, \cdot, \cdot) \notin \mathcal{Q}$
 return 0

Fig. 4: P3S Unforgeability

Immutability. The above unforgeability definition assumes that the holder of sk_{P3S} (the group manager) is honest. If this is not the case, however, the adversary can generate its own key pair for a sanitizer and can generate $\text{sk}_{\mathbb{S}}$ for any attribute-set it likes. Still, in such a case, we want to prohibit that an adversary generates any signatures which are outside the span the honest signer has endorsed for *any* combination of attributes. This is captured by the immutability definition — if a block is marked as non-admissible by a signer, no one must be able to change this block. This also includes that an adversary must not be able to redact or append a block. Clearly, we cannot limit the adversary to change admissible blocks, as it can grant sanitizing rights to itself.

This is modeled in such a way that the challenger draws pp_{P3S} honestly, along with a key-pair for the signer. The adversary only receives pp_{P3S} and $\text{pk}_{\text{P3S}}^{\text{Sig}}$. Then, the adversary gains adaptive access to signing-oracle (where the adversary can choose pk_{P3S} , m , \mathbb{A} , \mathbb{A} , but not $\text{sk}_{\text{P3S}}^{\text{Sig}}$), and access to a proof-oracle. We keep a set \mathcal{M} which contains all possible messages which can “legally” be derived by the adversary (bound to pk_{P3S} , also chosen by the adversary, and tracked by \mathcal{M} ; Again, this set may be exponential in size, but membership is trivial to decide). If, and only if, the adversary finds a valid signature σ^* w.r.t. $\text{pk}_{\text{P3S}}^{\text{Sig}}$ and pk^* , which could never been derived from any input, it wins.

```

Exp $\mathcal{A}, \text{P3S}$ Immutability( $\kappa$ )
   $\text{pp}_{\text{P3S}} \leftarrow_r \text{ParGen}_{\text{P3S}}(1^\kappa)$ 
   $(\text{sk}_{\text{P3S}}^{\text{Sig}}, \text{pk}_{\text{P3S}}^{\text{Sig}}) \leftarrow_r \text{KGenSig}_{\text{P3S}}(\text{pp}_{\text{P3S}})$ 
   $\mathcal{M} \leftarrow \emptyset$ 
   $(\text{pk}^*, \sigma^*, m^*) \leftarrow_r \mathcal{A}^{\text{Sign}'_{\text{P3S}}(\cdot, \text{sk}_{\text{P3S}}^{\text{Sig}}, \cdot, \cdot, \cdot), \text{Proof}_{\text{P3S}}(\cdot, \text{sk}_{\text{P3S}}^{\text{Sig}}, \cdot, \cdot)}(\text{pk}_{\text{P3S}}^{\text{Sig}})$ 
  where  $\text{Sign}'_{\text{P3S}}$  on input  $\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, m, \mathbb{A}, \mathbb{A}$ :
     $\sigma \leftarrow_r \text{Sign}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, m, \mathbb{A}, \mathbb{A})$ 
    return  $\perp$ , if  $\sigma = \perp$ 
     $\mathcal{M} \cup \{(\text{pk}_{\text{P3S}}, \text{M}(m)) \mid \text{M}(\mathbb{A}) = 1\}$ 
    return  $\sigma$ 
  return 1, if:
     $\text{Verify}_{\text{P3S}}(\text{pk}^*, \text{pk}_{\text{P3S}}^{\text{Sig}}, \sigma^*, m^*) = 1 \wedge (\text{pk}^*, m^*) \notin \mathcal{M}$ 
  return 0

```

Fig. 5: P3S Immutability

Definition 8 (P3S Immutability). *We say a P3S scheme is immutable, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:*

$$\Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{P3S}}^{\text{Immutability}}(\kappa) = 1 \right] \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 5.

Privacy. Privacy prohibits that an adversary can derive any useful information from a sanitized signature. We define a very strong version, where all values can be generated by the adversary, making our definition even stronger than existing ones [dMPPS14, FF15].

In more detail, the challenger draws a bit $b \leftarrow_r \{0, 1\}$, while the parameters pp_{P3S} are generated honestly. The adversary gains access to a LoRSanit -oracle, where it can input $\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, \text{sk}_{\text{P3S}}^{\text{San}}, \mathbb{A}, m_0, m_1, M_0, M_1, A$, and $\text{sk}_{\mathbb{S}}$ (b is input by the challenger). The oracle then signs m_b with A and \mathbb{A} . Then, the resulting signature is sanitized to $M_b(m_b)$, while $M_0(m_0) = M_1(m_1)$ must hold to prevent trivial attacks. The goal of the adversary is to guess the bit b .

We stress that this definition seems to be overly strong. However, it also preserves privacy in case of bad randomness at key generation, completely leaked keys, and even corrupt group managers.

```

Exp $\mathcal{A}, \text{P3S}$ Privacy( $\kappa$ )
   $\text{pp}_{\text{P3S}} \leftarrow_r \text{ParGen}_{\text{P3S}}(1^\kappa)$ 
   $b \leftarrow_r \{0, 1\}$ 
   $b^* \leftarrow_r \mathcal{A}^{\text{LoRSanit}(\dots, \dots, \dots, b)}(\text{pp}_{\text{P3S}})$ 
  where  $\text{LoRSanit}$  on input of  $\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, \text{sk}_{\text{P3S}}^{\text{San}}, \mathbb{A}, m_0, m_1, M_0, M_1, A, \text{sk}_{\mathbb{S}}, b$ :
     $\sigma \leftarrow_r \text{Sign}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, m_b, A, \mathbb{A})$ 
    for  $b' \in \{0, 1\}$ ,  $(\sigma'_{b'}, \cdot) \leftarrow_r \text{Sanitize}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \text{sk}_{\text{P3S}}^{\text{San}}, \text{sk}_{\mathbb{S}}, m_{b'}, \sigma, M_{b'})$ 
    return  $\perp$ , if  $\sigma'_0 = \perp \vee \sigma'_1 = \perp \vee A(m_0) = 0 \vee$ 
       $A(m_1) = 0 \vee M_0(m_0) \neq M_1(m_1)$ 
    return  $\sigma'_b$ 
  return 1, if  $b = b^*$ 
  return 0

```

Fig. 6: P3S Privacy

Definition 9 (P3S Privacy). *We say a P3S scheme is private, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:*

$$\left| \Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{P3S}}^{\text{Privacy}}(\kappa) = 1 \right] - \frac{1}{2} \right| \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 6.

Transparency. Transparency prohibits that an adversary can decide whether a signature is fresh or the result of a sanitization. As for privacy, we define a very strong version, where all values, but the signer’s key pair $(\text{sk}_{\text{P3S}}^{\text{Sig}}, \text{pk}_{\text{P3S}}^{\text{Sig}})$, can be generated by the adversary, making our definition even stronger than existing ones [dMPPS14, FF15, KSS15]. The reason why the signer’s key pair must be generated honestly is that the signer can always pinpoint the accountable party due to correctness.

In more detail, the challenger draws a bit $b \leftarrow_r \{0, 1\}$, while the parameters pp_{P3S} and the signer’s key pair $(\text{sk}_{\text{P3S}}^{\text{Sig}}, \text{pk}_{\text{P3S}}^{\text{Sig}})$ are generated honestly. The adversary gains access to three oracles: Sign_{P3S} , SignOrSanit , and $\text{Proof}'_{\text{P3S}}$. The Sign_{P3S} -oracle allows the adversary to generate new signatures; the only fixed input is $\text{sk}_{\text{P3S}}^{\text{Sig}}$. The SignOrSanit -oracle is the challenge oracle. It allows the adversary \mathcal{A} to input $\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{San}}, \mathbb{A}, m, \mathbb{M}, \mathbb{A}$, and $\text{sk}_{\mathbb{S}}$ (b and $\text{sk}_{\text{P3S}}^{\text{Sig}}$ are input by the challenger). The oracle then signs m with \mathbb{A} and \mathbb{A} . Then, the resulting signature is sanitized to $\mathbb{M}(m)$. If $b = 1$, however, a fresh signature on $\mathbb{M}(m)$ is generated. The resulting signature is returned to the adversary. However, we also log the signatures generated by this oracle in a list \mathcal{Q} . The list \mathcal{Q} is required to prohibit that the adversary \mathcal{A} can generate a proof using the $\text{Proof}'_{\text{P3S}}$ -oracle with signatures generated by the SignOrSanit -oracle, which directly returns the accountable party. Thus, the adversary can only input $\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, \sigma, m$ for which $(\text{pk}_{\text{P3S}}, \sigma, m)$ was never input/output to the SignOrSanit -oracle. The goal of the adversary is to guess the bit b .

We stress that this definition also seems to be overly strong. However, it also preserves transparency in case of bad randomness at key generation, leaked keys, and even corrupt group managers.

Definition 10 (P3S Transparency). *We say a P3S scheme is transparent, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:*

$$\left| \Pr \left[\text{Exp}_{\mathcal{A}, \text{P3S}}^{\text{Transparency}}(\kappa) = 1 \right] - \frac{1}{2} \right| \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 7.

Pseudonymity. Pseudonymity prohibits that an adversary can decide which sanitizer actually is responsible for a given signature, if it does not have access to $\text{sk}_{\text{P3S}}^{\text{Sig}}$. This is related to the anonymity of group signatures [CvH91]. We formalize it in the following way. The challenger draws a bit $b \leftarrow_r \{0, 1\}$, generates the public parameters pp_{P3S} and the

$\text{Exp}_{\mathcal{A}, \text{P3S}}^{\text{Transparency}}(\kappa)$
 $\text{pp}_{\text{P3S}} \leftarrow_r \text{ParGen}_{\text{P3S}}(1^\kappa)$
 $(\text{sk}_{\text{P3S}}^{\text{Sig}}, \text{pk}_{\text{P3S}}^{\text{Sig}}) \leftarrow_r \text{KGenSig}_{\text{P3S}}(\text{pp}_{\text{P3S}})$
 $b \leftarrow_r \{0, 1\}$
 $\mathcal{Q} \leftarrow \emptyset$
 $b^* \leftarrow_r \mathcal{A}^{\text{Sign}_{\text{P3S}}(\cdot, \text{sk}_{\text{P3S}}^{\text{Sig}}, \cdot, \cdot), \text{SignOrSanit}(\cdot, \text{sk}_{\text{P3S}}^{\text{Sig}}, \cdot, \cdot, \cdot, b), \text{Proof}'_{\text{P3S}}(\cdot, \text{sk}_{\text{P3S}}^{\text{Sig}}, \cdot, \cdot)}(\text{pk}_{\text{P3S}}^{\text{Sig}})$
 where SignOrSanit on input of $\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, \text{sk}_{\text{P3S}}^{\text{San}}, \mathbb{A}, m, \text{M}, \text{A}, \text{sk}_{\mathbb{S}}, b$:
 $\sigma \leftarrow_r \text{Sign}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, m, \text{A}, \mathbb{A})$
 $(\sigma', m') \leftarrow_r \text{Sanitize}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \text{sk}_{\text{P3S}}^{\text{San}}, \text{sk}_{\mathbb{S}}, m, \sigma, \text{M})$
 if $b = 1$:
 $\sigma' \leftarrow_r \text{Sign}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, m', \text{A}, \mathbb{A})$
 $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(\text{pk}_{\text{P3S}}, \sigma', m')\}$
 return σ'
 and $\text{Proof}'_{\text{P3S}}$ on input of $\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, \sigma, m$:
 return \perp , if $(\text{pk}_{\text{P3S}}, \sigma, m) \in \mathcal{Q}$
 return $\text{Proof}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, \sigma, m)$
 return 1, if $b = b^*$
 return 0

Fig. 7: P3S Transparency

signer's key pair $(\text{sk}_{\text{P3S}}^{\text{Sig}}, \text{pk}_{\text{P3S}}^{\text{Sig}})$ honestly. The adversary gains access to three oracles: Sign_{P3S} , LoRSanit , and $\text{Proof}'_{\text{P3S}}$. The Sign_{P3S} -oracle allows the adversary to generate new signatures; the only fixed input is $\text{sk}_{\text{P3S}}^{\text{Sig}}$. The LoRSanit -oracle is the challenge oracle. It allows the adversary \mathcal{A} to input $\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \text{sk}_{\text{P3S},0}^{\text{San}}, \text{sk}_{\text{P3S},1}^{\text{San}}, \text{sk}_{\mathbb{S}0}, \text{sk}_{\mathbb{S}1}, m$, and σ (b and $\text{sk}_{\text{P3S}}^{\text{Sig}}$ are input by the challenger). The oracle then signs m with A and \mathbb{A} . Then, the resulting signature is sanitized to $\text{M}(m)$, using keys $\text{sk}_{\text{P3S},b}^{\text{San}}$ and $\text{sk}_{\mathbb{S},b}$. The resulting signature is given to the adversary. As done for transparency, we also log the signatures generated by this oracle in a list \mathcal{Q} . The list \mathcal{Q} is required to prohibit that the adversary \mathcal{A} wants to generate a proof using the $\text{Proof}'_{\text{P3S}}$ -oracle with signatures generated by the LoRSanit -oracle, which clearly contradicts pseudonymity. Thus, the adversary can only input $\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, \sigma, m$ for which $(\text{pk}_{\text{P3S}}, \sigma, m)$ was never input/output to the LoRSanit -oracle. The goal of the adversary is to guess the bit b .

Again, we stress that this definition also seems to be overly strong. However, as also done for group signatures, secrets keys may leak over time. This definition protects even against bad randomness at key generation.

$\mathbf{Exp}_{\mathcal{A}, \text{P3S}}^{\text{Pseudonymity}}(\kappa)$
 $\text{pp}_{\text{P3S}} \leftarrow_r \text{ParGen}_{\text{P3S}}(1^\kappa)$
 $(\text{sk}_{\text{P3S}}^{\text{Sig}}, \text{pk}_{\text{P3S}}^{\text{Sig}}) \leftarrow_r \text{KGenSig}_{\text{P3S}}(\text{pp}_{\text{P3S}})$
 $\mathcal{Q} \leftarrow \emptyset$
 $b \leftarrow_r \{0, 1\}$
 $b^* \leftarrow_r \mathcal{A}^{\text{Sign}_{\text{P3S}}(\cdot, \text{sk}_{\text{P3S}}^{\text{Sig}}, \cdot, \cdot), \text{Proof}'_{\text{P3S}}(\cdot, \text{sk}_{\text{P3S}}^{\text{Sig}}, \cdot, \cdot), \text{LoRSanit}(\cdot, \text{pk}_{\text{P3S}}^{\text{Sig}}, \cdot, \cdot, \cdot, \cdot, \cdot, b)}(\text{pk}_{\text{P3S}}^{\text{Sig}})$
 where $\text{Proof}'_{\text{P3S}}$ on input of $\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, \sigma, m$:
 return \perp , if $(\text{pk}_{\text{P3S}}, \sigma, m) \in \mathcal{Q}$
 return $\text{Proof}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, \sigma, m)$
 and LoRSanit on input of $\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \text{sk}_{\text{P3S}, 0}^{\text{San}}, \text{sk}_{\text{P3S}, 1}^{\text{San}}, \text{sk}_{\mathbb{S}0}, \text{sk}_{\mathbb{S}1}, m, \sigma, M, b$:
 for $b' \in \{0, 1\}$, $(\sigma'_{b'}, m'_{b'}) \leftarrow_r \text{Sanitize}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \text{sk}_{\text{P3S}, b'}^{\text{San}}, \text{sk}_{\mathbb{S}, b'}, m, \sigma, M)$
 return \perp , if $\sigma'_0 = \perp \vee \sigma'_1 = \perp$
 $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(\text{pk}_{\text{P3S}}, \sigma'_{b'}, m'_{b'})\}$
 return $\sigma'_{b'}$
 return 1, if $b = b^*$
 return 0

Fig. 8: P3S Pseudonymity

Definition 11 (P3S Pseudonymity). *We say a P3S scheme is pseudonymous, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:*

$$\left| \Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{P3S}}^{\text{Pseudonymity}}(\kappa) = 1 \right] - \frac{1}{2} \right| \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 8.

Signer-Accountability. Signer-accountability prohibits that an adversary can generate a bogus proof that makes $\text{Judge}_{\text{P3S}}$ decide that a sanitizer is responsible for a given signature/message pair (m^*, σ^*) , but that sanitizer has never generated this pair. This is even true, if the adversary can generate the signer's key pair, the global group key pair, while receiving full adaptive access to a sanitization-oracle.

Definition 12 (P3S Signer-Accountability). *We say a P3S scheme is signer-accountable, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:*

$$\Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{P3S}}^{\text{Signer-Accountability}}(\kappa) = 1 \right] \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 9.

$\text{Exp}_{\mathcal{A}, \text{P3S}}^{\text{Signer-Accountability}}(\kappa)$
 $\text{pp}_{\text{P3S}} \leftarrow_r \text{ParGen}_{\text{P3S}}(1^\kappa)$
 $(\text{sk}_{\text{P3S}}^{\text{San}}, \text{pk}_{\text{P3S}}^{\text{San}}) \leftarrow_r \text{KGenSan}_{\text{P3S}}(\text{pp}_{\text{P3S}})$
 $b \leftarrow_r \{0, 1\}$
 $\mathcal{Q} \leftarrow \emptyset$
 $(\text{pk}_0^*, \text{pk}_1^*, \sigma^*, m^*, \pi^*) \leftarrow_r \mathcal{A}^{\text{Sanitize}'_{\text{P3S}}(\cdot, \cdot, \text{sk}_{\text{P3S}}^{\text{San}}, \cdot, \cdot, \cdot)}(\text{pk}_{\text{P3S}}^{\text{San}})$
 where $\text{Sanitize}'_{\text{P3S}}$ on input of $\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \text{sk}_{\text{P3S}}^{\text{San}}, \text{sk}_{\mathbb{S}}, m, \sigma, M$:
 $(\sigma', m') \leftarrow_r \text{Sanitize}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \text{sk}_{\text{P3S}}^{\text{San}}, \text{sk}_{\mathbb{S}}, m, \sigma, M)$
 if $\sigma \neq \perp$, $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \sigma', m')\}$
 return σ'
 return 1, if $\text{Judge}_{\text{P3S}}(\text{pk}_0^*, \text{pk}_1^*, \text{pk}_{\text{P3S}}^{\text{San}}, \pi^*, \sigma^*, m^*) = 1 \wedge (\text{pk}_0^*, \text{pk}_1^*, \sigma^*, m^*) \notin \mathcal{Q}$
 return 0

Fig. 9: P3S Signer-Accountability

Sanitizer-Accountability. Sanitizer-accountability prohibits that an adversary can generate a bogus signature/message pair (m^*, σ^*) that makes $\text{Proof}_{\text{P3S}}$ outputs a (honestly generated) generated proof π_{P3S} which points to the signer, but (m^*, σ^*) has never been generated by the signer. This is even true, if the adversary can generate all sanitizers key pairs, while receiving full adaptive access to a signing-oracle and a proof-oracle.

$\text{Exp}_{\mathcal{A}, \text{P3S}}^{\text{Sanitizer-Accountability}}(\kappa)$
 $\text{pp}_{\text{P3S}} \leftarrow_r \text{ParGen}_{\text{P3S}}(1^\kappa)$
 $(\text{sk}_{\text{P3S}}^{\text{Sig}}, \text{pk}_{\text{P3S}}^{\text{Sig}}) \leftarrow_r \text{KGenSig}_{\text{P3S}}(\text{pp}_{\text{P3S}})$
 $b \leftarrow_r \{0, 1\}$
 $\mathcal{Q} \leftarrow \emptyset$
 $(\text{pk}^*, \sigma^*, m^*, \pi^*) \leftarrow_r \mathcal{A}^{\text{Sign}'_{\text{P3S}}(\cdot, \text{sk}_{\text{P3S}}^{\text{Sig}}, \cdot, \cdot, \cdot), \text{Proof}_{\text{P3S}}(\cdot, \text{sk}_{\text{P3S}}^{\text{Sig}}, \cdot, \cdot)}(\text{pk}_{\text{P3S}}^{\text{Sig}})$
 where $\text{Sign}'_{\text{P3S}}$ on input of $\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, m, A, \mathbb{A}$:
 $\sigma \leftarrow_r \text{Sign}_{\text{P3S}}(\text{pk}_{\text{P3S}}, \text{sk}_{\text{P3S}}^{\text{Sig}}, m, A, \mathbb{A})$
 if $\sigma \neq \perp$, $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(\text{pk}_{\text{P3S}}, \sigma', m')\}$
 return σ'
 $(\pi_{\text{P3S}}, \text{pk}) \leftarrow_r \text{Proof}_{\text{P3S}}(\text{pk}^*, \text{sk}_{\text{P3S}}^{\text{Sig}}, \sigma^*, m^*)$
 return 1, if $\text{Judge}_{\text{P3S}}(\text{pk}^*, \text{pk}_{\text{P3S}}^{\text{Sig}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \pi_{\text{P3S}}, \sigma^*, m^*) = 1 \wedge (\text{pk}^*, \sigma^*, m^*) \notin \mathcal{Q}$
 return 0

Fig. 10: P3S Sanitizer-Accountability

Definition 13 (P3S Sanitizer-Accountability). We say a P3S scheme is sanitizer-accountable, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:

$$\Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{P3S}}^{\text{Sanitizer-Accountability}}(\kappa) = 1 \right] \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 10.

Proof-Soundness. Proof-soundness essentially only handles the case that a signature σ can only be opened in an unambiguous way. Thus, the adversary’s goal is to output two proofs which “prove” different statements for the same signature/message pair. It is related to the property of opening-soundness introduced by Sakai et al. [SSE⁺12] for group signatures.

```

Exp $\mathcal{A}, \text{P3S}$ Proof-Soundness( $\kappa$ )
   $\text{pp}_{\text{P3S}} \leftarrow_r \text{ParGen}_{\text{P3S}}(1^\kappa)$ 
   $((\text{pk}_i^*)_{0 \leq i \leq 5}, \sigma^*, m^*, \pi_0^*, \pi_1^*) \leftarrow_r \mathcal{A}(\text{pp}_{\text{P3S}})$ 
  return 1, if  $\text{Judge}_{\text{P3S}}(\text{pk}_0^*, \text{pk}_1^*, \text{pk}_2^*, \pi_0^*, \sigma^*, m^*) = 1 \wedge$ 
      $\text{Judge}_{\text{P3S}}(\text{pk}_3^*, \text{pk}_4^*, \text{pk}_5^*, \pi_1^*, \sigma^*, m^*) = 1 \wedge$ 
      $(\text{pk}_0^*, \text{pk}_1^*, \text{pk}_2^*) \neq (\text{pk}_3^*, \text{pk}_4^*, \text{pk}_5^*)$ 
  return 0

```

Fig. 11: P3S Proof-Soundness

Definition 14 (P3S Proof-Soundness). We say a P3S scheme is proof-sound, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:

$$\Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{P3S}}^{\text{Proof-Soundness}}(\kappa) = 1 \right] \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 11.

