

# Fully Invisible Protean Signature Schemes

Stephan Krenn<sup>1</sup>, Henrich C. Pöhls<sup>2</sup>, Kai Samelin<sup>3</sup>, and Daniel Slamanig<sup>1</sup>

<sup>1</sup> AIT Austrian Institute of Technology, Vienna, Austria

{[stephan.krenn](mailto:stephan.krenn@ait.ac.at),[daniel.slamanig](mailto:daniel.slamanig@ait.ac.at)}@ait.ac.at

<sup>2</sup> ISL & Chair of IT-Security, University of Passau, Passau, Germany

[hp@sec.uni-passau.de](mailto:hp@sec.uni-passau.de)

<sup>3</sup> TÜV Rheinland i-sec GmbH, Hallbergmoos, Germany

[kaispapers@gmail.com](mailto:kaispapers@gmail.com)

**Abstract.** Protean Signatures (PS), recently introduced by Krenn et al. (CANS’18), allow a semi-trusted third party (the *sanitizer*), to modify a signed message in a controlled way: the signer can define message parts to be arbitrarily editable by the sanitizer, as well as message parts which can be redacted (but not altered otherwise) by the sanitizer. Thus, PSs generalize both redactable signatures (RS) and sanitizable signatures (SS) into a single notion.

*Invisibility* for protean signatures guarantees that no outsider (i.e., any party not being signer or sanitizer) can decide which message parts can be edited. However, the current definition of invisibility does not prohibit that an outsider can decide which parts are *redactable* — only which parts can be edited are hidden. This negatively impacts on the privacy guarantees provided by this definition.

We extend PSs to be *fully* invisible. Our notion guarantees that an outsider can neither identify editable nor redactable parts. We therefore introduce the new notions of Invisible RSs and Invisible Non-Accountable SSs (SS’), along with a consolidated framework for aggregate signatures. Using those building blocks, our resulting construction is significantly more efficient than the original scheme by Krenn et al., which we demonstrate in a prototypical implementation.

## 1 Introduction

By definition, (plain) digital signatures prohibit any kind of modification of signed messages. That is, only the holder of the secret signing key  $sk_{\Sigma}$  can sign messages [1]. Still, there are a plethora of application scenarios where a later modification of signed messages by a (semi-trusted) third party has its merits [2], e.g., in the following use-case, based on the handling of patient data [3,4,5,6]:

Assume that a M.D. always signs complete patient records. Further assume that each of those records consists of the patient’s name, its insurance number and the treatments given. After the patient is released from the hospital, the responsible accountant receives the complete signed record corresponding to the to-be-released patient to be able to write a bill to the patient’s insurance company.

It is easy to see that this process is not very privacy-friendly from the patient’s point of view, as the accountant receives all information related to patient. However, most of the information is not relevant for writing the bill, e.g., the patient’s name. Thus, a solution is to only give the treatments and the insurance number to the accountant, anonymizing the paperwork. The major obstacle is that standard digital signatures prohibit any alterations, and thus the M.D. either needs to re-sign the document, or an additional trusted entity does need to sign on behalf of the M.D. Both solutions are not satisfactory, as both induce additional overhead, one might even be impossible, e.g., if the M.D. is no longer employed. We conclude that modifying signed messages in a controlled way does have its merits.

*Motivation and Contribution.* Strictly speaking, the above application scenario only requires that parts of a signed message can be redacted without invalidating the signature, which is achieved by redactable signatures (RS) [7,8]. Yet, also editing, but not redacting, parts of a message has many use-cases, including secure routing, document sanitization, and outsourcing of computation, cf. [2], which can be achieved using sanitizable signatures (SS) [3].

Bilzhaue et al. [2] asked whether SSs and RSs could be combined, enabling signatures which allow for editing and redacting blocks at the same time. This question was answered by Krenn et al. by introducing “protean signatures” (PS) [4], which allow subsequent editing and redacting simultaneously. They provide a formal security model, a provably secure construction, and present first implementation results.

However, their definition of invisibility only guarantees that an outsider cannot decide which parts of a message are editable, but not which are redactable. This negatively impacts on the privacy guarantees their construction provides. Moreover, their corresponding implementation shows that their construction cannot be considered *practically* efficient.