Traceability. Traceability requires that an adversary cannot generate a signature which cannot be opened, i.e., it can be seen as the “dual” to proof-soundness. In more detail, the adversary’s goal is to generate a verifying signature for which an honest signer cannot generate $(\pi_{\text{P3S}}, \text{pk})$ for which $\text{Judge}_{\text{P3S}}$ outputs correct.

Definition 15 (P3S Traceability). We say a P3S scheme is traceable, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:

$$\Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{P3S}}^{\text{Traceability}}(\kappa) = 1 \right] \leq \nu(\kappa).$$

$$\begin{aligned}
& \mathbf{Exp}_{\mathcal{A}, \text{P3S}}^{\text{Traceability}}(\kappa) \\
& \text{pp}_{\text{P3S}} \leftarrow_r \text{ParGen}_{\text{P3S}}(1^\kappa) \\
& (\text{sk}_{\text{P3S}}^{\text{Sig}}, \text{pk}_{\text{P3S}}^{\text{Sig}}) \leftarrow_r \text{KGenSig}_{\text{P3S}}(\text{pp}_{\text{P3S}}) \\
& (\text{pk}^*, \sigma^*, m^*) \leftarrow_r \mathcal{A}^{\text{Sig}_{\text{P3S}}(\cdot, \text{sk}_{\text{P3S}}^{\text{Sig}}, \cdot, \cdot), \text{Proof}_{\text{P3S}}(\cdot, \text{sk}_{\text{P3S}}^{\text{Sig}}, \cdot, \cdot)}(\text{pk}_{\text{P3S}}^{\text{Sig}}) \\
& \text{return } 0, \text{ if } \text{Verify}_{\text{P3S}}(\text{pk}^*, \text{pk}_{\text{P3S}}^{\text{Sig}}, \sigma^*, m^*) = 0 \\
& (\pi_{\text{P3S}}, \text{pk}) \leftarrow_r \text{Proof}_{\text{P3S}}(\text{pk}^*, \text{sk}_{\text{P3S}}^{\text{Sig}}, \sigma^*, m^*) \\
& \text{return } 1, \text{ if } \text{Judge}_{\text{P3S}}(\text{pk}^*, \text{pk}_{\text{P3S}}^{\text{Sig}}, \text{pk}, \pi_{\text{P3S}}, \sigma^*, m^*) = 0 \\
& \text{return } 0
\end{aligned}$$

Fig. 12: P3S Traceability

The corresponding experiment is depicted in Figure 12.

Relationship of Properties. All properties are independent of each other. The full theorems and proofs are given in App. D.

4 Construction

In this section we present our P3S construction. The key ingredients are our strengthened version of a policy-based chameleon-hash PCH, a labeled simulation-sound extractable non-interactive zero-knowledge proof system Ω (NIZK for short), a one-way function f as well as a key-verifiable IND-CCA2 secure public key encryption scheme³ Π and an eUNF-CMA-secure signature scheme Σ . The intuition behind our construction, given in Construction 1, is as follows.

The global parameters of the scheme are a one-way function f , the CRS of the NIZK, and the parameters for the encryption scheme, the signature scheme and the policy-based chameleon hash. The group setup generates the keys of the policy-based chameleon-hash, and a key pair of the signature scheme. The signer generates a signature key pair and publishes the public key together with an image y_1 of a random pre-image x_1 of the OWF f . The sanitizer chooses a random pre-image x_2 of the OWF as secret key and as public key $y_2 = f(x_2)$. If a sanitizer joins a group, i.e., obtains secret keys for a set of attributes \mathbb{S} , the group manager signs the sanitizer's public key and additionally issues a secret key for the PCH for attributes \mathbb{S} .

³ Although key-verifiability is no property often explicitly used within IND-CCA2 encryption schemes, most encryption schemes are key-verifiable. See App. C.

For signing, the signer hashes the message using the PCH and signs the hash (along with some additional information). Moreover, it computes a NIZK for the relation R (using as label ℓ some additional information like the admissible changes).

$$\begin{aligned} (y_1, c, y_2, \text{pk}_\Pi, \text{pk}_\Sigma), (x, r, \sigma_{\text{sk}_S}) \in R &\iff \\ (y_1 = f(x) \wedge c = \text{Enc}_\Pi(\text{pk}_\Pi, y_1; r)) \vee & \\ (y_2 = f(x) \wedge c = \text{Enc}_\Pi(\text{pk}_\Pi, y_2; r) \wedge \text{Verf}_\Sigma(\text{pk}_\Sigma, y_2, \sigma_{\text{sk}_S}) = 1). & \end{aligned}$$

Sanitizing amounts to computing a collision for the PCH hash, updating the respective message blocks, and again attaching a NIZK for relation R . Verification is straightforward. Relation R is used within signing and sanitizing to force the signer or the sanitizer to commit to having performed the action. Intuitively, when determining whether a signer or sanitizer has performed the action, the $\text{Proof}_{\text{P3S}}$ algorithm (having access to the signer's secret key) can simply decrypt c and prove correct decryption.

It may be tempting to think that the weaker notion of witness indistinguishability is sufficient for our construction, but it turns out that one requires zero-knowledge. Moreover, we stress that due to the underlying construction paradigm, we do not consider the strong privacy notion of unlinkability [BFLS10], i.e., that sanitized signatures cannot be linked to its origin, which seems to be very hard to achieve with the current construction paradigm. However, finding such a construction may have its merits. Formally, for our construction, we can show the following:

Theorem 1. *If f is a one-way function, Π is IND-CCA2 secure and key-verifiable, Σ is eUNF-CMA secure, Ω is zero-knowledge and simulation-sound extractable, while PCH is fully indistinguishable, insider collision-resistant, and unique, the construction of a P3S given in Construction 1 is unforgeable, immutable, private, transparent, pseudonymous, signer-accountable, sanitizer-accountable, proof-sound, and traceable. Likewise, the construction is correct, if the underlying primitives are correct (and sound, resp.).*

The full proof of Theorem 1 is given in App. D.

Instantiation. The description of Construction 1 is as compact as reasonable. For a concrete instantiation, there are some aspects which can be optimized. Currently, it seems to be advisable to stick to elliptic curves and in particular to the type-3 bilinear group setting (a setting where we assume the SXDH assumption to hold), due to the efficiency of the

<p>ParGen_{P3S}(1^κ) : On input a security parameter κ, let $\text{pp}_\Pi \leftarrow_r \text{PPGen}_\Pi(1^\kappa)$, $\text{crs}_\Omega \leftarrow_r \text{PPGen}_\Omega(1^\kappa)$.^a Finally, choose a one-way function f, let $\text{pp}_\Sigma \leftarrow_r \text{PPGen}_\Sigma(1^\kappa)$, and $\text{pp}_{\text{PCH}} \leftarrow_r \text{PPGen}_{\text{PCH}}(1^\kappa)$. Return $\text{pp}_{\text{P3S}} \leftarrow (\text{crs}_\Omega, \text{pp}_\Pi, \text{pp}_\Sigma, \text{pp}_{\text{PCH}}, f)$.</p> <p>Setup_{P3S}(pp_{P3S}) : Let $(\text{sk}_{\text{PCH}}, \text{pk}_{\text{PCH}}) \leftarrow_r \text{MKeyGen}_{\text{PCH}}(\text{pp}_{\text{PCH}})$ and $(\text{sk}_\Sigma, \text{pk}_\Sigma) \leftarrow_r \text{KGen}_\Sigma(\text{pp}_\Sigma)$. Return $(\text{sk}_{\text{P3S}}, \text{pk}_{\text{P3S}}) \leftarrow ((\text{sk}_{\text{PCH}}, \text{sk}_\Sigma), (\text{pk}_{\text{PCH}}, \text{pk}_\Sigma))$.</p> <p>KGenSig_{P3S}(pp_{P3S}) : Draw $x_1 \leftarrow_r D_f$, $(\text{sk}_\Pi, \text{pk}_\Pi) \leftarrow_r \text{KGen}_\Pi(\text{pp}_\Pi)$, let $y_1 \leftarrow f(x_1)$ and $(\text{sk}'_\Sigma, \text{pk}'_\Sigma) \leftarrow_r \text{KGen}_\Sigma(\text{pp}_\Sigma)$. Return $(\text{sk}_{\text{P3S}}^{\text{Sig}}, \text{pk}_{\text{P3S}}^{\text{Sig}}) \leftarrow ((x_1, \text{sk}'_\Sigma, \text{sk}_\Pi), (y_1, \text{pk}'_\Sigma, \text{pk}_\Pi))$.</p> <p>KGenSan_{P3S}(pp_{P3S}) : Draw $x_2 \leftarrow_r D_f$. Let $y_2 \leftarrow f(x_2)$. Return (x_2, y_2).</p> <p>Sign_{P3S}(pk_{P3S}, sk_{P3S}^{Sig}, m, A, A) : If $\mathbb{A} = \emptyset$, return \perp. Let $(h, r) \leftarrow_r \text{Hash}_{\text{PCH}}(\text{pk}_{\text{PCH}}, m, \mathbb{A})$, $\sigma_m \leftarrow_r \text{Sign}_\Sigma(\text{sk}'_\Sigma, (\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \mathbb{A}, m_{!A}, h, \mathbb{A}))$, and $c \leftarrow_r \text{Enc}_\Pi(\text{pk}_\Pi, y_1)$. Let $\pi \leftarrow_r \text{Prove}_\Omega\{(x_1, x_2, \sigma_{\text{sk}_\Sigma}) : (y_1 = f(x_1) \wedge c = \text{Enc}_\Pi(\text{pk}_\Pi, y_1)) \vee (y_2 = f(x_2) \wedge c = \text{Enc}_\Pi(\text{pk}_\Pi, y_2) \wedge \text{Verf}_\Sigma(\text{pk}_\Sigma, (y_2, \text{pk}_{\text{P3S}}), \sigma_{\text{sk}_\Sigma}) = 1)\}$($\ell$), where $\ell = (\text{pp}_{\text{P3S}}, \text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, h, r, m, \mathbb{A}, \mathbb{A}, m_A, m_{!A}, \sigma_m, c)$. Return $\sigma \leftarrow (h, r, \mathbb{A}, \sigma_m, \mathbb{A}, \pi, c)$.</p> <p>AddSan_{P3S}(sk_{P3S}, pk_{P3S}^{San}, S) : If $\mathbb{S} \notin 2^{\mathbb{U}}$, return \perp. Let $\sigma_{\text{sk}_\Sigma} \leftarrow_r \text{Sign}_\Sigma(\text{sk}_\Sigma, (\text{pk}_{\text{P3S}}^{\text{San}}, \text{pk}_{\text{P3S}}))$ and $\text{sk}'_\mathbb{S} \leftarrow_r \text{KGen}_{\text{PCH}}(\text{sk}_{\text{PCH}}, \mathbb{S})$. Return $\text{sk}_\mathbb{S} \leftarrow (\sigma_{\text{sk}_\Sigma}, \text{sk}'_\mathbb{S})$.</p> <p>Verify_{P3S}(pk_{P3S}, pk_{P3S}^{Sig}, σ, m) : If π or σ_m is not valid, return \perp. Check that $m_{!A}$ is contained in m in the correct sequence at the right positions (derivable from \mathbb{A}). If $\text{Verify}_{\text{PCH}}(\text{pk}_{\text{PCH}}, m, r, h) = 1$, return 1. Otherwise, return 0.</p> <p>Sanitize_{P3S}(pk_{P3S}, pk_{P3S}^{Sig}, sk_{P3S}^{San}, sk_S, m, σ, M) : If $\sigma_{\text{sk}_\Sigma}$ or σ is not valid, return \perp. Let $r' \leftarrow_r \text{Adapt}_{\text{PCH}}(\text{pk}_{\text{PCH}}, \text{sk}'_\mathbb{S}, m, \text{M}(m), h, r)$, $c' \leftarrow_r \text{Enc}_\Pi(\text{pk}_\Pi, y_2)$, and $\pi' \leftarrow_r \text{Prove}_\Omega\{(x_1, x_2, \sigma_{\text{sk}_\Sigma}) : (y_1 = f(x_1) \wedge c' = \text{Enc}_\Pi(\text{pk}_\Pi, y_1)) \vee (y_2 = f(x_2) \wedge c' = \text{Enc}_\Pi(\text{pk}_\Pi, y_2) \wedge \text{Verf}_\Sigma(\text{pk}_\Sigma, (y_2, \text{pk}_{\text{P3S}}), \sigma_{\text{sk}_\Sigma}) = 1)\}$($\ell$), where $\ell = (\text{pp}_{\text{P3S}}, \text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, h, r', \text{M}(m), \mathbb{A}, \mathbb{A}, m_A, m_{!A}, \sigma_m, c')$. Let $(\sigma', m') \leftarrow ((h, r', \mathbb{A}, \sigma_m, \mathbb{A}, \pi', c'), \text{M}(m))$. If (σ', m') is not valid, return \perp. Return (σ', m').</p> <p>Proof_{P3S}(pk_{P3S}, sk_{P3S}^{Sig}, σ, m) : If σ is not valid, return \perp. Let $\text{pk} \leftarrow \text{Dec}_\Pi(\text{sk}_\Pi, c)$. Let $\pi_{\text{P3S}} \leftarrow_r \text{Prove}_\Omega\{(\text{sk}_\Pi) : \text{pk} = \text{Dec}_\Pi(\text{sk}_\Pi, c) \wedge \text{KVrf}_\Pi(\text{sk}_\Pi, \text{pk}_\Pi) = 1\}$($\ell$), where $\ell = (\text{pp}_{\text{P3S}}, \text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \sigma, \text{pk}, m)$. Return $(\pi_{\text{P3S}}, \text{pk})$.</p> <p>Judge_{P3S}(pk_{P3S}, pk_{P3S}^{Sig}, pk, π_{P3S}, σ, m) : If σ or π_{P3S} is not valid, return 0. Return 1.</p>
<p>^a Note, we need a different CRS for each language L involved. However, we keep the description short, and thus do not make this explicit.</p>

Construction 1: Our P3S

CP-ABE schemes in this setting (used by the PCH). Consequently, we consider the OWF f to be simply the function $f(x) = g^x$ for $x \in \mathbb{Z}_q$ and g being a generator of a group \mathbb{G} of prime order q (and in particular one of the base groups of a bilinear group). Then, as an encryption scheme to encrypt images under f and that is key-verifiable, we can use Cramer-Shoup encryption in either of the two base groups. For completeness, we show key-verifiability of CS-encryption where keys are generated with respect to a common group description (including both generators) in App. C). Now, the signature keys $(\text{sk}'_\Sigma, \text{pk}'_\Sigma)$ used by signer to produce

signatures can be any arbitrary eUNF-CMA-secure scheme. In contrast, the signature scheme associated to keys $(\text{sk}_\Sigma, \text{pk}_\Sigma)$ used by the group manager in `AddSanP35` to certify the y_2 values of sanitizers need to be chosen with care: we need a signature scheme with message space being one of the base groups of the bilinear group and thus the natural choice is a structure preserving signature scheme [AFG⁺10]. Moreover, the SPS (e.g., Groth [Gro15]) needs to be compatible with efficient labeled NIZK; the latter can be instantiated from standard Σ -protocols using the compiler by Faust et al. [FKMV12] and supporting labels is straightforward (cf. [ABM15]). As PCH instantiation we can use a strengthened version of the PCH by Derler et al. [DSSS19]. See App. C. To make the public key of the PCH compatible with the NIZK and Σ , it can simply be hashed using a collision-resistant hash-function.

Efficiency. Our scheme is reasonably efficient. The group manager only needs to create a key-pair for a PCH, while the sanitizer only needs to evaluate a one-way functions (the signer additionally needs to draw a key-pair for an encryption scheme Π). For signing, the signer needs to generate a hash, a signature, an encryption, and a simple NIZK. For sanitizing, the sanitizer has to create an encryption, adapt a hash, and attaches a simple NIZK. Granting sanitizing rights boils down to creating a signature and creating a key for the PCH. Verification is also straightforward: A verifier checks a signature and the NIZK. Likewise, proof-generation is a simple decryption and a NIZK proving that decryption was done honestly. Checking a proof is verifying a proof and a signature. Thus, ignoring the NIZK and the encryptions, our scheme is comparable to existing, way less expressive, constructions.

5 Conclusion

We have introduced the notion of policy-based sanitizable signatures, which are an extension to standard sanitizable signature schemes, along with a provably secure construction. Our construction features, for the first time, full accountability. In our new primitive, a sanitizer is no longer appointed by the signer at signature generation, but rather can sanitize based on a set attributes it has.

References

- ABC⁺15. Jae Hyun Ahn, Dan Boneh, Jan Camenisch, Susan Hohenberger, Abhi Shelat, and Brent Waters. Computing on authenticated data. *J. Cryptology*,

- 28(2):351–395, 2015.
- ABM15. Michel Abdalla, Fabrice Benhamouda, and Philip MacKenzie. Security of the J-PAKE password-authenticated key exchange protocol. In *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*, pages 571–587. IEEE Computer Society, 2015.
- ACdMT05. Giuseppe Ateniese, Daniel H. Chou, Breno de Medeiros, and Gene Tsudik. Sanitizable signatures. In Sabrina De Capitani di Vimercati, Paul F. Syverson, and Dieter Gollmann, editors, *Computer Security - ESORICS 2005, 10th European Symposium on Research in Computer Security, Milan, Italy, September 12-14, 2005, Proceedings*, volume 3679 of *Lecture Notes in Computer Science*, pages 159–177. Springer, 2005.
- ADK⁺13. Masayuki Abe, Bernardo David, Markulf Kohlweiss, Ryo Nishimaki, and Miyako Ohkubo. Tagged one-time signatures: Tight security and optimal tag size. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *Public-Key Cryptography - PKC 2013 - 16th International Conference on Practice and Theory in Public-Key Cryptography, Nara, Japan, February 26 - March 1, 2013. Proceedings*, volume 7778 of *Lecture Notes in Computer Science*, pages 312–331. Springer, 2013.
- AdM04. Giuseppe Ateniese and Breno de Medeiros. On the key exposure problem in chameleon hashes. In Carlo Blundo and Stelvio Cimato, editors, *Security in Communication Networks, 4th International Conference, SCN 2004, Amalfi, Italy, September 8-10, 2004, Revised Selected Papers*, volume 3352 of *Lecture Notes in Computer Science*, pages 165–179. Springer, 2004.
- AFG⁺10. Masayuki Abe, Georg Fuchsbauer, Jens Groth, Kristiyan Haralambiev, and Miyako Ohkubo. Structure-preserving signatures and commitments to group elements. In Tal Rabin, editor, *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings*, volume 6223 of *Lecture Notes in Computer Science*, pages 209–236. Springer, 2010.
- AMVA17. Giuseppe Ateniese, Bernardo Magri, Daniele Venturi, and Ewerton R. Andrade. Redactable blockchain - or - rewriting history in bitcoin and friends. In *2017 IEEE European Symposium on Security and Privacy, EuroS&P 2017, Paris, France, April 26-28, 2017*, pages 111–126. IEEE, 2017.
- BCD⁺17. Michael Till Beck, Jan Camenisch, David Derler, Stephan Krenn, Henrich C. Pöhls, Kai Samelin, and Daniel Slamanig. Practical strongly invisible and strongly accountable sanitizable signatures. In Josef Pieprzyk and Suriadi Suriadi, editors, *Information Security and Privacy - 22nd Australasian Conference, ACISP 2017, Auckland, New Zealand, July 3-5, 2017, Proceedings, Part I*, volume 10342 of *Lecture Notes in Computer Science*, pages 437–452. Springer, 2017.
- BFF⁺09. Christina Brzuska, Marc Fischlin, Tobias Freudenreich, Anja Lehmann, Marcus Page, Jakob Schelbert, Dominique Schröder, and Florian Volk. Security of sanitizable signatures revisited. In Jarecki and Tsudik [JT09], pages 317–336.
- BFKW09. Dan Boneh, David Mandell Freeman, Jonathan Katz, and Brent Waters. Signing a linear subspace: Signature schemes for network coding. In Jarecki and Tsudik [JT09], pages 68–87.
- BFLS09. Christina Brzuska, Marc Fischlin, Anja Lehmann, and Dominique Schröder. Sanitizable signatures: How to partially delegate control for authenticated

- data. In Arslan Brömme, Christoph Busch, and Detlef Hühnlein, editors, *BIOSIG 2009 - Proceedings of the Special Interest Group on Biometrics and Electronic Signatures, 17.-18. September 2009 in Darmstadt, Germany*, volume P-155 of *LNI*, pages 117–128. GI, 2009.
- BFSL10. Christina Brzuska, Marc Fischlin, Anja Lehmann, and Dominique Schröder. Unlinkability of sanitizable signatures. In Phong Q. Nguyen and David Pointcheval, editors, *Public Key Cryptography - PKC 2010, 13th International Conference on Practice and Theory in Public Key Cryptography, Paris, France, May 26-28, 2010. Proceedings*, volume 6056 of *Lecture Notes in Computer Science*, pages 444–461. Springer, 2010.
- BL17. Xavier Bultel and Pascal Lafourcade. Unlinkable and strongly accountable sanitizable signatures from verifiable ring signatures. In Srdjan Capkun and Sherman S. M. Chow, editors, *Cryptology and Network Security - 16th International Conference, CANS 2017, Hong Kong, China, November 30 - December 2, 2017, Revised Selected Papers*, volume 11261 of *Lecture Notes in Computer Science*, pages 203–226. Springer, 2017.
- BLL⁺19. Xavier Bultel, Pascal Lafourcade, Russell W. F. Lai, Giulio Malavolta, Dominique Schröder, and Sri Aravinda Krishnan Thyagarajan. Efficient invisible and unlinkable sanitizable signatures. In Dongdai Lin and Kazue Sako, editors, *Public-Key Cryptography - PKC 2019 - 22nd IACR International Conference on Practice and Theory of Public-Key Cryptography, Beijing, China, April 14-17, 2019, Proceedings, Part I*, volume 11442 of *Lecture Notes in Computer Science*, pages 159–189. Springer, 2019.
- BNPS03. Mihir Bellare, Chanathip Namprempre, David Pointcheval, and Michael Semanko. The one-more-rsa-inversion problems and the security of chaum’s blind signature scheme. *J. Cryptology*, 16(3):185–215, 2003.
- BPS12. Christina Brzuska, Henrich Christopher Pöhls, and Kai Samelin. Non-interactive public accountability for sanitizable signatures. In Sabrina De Capitani di Vimercati and Chris J. Mitchell, editors, *Public Key Infrastructures, Services and Applications - 9th European Workshop, EuroPKI 2012, Pisa, Italy, September 13-14, 2012, Revised Selected Papers*, volume 7868 of *Lecture Notes in Computer Science*, pages 178–193. Springer, 2012.
- BPS13. Christina Brzuska, Henrich Christopher Pöhls, and Kai Samelin. Efficient and perfectly unlinkable sanitizable signatures without group signatures. In Sokratis K. Katsikas and Isaac Agudo, editors, *Public Key Infrastructures, Services and Applications - 10th European Workshop, EuroPKI 2013, Egham, UK, September 12-13, 2013, Revised Selected Papers*, volume 8341 of *Lecture Notes in Computer Science*, pages 12–30. Springer, 2013.
- BPS17. Arne Bilzhause, Henrich C. Pöhls, and Kai Samelin. Position paper: The past, present, and future of sanitizable and redactable signatures. In *Proceedings of the 12th International Conference on Availability, Reliability and Security, Reggio Calabria, Italy, August 29 - September 01, 2017*, pages 87:1–87:9. ACM, 2017.
- BR93. Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *CCS '93, Proceedings of the 1st ACM Conference on Computer and Communications Security, Fairfax, Virginia, USA, November 3-5, 1993*, pages 62–73. ACM, 1993.

- BSW07. John Bethencourt, Amit Sahai, and Brent Waters. Ciphertext-policy attribute-based encryption. In *2007 IEEE Symposium on Security and Privacy (S&P 2007), 20-23 May 2007, Oakland, California, USA*, pages 321–334. IEEE Computer Society, 2007.
- BSZ05. Mihir Bellare, Haixia Shi, and Chong Zhang. Foundations of group signatures: The case of dynamic groups. In Alfred Menezes, editor, *Topics in Cryptology - CT-RSA 2005, The Cryptographers' Track at the RSA Conference 2005, San Francisco, CA, USA, February 14-18, 2005, Proceedings*, volume 3376 of *Lecture Notes in Computer Science*, pages 136–153. Springer, 2005.
- CDK⁺17. Jan Camenisch, David Derler, Stephan Krenn, Henrich C. Pöhls, Kai Samelin, and Daniel Slamanig. Chameleon-hashes with ephemeral trapdoors - and applications to invisible sanitizable signatures. In Serge Fehr, editor, *Public-Key Cryptography - PKC 2017 - 20th IACR International Conference on Practice and Theory in Public-Key Cryptography, Amsterdam, The Netherlands, March 28-31, 2017, Proceedings, Part II*, volume 10175 of *Lecture Notes in Computer Science*, pages 152–182. Springer, 2017.
- CJ10. Sébastien Canard and Amandine Jambert. On extended sanitizable signature schemes. In Josef Pieprzyk, editor, *Topics in Cryptology - CT-RSA 2010, The Cryptographers' Track at the RSA Conference 2010, San Francisco, CA, USA, March 1-5, 2010. Proceedings*, volume 5985 of *Lecture Notes in Computer Science*, pages 179–194. Springer, 2010.
- CJL12. Sébastien Canard, Amandine Jambert, and Roch Lescuyer. Sanitizable signatures with several signers and sanitizers. In Aikaterini Mitrokotsa and Serge Vaudenay, editors, *Progress in Cryptology - AFRICACRYPT 2012 - 5th International Conference on Cryptology in Africa, Ifrance, Morocco, July 10-12, 2012. Proceedings*, volume 7374 of *Lecture Notes in Computer Science*, pages 35–52. Springer, 2012.
- CLM08. Sébastien Canard, Fabien Laguillaumie, and Michel Milhau. Trapdoorsanitizable signatures and their application to content protection. In Steven M. Bellovin, Rosario Gennaro, Angelos D. Keromytis, and Moti Yung, editors, *Applied Cryptography and Network Security, 6th International Conference, ACNS 2008, New York, NY, USA, June 3-6, 2008. Proceedings*, volume 5037 of *Lecture Notes in Computer Science*, pages 258–276, 2008.
- CS97. Jan Camenisch and Markus Stadler. Efficient group signature schemes for large groups (extended abstract). In Burton S. Kaliski Jr., editor, *Advances in Cryptology - CRYPTO '97, 17th Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 1997, Proceedings*, volume 1294 of *Lecture Notes in Computer Science*, pages 410–424. Springer, 1997.
- CS98. Ronald Cramer and Victor Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In Hugo Krawczyk, editor, *Advances in Cryptology - CRYPTO '98, 18th Annual International Cryptology Conference, Santa Barbara, California, USA, August 23-27, 1998, Proceedings*, volume 1462 of *Lecture Notes in Computer Science*, pages 13–25. Springer, 1998.
- CvH91. David Chaum and Eugène van Heyst. Group signatures. In Donald W. Davies, editor, *Advances in Cryptology - EUROCRYPT '91, Workshop on the Theory and Application of of Cryptographic Techniques, Brighton, UK*,

- April 8-11, 1991, *Proceedings*, volume 547 of *Lecture Notes in Computer Science*, pages 257–265. Springer, 1991.
- DDH⁺15. Denise Demirel, David Derler, Christian Hanser, Henrich C. Pöhls, Daniel Slamanig, and Giulia Traverso. PRISMACLOUD D4.4: Overview of Functional and Malleable Signature Schemes. Technical report, H2020 Prismacloud, www.prismacloud.eu, 2015.
- DHLW10. Yevgeniy Dodis, Kristiyan Haralambiev, Adriana López-Alt, and Daniel Wichs. Efficient public-key cryptography in the presence of key leakage. In Masayuki Abe, editor, *Advances in Cryptology - ASIACRYPT 2010 - 16th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 5-9, 2010. Proceedings*, volume 6477 of *Lecture Notes in Computer Science*, pages 613–631. Springer, 2010.
- dMPPS14. Hermann de Meer, Henrich Christopher Pöhls, Joachim Posegga, and Kai Samelin. On the relation between redactable and sanitizable signature schemes. In Jan Jürjens, Frank Piessens, and Nataliia Bielova, editors, *Engineering Secure Software and Systems - 6th International Symposium, ESSoS 2014, Munich, Germany, February 26-28, 2014, Proceedings*, volume 8364 of *Lecture Notes in Computer Science*, pages 113–130. Springer, 2014.
- DPSS15. David Derler, Henrich C. Pöhls, Kai Samelin, and Daniel Slamanig. A general framework for redactable signatures and new constructions. In Soonhak Kwon and Aaram Yun, editors, *Information Security and Cryptology - ICISC 2015 - 18th International Conference, Seoul, South Korea, November 25-27, 2015, Revised Selected Papers*, volume 9558 of *Lecture Notes in Computer Science*, pages 3–19. Springer, 2015.
- DS15. David Derler and Daniel Slamanig. Rethinking privacy for extended sanitizable signatures and a black-box construction of strongly private schemes. In Man Ho Au and Atsuko Miyaji, editors, *Provable Security - 9th International Conference, ProvSec 2015, Kanazawa, Japan, November 24-26, 2015, Proceedings*, volume 9451 of *Lecture Notes in Computer Science*, pages 455–474. Springer, 2015.
- DS19. David Derler and Daniel Slamanig. Key-homomorphic signatures: definitions and applications to multiparty signatures and non-interactive zero-knowledge. *Des. Codes Cryptogr.*, 87(6):1373–1413, 2019.
- DSSS19. David Derler, Kai Samelin, Daniel Slamanig, and Christoph Striecks. Fine-grained and controlled rewriting in blockchains: Chameleon-hashing gone attribute-based. In *26th Annual Network and Distributed System Security Symposium, NDSS 2019, San Diego, California, USA, February 24-27, 2019*. The Internet Society, 2019.
- FF15. Victoria Fehr and Marc Fischlin. Sanitizable signcryption: Sanitization over encrypted data (full version). *IACR Cryptology ePrint Archive*, 2015:765, 2015.
- FH18. Marc Fischlin and Patrick Harasser. Invisible sanitizable signatures and public-key encryption are equivalent. In Bart Preneel and Frederik Vercauteren, editors, *Applied Cryptography and Network Security - 16th International Conference, ACNS 2018, Leuven, Belgium, July 2-4, 2018, Proceedings*, volume 10892 of *Lecture Notes in Computer Science*, pages 202–220. Springer, 2018.
- FKM⁺16. Nils Fleischhacker, Johannes Krupp, Giulio Malavolta, Jonas Schneider, Dominique Schröder, and Mark Simkin. Efficient unlinkable sanitizable signatures from signatures with re-randomizable keys. In Chen-Mou Cheng,

- Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang, editors, *Public-Key Cryptography - PKC 2016 - 19th IACR International Conference on Practice and Theory in Public-Key Cryptography, Taipei, Taiwan, March 6-9, 2016, Proceedings, Part I*, volume 9614 of *Lecture Notes in Computer Science*, pages 301–330. Springer, 2016.
- FKMV12. Sebastian Faust, Markulf Kohlweiss, Giorgia Azzurra Marson, and Daniele Venturi. On the non-malleability of the fiat-shamir transform. In Steven D. Galbraith and Mridul Nandi, editors, *Progress in Cryptology - INDOCRYPT 2012, 13th International Conference on Cryptology in India, Kolkata, India, December 9-12, 2012. Proceedings*, volume 7668 of *Lecture Notes in Computer Science*, pages 60–79. Springer, 2012.
- GGOT16. Esha Ghosh, Michael T. Goodrich, Olga Ohrimenko, and Roberto Tamassia. Verifiable zero-knowledge order queries and updates for fully dynamic lists and trees. In Vassilis Zikas and Roberto De Prisco, editors, *Security and Cryptography for Networks - 10th International Conference, SCN 2016, Amalfi, Italy, August 31 - September 2, 2016, Proceedings*, volume 9841 of *Lecture Notes in Computer Science*, pages 216–236. Springer, 2016.
- GQZ10. Junqing Gong, Haifeng Qian, and Yuan Zhou. Fully-secure and practical sanitizable signatures. In Xuejia Lai, Moti Yung, and Dongdai Lin, editors, *Information Security and Cryptology - 6th International Conference, Inscrypt 2010, Shanghai, China, October 20-24, 2010, Revised Selected Papers*, volume 6584 of *Lecture Notes in Computer Science*, pages 300–317. Springer, 2010.
- Gro06. Jens Groth. Simulation-sound NIZK proofs for a practical language and constant size group signatures. In Xuejia Lai and Kefei Chen, editors, *Advances in Cryptology - ASIACRYPT 2006, 12th International Conference on the Theory and Application of Cryptology and Information Security, Shanghai, China, December 3-7, 2006, Proceedings*, volume 4284 of *Lecture Notes in Computer Science*, pages 444–459. Springer, 2006.
- Gro15. Jens Groth. Efficient fully structure-preserving signatures for large messages. In Tetsu Iwata and Jung Hee Cheon, editors, *Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application of Cryptology and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part I*, volume 9452 of *Lecture Notes in Computer Science*, pages 239–259. Springer, 2015.
- JMSW02. Robert Johnson, David Molnar, Dawn Xiaodong Song, and David A. Wagner. Homomorphic signature schemes. In Bart Preneel, editor, *Topics in Cryptology - CT-RSA 2002, The Cryptographer’s Track at the RSA Conference, 2002, San Jose, CA, USA, February 18-22, 2002, Proceedings*, volume 2271 of *Lecture Notes in Computer Science*, pages 244–262. Springer, 2002.
- JT09. Stanislaw Jarecki and Gene Tsudik, editors. *Public Key Cryptography - PKC 2009, 12th International Conference on Practice and Theory in Public Key Cryptography, Irvine, CA, USA, March 18-20, 2009. Proceedings*, volume 5443 of *Lecture Notes in Computer Science*. Springer, 2009.
- KPSS18a. Stephan Krenn, Henrich C. Pöhls, Kai Samelin, and Daniel Slamanig. Chameleon-hashes with dual long-term trapdoors and their applications. In Antoine Joux, Abderrahmane Nitaj, and Tajjeeddine Rachidi, editors, *Progress in Cryptology - AFRICACRYPT 2018 - 10th International Conference on Cryptology in Africa, Marrakesh, Morocco, May 7-9, 2018*,

- Proceedings*, volume 10831 of *Lecture Notes in Computer Science*, pages 11–32. Springer, 2018.
- KPSS18b. Stephan Krenn, Henrich C. Pöhls, Kai Samelin, and Daniel Slamanig. Protean signature schemes. In Jan Camenisch and Panos Papadimitratos, editors, *Cryptology and Network Security - 17th International Conference, CANS 2018, Naples, Italy, September 30 - October 3, 2018, Proceedings*, volume 11124 of *Lecture Notes in Computer Science*, pages 256–276. Springer, 2018.
- KPSS19. Stephan Krenn, Henrich C. Pöhls, Kai Samelin, and Daniel Slamanig. Fully invisible protean signatures schemes. *IACR Cryptology ePrint Archive*, 2019:39, 2019.
- KR00. Hugo Krawczyk and Tal Rabin. Chameleon signatures. In *Proceedings of the Network and Distributed System Security Symposium, NDSS 2000, San Diego, California, USA*, pages 143–154. The Internet Society, 2000.
- KSS15. Stephan Krenn, Kai Samelin, and Dieter Sommer. Stronger security for sanitizable signatures. In Joaquín García-Alfaro, Guillermo Navarro-Arribas, Alessandro Aldini, Fabio Martinelli, and Neeraj Suri, editors, *Data Privacy Management, and Security Assurance - 10th International Workshop, DPM 2015, and 4th International Workshop, QASA 2015, Vienna, Austria, September 21-22, 2015. Revised Selected Papers*, volume 9481 of *Lecture Notes in Computer Science*, pages 100–117. Springer, 2015.
- LDW13. Junzuo Lai, Xuhua Ding, and Yongdong Wu. Accountable trapdoor sanitizable signatures. In Robert H. Deng and Tao Feng, editors, *Information Security Practice and Experience - 9th International Conference, ISPEC 2013, Lanzhou, China, May 12-14, 2013. Proceedings*, volume 7863 of *Lecture Notes in Computer Science*, pages 117–131. Springer, 2013.
- LOS⁺10. Allison B. Lewko, Tatsuki Okamoto, Amit Sahai, Katsuyuki Takashima, and Brent Waters. Fully secure functional encryption: Attribute-based encryption and (hierarchical) inner product encryption. In Henri Gilbert, editor, *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Monaco / French Riviera, May 30 - June 3, 2010. Proceedings*, volume 6110 of *Lecture Notes in Computer Science*, pages 62–91. Springer, 2010.
- SBZ01. Ron Steinfeld, Laurence Bull, and Yuliang Zheng. Content extraction signatures. In Kwangjo Kim, editor, *Information Security and Cryptology - ICISC 2001, 4th International Conference Seoul, Korea, December 6-7, 2001, Proceedings*, volume 2288 of *Lecture Notes in Computer Science*, pages 285–304. Springer, 2001.
- 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 Marc Fischlin, Johannes A. Buchmann, and Mark Manulis, editors, *Public Key Cryptography - PKC 2012 - 15th International Conference on Practice and Theory in Public Key Cryptography, Darmstadt, Germany, May 21-23, 2012. Proceedings*, volume 7293 of *Lecture Notes in Computer Science*, pages 715–732. Springer, 2012.
- YAHK11. Shota Yamada, Nuttapong Attrapadung, Goichiro Hanaoka, and Noboru Kunihiro. Generic constructions for chosen-ciphertext secure attribute based encryption. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *Public Key Cryptography - PKC 2011 - 14th In-*

ternational Conference on Practice and Theory in Public Key Cryptography, Taormina, Italy, March 6-9, 2011. Proceedings, volume 6571 of *Lecture Notes in Computer Science*, pages 71–89. Springer, 2011.

- YSL10. Dae Hyun Yum, Jae Woo Seo, and Pil Joong Lee. Trapdoor sanitizable signatures made easy. In Jianying Zhou and Moti Yung, editors, *Applied Cryptography and Network Security, 8th International Conference, ACNS 2010, Beijing, China, June 22-25, 2010. Proceedings*, volume 6123 of *Lecture Notes in Computer Science*, pages 53–68, 2010.

A Additional Preliminaries

Definition 16 (One-Way Functions). *A function $f : D_f \rightarrow R_f$ is κ -one-way, if for every PPT adversary \mathcal{A} there exists a negligible function ν such that:*

$$\Pr[x \leftarrow_r D_f, x' \leftarrow_r \mathcal{A}(f(x)) : f(x) = f(x')] \leq \nu(\kappa)$$

We assume that D_f and R_f are implicitly defined by f .

Definition 17 (Digital Signatures). *A digital signature scheme Σ consists of four algorithms $\{\text{PPGen}_\Sigma, \text{KGen}_\Sigma, \text{Sign}_\Sigma, \text{Verf}_\Sigma\}$ such that:*

PPGen_Σ . *The algorithm PPGen_Σ outputs the public parameters*

$$\text{pp}_\Sigma \leftarrow_r \text{PPGen}_\Sigma(1^\kappa)$$

We assume that pp_Σ contains 1^κ and is implicit input to all other algorithms.

KGen_Σ . *The algorithm KGen_Σ outputs the public and private key of the signer, where κ is the security parameter:*

$$(\text{sk}_\Sigma, \text{pk}_\Sigma) \leftarrow_r \text{KGen}_\Sigma(\text{pp}_\Sigma)$$

Sign_Σ . *The algorithm Sign_Σ gets as input the secret key sk_Σ and the message $m \in \mathcal{M}$ to sign. It outputs a signature:*

$$\sigma \leftarrow_r \text{Sign}_\Sigma(\text{sk}_\Sigma, m)$$

Verf_Σ . *The deterministic algorithm Verf_Σ outputs a decision bit $d \in \{0, 1\}$, indicating if the signature σ is valid, w.r.t. pk_Σ and m :*

$$d \leftarrow \text{Verf}_\Sigma(\text{pk}_\Sigma, m, \sigma)$$

For each Σ it is required that the correctness properties hold. In particular, it is required that for all $\kappa \in \mathbb{N}$, for all $\text{pp}_\Sigma \leftarrow_r \text{PPGen}_\Sigma(1^\kappa)$, for all $(\text{sk}_\Sigma, \text{pk}_\Sigma) \leftarrow_r \text{KGen}_\Sigma(\text{pp}_\Sigma)$, for all $m \in \mathcal{M}$, $\text{Verf}_\Sigma(\text{pk}_\Sigma, m, \text{Sign}_\Sigma(\text{sk}_\Sigma, m)) = 1$ is true. This definition captures perfect correctness.

We require existential unforgeability (eUNF-CMA) of digital signature schemes. In a nutshell, unforgeability requires that an adversary \mathcal{A} cannot (except with negligible probability) come up with a signature for a message m^* for which the adversary did not see any signature before. As usual, the adversary \mathcal{A} can adaptively query for signatures on messages of its own choice.

```

Experiment  $\text{eUNF-CMA}_{\mathcal{A}}^\Sigma(\kappa)$ 
 $\text{pp}_\Sigma \leftarrow_r \text{PPGen}_\Sigma(1^\kappa)$ 
 $(\text{sk}_\Sigma, \text{pk}_\Sigma) \leftarrow_r \text{KGen}_\Sigma(\text{pp}_\Sigma)$ 
 $\mathcal{Q} \leftarrow \emptyset$ 
 $(m^*, \sigma^*) \leftarrow_r \mathcal{A}^{\text{Sign}'_\Sigma(\text{sk}_\Sigma, \cdot)}(\text{pk}_\Sigma)$ 
  where  $\text{Sign}'_\Sigma$  on input  $\text{sk}_\Sigma$  and  $m$ :
     $\sigma \leftarrow_r \text{Sign}_\Sigma(\text{sk}_\Sigma, m)$ 
    set  $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{m\}$ 
    return  $\sigma$ 
return 1, if  $\text{Verf}_\Sigma(\text{pk}_\Sigma, m^*, \sigma^*) = 1 \wedge m^* \notin \mathcal{Q}$ 
return 0

```

Fig. 13: Σ Unforgeability

Definition 18 (Σ Unforgeability). *We say a Σ scheme is unforgeable, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:*

$$\Pr \left[\mathbf{Exp}_{\mathcal{A}, \Sigma}^{\text{eUNF-CMA}}(\kappa) = 1 \right] \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 13.

For our definition of public key encryption we need an additional algorithm KVrf_Π verifying if a given key pair is valid along with a corresponding security notion requiring that even for adversarially chosen public keys one can find at most one corresponding secret key.

Definition 19 (Public-Key Encryption). *A public-key encryption-scheme Π consists of five algorithms $\{\text{PPGen}_\Pi, \text{KGen}_\Pi, \text{Enc}_\Pi, \text{Dec}_\Pi, \text{KVrf}_\Pi\}$*

PPGen_Π . The algorithm PPGen_Π outputs the public parameters of the scheme:

$$\text{pp}_\Pi \leftarrow_r \text{PPGen}_\Pi(1^\kappa)$$

It is assumed that pp_Π is implicit input to all other algorithms. Also, this algorithm may be omitted, if it is clear from the context.

KGen_Π . The algorithm KGen_Π outputs the public and private key, on input pp_Π :

$$(\text{sk}_\Pi, \text{pk}_\Pi) \leftarrow_r \text{KGen}_\Pi(\text{pp}_\Pi)$$

Enc_Π . The algorithm Enc_Π gets as input the public key pk_Π , and a message $m \in \mathcal{M}$ to encrypt. It outputs a ciphertext:

$$c \leftarrow_r \text{Enc}_\Pi(\text{pk}_\Pi, m)$$

Dec_Π . The deterministic algorithm Dec_Π outputs a message m (or \perp , if the ciphertext is invalid) on input sk_Π , and a ciphertext c :

$$m \leftarrow \text{Dec}_\Pi(\text{sk}_\Pi, c)$$

KVrf_Π . The deterministic algorithm KVrf_Π decides whether a given secret key sk_Π belongs to pk_Π , outputting a decision bit $b \in \{1, 0\}$.

$$b \leftarrow \text{KVrf}_\Pi(\text{sk}_\Pi, \text{pk}_\Pi)$$

For each Π , the usual correctness properties must hold. In particular, it is required that for all $\kappa \in \mathbb{N}$, for all $\text{pp}_\Pi \leftarrow_r \text{PPGen}_\Pi(1^\kappa)$, for all $(\text{sk}_\Pi, \text{pk}_\Pi) \leftarrow_r \text{KGen}_\Pi(\text{pp}_\Pi)$, for all $m \in \mathcal{M}$, it holds that $\text{Dec}_\Pi(\text{sk}_\Pi, \text{Enc}_\Pi(\text{pk}_\Pi, m)) = m$ and $\text{KVrf}_\Pi(\text{sk}_\Pi, \text{pk}_\Pi) = 1$ are true.

Moreover, we require that the encryption scheme is Π is IND-CCA2 secure and key-verifiable.

Definition 20 (Π IND-CCA2 Security). An encryption scheme Π is IND-CCA2 secure, if for any PPT adversary \mathcal{A} there exists a negligible function ν such that:

$$\left| \Pr \left[\mathbf{Exp}_{\mathcal{A}, \Pi}^{\text{IND-CCA2}}(\kappa) = 1 \right] - \frac{1}{2} \right| \leq \nu(\kappa)$$

The corresponding experiment is depicted in Figure 14.

Definition 21 (Π Key-Verifiability). An encryption scheme Π is key-verifiable, if for any PPT adversary \mathcal{A} there exists a negligible function ν such that:

$$\Pr \left[\mathbf{Exp}_{\mathcal{A}, \Pi}^{\text{Key-Verifiability}}(\kappa) = 1 \right] \leq \nu(\kappa)$$

The corresponding experiment is depicted in Figure 15.

```

Exp $\mathcal{A}, \Pi$ IND-CCA2( $\kappa$ )
   $\text{pp}_\Pi \leftarrow_r \text{PPGen}_\Pi(1^\kappa)$ 
   $(\text{sk}_\Pi, \text{pk}_\Pi) \leftarrow_r \text{KGen}_\Pi(\text{pp}_\Pi)$ 
   $b \leftarrow_r \{0, 1\}$ 
   $((m_0^*, m_1^*), \text{state}_\mathcal{A}) \leftarrow_r \mathcal{A}^{\text{Dec}_\Pi(\text{sk}_\Pi, \cdot)}(\text{pk}_\Pi)$ 
  If  $|m_0^*| \neq |m_1^*| \vee m_0^* \notin \mathcal{M} \vee m_1^* \notin \mathcal{M}$ :
     $c^* \leftarrow \perp$ 
  Else:
     $c^* \leftarrow_r \text{Enc}_\Pi(\text{pk}_\Pi, m_b^*)$ 
     $b^* \leftarrow_r \mathcal{A}^{\text{Dec}'_\Pi(\text{sk}_\Pi, \cdot)}(\text{state}_\mathcal{A}, c^*)$ 
    where  $\text{Dec}'_\Pi$  on input  $\text{sk}_\Pi$  and  $c$ :
      return  $\perp$ , if  $c = c^*$ 
      return  $\text{Dec}_\Pi(\text{sk}_\Pi, c)$ 
  return 1, if  $b^* = b$ 
  return 0

```

Fig. 14: Π IND-CCA2 Security

```

Exp $\mathcal{A}, \Pi$ Key-Verifiability( $\kappa$ )
   $\text{pp}_\Pi \leftarrow_r \text{PPGen}_\Pi(1^\kappa)$ 
   $(\text{sk}_0^*, \text{sk}_1^*, \text{pk}^*) \leftarrow_r \mathcal{A}(\text{pp}_\Pi)$ 
  return 0, if  $\text{KVrf}_\Pi(\text{sk}_0^*, \text{pk}^*) = 0 \vee \text{KVrf}_\Pi(\text{sk}_1^*, \text{pk}^*) = 0$ 
  return 1, if  $\text{sk}_0^* \neq \text{sk}_1^*$ 
  return 0

```

Fig. 15: Π Key-Verifiability

Non-Interactive Proof Systems. Let L be an NP-language with associated witness relation R , i.e., such that $L = \{x \mid \exists w : R(x, w) = 1\}$. A non-interactive proof system allows to prove membership of some statement x in the language L . More formally, such a system is defined as follows.

Definition 22 (Non-Interactive Proof System). A non-interactive proof system Ω for language L consists of three algorithms $\{\text{PPGen}_\Omega, \text{Prove}_\Omega, \text{Verify}_\Omega\}$, such that:

PPGen_Ω . The algorithm PPGen_Ω outputs public parameters of the scheme, where κ is the security parameter:

$$\text{crs}_\Omega \leftarrow_r \text{PPGen}_\Omega(1^\kappa)$$

Prove_Ω . The algorithm Prove_Ω outputs the proof π , on input of the CRS crs_Ω , statement x to be proven, and the corresponding witness w :

$$\pi \leftarrow_r \text{Prove}_\Omega(\text{crs}_\Omega, x, w)$$

Verify_Ω . The deterministic algorithm Verify_Ω verifies the proof π by outputting a bit $d \in \{0, 1\}$, w.r.t. to some CRS crs_Ω and some statement x :

$$d \leftarrow \text{Verify}_\Omega(\text{crs}_\Omega, x, \pi)$$

Definition 23 (Completeness). A non-interactive proof system is called complete, if for all $\kappa \in \mathbb{N}$, for all $\text{crs}_\Omega \leftarrow_r \text{PPGen}_\Omega(1^\kappa)$, for all $x \in L$, for all w such that $R(x, w) = 1$, for all $\pi \leftarrow_r \text{Prove}_\Omega(\text{crs}_\Omega, x, w)$, it holds that $\text{Verify}_\Omega(\text{crs}_\Omega, x, \pi) = 1$.

In addition to completeness, we require two standard security notions for zero-knowledge proofs of knowledge: zero-knowledge and simulation-sound extractability. We define them analogous to the definitions given in [DS19].

Informally speaking, zero-knowledge says that the receiver of the proof π does not learn anything except the validity of the statement.

Exp $_{\mathcal{A}, \Omega, \text{SIM}}^{\text{Zero-Knowledge}}(\kappa)$
 $(\text{crs}_\Omega, \tau) \leftarrow_r \text{SIM}_1(1^\kappa)$
 $b \leftarrow_r \{0, 1\}$
 $b^* \leftarrow_r \mathcal{A}^{P_b(\cdot, \cdot)}(\text{crs}_\Omega)$
 where P_0 on input x and w :
 return $\pi \leftarrow_r \text{Prove}_\Omega(\text{crs}_\Omega, x, w)$, if $R(x, w) = 1$
 return \perp
 and P_1 on input (x, w) :
 return $\pi \leftarrow_r \text{SIM}_2(\text{crs}_\Omega, \tau, x)$, if $R(x, w) = 1$
 return \perp
 return 1, if $b^* = b$
 return 0

Fig. 16: Ω Zero-Knowledge

Definition 24 (Zero-Knowledge). A non-interactive proof system Ω for language L is zero-knowledge, if for any PPT adversary \mathcal{A} , there exists an PPT simulator $\text{SIM} = (\text{SIM}_1, \text{SIM}_2)$ such that there exist negligible

functions ν_1 and ν_2 such that

$$\left| \Pr [\text{crs}_\Omega \leftarrow_r \text{PPGen}_\Omega(1^\kappa) : \mathcal{A}(\text{crs}_\Omega) = 1] - \Pr [(\text{crs}_\Omega, \tau) \leftarrow_r \text{SIM}_1(1^\kappa) : \mathcal{A}(\text{crs}_\Omega) = 1] \right| \leq \nu_1(\kappa),$$

and that

$$\left| \Pr [\text{Exp}_{\mathcal{A}, \Omega, \text{SIM}}^{\text{Zero-Knowledge}}(\kappa) = 1] - \frac{1}{2} \right| \leq \nu_2(\kappa),$$

where the corresponding experiment is depicted in Figure 16.

Simulation-sound extractability says every adversary which is able to come up with a proof π^* for a statement must know the witness, even when seeing proofs for statements potentially not in L . Clearly, this implies that the proofs output by a simulation-sound extractable proof-systems are non-malleable. Note that the definition of simulation-sound extractability

```

Exp $\mathcal{A}, \Omega, \mathcal{E}$ SimSoundExt( $\kappa$ )
  ( $\text{crs}_\Omega, \tau, \zeta$ )  $\leftarrow_r \mathcal{E}_1(1^\kappa)$ 
  ( $x^*, \pi^*$ )  $\leftarrow_r \mathcal{A}^{\text{SIM}(\cdot)}(\text{crs}_\Omega)$ 
   $\mathcal{Q} \leftarrow \emptyset$ 
  where SIM on input  $x$ :
    obtain  $\pi \leftarrow_r \text{SIM}_2(\text{crs}_\Omega, \tau, x)$ 
     $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(x, \pi)\}$ 
    return  $\pi$ 
   $w^* \leftarrow_r \mathcal{E}_2(\text{crs}_\Omega, \zeta, x^*, \pi^*)$ 
  return 1, if  $\text{Verify}_\Omega(x^*, \pi^*) = 1 \wedge R(x^*, w^*) = 0 \wedge (x^*, \pi^*) \notin \mathcal{Q}$ 
  return 0

```

Fig. 17: Ω Simulation-Sound Extractability

of [Gro06] is stronger than ours in the sense that the adversary also gets the trapdoor ζ as input. However, in our context this weaker notion (previously also used [ADK⁺13, DHLW10, DS19]) suffices.

Definition 25 (Simulation-Sound Extractability). *A zero-knowledge non-interactive proof system Ω for language L is said to be simulation-sound extractable, if for any PPT adversary \mathcal{A} , there exists a PPT extractor*

$\mathcal{E} = (\mathcal{E}_1, \mathcal{E}_2)$, such that

$$\left| \Pr[(\text{crs}_\Omega, \tau) \leftarrow_r \text{SIM}_1(1^\kappa) : \mathcal{A}(\text{crs}_\Omega, \tau) = 1] - \Pr[(\text{crs}_\Omega, \tau, \zeta) \leftarrow_r \mathcal{E}_1(1^\kappa) : \mathcal{A}(\text{crs}_\Omega, \tau) = 1] \right| = 0,$$

and that there exist a negligible function ν so that

$$\Pr[\mathbf{Exp}_{\mathcal{A}, \Omega, \mathcal{E}}^{\text{SimSoundExt}}(\kappa)] = 1 \leq \nu(\kappa),$$

where the corresponding experiment is depicted in Figure 17.

Supporting Labels. We note that to support labels, the definition of the Prove_Ω , Verify_Ω , SIM_2 , and \mathcal{E}_2 algorithms also take a public label ℓ as input, and the completeness, soundness, and zero-knowledge properties are updated accordingly (cf. [DHLW10]). Achieving this for SSE NIZK proofs obtained via the Fiat-Shamir transform from Σ -protocol [FKMV12] can be efficiently done by including the label into the hash computation (cf. [ABM15]).

B More Preliminaries and Building Blocks

This section is devoted to give additional background on the building blocks.

B.1 Additional Preliminaries

We first give some additional preliminaries required to understand the concrete constructions given in Appendix C.

Known-Order Group Definitions and Assumptions.

Group Generator. Let $(\mathbb{G}, g, q) \leftarrow_r \text{DLGen}(1^\kappa)$ be a group-generator, where \mathbb{G} is multiplicatively written and of prime-order q , where g is a generator.

Discrete Logarithm Assumption. Let $(\mathbb{G}, g, q) \leftarrow_r \text{DLGen}(1^\kappa)$ be as defined above. The discrete logarithm assumption states that given g^x for some random $x \leftarrow_r \mathbb{Z}_q$, it is hard to find that x .

Definition 26 (Discrete Logarithm Assumption). *The discrete logarithm assumption holds, if for every PPT adversary \mathcal{A} there exists a negligible function ν such that:*

$$\Pr[(\mathbb{G}, g, q) \leftarrow_r \text{DLGen}(1^\kappa), x \leftarrow_r \mathbb{Z}_q, x' \leftarrow_r \mathcal{A}(\mathbb{G}, g, q, g^x) : x = x'] \leq \nu(\kappa)$$

Decisional Diffie-Hellman Assumption. Let $(\mathbb{G}, g, q) \leftarrow_r \text{DLGen}(1^\kappa)$ be as defined above. The decisional Diffie-Hellman (DDH) assumption states that given (g^x, g^y, g^{xy}) is computationally indistinguishable from (g^x, g^y, g^z) for some random $x, y, z \leftarrow_r \mathbb{Z}_q^3$.

Definition 27 (Decisional Diffie-Hellman Assumption). *The decisional Diffie-Hellman assumption holds, if for every PPT adversary \mathcal{A} there exists a negligible function ν such that:*

$$\left| \Pr[(\mathbb{G}, g, q) \leftarrow_r \text{DLGen}(1^\kappa), (x, y, z) \leftarrow_r \mathbb{Z}_q^3, b \leftarrow_r \{0, 1\} \right. \\ \left. b' \leftarrow_r \mathcal{A}(\mathbb{G}, g, q, g^x, g^y, g^{bxy+(1-b)z}) : b = b' \right] - 1/2 \right| \leq \nu(\kappa)$$

Unknown-Order Group Definitions and Assumptions.

RSA Key-Generator. Let $(N, p, q, e, d) \leftarrow_r \text{RSAGen}(1^\kappa)$ be an instance generator which returns an RSA modulus $N = pq$, where p and q are distinct primes, $e > 1$ an integer co-prime to $\varphi(n)$, and $de \equiv 1 \pmod{\varphi(n)}$. We require that RSAGen always outputs moduli with the same bit-length, based on κ .

The One-More-RSA Inversion Assumption [BNPS03]. Let $(n, e, d, p, q) \leftarrow_r \text{RSAGen}(1^\kappa)$ be an RSA-key generator returning an RSA modulus $n = pq$, where p and q are random distinct primes, $e > 1$ an integer co-prime to $\varphi(n)$, and $d \equiv e^{-1} \pmod{\varphi(n)}$. The one-more-RSA-assumption associated to RSAGen is provided an inversion oracle \mathcal{I} , which inverts any element $x \in \mathbb{Z}_n^*$ w.r.t. e , and a challenge oracle \mathcal{C} , which at each call returns a random element $y_i \in \mathbb{Z}_n^*$, it is hard to invert more challenges than calls to the inversion oracle.

Definition 28 (One-More-RSA Inversion Assumption). *The one-more RSA inversion assumption holds, if for every PPT adversary \mathcal{A} there exists a negligible function ν such that:*

$$\Pr[(n, p, q, e, d) \leftarrow_r \text{RSAGen}(1^\kappa), X \leftarrow_r \mathcal{A}(n, e)^{\mathcal{C}(n), \mathcal{I}(d, n, \cdot)} : \\ \text{more values returned by } \mathcal{C} \text{ are inverted than queries to } \mathcal{I}] \leq \nu(\kappa)$$

Here, X is the set of inverted challenges.

We require that e is larger than any possible n w.r.t. κ and that it is prime. Re-stating the assumption with this condition is straightforward. In this case, it is also required that e is drawn independently from p , q , or n (and d is then calculated from e , and not vice versa). This can, e.g., be

achieved by demanding that e is drawn uniformly from $[n'+1, \dots, 2n'] \cap \{p \mid p \text{ is prime}\}$, where n' is the largest RSA modulus possible w.r.t. to κ . The details are left to the concrete instantiation of `RSAGen`.

B.2 Additional Building Blocks

We now present our additional building blocks.

Standard Chameleon-Hashes. Chameleon-hashes behave similar to standard collision-resistant hash-functions, but allow to find arbitrary collisions, if a trapdoor is known [KR00].

The following framework is derived from Camenisch et al. [CDK⁺17].

Definition 29 (Chameleon-Hashes). *A chameleon-hash CH consists of five algorithms $(\text{PPGen}_{\text{CH}}, \text{KGen}_{\text{CH}}, \text{Hash}_{\text{CH}}, \text{Verify}_{\text{CH}}, \text{Adapt}_{\text{CH}})$, such that:*

PPGen_{CH}. *The algorithm PPGen_{CH} on input security parameter κ outputs public parameters pp_{CH} of the scheme. For brevity, we assume that pp_{CH} is implicit input to all other algorithms:*

$$\text{pp}_{\text{CH}} \leftarrow_r \text{PPGen}_{\text{CH}}(1^\kappa)$$

KGen_{CH}. *The algorithm KGen_{CH} , given the public parameters pp_{CH} , outputs the private (sk_{CH}) and public key (pk_{CH}) of the scheme*

$$(\text{sk}_{\text{CH}}, \text{pk}_{\text{CH}}) \leftarrow_r \text{KGen}_{\text{CH}}(\text{pp}_{\text{CH}})$$

Hash_{CH}. *The algorithm Hash_{CH} gets as input the public key pk_{CH} , and a message m to hash. It outputs a hash h , and some randomness r :*

$$(h, r) \leftarrow_r \text{Hash}_{\text{CH}}(\text{pk}_{\text{CH}}, m)$$

Verify_{CH}. *The deterministic algorithm $\text{Verify}_{\text{CH}}$ gets as input the public key pk_{CH} , a message m , randomness r , and a hash h . It outputs a decision $d \in \{0, 1\}$ indicating whether the hash h is valid:*

$$d \leftarrow \text{Verify}_{\text{CH}}(\text{pk}_{\text{CH}}, m, h, r)$$

Adapt_{CH}. *The algorithm Adapt_{CH} on input of secret key sk , the old message m , the old randomness r , hash h , and a new message m' outputs new randomness r' :*

$$r' \leftarrow_r \text{Adapt}_{\text{CH}}(\text{sk}_{\text{CH}}, m, m', r, h)$$

Note that we assume that the Adapt_{CH} algorithm always verifies if the hash it is given is valid, and outputs \perp otherwise.

For a CH we require the correctness property to hold. In particular, we require that for all $\kappa \in \mathbb{N}$, for all $\text{pp}_{\text{CH}} \leftarrow_r \text{PPGen}_{\text{CH}}(1^\kappa)$, for all $(\text{sk}_{\text{CH}}, \text{pk}_{\text{CH}}) \leftarrow_r \text{KGen}_{\text{CH}}(\text{pp}_{\text{CH}})$, for all $m \in \mathcal{M}$, for all $(h, r) \leftarrow_r \text{Hash}_{\text{CH}}(\text{pk}_{\text{CH}}, m)$, for all $m' \in \mathcal{M}$, we have for all for all $r' \leftarrow_r \text{Adapt}_{\text{CH}}(\text{sk}_{\text{CH}}, m, m', r, h)$, that $1 = \text{Verify}_{\text{CH}}(\text{pk}_{\text{CH}}, m, h, r) = \text{Verify}_{\text{CH}}(\text{pk}_{\text{CH}}, m', h, r')$. This definition captures perfect correctness.

The randomness is drawn by Hash_{CH} , and not outside. This was done to capture “private-coin” constructions [AMVA17].

Next, we present security notions of CHs.

Full Indistinguishability. Indistinguishability requires that the randomnesses r does not reveal if it was obtained through $\text{Hash}_{\text{CHET}}$ or $\text{Adapt}_{\text{PCH}}$, which is captured by the HashOrAdapt -oracle. The messages are chosen by the adversary.

We relax the indistinguishability definition by Brzuska et al. [BFF⁺09] to a computational version, which is enough for most use-cases, including ours. However, compared to the existing definitions in [BCD⁺17, CDK⁺17, DSSS19, KPSS18a], the adversary is now also allowed to *generate* the secret keys involved.

```

Exp $\mathcal{A}, \text{CH}$ FullIndistinguishability( $\kappa$ )
   $\text{pp}_{\text{CH}} \leftarrow_r \text{PPGen}_{\text{CH}}(1^\kappa)$ 
   $b \leftarrow_r \{0, 1\}$ 
   $b^* \leftarrow_r \mathcal{A}^{\text{HashOrAdapt}(\cdot, \cdot, \cdot, b)}(\text{pp}_{\text{CH}})$ 
  where oracle  $\text{HashOrAdapt}$  on input  $\text{sk}_{\text{CH}}, \text{pk}_{\text{CH}}, m, m', b$ :
     $(h, r) \leftarrow_r \text{Hash}_{\text{CH}}(\text{pk}_{\text{CH}}, m')$ 
     $(h', r') \leftarrow_r \text{Hash}_{\text{CH}}(\text{pk}_{\text{CH}}, m)$ 
     $r'' \leftarrow_r \text{Adapt}_{\text{CH}}(\text{sk}_{\text{CH}}, m, m', r', h')$ 
    return  $\perp$ , if  $r'' = \perp \vee r' = \perp \vee r = \perp$ 
    if  $b = 0$ , return  $(h, r)$ 
    if  $b = 1$ , return  $(h', r'')$ 
  return 1, if  $b^* = b$ 
  return 0

```

Fig. 18: CH Full Indistinguishability

We return \perp in the HashOrAdapt oracle (in case of an error), as the adversary \mathcal{A} may try to enter a message $m \notin \mathcal{M}$, even if $\mathcal{M} = \{0, 1\}^*$,

which makes the algorithm output \perp . If we would not do this, the adversary could trivially decide which case it sees. For similar reasons these checks are also included in other definitions.

Definition 30 (CH Full Indistinguishability). *We say a CH scheme is fully indistinguishable, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:*

$$\Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{CH}}^{\text{FullIndistinguishability}}(\kappa) = 1 \right] \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 18.

Collision-Resistance. Collision-resistance says, that even if an adversary has access to an adapt oracle, it cannot find any collisions for messages other than the ones queried to the adapt oracle. Note, this is an even stronger definition than key-exposure freeness [AdM04]: key-exposure freeness only requires that one cannot find a collision for some new “tag”, i.e., for some auxiliary value for which the adversary has never seen a collision.

```

Exp $\mathcal{A}, \text{CH}$ Collision-Resistance( $\kappa$ )
  ppCH  $\leftarrow_r$  PPGenCH( $1^\kappa$ )
  (skCH, pkCH)  $\leftarrow_r$  KGenCH(ppCH)
  Q  $\leftarrow \emptyset$ 
  ( $m^*, r^*, m'^*, r'^*, h^*$ )  $\leftarrow_r$   $\mathcal{A}^{\text{Adapt}'_{\text{CH}}(\text{sk}_{\text{CH}}, \cdot, \cdot, \cdot)}(\text{pk}_{\text{CH}})$ 
  where  $\text{Adapt}'_{\text{CH}}$  on input skCH,  $m, m', r, h$ :
     $r' \leftarrow_r \text{Adapt}_{\text{CH}}(\text{sk}_{\text{CH}}, m, m', r, h)$ 
    Q  $\leftarrow Q \cup \{m, m'\}$ 
    return  $r'$ 
  return 1, if  $\text{Verify}_{\text{CH}}(\text{pk}_{\text{CH}}, m^*, h^*, r^*) = \text{Verify}_{\text{CH}}(\text{pk}_{\text{CH}}, m'^*, h^*, r'^*) = 1 \wedge$ 
     $m^* \notin Q \wedge m^* \neq m'^*$ 
  return 0

```

Fig. 19: CH Collision-resistance

Definition 31 (CH Collision-Resistance). *We say a CH scheme is collision-resistant, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:*

$$\Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{CH}}^{\text{Collision-Resistance}}(\kappa) = 1 \right] \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 19.

Uniqueness. Uniqueness requires that it be hard to come up with two different randomness values for the same message m^* such that the hashes are equal, for the same adversarially chosen pk^* .

```

ExpA,CHUniqueness( $\kappa$ )
   $\text{pp}_{\text{CH}} \leftarrow_r \text{PPGen}_{\text{CH}}(1^\kappa)$ 
   $(\text{pk}^*, m^*, r^*, r'^*, h^*) \leftarrow_r \mathcal{A}(\text{pp}_{\text{CH}})$ 
  return 1, if  $\text{Verify}_{\text{CH}}(\text{pk}^*, m^*, h^*, r^*) = \text{Verify}_{\text{CH}}(\text{pk}^*, m^*, h^*, r'^*) = 1$ 
      $\wedge r^* \neq r'^*$ 
  return 0

```

Fig. 20: CH Uniqueness

Definition 32 (CH Uniqueness). *We say a CH scheme is unique, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:*

$$\Pr \left[\mathbf{Exp}_{\mathcal{A},\text{CH}}^{\text{Uniqueness}}(\kappa) = 1 \right] \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 20.

We do not consider uniqueness as a fundamental security property, as it depends on the concrete use-case whether this notion is required.

Chameleon-Hashes with Ephemeral Trapdoors. We recall the notion of chameleon-hashes with ephemeral trapdoors (CHET) from [CDK⁺17]. This primitive is a variant of a chameleon-hash where, in addition to the long-term trapdoor, another ephemeral trapdoor etd (chosen freshly during hashing) is required to compute collisions.

Definition 33 (Chameleon-Hashes with Ephemeral Trapdoors). *A chameleon-hash with ephemeral trapdoors CHET is a tuple of five algorithms $(\text{PPGen}_{\text{CHET}}, \text{KGen}_{\text{CHET}}, \text{Hash}_{\text{CHET}}, \text{Verify}_{\text{CHET}}, \text{Adapt}_{\text{CHET}})$, such that:*

$\text{PPGen}_{\text{CHET}}$: *On input security parameter κ , this algorithm outputs the public parameters pp_{CHET} .*

$$\text{pp}_{\text{CHET}} \leftarrow_r \text{PPGen}_{\text{CHET}}(1^\kappa)$$

We assume that pp_{CHET} implicitly defines the message space \mathcal{M} .

$\text{KGen}_{\text{CHET}}$: On input the public parameters pp_{CHET} , this algorithm outputs the long-term key pair $(\text{sk}_{\text{CHET}}, \text{pk}_{\text{CHET}})$:

$$(\text{sk}_{\text{CHET}}, \text{pk}_{\text{CHET}}) \leftarrow_r \text{KGen}_{\text{CHET}}(\text{pp}_{\text{CHET}})$$

$\text{Hash}_{\text{CHET}}$: On input the public key pk_{CHET} and a message m , this algorithm outputs a hash h , corresponding randomness r , as well as the ephemeral trapdoor etd :

$$(h, r, \text{etd}) \leftarrow_r \text{Hash}_{\text{CHET}}(\text{pk}_{\text{CHET}}, m)$$

$\text{Verify}_{\text{CHET}}$: On input the public key pk_{CHET} , a message m , a hash h , and randomness r , this algorithm outputs a bit $b \in \{1, 0\}$:

$$b \leftarrow \text{Verify}_{\text{CHET}}(\text{pk}_{\text{CHET}}, m, h, r)$$

$\text{Adapt}_{\text{CHET}}$: On input secret key sk_{CHET} , ephemeral trapdoor etd , a message m , a message m' , hash h , randomness r , and trapdoor information etd , this algorithm outputs randomness r' :

$$r' \leftarrow_r \text{Adapt}_{\text{CHET}}(\text{sk}_{\text{CHET}}, \text{etd}, m, m', h, r)$$

Note that we assume that the $\text{Adapt}_{\text{CHET}}$ algorithm always verifies if the hash it is given is valid, and output \perp otherwise.

For correctness, we require that for all $\kappa \in \mathbb{N}$, all $\text{pp}_{\text{CHET}} \leftarrow_r \text{PPGen}_{\text{CHET}}(1^\kappa)$, all $(\text{sk}_{\text{CHET}}, \text{pk}_{\text{CHET}}) \leftarrow_r \text{KGen}_{\text{CHET}}(\text{pp}_{\text{CHET}})$, all $m, m' \in \mathcal{M}$, all $(h, r, \text{etd}) \leftarrow_r \text{Hash}_{\text{CHET}}(\text{pk}_{\text{CHET}}, m)$, all $r' \leftarrow_r \text{Adapt}_{\text{CHET}}(\text{sk}_{\text{CHET}}, \text{etd}, m, m', h, r)$, we have that $\text{Verify}_{\text{CHET}}(\text{pk}_{\text{CHET}}, m, h, r) = \text{Verify}_{\text{CHET}}(\text{pk}_{\text{CHET}}, m', h, r') = 1$.

Full Indistinguishability. Full indistinguishability requires that it be intractable for outsiders to distinguish whether a given randomness corresponds to an output of $\text{Hash}_{\text{CHET}}$ or $\text{Adapt}_{\text{CHET}}$. This is captured within the HashOrAdapt -oracle. Note, however, that—when compared to the definitions in [BCD⁺17, CDK⁺17, DSSS19]—the adversary can additionally generate all secret keys.

Definition 34 (CHET Full Indistinguishability). *We say a CHET scheme is fully indistinguishable, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:*

$$\Pr \left[\text{Exp}_{\mathcal{A}, \text{CHET}}^{\text{FullIndistinguishability}}(\kappa) = 1 \right] \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 21.

```

Exp $\mathcal{A}, \text{CHET}$ Findistinguishability( $\kappa$ )
   $\text{pp}_{\text{CHET}} \leftarrow_r \text{PPGen}_{\text{CHET}}(1^\kappa)$ 
   $b \leftarrow_r \{0, 1\}$ 
   $b^* \leftarrow_r \mathcal{A}^{\text{HashOrAdapt}(\cdot, \cdot, \cdot, b)}(\text{pp}_{\text{CHET}})$ 
  where  $\text{HashOrAdapt}$  on input  $\text{sk}_{\text{CHET}}, \text{pk}_{\text{CHET}}, m, m', b$ :
    let  $(h_0, r_0, \text{etd}_0) \leftarrow_r \text{Hash}_{\text{CHET}}(\text{pk}, m')$ 
    let  $(h_1, r_1, \text{etd}_1) \leftarrow_r \text{Hash}_{\text{CHET}}(\text{pk}, m)$ 
    let  $r_1 \leftarrow_r \text{Adapt}_{\text{CHET}}(\text{sk}_{\text{CHET}}, \text{etd}_1, m, m', h_1, r_1)$ 
    return  $\perp$ , if  $r_0 = \perp \vee r_1 = \perp$ 
    return  $(h_b, r_b, \text{etd}_b)$ 
  return  $b = b^*$ 

```

Fig. 21: CHET Full Indistinguishability

Public Collision-Resistance. Public collision-resistance grants the adversary access to an $\text{Adapt}_{\text{PCH}}$ oracle. It requires that it is intractable to produce collisions, other than the ones produced by the $\text{Adapt}_{\text{PCH}}$ oracle. Thus, the adversary gains access to a $\text{Adapt}'_{\text{CHET}}$ -oracle, which also keeps track of the produced collisions, which we need to exclude to have a meaningful definition.

```

Exp $\mathcal{A}, \text{CHET}$ Public Collision-Resistance( $\kappa$ )
   $\text{pp}_{\text{CHET}} \leftarrow_r \text{PPGen}_{\text{CHET}}(1^\kappa)$ 
   $(\text{sk}_{\text{CHET}}, \text{pk}_{\text{CHET}}) \leftarrow_r \text{KGen}_{\text{CHET}}(\text{PPGen}_{\text{CHET}})$ 
   $\mathcal{Q} \leftarrow \emptyset$ 
   $(m^*, r^*, m'^*, r'^*, h^*) \leftarrow_r \mathcal{A}^{\text{Adapt}'_{\text{CHET}}(\text{sk}_{\text{CHET}}, \cdot, \cdot, \cdot)}(\text{pk}_{\text{CHET}})$ 
  where  $\text{Adapt}'_{\text{CHET}}$  on input  $\text{etd}, m, m', h, r$ :
     $r' \leftarrow_r \text{Adapt}_{\text{CHET}}(\text{sk}, \text{etd}, m, m', h, r)$ 
    if  $r' \neq \perp$ , let  $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{m, m'\}$ 
    return  $r'$ 
  return 1, if  $\text{Verify}_{\text{CHET}}(\text{pk}_{\text{CHET}}, m^*, h^*, r^*) = 1 \wedge$ 
     $\text{Verify}_{\text{CHET}}(\text{pk}_{\text{CHET}}, m'^*, h^*, r'^*) = 1 \wedge$ 
     $m^* \notin \mathcal{Q} \wedge m'^* \neq m^*$ 
  return 0

```

Fig. 22: CHET Public Collision-Resistance

Definition 35 (CHET Public Collision-Resistance). *We say a CHET scheme is publicly collision-resistant, if for every PPT adversary \mathcal{A} , there*

exists a negligible function ν such that:

$$\Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{CHET}}^{\text{Public Collision-Resistance}}(\kappa) = 1 \right] \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 22.

Strong Private Collision-Resistance. Strong private collision-resistance requires that it is even intractable for the holder of the secret key sk to find collisions without knowledge of etd . Note, the adversary can obtain arbitrary collisions.

```

Exp $\mathcal{A}, \text{CHET}$ Strong Private Collision-Resistance( $\kappa$ )
   $\text{PP}_{\text{CHET}} \leftarrow_r \text{PPGen}_{\text{CHET}}(1^\kappa)$ 
   $\mathcal{Q} \leftarrow \emptyset$ 
   $i \leftarrow 0$ 
   $(\text{pk}^*, m^*, r^*, m'^*, r'^*, h^*) \leftarrow_r \mathcal{A}^{\text{Hash}'_{\text{CHET}}(\cdot, \cdot), \text{Adapt}'_{\text{CHET}}(\cdot, \cdot, \cdot, \cdot, \cdot)}(\text{PP}_{\text{CHET}})$ 
  where  $\text{Hash}'_{\text{CHET}}$  on input  $\text{pk}, m$ :
     $(h, r, \text{etd}) \leftarrow_r \text{Hash}_{\text{CHET}}(\text{pk}, m)$ 
    return  $\perp$ , if  $r = \perp$ 
     $i \leftarrow i + 1$ 
    let  $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(\text{pk}, h, m, \text{etd}, i)\}$ 
    return  $(h, r)$ 
  and  $\text{Adapt}'_{\text{CHET}}$  on input  $\text{sk}, \text{pk}, h, r, m, m', i$ :
    return  $\perp$ , if  $(\text{pk}, h', m'', \text{etd}, i) \notin \mathcal{Q}$  for some  $h', m'', \text{etd}, \text{pk}$ 
     $r' \leftarrow_r \text{Adapt}_{\text{CHET}}(\text{sk}, \text{etd}, m, m', h, r)$ 
    if  $r' \neq \perp$ , let  $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(\text{pk}, h', m, \text{etd}, i), (\text{pk}, h', m', \text{etd}, i)\}$ 
    return  $r'$ 
  return 1, if  $\text{Verify}_{\text{CHET}}(\text{pk}^*, m^*, h^*, r^*) = 1 \wedge$ 
     $\text{Verify}_{\text{CHET}}(\text{pk}^*, m'^*, h^*, r'^*) = 1 \wedge m^* \neq m'^* \wedge$ 
     $(\text{pk}^*, h^*, m^*, \cdot, \cdot) \notin \mathcal{Q} \wedge (\text{pk}^*, h^*, \cdot, \cdot, \cdot) \in \mathcal{Q}$ 
  return 0

```

Fig. 23: CHET Strong Private Collision-Resistance

Definition 36 (CHET Strong Private Collision-Resistance). We say a CHET scheme is strongly privately collision-resistant, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:

$$\Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{CHET}}^{\text{Strong Private Collision-Resistance}}(\kappa) = 1 \right] \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 23.

Uniqueness. Uniqueness requires that it be hard to come up with two different randomness values for the same message m^* such that the hashes are equal, for the same adversarially chosen pk^* .

```

ExpA,CHETUniqueness( $\kappa$ )
   $\text{pp}_{\text{CHET}} \leftarrow_r \text{PPGen}_{\text{CHET}}(1^\kappa)$ 
   $(\text{pk}^*, m^*, r^*, r'^*, h^*) \leftarrow_r \mathcal{A}(\text{pp}_{\text{CHET}})$ 
  return 1, if  $\text{Verify}_{\text{CHET}}(\text{pk}^*, m^*, h^*, r^*) = \text{Verify}_{\text{CHET}}(\text{pk}^*, m^*, h^*, r'^*) = 1$ 
      $\wedge r^* \neq r'^*$ 
  return 0

```

Fig. 24: CHET Uniqueness

Definition 37 (CHET Uniqueness). *We say a CHET scheme is unique, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:*

$$\Pr \left[\mathbf{Exp}_{\mathcal{A},\text{CHET}}^{\text{Uniqueness}}(\kappa) = 1 \right] \leq \nu(\kappa).$$

The corresponding experiment is depicted in Figure 24.

Attribute-Based Encryption. Let us recall the description of a ciphertext-policy attribute encryption scheme (ABE henceforth) [BSW07].

Definition 38 (Ciphertext-Policy Attribute-Based Encryption). *A ABE scheme is a tuple of PPT algorithms $(\text{PPGen}_{\text{ABE}}, \text{KGen}_{\text{ABE}}, \text{Enc}_{\text{ABE}}, \text{Dec}_{\text{ABE}})$ such that:*

$\text{PPGen}_{\text{ABE}}(1^\kappa)$: *Takes as input a security parameter κ and outputs a master secret and public key $(\text{msk}_{\text{ABE}}, \text{mpk}_{\text{ABE}})$:*

$$(\text{msk}_{\text{ABE}}, \text{mpk}_{\text{ABE}}) \leftarrow_r \text{PPGen}_{\text{ABE}}(1^\kappa)$$

We assume that all subsequent algorithms will implicitly receive the master public key mpk_{ABE} (public parameters) as input which implicitly fixes a message space \mathcal{M} .

$\text{KGen}_{\text{ABE}}(\text{msk}_{\text{ABE}}, \mathbb{S})$: *Takes as input the master secret key msk_{ABE} and a set of attributes \mathbb{S} and outputs a secret key ssk :*

$$\text{ssk} \leftarrow_r \text{KGen}_{\text{ABE}}(\text{msk}_{\text{ABE}}, \mathbb{S})$$

$\text{Enc}_{\text{ABE}}(m, \mathbb{A})$: Takes as input a message $m \in \mathcal{M}$ and an access structure \mathbb{A} . It outputs a ciphertext c :

$$c \leftarrow_r \text{Enc}_{\text{ABE}}(m, \mathbb{A})$$

$\text{Dec}_{\text{ABE}}(\text{ssk}, c)$: Takes as input a secret key ssk and a ciphertext c and outputs a message m or \perp in case decryption does not work:

$$m \leftarrow \text{Dec}_{\text{ABE}}(\text{ssk}, c)$$

Correctness of a ABE scheme requires that for all $\kappa \in \mathbb{N}$, for all access structures \mathbb{A} , all $(\text{msk}_{\text{ABE}}, \text{mpk}_{\text{ABE}}) \leftarrow_r \text{PPGen}_{\text{ABE}}(1^\kappa)$, all $m \in \mathcal{M}$, all $\mathbb{S} \in \mathbb{A}$, all $\text{ssk} \leftarrow_r \text{KGen}_{\text{ABE}}(\text{msk}_{\text{ABE}}, \mathbb{S})$ we have that $\text{Dec}_{\text{ABE}}(\text{ssk}, \text{Enc}_{\text{ABE}}(m, \mathbb{A})) = m$.

Security of ABE. In the following, we recall adaptive IND-CCA2 security, for ABE. It is derived from the definition given by Lewko et al. [LOS⁺10] and Derler et al. [DSSS19], but altered for our used notation. Refer, e.g., to [YAHK11] for how to construct chosen-ciphertext secure ABEs from CPA-secure ones.

Definition 39 (ABE IND-CCA2-Security). *An ABE scheme is IND-CCA2-secure, if for any PPT adversary \mathcal{A} there exists a negligible function ν such that:*

$$\left| \Pr \left[\mathbf{Exp}_{\mathcal{A}, \text{ABE}}^{\text{IND-CCA2}}(\kappa) = 1 \right] - \frac{1}{2} \right| \leq \nu(\kappa)$$

The corresponding experiment is depicted in Figure 25.

C Concrete Instantiations of Primitives

We now present the instantiations of our building blocks.

Instantiation of Secure CHs. We recall a construction from [CDK⁺17] in Construction 2.

Theorem 2. *If the one-more-RSA inversion assumption [BNPS03] holds, then the construction of a CH given in Construction 2 is fully indistinguishable, correct, unique and collision-resistant, in the random-oracle model [BR93].*

Proof. All properties, but full indistinguishability, have already been proven by Camenisch et al. [CDK⁺17]. Thus, it remains to prove full indistinguishability.

$\mathbf{Exp}_{\mathcal{A}, \text{ABE}}^{\text{IND-CCA2}}(\kappa)$:
 $(\text{msk}_{\text{ABE}}, \text{mpk}_{\text{ABE}}) \leftarrow_r \text{PPGen}_{\text{ABE}}(1^\kappa)$
 $b \leftarrow_r \{0, 1\}$
 $\mathcal{Q} \leftarrow \emptyset$
 $\mathcal{S} \leftarrow \emptyset$
 $i \leftarrow 0$
 $(m_0, m_1, \mathbb{A}^*, \text{state}) \leftarrow_r \mathcal{A}^{\text{KGen}'_{\text{ABE}}(\text{msk}_{\text{ABE}}, \cdot), \text{KGen}''_{\text{ABE}}(\text{msk}_{\text{ABE}}, \cdot), \text{Dec}'_{\text{ABE}}(\cdot, \cdot)}(\text{mpk}_{\text{ABE}})$
 where $\text{KGen}'_{\text{ABE}}$ on input $\text{msk}_{\text{ABE}}, \mathbb{S}$:
 return $\text{KGen}_{\text{ABE}}(\text{msk}_{\text{ABE}}, \mathbb{S})$ and set $\mathcal{S} \leftarrow \mathcal{S} \cup \mathbb{S}$
 and $\text{KGen}''_{\text{ABE}}$ on input j, \mathbb{S} :
 let $\text{ssk} \leftarrow_r \text{KGen}_{\text{ABE}}(\text{msk}_{\text{ABE}}, \mathbb{S})$ and set $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(j, \text{ssk})\}$
 $i \leftarrow i + 1$
 and Dec'_{ABE} on input j, c :
 return \perp , if $(j, \text{ssk}) \notin \mathcal{Q}$ for some ssk
 return $\text{Dec}_{\text{ABE}}(\text{ssk}, c)$
 if $m_0, m_1 \notin \mathcal{M} \vee |m_0| \neq |m_1| \vee \mathbb{A}^* \cap \mathcal{S} \neq \emptyset$, let $c^* \leftarrow \perp$
 else let $c^* \leftarrow_r \text{Enc}_{\text{ABE}}(m_b, \mathbb{A}^*)$
 $b^* \leftarrow_r \mathcal{A}^{\text{KGen}'''_{\text{ABE}}(\text{msk}_{\text{ABE}}, \cdot), \text{KGen}''''_{\text{ABE}}(\text{msk}_{\text{ABE}}, \cdot), \text{Dec}''_{\text{ABE}}(\cdot, \cdot)}(c^*, \text{state})$
 where $\text{KGen}''''_{\text{ABE}}$ on input $\text{msk}_{\text{ABE}}, \mathbb{S}$:
 return \perp , if $\mathbb{S} \in \mathbb{A}^*$
 return $\text{KGen}_{\text{ABE}}(\text{msk}_{\text{ABE}}, \mathbb{S})$
 and $\text{KGen}'''_{\text{ABE}}$ on input j, \mathbb{S} :
 let $\text{ssk} \leftarrow_r \text{KGen}_{\text{ABE}}(\text{msk}_{\text{ABE}}, \mathbb{S})$ and set $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(j, \text{ssk})\}$
 $i \leftarrow i + 1$
 and $\text{Dec}''_{\text{ABE}}$ on input j, c :
 return \perp , if $(j, \text{ssk}) \notin \mathcal{Q}$ for some $\text{ssk} \vee c = c^*$
 return $\text{Dec}_{\text{ABE}}(\text{ssk}, c)$
 if $b^* = b$ return 1 else return 0

Fig. 25: ABE IND-CCA2 Security

Full Indistinguishability. We prove full indistinguishability by a sequence of games.

Game 0: The original full indistinguishability game in the case $b = 0$.

Game 1: As Game 0, but we now make the transition to $b = 1$.

Transition - Game 0 \rightarrow Game 1: As there is exactly one secret key (up to the group order, which can be ignored), which makes adaption work correctly, which we explicitly check, while r is always chosen randomly, the distributions are exactly equal and thus $|\Pr[S_0] - \Pr[S_1]| = 0$ follows.

As the adversary now has to other way to win the full indistinguishability game and each hop only changes the view of the adversary negligibly, full indistinguishability is proven.

<p>$\text{PPGen}_{\text{CH}}(1^\kappa)$: On input a security parameter κ it outputs the public parameters $\text{pp}_{\text{CH}} \leftarrow (1^\kappa, e)$, where e is prime and $e > N'$, and</p> $N' = \max_r \{N \in \mathbb{N} : (N, \cdot, \cdot, \cdot, \cdot) \leftarrow_r \text{RSAGen}(1^\kappa; r)\}$
<p>$\text{KGen}_{\text{CH}}(\text{pp}_{\text{CH}})$: On input $\text{pp}_{\text{CH}} = (1^\kappa, e)$ run $(N, p, q, \cdot, \cdot) \leftarrow_r \text{RSAGen}(1^\kappa)$, choose a hash-function $H : \{0, 1\}^* \rightarrow \mathbb{Z}_N^*$ (modeled as a random-oracle), compute d s.t. $ed \equiv 1 \pmod{\varphi(N)}$, set $\text{sk}_{\text{CH}} \leftarrow d$, $\text{pk}_{\text{CHET}} \leftarrow (N, H)$, and return $(\text{sk}_{\text{CH}}, \text{pk}_{\text{CH}})$.</p>
<p>$\text{Hash}_{\text{CH}}(\text{pk}_{\text{CH}}, m)$: On input a public key $\text{pk}_{\text{CH}} = (N, H)$ and a message m, choose $r \leftarrow_r \mathbb{Z}_N^*$, compute $h \leftarrow H(m)r^e \pmod N$ and output (h, r).</p>
<p>$\text{Verify}_{\text{CH}}(\text{pk}_{\text{CH}}, m, h, r)$: On input public key $\text{pk}_{\text{CH}} = (N, H)$, a message m, a hash h, and a randomness $r \in \mathbb{Z}_N^*$, it computes $h' \leftarrow H(m)r^e \pmod N$ and outputs 1 if $h' = h$ and 0 otherwise.</p>
<p>$\text{Adapt}_{\text{CH}}(\text{sk}_{\text{CH}}, m, m', h, r)$: On input a secret key $\text{sk}_{\text{CH}} = d$, messages m and m', a hash h, and randomness values r and r', the adaptation algorithm outputs \perp if $\text{Verify}_{\text{CH}}(\text{pk}_{\text{CH}}, m, h, r) \neq 1$. Otherwise, let $x \leftarrow H(m)$, $x' \leftarrow H(m')$, $y \leftarrow xr^e \pmod N$. Output \perp, if $\text{Verify}_{\text{CH}}(\text{pk}_{\text{CH}}, m', h, r') \neq 1$. Return $r' \leftarrow (y(x'^{-1}))^d \pmod N$.</p>

Construction 2: RSA-based CH

Instantiation of Secure CHETs. The generic construction is given in Construction 3. This construction is essentially the one given by Krenn et al. [KPSS18a], but we additionally check whether a hash h is valid after adaption, and use the stronger CH introduced above.

Theorem 3. *If CH is fully indistinguishable, collision-resistant, unique, and correct, then the construction of a CHET given in Construction 3 is fully indistinguishable, publicly collision-resistant, strongly private collision-resistant, unique, and correct.*

Proof. All properties, but full indistinguishability and uniqueness, have already been proven [DSSS19, KPSS18a]. We thus prove each remaining property on its own.

Full Indistinguishability. First, we prove full indistinguishability by a sequence of games.

Game 0: The original full indistinguishability game in the case $b = 1$.

Game 1: As Game 0, but instead of calculating the hash h^1 as in the game, directly hash.

Transition - Game 0 \rightarrow Game 1: We claim that Game 0 and Game 1 are indistinguishable under the full indistinguishability of CH. More formally, assume that the adversary \mathcal{A} can distinguish this hop. We can then construct an adversary \mathcal{B} which breaks the indistinguishability of

$\text{PPGen}_{\text{CHET}}(1^\kappa)$: On input a security parameter κ , let $\text{pp}_{\text{CH}} \leftarrow_r \text{PPGen}_{\text{CH}}(1^\kappa)$. Return $\text{pp}_{\text{CHET}} \leftarrow \text{pp}_{\text{CH}}$.
$\text{KGen}_{\text{CHET}}(\text{pp}_{\text{CHET}})$: On input $\text{pp}_{\text{CHET}} = \text{pp}_{\text{CH}}$ run $(\text{sk}_{\text{CH}}^1, \text{pk}_{\text{CH}}^1) \leftarrow_r \text{KGen}_{\text{CH}}(\text{pp}_{\text{CH}})$. Return $(\text{sk}_{\text{CH}}^1, \text{pk}_{\text{CH}}^1)$.
$\text{Hash}_{\text{CHET}}(\text{pk}_{\text{CHET}}, m)$: On input of $\text{pk}_{\text{CHET}} = \text{pk}_{\text{CH}}^1$ and m , let: $(\text{etd}, \text{pk}_{\text{CH}}^2) \leftarrow_r \text{KGen}_{\text{CH}}(\text{pp}_{\text{CH}})$. Let $(h^1, r^1) \leftarrow_r \text{Hash}_{\text{CH}}(\text{pk}_{\text{CH}}^1, (m, \text{pk}_{\text{CH}}^1, \text{pk}_{\text{CH}}^2))$ and $(h^2, r^2) \leftarrow_r \text{Hash}_{\text{CH}}(\text{pk}_{\text{CH}}^2, (m, \text{pk}_{\text{CH}}^1, \text{pk}_{\text{CH}}^2))$. Return
$((h^1, h^2, \text{pk}_{\text{CH}}^1, \text{pk}_{\text{CH}}^2), (r^1, r^2), \text{etd})$
$\text{Verify}_{\text{CHET}}(\text{pk}_{\text{CHET}}, m, h, r)$: On input of $\text{pk}_{\text{CHET}} = \text{pk}_{\text{CH}}^1$, m , $h = (h^1, h^2, \text{pk}_{\text{CH}}^1, \text{pk}_{\text{CH}}^2)$ and $r = (r^1, r^2)$, return 1, if
$\text{Verify}_{\text{CH}}(\text{pk}_{\text{CH}}^1, (m, \text{pk}_{\text{CH}}^1, \text{pk}_{\text{CH}}^2), h^1, r^1) = 1$
and
$\text{Verify}_{\text{CH}}(\text{pk}_{\text{CH}}^2, (m, \text{pk}_{\text{CH}}^1, \text{pk}_{\text{CH}}^2), h^2, r^2) = 1$
Otherwise, return 0.
$\text{Adapt}_{\text{CHET}}(\text{sk}_{\text{CHET}}, \text{etd}, m, m', h, r)$: On input a secret key $\text{sk}_{\text{CHET}} = \text{sk}_{\text{CH}}^1$, etd , messages m and m' , a hash $h = (h^1, h^2, \text{pk}_{\text{CH}}^1, \text{pk}_{\text{CH}}^2)$ and $r = (r^1, r^2)$, first check that $\text{Verify}_{\text{CHET}}(\text{pk}_{\text{CHET}}, m, h, r) = 1$. Otherwise, return \perp . Let
$r'^1 \leftarrow_r \text{Adapt}_{\text{CH}}(\text{sk}_{\text{CHET}}, (m, \text{pk}_{\text{CH}}^1, \text{pk}_{\text{CH}}^2), (m', \text{pk}_{\text{CH}}^1, \text{pk}_{\text{CH}}^2), r^1, h^1)$
and
$r'^2 \leftarrow_r \text{Adapt}_{\text{CH}}(\text{etd}, (m, \text{pk}_{\text{CH}}^1, \text{pk}_{\text{CH}}^2), (m', \text{pk}_{\text{CH}}^1, \text{pk}_{\text{CH}}^2), r^2, h^2)$
Let $r' \leftarrow (r'^1, r'^2)$. If $\text{Verify}_{\text{CHET}}(\text{pk}_{\text{CHET}}, m', h, r') = 0$, return \perp . Return r' .

Construction 3: Construction of a CHET

CH. In particular, the reduction works as follows. \mathcal{B} receives pp_{CH} as its own challenge, passing them through to \mathcal{A} within pp_{PCH} (generating the rest honestly), and proceeds as in the prior hop, with the exception that it uses the HashOrAdapt oracle to generate h^1 . Then, whatever \mathcal{A} outputs, is also output by \mathcal{B} . Clearly, the simulation is perfect from \mathcal{A} 's point of view. Note, the HashOrAdapt always checks if the adaption was successful, and thus so does \mathcal{B} , making the distributions equal. $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{CH-sind}}(\kappa)$ follows.

Game 2: As Game 1, but instead of calculating the hash h^2 as in the game, directly hash.

Transition - Game 1 \rightarrow Game 2: We claim that Game 1 and Game 2 are indistinguishable under the full indistinguishability of CH. More formally, assume that the adversary \mathcal{A} can distinguish this hop. We can then construct an adversary \mathcal{B} which breaks the indistinguishability of

CH. In particular, the reduction works as follows. \mathcal{B} receives pp_{CH} as its own challenge, passing them through to \mathcal{A} within pp_{PCH} (generating the rest honestly), and proceeds as in the prior hop, with the exception that it uses the `HashOrAdapt` oracle to generate h^2 . Then, whatever \mathcal{A} outputs, is also output by \mathcal{B} . Clearly, the simulation is perfect from \mathcal{A} 's point of view. Note, the `HashOrAdapt` always checks if the adaption was successful, and thus so does \mathcal{B} , making the distributions equal. $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{CH-sind}}(\kappa)$ follows.

We are now in the case $b = 0$. However, as the adversary only sees negligible changes, full indistinguishability is proven.

Uniqueness. Finally, we prove uniqueness by a sequence of games.

Game 0: The original strong private collision-resistance game.

Game 1: As Game 0, but we abort if the adversary outputs $(\text{pk}^*, m^*, r^*, r'^*, h^*)$ such that the winning conditions are fulfilled. Let this event be E_1 .

Transition - Game 0 \rightarrow Game 1: Assume that event E_1 happens. We can then construct an adversary \mathcal{B} which breaks the uniqueness of the underlying CH.

The reduction works as follows. It receives pp_{CH} from its own challenger and embeds it into pp_{CHET} . Then, when the adversary outputs $(\text{pk}^*, m^*, r^*, r'^*, h^*)$ such that the winning conditions are fulfilled, we know that $r_i^* \neq r_i'^*$ must hold for either $i = 1$ or $i = 2$ (or even both). Thus, the adversary can return $(\text{pk}'^*, (m^*, \text{pk}_{\text{CH}}^1, \text{pk}_{\text{CH}}^2), r_i^*, r_i'^*, h_i^*)$, where $\text{pk}'^* = \text{pk}_{\text{CH}}^1$ if $i = 1$ and $\text{pk}'^* = \text{pk}_{\text{CH}}^2$ otherwise, while for the hash $h^* = (h_1^*, h_2^*)$ holds. $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{CH-unique}}(\kappa)$ follows.

As now the adversary has no longer the possibility to win the uniqueness game, while each hop changes the view only negligibly, uniqueness is proven.

Instantiation of Secure PCHs. Our generic construction is depicted in Construction 4. This construction is taken from [DSSS19], but we also check whether an adaption was successful.

Theorem 4. *If ABE is IND-CCA2-secure and correct, while CHET is fully indistinguishable, strongly private collision-resistant, unique, and correct, then the construction of a PCH given in Construction 4 is fully indistinguishable, insider collision-resistant, unique, and correct.*

$\text{PPGen}_{\text{PCH}}(1^\kappa)$: Return $\text{pp}_{\text{PCH}} \leftarrow_r \text{PPGen}_{\text{CHET}}(1^\kappa)$.
$\text{MKeyGen}_{\text{PCH}}(\text{pp}_{\text{PCH}})$: Return $\text{sk}_{\text{PCH}} \leftarrow (\text{msk}_{\text{ABE}}, \text{sk}_{\text{CHET}})$ and $\text{pk}_{\text{PCH}} \leftarrow (\text{mpk}_{\text{ABE}}, \text{pk}_{\text{CHET}})$, where $(\text{sk}_{\text{CHET}}, \text{pk}_{\text{CHET}}) \leftarrow_r \text{KGen}_{\text{CHET}}(\text{pp}_{\text{PCH}})$, and $(\text{msk}_{\text{ABE}}, \text{mpk}_{\text{ABE}}) \leftarrow_r \text{PPGen}_{\text{ABE}}(1^\kappa)$.
$\text{KGen}_{\text{PCH}}(\text{sk}_{\text{PCH}}, \mathbb{S})$: Parse sk_{PCH} as $(\text{msk}_{\text{ABE}}, \text{sk}_{\text{CHET}})$ and return $\text{ssk} \leftarrow (\text{sk}_{\text{CHET}}, \text{ssk}')$, where $\text{ssk}' \leftarrow_r \text{KGen}_{\text{ABE}}(\text{msk}_{\text{ABE}}, \mathbb{S})$.
$\text{Hash}_{\text{PCH}}(\text{pk}_{\text{PCH}}, m, \mathbb{A})$: Parse pk_{PCH} as $(\text{mpk}_{\text{ABE}}, \text{pk}_{\text{CHET}})$ and return $(h, r) \leftarrow ((h_{\text{CHET}}, c), r_{\text{CHET}})$, where $(h_{\text{CHET}}, r_{\text{CHET}}, \text{etd}) \leftarrow_r \text{Hash}_{\text{CHET}}(\text{pk}_{\text{CHET}}, m)$, and $c \leftarrow_r \text{Enc}_{\text{ABE}}(\text{etd}, \mathbb{A})$.
$\text{Verify}_{\text{PCH}}(\text{pk}_{\text{PCH}}, m, h, r)$: Parse pk_{PCH} as $(\text{mpk}_{\text{ABE}}, \text{pk}_{\text{CHET}})$, h as (h_{CHET}, c) , and r as r_{CHET} . Return 1, if the following check holds and 0 otherwise: $\text{Verify}_{\text{CHET}}(\text{pk}_{\text{CHET}}, (m, c), h_{\text{CHET}}, r_{\text{CHET}}) = 1$
$\text{Adapt}_{\text{PCH}}(\text{ssk}, m, m', h, r)$: Parse ssk as $(\text{sk}_{\text{CHET}}, \text{ssk}')$ and h as (h_{CHET}, c) , and r as r_{CHET} . Check whether $\text{Verify}_{\text{PCH}}(\text{pk}_{\text{PCH}}, m, h, r) = 1$ and return \perp otherwise. Compute $\text{etd} \leftarrow \text{Dec}_{\text{ABE}}(\text{ssk}', c)$ and return \perp if $\text{etd} = \perp$. Let $r' \leftarrow r_{\text{CHET}}$, where $r'_{\text{CHET}} \leftarrow_r \text{Adapt}_{\text{CHET}}(\text{sk}_{\text{CHET}}, \text{etd}, m, m', h, r_{\text{CHET}})$. Return \perp , if $\text{Verify}_{\text{PCH}}(\text{pk}_{\text{PCH}}, m', h, r') = 0$. Return r' .

Construction 4: Black-box construction of a PCH scheme

Note, we do not require outsider collision-resistance. However, this property was already proven by Derler et al. [DSSS19].

Proof. Due to our strengthened notions, we need to prove each property on its own.

Uniqueness. First, we prove uniqueness by a sequence of games.

Game 0: The original uniqueness game.

Game 1: As Game 0, but we abort, if the adversary found $(\text{pk}^*, m^*, r^*, r'^*, h^*)$ such that it wins the uniqueness game. Let this event be E_1 .

Transition - Game 0 \rightarrow Game 1: Assume towards contradiction that event E_1 happens, we can build an adversary \mathcal{B} which breaks uniqueness of the underlying CHET. Our reduction receives pp_{CHET} and embeds it into pp_{PCH} . Then, by assumption, \mathcal{B} can directly return $(\text{pk}_2^*, m^*, r^*, r'^*, h_0^*)$, where $\text{pk}^* = (\text{pk}_0^*, \text{pk}_1^*)$ and $h^* = (h_0^*, c^*)$. $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{CHET-uniq}}(\kappa)$ follows, as c^* is part of the hash, while the randomness only applies to the CHET.

As the adversary now has no way to win the uniqueness game and the hop only changes the view of the adversary negligibly, uniqueness is proven.

Full Indistinguishability. Now, we prove full indistinguishability by a sequence of games.

Game 0: The original full indistinguishability game in the case $b = 1$.

Game 1: As Game 0, but instead of calculating the hash h as in the game, directly hash.

Transition - Game 0 \rightarrow Game 1: We claim that Game 0 and Game 1 are indistinguishable under the full indistinguishability of CHET. More formally, assume that the adversary \mathcal{A} can distinguish this hop. We can then construct an adversary \mathcal{B} which breaks the full indistinguishability of CHET. In particular, the reduction works as follows. \mathcal{B} receives pp_{CHET} as it's own challenge, passing them through to \mathcal{A} within pp_{PCH} (generating the rest honestly), and proceeds as in the prior game, with the exception that it uses the `HashOrAdapt` oracle to generate h_{CHET} . Then, whatever \mathcal{A} outputs, is also output by \mathcal{B} . Clearly, the simulation is perfect from \mathcal{A} 's point of view. Note, the `HashOrAdapt` always checks if the adaption was successful, and thus so does \mathcal{B} , making the output behave the same. $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{CHET-sInd}}(\kappa)$ follows.

We are now in the case $b = 0$. However, as the adversary only sees negligible changes, full indistinguishability is proven. Note, the ciphertext is distributed equally in all cases.

Insider Collision-Resistance. Finally, we prove insider collision-resistance by a sequence of games.

Game 0: The original insider collision-resistance game.

Game 1: As Game 0, but we abort, if the adversary makes a query (m, m', h, r, j) , for which h verifies, to the adaption oracle, for a h returned by the hashing oracle, but m has never been input to the hashing oracle or the adaption oracle, and \mathcal{A} does not have enough attributes to find a collision all by itself. Let this event be E_1 .

Transition - Game 0 \rightarrow Game 1: Assume that event E_1 happens with non-negligible probability. We can then construct a reduction \mathcal{B} which breaks the strong private collision-resistance of the underlying CHET. Our reduction \mathcal{B} works as follows. Let q be an upper bound on the queries to the hashing oracle. The adversary \mathcal{B} then makes a guess $i \leftarrow_r \{1, 2, \dots, q\}$. All queries, but the i th one, are answered as in the prior game. On the i th query, however, \mathcal{B} encrypts 0 instead of the real etd. If, at some point, the adversary has asked or asks to receive `ssk` which would allow to decrypt that c , we abort. However, by assumption, this does not happen in at least one case, thus we at

most lose a factor of q . Further assume, towards contradiction, that \mathcal{B} guessed right, but \mathcal{A} behaves noticeably different now. Our reduction \mathcal{B} can then use \mathcal{A} to break the IND-CCA2 security of the used ABE. The reduction proceeds as follows. It receives mpk_{ABE} as its own challenge, and embeds it accordingly. The oracles are simulated as follows:

Before the challenge ciphertext is embedded on the i th query (see below), every query to $\text{KGen}'_{\text{PCH}}$ is answered by the $\text{KGen}'_{\text{ABE}}$ -oracle provided. However, calls to $\text{KGen}''_{\text{PCH}}$ are simply stored as (j, \mathbb{S}) by \mathcal{B} . Hashing is done honestly for all queries except for the i th query, where the reduction queries its own challenger with either 0 or the correct etd , embedding the response c in the returned h . All following queries are performed honestly. After this embedding, all queries to the $\text{KGen}''_{\text{PCH}}$ -oracle are redirected to the $\text{KGen}'''_{\text{ABE}}$ -oracle, while queries to the $\text{KGen}''_{\text{PCH}}$ -oracle are again stored as (j, \mathbb{S}) . Note, by assumption \mathcal{A} never queries for keys which would allow decrypting that ciphertext. Adaption is done in such a way that if h was generated by the hashing-oracle, then we only continue if (j, \mathbb{S}) is sufficient to decrypt (note, h is known to \mathcal{B} , including the access structure \mathbb{A} used to generate that hash). Finally, for every decryption necessary during adaption, i.e., for ciphertexts not generated by the reduction (and ssk_j , defined by the index j , is actually sufficient to adapt; c , as part of h , never needs to be decrypted, even if it is re-used in another hash), \mathcal{B} uses the provided decryption oracle to receive each etd , and proceeds like in the game. Note, adaption can still be performed honestly, as all etds are thus known. Then, whatever \mathcal{A} outputs, is also output by \mathcal{B} .

We are now in the case that etd is no longer given to the adversary \mathcal{A} . However, this now also means that the adversary \mathcal{A} was able to find a collision, without ever having grasp on the valuable information of etd . Thus, \mathcal{B} can finally use this adversary to break the strong private collision-resistance of CHET. Consider the following reduction \mathcal{B} : it receives pp_{CHET} and embeds it into pp_{PCH} . $(\text{sk}_{\text{CHET}}, \text{pk}_{\text{CHET}}) \leftarrow_r \text{KGen}_{\text{CHET}}(\text{PPGen}_{\text{CHET}})$ is generated honestly. It then uses those to initialize the adversary \mathcal{A} . The ABE-part is done as before. The reduction \mathcal{B} now proceeds as follows: every hash is generated honestly, but the i th one; here, the oracle $\text{Hash}'_{\text{CHET}}$ is queried. All adaptations, but the challenge one, can be performed honestly (as described above with the decryption oracle provided). For the challenge one, however, \mathcal{B} uses its own oracle to find the collision. Then, if E_1 happens, \mathcal{B} can return $((m^*, c^*), r^*, (m'^*, c^*), r'^*, h_0^*)$ by assumption, where $h^* = (h_0^*, c^*)$ by construction.

$|\Pr[S_0] - \Pr[S_1]| \leq q(\nu_{\text{ABE-CCA2}}(\kappa) + \nu_{\text{CHET-SPrivColl}}(\kappa))$ follows, where q is the number of queries to the hashing oracle.

Game 2: As Game 1, but we abort, if the adversary outputs $(m^*, r^*, m'^*, r'^*, h'^*)$, such that the winning conditions are fulfilled. Let this event be E_2 .

Transition - Game 1 \rightarrow Game 2: Assume that event E_2 happens with non-negligible probability. We can then construct a reduction \mathcal{B} which breaks the strong private collision-resistance of the underlying CHET. Our reduction \mathcal{B} works as follows. Let q be an upper bound on the queries to hashing oracle. The adversary \mathcal{B} then makes a guess $i \leftarrow_r \{1, 2, \dots, q\}$. All queries, but the i th one, are answered as in the prior game. On the i th query, however, \mathcal{B} encrypts 0 instead of the real etd. If, at some point, the adversary has asked or asks to receive ssk which would allow to decrypt that c , we abort. However, by assumption, this does not happen in at least one case, thus we at most lose a factor of q . Further assume, towards contradiction, that \mathcal{B} guessed right, but \mathcal{A} behaves noticeably different now. Our reduction \mathcal{B} can then use \mathcal{A} to break the IND-CCA2 security of the used ABE. The reduction proceeds as follows. It receives mpk_{ABE} as its own challenge, and embeds it accordingly. The oracles are simulated as follows:

Before the challenge ciphertext is embedded on the i th query (see below), every query to $\text{KGen}'_{\text{PCH}}$ is answered by the $\text{KGen}'_{\text{ABE}}$ -oracle provided. However, calls to $\text{KGen}''_{\text{PCH}}$ are simply stored as (j, \mathbb{S}) by \mathcal{B} . Hashing is done honestly for all queries except for the i th query, where the reduction queries its own challenger with either 0 or the correct etd, embedding the response c in the returned h . All following queries are performed honestly. After this embedding, all queries to the $\text{KGen}''_{\text{PCH}}$ -oracle are redirected to the $\text{KGen}'''_{\text{ABE}}$ -oracle, while queries to the $\text{KGen}''_{\text{PCH}}$ -oracle are again stored as (j, \mathbb{S}) . Note, by assumption, i.e., \mathcal{A} never queries for key which would allow decrypting that ciphertext. Adaption is done in such a way that if h was generated by the hashing-oracle, then we only continue if (j, \mathbb{S}) is sufficient to decrypt (note, h is known to \mathcal{B} , including the access structure \mathbb{A} used to generate that hash). Finally, for every decryption necessary during adaption, i.e., for ciphertexts not generated by the reduction (and ssk_j , defined by the index j , is actually sufficient to adapt; c , as part of h , never needs to be decrypted, even if it is re-used in another hash), \mathcal{B} uses the provided decryption oracle to receive each etd, and proceeds like in the game. Note, adaption can still be performed honestly, as all etds are thus known. Then, whatever \mathcal{A} outputs, is also output by \mathcal{B} .

We are now in the case that etd is no longer given to the adversary \mathcal{A} . However, this now also means that the adversary \mathcal{A} was able to find a collision, without ever having grasp on the valuable information of etd . Thus, \mathcal{B} can finally use this adversary to break the strong private collision-resistance of CHET. Consider the following reduction \mathcal{B} : it receives pp_{CHET} and embeds it into pp_{PCH} . $(\text{sk}_{\text{CHET}}, \text{pk}_{\text{CHET}}) \leftarrow_r \text{KGen}_{\text{CHET}}(\text{PPGen}_{\text{CHET}})$ is generated honestly. It then uses those to initialize the adversary \mathcal{A} . The ABE-part is done as before. The reduction \mathcal{B} now proceeds as follows: every hash is generated honestly, but the i th one; here, the oracle $\text{Hash}'_{\text{CHET}}$ is queried. All adaptations, but the challenge one, can be performed honestly (as described above with the decryption oracle provided). For the challenge one, however, \mathcal{B} uses its own oracle to find the collision. Then, if E_2 happens, \mathcal{B} can return $((m^*, c^*), r^*, (m'^*, c^*), r'^*, h_0^*)$ by assumption, where $h^* = (h_0^*, c^*)$ by construction.

$|\Pr[S_1] - \Pr[S_2]| \leq q(\nu_{\text{ABE-CCA2}}(\kappa) + \nu_{\text{CHET-SPrivColl}}(\kappa))$ follows, where q is the number of queries to the hashing oracle.

As now the adversary \mathcal{A} has no additional way to win this game, our statement is proven.

Instantiation of a Key-Verifiable Π . We recall a construction from Cramer and Shoup [CS98] in Construction 5, with the alteration that g_1 and g_2 are part of the parameters. We require this alteration to prove key-verifiability.

<p>$\text{PPGen}_{\Pi}(1^\kappa)$: On input a security parameter κ, it outputs the public parameters $\text{pp}_{\Pi} = (\mathbb{G}, g_1, g_2, q, \text{H})$, where $(\mathbb{G}, g, q) \leftarrow_r \text{DLGen}(1^\kappa)$, where g is some generator of \mathbb{G}. Draw $x \leftarrow_r \mathbb{Z}_q$, and let $g_1 \leftarrow g$ and $g_2 \leftarrow g^x$. H is some universal hash-function.</p> <p>$\text{KGen}_{\Pi}(\text{pp}_{\Pi})$: On input $\text{pp}_{\Pi} = (\mathbb{G}, g_1, g_2, q, \text{H})$, draw $(x_1, x_2, y_1, y_2, z) \leftarrow_r \mathbb{Z}_q^5$. Let $c \leftarrow g_1^{x_1} g_2^{x_2}$, $d \leftarrow g_1^{y_1} g_2^{y_2}$, and $h \leftarrow g_1^z$. Set $\text{pk}_{\Pi} \leftarrow (g_1, g_2, c, d, h)$, and $\text{sk}_{\Pi} \leftarrow (x_1, x_2, y_1, y_2, z)$. Return $(\text{sk}_{\Pi}, \text{pk}_{\Pi})$.</p> <p>$\text{Enc}_{\Pi}(\text{pk}_{\Pi}, m)$: On input a public key $\text{pk}_{\Pi} = (g_1, g_2, c, d, h)$ and a message m, draw $r \leftarrow_r \mathbb{Z}_q$. Compute $u_1 \leftarrow g_1^r$, $u_2 \leftarrow g_2^r$, $e \leftarrow h^r m$, $\alpha \leftarrow \text{H}(u_1, u_2, e)$, and $v \leftarrow c^r d^{r\alpha}$. Return (u_1, u_2, e, v).</p> <p>$\text{Dec}_{\Pi}(\text{sk}_{\Pi}, c)$: On input secret key $\text{sk}_{\Pi} = (x_1, x_2, y_1, y_2, z)$, and a ciphertext $c = (u_1, u_2, e, v)$, compute $\alpha \leftarrow \text{H}(u_1, u_2, e)$. If $u_1^{x_1 + y_1 \alpha} u_2^{x_2 + y_2 \alpha} \neq v$, return \perp. Return e/u_1^z.</p> <p>$\text{KVrf}_{\Pi}(\text{sk}_{\Pi}, \text{pk}_{\Pi})$: On input a secret key $\text{sk}_{\Pi} = (x_1, x_2, y_1, y_2, z)$ and a public key $\text{pk}_{\Pi} = (g_1, g_2, c, d, h)$, output 1, if $\text{pk}_{\Pi} = (g_1^{x_1} g_2^{x_2}, g_1^{y_1} g_2^{y_2}, g_1^z)$, and 0 otherwise.</p>
--

Construction 5: Cramer Shoup Π

Theorem 5. *If the DDH-Assumption holds in \mathbb{G} , then the above construction is correct, IND-CCA2 secure, and key-verifiable.*

Proof. Correctness and IND-CCA2 security have already been proven by Cramer and Shoup [CS98].

Thus, it remains to prove key-verifiability.

Key-Verifiability. We prove key-verifiability by a sequence of games.

Game 0: The original key-verifiability game.

Game 1: As Game 0, but abort, if the adversary \mathcal{A} wins the game as defined. Let this event be E_1 .

Transition - Game 0 \rightarrow Game 1: Assume, towards contradiction, that E_1 happens. We can then construct an adversary \mathcal{B} which breaks the discrete logarithm assumption. \mathcal{B} proceeds as follows. It receives (\mathbb{G}, g, q) and g^x . It embeds g as g_1 and g^x as g_2 . Whenever \mathcal{A} outputs $(x_1, x_2, y_1, y_2, z) \neq (x'_1, x'_2, y'_1, y'_2, z')$. We know that $z = z'$ must hold, as we are working in prime-order groups. Assume that $x_1 \neq x'_1$. It follows that $x_2 \neq x'_2$. \mathcal{B} can extract x by calculating $x \leftarrow (x_1 - x'_1)/(x'_2 - x_2)$. The case that $y_1 \neq y'_1$ is similar. $|\Pr[S_0] - \Pr[S_1]| \leq 2\nu_{\text{dlog}}(\kappa)$ follows.

As the adversary has no more possibilities to win the game, key-verifiability is proven.

D Proof of Theorem 1

We now present the proof of the Theorem 1.

Proof. We prove each property on its own, while correctness follows from inspection.

Unforgeability. To prove unforgeability, we use a sequence of games:

Game 0: The original unforgeability game.

Game 1: As Game 0, but we replace crs_Ω with the one generated by $(\text{crs}_\Omega, \tau) \leftarrow_r \text{SIM}_1(1^\kappa)$, keep the trapdoor τ , and start simulating all proofs.

Transition - Game 0 \rightarrow Game 1: Assume towards contradiction that the adversary behaves differently. We can then build an adversary \mathcal{B} which breaks the zero-knowledge property of the underlying proof-system. The reduction works as follows. Our adversary \mathcal{B} receives crs_Ω from its own challenger and embeds it into pp_{P3S} and generates all other values

honestly. All proofs are then generated using the oracle P provided and embedded honestly. Then, whatever \mathcal{A} outputs, is also output by \mathcal{B} . $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{nizk-zk}}(\kappa)$ follows. Note, this also means that all proofs are now simulated, even though they still prove valid statements.

Game 2: As Game 1, but we replace crs_Ω with the one generated by $(\text{crs}_\Omega, \tau, \xi) \leftarrow_{\mathcal{r}} \mathcal{E}_1(1^\kappa)$ and keep the trapdoors τ and ξ . Let E_2 be the event that \mathcal{A} can distinguish this replacement with non-negligible probability. Moreover, note that by definition crs_Ω is exactly distributed as in the prior hop.

Transition - Game 1 \rightarrow Game 2: As we only keep one additional value, i.e., ξ , this is only an internal change. $|\Pr[S_1] - \Pr[S_2]| = 0$ immediately follows.

Game 3: As Game 2, but we abort, if the adversary was able to generate a signature σ_m^* on a string never generated by the signing-oracle. Let this event be E_3 .

Transition - Game 2 \rightarrow Game 3: Assume, towards contradiction, that event E_3 happens. We can then construct an adversary \mathcal{B} which breaks the unforgeability of the underlying signature scheme. Namely, \mathcal{B} receives pk of the signature scheme. This is embedded in pk'_Σ , while all other values are generated as in Game 2. All oracles are simulated honestly, but $\text{Sign}'_{\text{P3S}}$. The only change is, however, that the generation of each σ_m is outsourced to the signature-generation oracle. Then, whenever E_3 happens, \mathcal{B} can return $((\text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, \mathbf{A}, m_{\text{!A}}, h, \mathbb{A}), \sigma_m^*)$. These values can easily be compiled using \mathcal{A} 's output, i.e., (m^*, σ^*) . Note, this already includes that the adversary cannot temper with \mathbf{A} . $|\Pr[S_2] - \Pr[S_3]| \leq \nu_{\text{eUNF-CMA}}(\kappa)$ follows.

Game 4: As Game 3, but we abort, if the adversary was able to generate (m^*, σ^*) for which m^* should not have been derivable. Let this event be E_4 .

Transition - Game 3 \rightarrow Game 4: Assume, towards contradiction, that event E_4 happens. We can then construct an adversary \mathcal{B} which breaks the strong insider collision-resistance of the used PCH. Namely, \mathcal{B} receives pk_{PCH} of the PCH. This is embedded in pk_{P3S} , while all other values are generated as in Game 3. The GetSan -oracle is simulated honestly. Calls to $\text{Sign}'_{\text{P3S}}$ -oracle are done honestly, but the hash is generated using the $\text{Hash}'_{\text{PCH}}$ -oracle. Calls to the $\text{AddSan}'_{\text{P3S}}$ -oracle are simulated as follows. If a key for a simulated sanitizer (obtained by a call to the GetSan -oracle) is to be generated, it is rerouted to $\text{KGen}''_{\text{PCH}}$. If the adversary wants to get a key for itself, it is re-routed to the $\text{KGen}'_{\text{PCH}}$ -oracle and the answer embedded honestly in the re-

sponse. Sanitization requests are performed honestly (but simulated proofs), with the exception that adaptations for simulated sanitizers are done using the $\text{Adapt}'_{\text{PCH}}$ -oracle. So far, the distributions are equal. Then, whenever the adversary outputs (m^*, σ^*) such that the winning-conditions are fulfilled, our reduction \mathcal{B} can return $(m^*, r^*, m'^*, r'^*, h^*)$. The values can be compiled from (m^*, σ^*) and the transcript from the signing-oracle (note, we already excluded that the adversary can temper with the hash h). $|\Pr[S_3] - \Pr[S_4]| \leq \nu_{\text{PBCH-SInsider-CollRes}}(\kappa)$ follows.

Game 5: As Game 4, but we abort, if the adversary was able to generate (m^*, σ^*) for which σ^* is fresh. Let this event be E_5 .

Transition - Game 4 \rightarrow Game 5: Assume, towards contradiction, that event E_5 happens. We can then construct an adversary \mathcal{B} which breaks the unforgeability of the used Σ or the one-wayness of the used one-way function f . Namely, \mathcal{B} receives pk_Σ of the Σ and f , and $f(x) = y$ from its own challenger. This is embedded in pk_{P3S} (and, of course, the public parameters), while all other values are generated as in Game 4. y is embedded in $\text{pk}_{\text{P3S}}^{\text{Sig}}$. For signing, the proofs are already simulated, and thus x is not required to be known. Each call to $\text{AddSan}_{\text{P3S}}$ for keys for which the adversary knows the corresponding secret keys, \mathcal{B} calls its signature oracle to obtain such a key. For simulated sanitizers, those signature do not need to be obtained, as the proofs are already simulated. Then, whenever the adversary outputs (m^*, σ^*) , \mathcal{B} extracts values (x_1, x_2, σ') . If $f(x_1) = y$, \mathcal{B} can return x_1 to break the one-wayness of f . In the other case, \mathcal{B} can return $(m^*, \sigma^*) = (f(x_2), \text{pk}_{\text{P3S}})$ as its own forgery attempt for Σ . If extraction fails or a wrong statement was proven, SSE does not hold. A reduction is straightforward.

$|\Pr[S_4] - \Pr[S_5]| \leq \nu_{\text{eUNF-CMA}}(\kappa) + \nu_{\text{ow}}(\kappa) + \nu_{\text{nizk-sse}}(\kappa)$ follows.

Now, the adversary can no longer win the unforgeability game; this game is computationally indistinguishable from the original game, which concludes the proof.

Immutability. To prove immutability, we use a sequence of games:

Game 0: The original immutability game.

Game 1: As Game 0, we abort if the adversary outputs $(\text{pk}^*, \sigma^*, m^*)$ such that the winning conditions are met. Let this event be E_1 .

Transition - Game 0 \rightarrow Game 1: Assume, towards contradiction, that event E_1 happens. We can then build an adversary \mathcal{B} which breaks the unforgeability of the used signature scheme. Namely, we know that A

(which also contains the length of the message and all non-modifiable blocks along with their location), along with pk_{PCH} , is signed. As, however, by definition, the message m^* must be different from any derivable message, \mathcal{A} w.r.t. pk_{PCH} was never signed in this regard. Thus, $(\text{pk}^*, \text{pk}_{\text{P3S}}^{\text{Sig}}, \mathbb{A}^*, m_{\mathcal{A}}^*, h^*, \mathbb{A}^*)$ was never signed by the signer. Constructing a reduction \mathcal{B} is now straightforward. Our reduction \mathcal{B} receives the public key pk'_{Σ} (along with the public parameters) from its own challenger. This public key is embedded as pk'_{Σ} . All other values are generated honestly. If a signature σ_m is to be generated, \mathcal{B} asks its own oracle to generate that signature, embedding it into the response \mathcal{A} receives. At some point, \mathcal{A} returns $(\text{pk}^*, \sigma^*, m^*)$. The forgery can be extracted as described above. $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{eUNF-CMA}}(\kappa)$ follows. We stress that, by construction, a sanitizer always exists. Now, the adversary can no longer win the immutability game; this game is computationally indistinguishable from the original game, which concludes the proof.

Privacy. To prove privacy, we use a sequence of games:

Game 0: The original privacy game.

Game 1: As Game 0, but we replace crs_{Ω} with the one generated by $(\text{crs}_{\Omega}, \tau) \leftarrow_r \text{SIM}_1(1^{\kappa})$, keep the trapdoor τ , and start simulating all proofs.

Transition - Game 0 \rightarrow Game 1: Assume towards contradiction that the adversary behaves differently. We can then build an adversary \mathcal{B} which breaks the zero-knowledge property of the underlying proof-system. The reduction works as follows. Our adversary \mathcal{B} receives crs_{Ω} from its own challenger and embeds it into pp_{P3S} and generates all other values honestly. All proofs are then generated using the oracle P provided and embedded honestly. Then, whatever \mathcal{A} outputs, is also output by \mathcal{B} . $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{nik-zk}}(\kappa)$ follows. Note, this also means that all proofs are now simulated, even though they still prove valid statements.

Game 2: As Game 1, but we abort if (σ'_0, m) and (σ'_1, m) contain different randomness $r'_0 \neq r'_1$ if generated inside LoRSanit . Let this event be E_2 .

Transition - Game 1 \rightarrow Game 2: Assume, towards contradiction, that event E_2 happens. We can then construct an adversary \mathcal{B} which breaks the uniqueness of PCH. In particular, it receives pp_{PCH} and embeds it accordingly. All other values are generated as in Game 2. Then, when \mathcal{A} was able to generate $r'_0 \neq r'_1$, the reduction \mathcal{B} can directly return $(\text{pk}^*, m, r'_0, r'_1, h^*)$, where pk^* is contained in pk_{P3S} .

$|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{PCH-uniq}}(\kappa)$ follows, while the signature does not matter, as it is already hidden behind a simulated zero-knowledge proof, making the distributions equal.

Game 3: As Game 2, but we directly generate $(\sigma, M_0(m_0))$ without using sanitizing, i.e., we freshly hash with $M_0(m_0)$ (if the oracle would return a signature). Note, the proofs are already simulated, but we also need to encrypt $\text{pk}_{\text{P3S}}^{\text{San}}$, as it would be done at sanitization anyway. Moreover, the adversary never sees a non-sanitized signature from that oracle, while all proofs are already simulated.

Transition - Game 2 \rightarrow Game 3: If the adversary behaves noticeably different, we can build an adversary \mathcal{B} which breaks the strong indistinguishability of the used PCH. The reduction works as follows. \mathcal{B} receives pp_{PCH} and embeds is honestly. All other values are generated according to Game 3. Then, for every hash generated in the LoRSanit-oracle, the challenge oracle is queried and the answer embedded into the response. Whatever \mathcal{A} then outputs, is also output by \mathcal{B} . $|\Pr[S_2] - \Pr[S_3]| \leq \nu_{\text{PCH-SInd}}(\kappa)$ immediately follows.

We stress that, by construction, a sanitizer always exists, because $\mathbb{A} \neq \emptyset$ must hold. Thus, sanitization is always possible from any generated signature, even in the case $\mathbb{A} = (\emptyset, m_\ell)$, i.e., where a sanitizer only claims accountability, but does not modify the message itself.

Now, the privacy game is independent of the bit b , proving privacy.

Transparency. To prove transparency, we use a sequence of games:

Game 0: The original transparency game, where $b = 0$.

Game 1: As Game 0, but we replace crs_Ω with the one generated by $(\text{crs}_\Omega, \tau) \leftarrow_r \text{SIM}_1(1^\kappa)$, keep the trapdoor τ , and start simulating all proofs.

Transition - Game 0 \rightarrow Game 1: Assume towards contradiction that the adversary behaves differently. We can then build an adversary \mathcal{B} which breaks the zero-knowledge property of the underlying proof-system. The reduction works as follows. Our adversary \mathcal{B} receives crs_Ω from its own challenger and embeds it into pp_{P3S} and generates all other values honestly. All proofs are then generated using the oracle P provided and embedded honestly. Then, whatever \mathcal{A} outputs, is also output by \mathcal{B} . $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{nik-zk}}(\kappa)$ follows. Note, this also means that all proofs are now simulated, even though they still prove valid statements.

Game 2: As Game 1, but we replace the contents of c (or c' resp.) with a 0, if generated by the SignOrSanit-oracle.

Transition - Game 1 \rightarrow Game 2: Assume, towards contradiction, that the adversary behaves noticeably different. We can then construct an adversary \mathcal{B} which breaks the IND-CCA2 security of the used encryption scheme. Namely, we use a series of hybrids. Our reduction \mathcal{B} proceeds as follows. It receives pk_H and (and the corresponding parameters) from its own challenger and embeds them correctly. All other values are generated as in Game 1. For the first i ciphertexts generated, encrypt a 0. If, however, the i th ciphertext is generated, \mathcal{B} asks its own challenge oracle to either encrypt 0 or the correct value. The response is embedded to \mathcal{B} 's response to \mathcal{A} . All following ciphertexts are generated honestly. Thus, Game 2.0 is the same as Game 1, while in Game 2.1., however, we make the first replacement. Then, whatever \mathcal{A} outputs in Game 3.i is also output by \mathcal{B} . Note, if a ciphertext is to be decrypted (e.g., for proof-generation of signatures not generated by the SignOrSanit-oracle), \mathcal{B} uses the provided decryption oracle provided. $|\Pr[S_1] - \Pr[S_2]| \leq q\nu_{\text{ind-cca2}}(\kappa)$ follows, where q is the number of ciphertexts generated. We stress that we do not need to “cheat” during proof-generation, as the adversary is not allowed to query such signatures to the $\text{Proof}'_{\text{P3S}}$ -oracle.

Game 3: As Game 2, but we directly generate $(\sigma, \text{M}(m))$ without using sanitizing, i.e., we always freshly hash with $\text{M}(m)$ (if the oracle would return a signature). Note, the proofs are already simulated. Moreover, the adversary never sees a non-sanitized signature from that oracle, while all proofs are already simulated.

Transition - Game 2 \rightarrow Game 3: Assume, towards contradiction, that the adversary behaves noticeably different. We can build an adversary \mathcal{B} which breaks the strong indistinguishability of the used PCH. The reduction works as follows. \mathcal{B} receives pp_{PCH} and embeds is honestly. All other values are generated according to Game 2. Then, for every hash generated in the SignOrSanit oracle the challenge oracle is queried and the answer embedded into the response. Whatever \mathcal{A} then outputs, is also output by \mathcal{B} . $|\Pr[S_2] - \Pr[S_3]| \leq \nu_{\text{PCH-SInd}}(\kappa)$ immediately follows.

We stress that, by construction, a sanitizer always exists, because $\mathbb{A} \neq \emptyset$ must hold. Thus, sanitization is always possible from any generated signature, even in the case $\mathbb{A} = (\emptyset, m_\ell)$, i.e., where a sanitizer only claims accountability.

Now, we are in the case that a signature is freshly generated ($b = 1$). Thus, transparency is proven, as each hop only changes the view negligibly.

Pseudonymity. To prove pseudonymity, we use a sequence of games:

Game 0: The original transparency game.

Game 1: As Game 0, but we replace crs_Ω with the one generated by $(\text{crs}_\Omega, \tau) \leftarrow_r \text{SIM}_1(1^\kappa)$, keep the trapdoor τ , and start simulating all proofs.

Transition - Game 0 \rightarrow Game 1: Assume towards contradiction that the adversary behaves differently. We can then build an adversary \mathcal{B} which breaks the zero-knowledge property of the underlying proof-system. The reduction works as follows. Our adversary \mathcal{B} receives crs_Ω from its own challenger and embeds it into pp_{P3S} and generates all other values honestly. All proofs are then generated using the oracle P provided and embedded honestly. Then, whatever \mathcal{A} outputs, is also output by \mathcal{B} . $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{nik-zk}}(\kappa)$ follows. Note, this also means that all proofs are now simulated, even though they still prove valid statements.

Game 2: As Game 1, but we replace crs_Ω with the one generated by $(\text{crs}_\Omega, \tau, \xi) \leftarrow_r \mathcal{E}_1(1^\kappa)$ and keep the trapdoors τ and ξ . Let E_2 be the event that \mathcal{A} can distinguish this replacement with non-negligible probability. Moreover, note that by definition crs_Ω is exactly distributed as in the prior hop.

Transition - Game 1 \rightarrow Game 2: As we only keep one additional value, i.e., ξ , this is only an internal change. $|\Pr[S_1] - \Pr[S_2]| = 0$ immediately follows.

Game 3: As Game 2, but we abort if (σ'_0, m) and (σ'_1, m) contain different randomness $r'_0 \neq r'_1$ if generated inside LoRSanit . Let this event be E_3 .

Transition - Game 2 \rightarrow Game 3: Assume, towards contradiction, that event E_3 happens. We can then construct an adversary \mathcal{B} which breaks the uniqueness of PCH. In particular, it receives pp_{PCH} and embeds it accordingly. All other values are generated as in Game 2. Then, when \mathcal{A} was able to generate $r'_0 \neq r'_1$, the reduction \mathcal{B} can directly return $(\text{pk}^*, m, r'_0, r'_1, h^*)$, where pk^* is contained in pk_{P3S} . $|\Pr[S_2] - \Pr[S_3]| \leq \nu_{\text{PCH-uniq}}(\kappa)$ follows, while the signature does not matter, as it is already hidden behind a simulated zero-knowledge proof, making the distributions equal.

Game 4: As Game 3, but we replace the contents of c' with a 0, if generated by the LoRSanit -oracle.

Transition - Game 3 \rightarrow Game 4: Assume, towards contradiction, that the adversary behaves noticeably different. We can then construct an adversary \mathcal{B} which breaks the IND-CCA2 security of the used encryption scheme. Namely, we use a series of hybrids. Our reduction

\mathcal{B} proceeds as follows. It receives pk_H and (and the corresponding parameters) from its own challenger and embeds them correctly. All other values are generated as in Game 3. For the first i ciphertexts generated, encrypt a 0. If, however, the i th ciphertext is generated, \mathcal{B} asks its own challenge oracle to either encrypt 0 or the correct value. The response is embedded to \mathcal{B} 's response to \mathcal{A} . All following ciphertexts are generated honestly. Thus, Game 4.0 is the same as Game 4, while in Game 4.1., however, we make the first replacement. Then, whatever \mathcal{A} outputs in Game 4.i is also output by \mathcal{B} . All decryption queries required can be obtained by the decryption oracle provided.

$|\Pr[S_3] - \Pr[S_4]| \leq q\nu_{\text{ind-cca2}}(\kappa)$ follows, where q is the number of queries to the LoRSanit-oracle. We stress that we do not need to “cheat” during proof-generation, as the adversary is not allowed to query such signatures to the $\text{Proof}'_{\text{p35}}$ -oracle.

Now, the game is independent of the bit b , proving the theorem.

Signer-Accountability. To prove signer-accountability, we use a sequence of games:

Game 0: The original signer-accountability game.

Game 1: As Game 0, but we replace crs_Ω with the one generated by $(\text{crs}_\Omega, \tau) \leftarrow_r \text{SIM}_1(1^\kappa)$, keep the trapdoor τ , and start simulating all proofs.

Transition - Game 0 \rightarrow Game 1: Assume towards contradiction that the adversary behaves differently. We can then build an adversary \mathcal{B} which breaks the zero-knowledge property of the underlying proof-system. The reduction works as follows. Our adversary \mathcal{B} receives crs_Ω from its own challenger and embeds it into pp_{p35} and generates all other values honestly. All proofs are then generated using the oracle P provided and embedded honestly. Then, whatever \mathcal{A} outputs, is also output by \mathcal{B} . $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{nik-zk}}(\kappa)$ follows. Note, this also means that all proofs are now simulated, even though they still prove valid statements.

Game 2: As Game 1, but we replace crs_Ω with the one generated by $(\text{crs}_\Omega, \tau, \xi) \leftarrow_r \mathcal{E}_1(1^\kappa)$ and keep the trapdoors τ and ξ . Let E_2 be the event that \mathcal{A} can distinguish this replacement with non-negligible probability. Moreover, note that by definition crs_Ω is exactly distributed as in the prior hop.

Transition - Game 1 \rightarrow Game 2: As we only keep one additional value, i.e., ξ , this is only an internal change. $|\Pr[S_1] - \Pr[S_2]| = 0$ immediately follows.

Game 3: As Game 2, but we abort, if the adversary outputs $(pk_0^*, pk_1^*, \sigma^*, m^*, \pi^*)$ such that the winning conditions are met. Let this event be E_3 .

Transition - Game 2 \rightarrow Game 3: Assume, towards contradiction, that E_3 happened. We can then construct an adversary \mathcal{B} against the one-wayness of f . The reduction works as follows. It receives f and $f(x)$. It embeds both accordingly. Note, the proofs are simulated, and thus x is not needed to be known for π_{pk} . Every sanitization is done honestly, with the exception of simulated proofs. Then, as we know that the adversary wins its game, \mathcal{B} can extract a pre-image x' such that $f(x') = f(x)$ (if extraction fails, this adversary breaks SSE using a straightforward reduction), and can return it to its own challenger. $|\Pr[S_2] - \Pr[S_3]| \leq \nu_{ow}(\kappa) + \nu_{nizk-sse}(\kappa)$ follows.

As now the adversary has no more possibilities to win the signer-accountability game, the theorem is proven.

Sanitizer-Accountability. To prove sanitizer-accountability, we use a sequence of games:

Game 0: The original sanitizer-accountability game.

Game 1: As Game 0, but we replace crs_Ω with the one generated by $(crs_\Omega, \tau) \leftarrow_r \text{SIM}_1(1^\kappa)$, keep the trapdoor τ , and start simulating all proofs.

Transition - Game 0 \rightarrow Game 1: Assume towards contradiction that the adversary behaves differently. We can then build an adversary \mathcal{B} which breaks the zero-knowledge property of the underlying proof-system. The reduction works as follows. Our adversary \mathcal{B} receives crs_Ω from its own challenger and embeds it into pp_{P3S} and generates all other values honestly. All proofs are then generated using the oracle P provided and embedded honestly. Then, whatever \mathcal{A} outputs, is also output by \mathcal{B} . $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{nizk-zk}(\kappa)$ follows. Note, this also means that all proofs are now simulated, even though they still prove valid statements.

Game 2: As Game 1, but we replace crs_Ω with the one generated by $(crs_\Omega, \tau, \xi) \leftarrow_r \mathcal{E}_1(1^\kappa)$ and keep the trapdoors τ and ξ . Let E_2 be the event that \mathcal{A} can distinguish this replacement with non-negligible probability. Moreover, note that by definition crs_Ω is exactly distributed as in the prior hop.

Transition - Game 1 \rightarrow Game 2: As we only keep one additional value, i.e., ξ , this is only an internal change. $|\Pr[S_1] - \Pr[S_2]| = 0$ immediately follows.

Game 3: As Game 2, but we abort, if the adversary outputs $(\text{pk}^*, \sigma^*, m^*, \pi^*)$ such that the winning conditions are met. Let this event be E_3 .

Transition - Game 2 \rightarrow Game 3: Assume, towards contradiction, that E_3 happened. We can then construct an adversary \mathcal{B} against the one-wayness of f . The reduction works as follows. It receives f and $f(x)$. It embeds both accordingly. Every signing and proof-generation is done honestly, with the exception of simulated proofs. Then, as we know that the adversary wins its game, \mathcal{B} can extract a pre-image x' such that $f(x') = f(x)$ (if extraction fails, this adversary breaks SSE using a straightforward reduction), and can return it to its own challenger. $|\Pr[S_2] - \Pr[S_3]| \leq \nu_{\text{ow}}(\kappa) + \nu_{\text{nik-sse}}(\kappa)$ follows.

As now the adversary has no more possibilities to win the sanitizer-accountability game, the theorem is proven.

Proof-Soundness. First, we prove proof-soundness by a sequence of games.

Game 0: The original proof-soundness game.

Game 1: As Game 0, but we replace crs_Ω with the one generated by $(\text{crs}_\Omega, \tau) \leftarrow_r \text{SIM}_1(1^\kappa)$, keep the trapdoor τ .

Transition - Game 0 \rightarrow Game 1: Assume towards contradiction that the adversary behaves differently. We can then build an adversary \mathcal{B} which breaks the zero-knowledge property of the underlying proof-system. The reduction works as follows. Our adversary \mathcal{B} receives crs_Ω from its own challenger and embeds it into pp_{P3S} and generates all other values honestly. Note, in this case no proofs need to be simulated, as we do not have any oracles. $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{nik-zk}}(\kappa)$ follows.

Game 2: As Game 1, but we replace crs_Ω with the one generated by $(\text{crs}_\Omega, \tau, \xi) \leftarrow_r \mathcal{E}_1(1^\kappa)$ and keep the trapdoors τ and ξ . Let E_2 be the event that \mathcal{A} can distinguish this replacement with non-negligible probability. Moreover, note that by definition crs_Ω is exactly distributed as in the prior hop.

Transition - Game 1 \rightarrow Game 2: As we only keep one additional value, i.e., ξ , this is only an internal change. $|\Pr[S_1] - \Pr[S_2]| = 0$ immediately follows.

Game 3: As Game 2, but we abort if the adversary outputs $((\text{pk}_i^*)_{0 \leq i \leq 5}, \sigma^*, m^*, \pi_0^*, \pi_1^*)$ such that $\text{pk}_2 \neq \text{pk}_5$, and the winning conditions are met (Note, decryption is deterministic). Let this event be E_3 .

Transition - Game 2 \rightarrow Game 3: Assume, towards contradiction, that event E_3 happens. We can then construct an adversary \mathcal{B} against the key-verifiability of the used encryption scheme. The reduction works as follows. It receives pp_Π , and once the adversary outputs

$((\mathbf{pk}_i^*)_{0 \leq i \leq 5}, \sigma^*, m^*, \pi_0^*, \pi_1^*)$, \mathcal{B} extracts \mathbf{sk}_0^* from π_0^* and \mathbf{sk}_1^* from π_1^* (if extraction fails, this adversary breaks SSE; a reduction is straightforward). Then, it can return $(\mathbf{sk}_0^*, \mathbf{sk}_1^*, \mathbf{pk}_1^*)$ as its own forgery.

$|\Pr[S_2] - \Pr[S_3]| \leq \nu_{\text{enc-key-verf}}(\kappa) + 2\nu_{\text{nik-sse}}(\kappa)$ immediately follows. This hop essentially rules out the possibility an adversary can create more than one secret key w.r.t. to its \mathbf{pk} such that decryption points to a different sanitizer.

Game 4: As Game 4, but we abort if the adversary outputs $((\mathbf{pk}_i^*)_{0 \leq i \leq 5}, \sigma^*, m^*, \pi_0^*, \pi_1^*)$ for which the winning conditions are met. Let this event be E_4 .

Transition - Game 3 \rightarrow Game 4: If this event (E_4) happens, either π_0^* or π_1^* is a bogus proof, as at least one proves a false statement. For the reduction, \mathcal{B} proceeds as in the prior game (doing everything honestly, but using crs_Ω received from \mathcal{B} 's own challenger) and randomly selects either the first statement (concerning $(\mathbf{pk}_i^*)_{0 \leq i \leq 2}$) or the second statement (concerning $(\mathbf{pk}_i^*)_{3 \leq i \leq 5}$), with the proof contained in σ^* . As \mathcal{B} needs to randomly guess (one may still be true), we lose a factor of 2 in the reduction. $|\Pr[S_3] - \Pr[S_4]| \leq 2\nu_{\text{nik-sse}}(\kappa)$ follows.

As the adversary now has to other way to win the proof-soundness game and each hop only changes the view of the adversary negligibly, proof-soundness is proven.

Traceability. Next, we prove traceability by a sequence of games.

Game 0: The original traceability game.

Game 1: As Game 0, but we replace crs_Ω with the one generated by $(\text{crs}_\Omega, \tau) \leftarrow_r \text{SIM}_1(1^\kappa)$, keep the trapdoor τ , and start simulating all proofs.

Transition - Game 0 \rightarrow Game 1: Assume towards contradiction that the adversary behaves differently. We can then build an adversary \mathcal{B} which breaks the zero-knowledge property of the underlying proof-system. The reduction works as follows. Our adversary \mathcal{B} receives crs_Ω from its own challenger and embeds it into $\text{pp}_{\text{P}_{35}}$ and generates all other values honestly. All proofs are then generated using the oracle P provided and embedded honestly. Then, whatever \mathcal{A} outputs, is also output by \mathcal{B} . $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{nik-zk}}(\kappa)$ follows. Note, this also means that all proofs are now simulated, even though they still prove valid statements.

Game 2: As Game 1, but we replace crs_Ω with the one generated by $(\text{crs}_\Omega, \tau, \xi) \leftarrow_r \mathcal{E}_1(1^\kappa)$ and keep the trapdoors τ and ξ . Let E_2 be the event that \mathcal{A} can distinguish this replacement with non-negligible

probability. Moreover, note that by definition crs_Ω is exactly distributed as in the prior hop.

Transition - Game 1 \rightarrow Game 2: As we only keep one additional value, i.e., ξ , this is only an internal change. $|\Pr[S_1] - \Pr[S_2]| = 0$ immediately follows.

Game 3: As Game 2, but we abort if the adversary outputs a valid $(\text{pk}^*, \sigma^*, m^*)$ for which we cannot (as the holder of $\text{sk}_{\text{P3S}}^{\text{Sig}}$) calculate a pk which makes $\text{Judge}_{\text{P3S}}(\text{pk}^*, \text{pk}_{\text{P3S}}^{\text{Sig}}, \text{pk}, \pi_{\text{P3S}}, \sigma^*, m^*)$ output 0. Let this event be E_3 .

Transition - Game 2 \rightarrow Game 3: If this event (E_3) happens, we have a bogus proof π contained in σ^* , as it proves a false statement. Thus, \mathcal{B} proceeds as in the prior game (doing everything honestly, but using simulated proofs and the simulated crs_Ω), and can simply return the statement claimed to be proven by π , and π itself. $|\Pr[S_2] - \Pr[S_3]| \leq \nu_{\text{nizk-sse}}(\kappa)$ directly follows.

Relations of Security Properties. We now show several relations among the security properties defined. These relations may only hold relative to the assumptions we use in our construction.

Theorem 6 (Unforgeability is independent). *There exists a P3S which offers all security properties, but unforgeability.*

Proof. A counter-example is simple: Alter $\text{AddSan}_{\text{P3S}}$ in such a way, that a sanitizer receives a sk_S not only for the asked for attributes, but for all attributes. Clearly, all other properties, including correctness, are still preserved, but now a sanitizer can alter more signatures than it should be allowed to, as it holds a sk_S for all attributes. Moreover, it still cannot blame a signer or sanitizer for the signatures it creates.

Theorem 7 (Transparency is independent). *There exists a P3S which offers all security properties, but transparency.*

Proof. This holds by altering our construction. Namely, at signing, the label to the proof system is no longer $\ell = (\text{pp}_{\text{P3S}}, \text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, h, r, m, \mathbb{A}, \mathbb{A}, m_{\mathbb{A}}, m_{\mathbb{A}}, \sigma_m, c)$, but $\ell = (\text{pp}_{\text{P3S}}, \text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, h, r, m, \mathbb{A}, \mathbb{A}, m_{\mathbb{A}}, m_{\mathbb{A}}, \sigma_m, c, 0)$. For sanitization, the label is changed to $\ell = (\text{pp}_{\text{P3S}}, \text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, h, r, m, \mathbb{A}, \mathbb{A}, m_{\mathbb{A}}, m_{\mathbb{A}}, \sigma_m, c, 1)$. For verification, both possibilities (last bit equal to 0 or equal to 1; Both values are distinct, and are neither derived from any secrets or message) are tested, and only returns 1, if one of the verification procedures return 1. Clearly, all other properties still hold,

while an adversary can use the last bit to decide whether a sanitization was performed or not.

Theorem 8 (Privacy is independent). *There exists a P3S which offers all security properties, but privacy.*

Proof. We prove this by slightly altering our construction. First note that, in the privacy experiment, the adversary \mathcal{A} is allowed to generate $\text{sk}_{\text{P3S}}^{\text{Sig}}$, and thus obviously knows it. We now alter our construction in the following way; At signing, the original message (if the message space is not compatible, one can use a hash-function) is also encrypted to the signer itself as c' (note, the signer already owns an encryption key-pair), and appended to the label $\ell = (\text{pp}_{\text{P3S}}, \text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, h, r, m, \mathbb{A}, \mathbb{A}, m_{\mathbb{A}}, m_{!_{\mathbb{A}}}, \sigma_m, c, c')$, and is also part of the signature $\sigma' = (\sigma, c')$. Verification works as expected. Sanitization remains the same, also using c' in the augmented label, but returns c' as part of the sanitized signature. Proof-generation and the judge now also take c' into account in a straightforward manner. All properties, but privacy, remain to hold, as we only add an additional value to the label ℓ of the proof-system. However, an adversary \mathcal{A} can use its secret decryption key to decrypt the *original* message (or its hash), directly contradicting the privacy requirements. Transparency continues to hold, as the message is encrypted. Note, in the altered construction IND-CPA is sufficient, as this value is never decrypted by an honest party.

Theorem 9 (Immutability is independent). *There exists a P3S which offers all security properties, but immutability.*

Proof. We alter the construction in the following way: An honestly generated pk_{PCH} is augmented by appending a 0. For usage outside of ℓ for the proof-system, this bit is dropped. However, if the appended bit is a 1, the verification algorithm now also accepts, if \mathbb{A} , $m_{\mathbb{A}}$, and $m_{!_{\mathbb{A}}}$ are not consistent, i.e., arbitrary. Thus, an adversary can sanitize a seen signature to arbitrary ones. Again, all properties, but immutability, remain to hold: An adversary \mathcal{A} simply needs to generate a bogus public key (which is never generated in the honest case), and can then alter immutable blocks.

Theorem 10 (Pseudonymity is independent). *There exists a P3S which offers all security properties, but pseudonymity.*

Proof. We first want to remind the reader that, in the pseudonymity experiment, the adversary \mathcal{A} is allowed to input arbitrary signatures, while in the transparency experiment the adversary never sees a signature from the signer in the case $b = 0$ from the LeftOrRight-oracle.

We use this gap to encode the sanitizer’s identity such that it can only be noticed, if a sanitized and the original signatures are available. Let l be an upper bound on the bit-length of the output of the one-way function f . Let e be an additional security parameter. We alter signing as follows: At signing, the signer chooses a random integer $i \leftarrow_r \{0, 1\}^{l+e}$, and attaches it to the label ℓ' for the NIZK, i.e., $\ell' = (\text{pp}_{\text{P3S}}, \text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, h, r, m, \mathbb{A}, \mathbb{A}, m_{\mathbb{A}}, m_{\mathbb{A}}, \sigma_m, c, i)$, yet also to the signature, i.e., $\sigma' = (\sigma, i)$. Verification simply also takes the altered values into account. At sanitization, however, i is altered by setting $i' \leftarrow i + x'_2$, where x'_2 is the binary representation of x_2 . The variable i' then becomes part of the used label for the new NIZK and the sanitized signatures. If e is chosen large enough, while l is a constant, the distributions remain indistinguishable in the transparency experiment. Note, all attached values are independent of the messages, thus privacy still holds.

Clearly, all properties, but pseudonymity, hold. Namely, the adversary \mathcal{A} simply checks whether a chosen x_2 and i (note, the adversary \mathcal{A} also chooses the corresponding secret keys in the pseudonymity experiment, and thus knows the corresponding public keys) match by checking whether i' (generated by the challenger) equals $x'_2 + i$ or not.

Theorem 11 (Signer-Accountability is independent). *There exists a P3S which offers all security properties, but signer-accountability.*

Proof. The idea is similar to the proof for showing that immutability is independent. Namely, we alter our construction as follows. At key-generation for the signer, a 0 is appended to $\text{pk}_{\text{P3S}}^{\text{Sig}}$. If some of the inner keys of $\text{pk}_{\text{P3S}}^{\text{Sig}}$ are used, the last bit is simply dropped for the underlying algorithms. For the judge, however, if $\text{pk}_{\text{P3S}}^{\text{Sig}}$ has a trailing 1, it always outputs 1, if the key to the checked is the corresponding sanitizer one, if signature verification passes (in other words, the generated proof is ignored, but only the validity is checked). Otherwise, the original algorithm is executed.

Now, if the signer generates a key with a trailing 1, it can make the sanitizer accountable for any signature it wants. All other properties are, however, still preserved, as all keys are part of the label, which still preserves proof-soundness.

Theorem 12 (Sanitizer-Accountability is independent). *There exists a P3S which offers all security properties, but sanitizer-accountability.*

Proof. The proof follows the same line as for proving the independence of signer-accountability. Namely, we alter our construction as follows. At

key-generation for the sanitizer, a 0 is appended to $\text{pk}_{\text{P3S}}^{\text{San}}$. If some of the inner keys of $\text{pk}_{\text{P3S}}^{\text{San}}$ are used, the last bit is simply dropped for the underlying algorithms. For the judge, however, if $\text{pk}_{\text{P3S}}^{\text{San}}$ has a trailing 1, it always outputs 1, if the key to the checked is the corresponding signer one, if signature verification passes (in other words, the generated proof is ignored, but only the validity is checked). Otherwise, the original algorithm is executed.

Now, if the sanitizer generates a key with a trailing 1, it can make the signer accountable for any signature it wants. All other properties are, however, still preserved, as all keys are part of the label, which still preserves proof-soundness.

Theorem 13 (Proof-Soundness is independent). *There exists a P3S which offers all security properties, but proof-soundness.*

Proof. We alter our construction as follows. At key-generation, all keys (group, signer, and sanitizer) are appended with a 0. If an algorithm uses an inner key, that bit is ignored. Judge, however, outputs also 1 (if the corresponding signature verifies), if *all* public keys have a trailing 1. This allows the adversary to easily win the proof-soundness experiment. All other properties are still preserved, as the adversary need to control all three key-pairs to win, which is not the case in the other definitions, but privacy. Privacy, however, still continues to hold, as the message is not input to the changes in our contrived scheme.

Theorem 14 (Traceability is independent). *There exists a P3S which offers all security properties, but traceability.*

Proof. This holds by altering our construction. Namely, at signing, the label to the proof system is no longer $\ell = (\text{pp}_{\text{P3S}}, \text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, h, r, m, \mathbb{A}, \mathbb{A}, m_{\mathbb{A}}, m_{\mathbb{A}}, \sigma_m, c)$, but $\ell = (\text{pp}_{\text{P3S}}, \text{pk}_{\text{P3S}}, \text{pk}_{\text{P3S}}^{\text{Sig}}, h, r, m, \mathbb{A}, \mathbb{A}, m_{\mathbb{A}}, m_{\mathbb{A}}, \sigma_m, c, 0)$. For sanitization, the label remains the same. If, however, the last bit is a 1, judge outputs 0.

Again, all properties, but traceability, are preserved, as an adversary can simply append a 1 to the label, which an honest player would never do.