We extend their work in the following areas: (i) As our main contribution we introduce a stronger invisibility definition for PSs. In more detail, in our definition, an outsider can neither decide which parts of a message are redactable nor which parts are editable. (ii) To show that our strengthened definition is actually achievable, we provide an altered (black-box) construction of a PS, derived from the original one by Krenn et al. [4]. (iii) We provide a new framework for aggregate signatures [9], which may be useful in other contexts as well. Namely, we introduce the new notions of strict aggregate uniqueness, a black-box “no-extraction” notion, and explicitly allow for de-aggregation. (iv) Our construction makes use of two new primitives which we also introduce: An invisible designated redactor RS, and a non-accountable invisible SS'. In an invisible designated redactor RS, only a signer-chosen semi-trusted third party can redact, while an outsider cannot decide which parts are redactable. Likewise, non-accountable invisible SS's behave as standard invisible SSs, but do not offer any form of accountability. This allows for a far more efficient instantiation. For both new notions, we provide formal frameworks, formal security models, and provably secure instantiations. (v) We have implemented our scheme in a prototypical way. Using our new primitives, the resulting construction is an order of magnitude more efficient than the one given by Krenn et al. [4], and can be considered really practical. (vi) Finally, on a very technical level, we provide a slightly stronger definition of invisibility for standard sanitizable signatures, where the adversary is now able to query *arbitrary* messages to the sanitization oracle instead of only honestly generated ones.

*Related Work.* Signatures allowing for subsequent alterations received a lot of attention in the recent past, as it became apparent that there are many application scenarios where signed messages need to be modified in a controlled way [10,2,11,12]. This weakens the standard unforgeability definition, where the messages protected by signatures cannot be altered at all, which is clearly not avoidable, if one wants to allow for modifications or derivations.

From our perspective, existing work can be grouped into three, not always distinct, directions. The first direction are homomorphic signatures [10,13,7,14], and some other closely related concepts [15,16]. Homomorphic signatures take several (signed) messages as input and can be used to compute functions on authenticated data-sets. In such schemes, an entity not holding any secrets can derive a new (valid) signature  $\sigma'$  on  $f(m)$ , where the function  $f$  is public.

Related are RSs, where anyone (i.e., no secrets are required) can publish a subset of signed data, along with a new signature  $\sigma'$ . To illustrate this, let  $m = (\text{I, do, not, like, fish})$  along with a valid redactable signature  $\sigma$ . Anyone can then derive a signature  $\sigma'$  on  $m' = (\text{I, like, fish})$ , i.e., redact the second and third block  $m^2 = \text{do}$  and  $m^3 = \text{not}$ , if both blocks are marked as redactable. The original ideas of RSs [7,8] were later formalized [17,18], including adding new values after signature generation [19]. Then, RSs have been extended to allow for additional use-cases, including adding accountability [20], discussing their relation to SSs [21], allowing for redactable structure [22], prohibiting additional redactions [23,24,25,26], yet also defining dependencies between different parts of a message [5]. Moreover, there are also some real-world implementations of this primitive proving that they are practical [27,6]. All these approaches (but accountability) have later been unified into a generalized framework by Derler et al. [28]. We stress that the work by Izu et al. [24] addresses the case of “sanitizable and deletable signatures”. However, they actually address the case of RSs and not SSs. In particular, in their scheme, a third party can decide whether a redaction is visible or not,

but does not allow for any other alterations. We follow the nomenclature clarified by Bilzhouse et al. [2], and thus classify the work by Izu et al. [24] as an RS.

In contrast, SSs allow editing of signer-chosen blocks of signed messages by a semi-trusted entity named the sanitizer [3]. In particular, the sanitizer holds its own secret key and can derive new messages, along with the corresponding signatures, but cannot completely redact blocks. For example, if  $m = (\text{I, do, not, like, fish})$  (and  $m^5$  is admissible, i.e., modifiable), then the sanitizer can, e.g., derive a new signature  $\sigma'$  on the message  $m' = (\text{I, do, not, like, meat})$ . Even though this seems to be off the limits, it turned out that this primitive has many real-life application scenarios, see, e.g., Bilzhouse et al. [2]. After the initial ideas by Ateniese et al. [3], SSs also received a lot of attention in the recent past. Namely, the first thorough security model was given by Brzuska et al. [29] (later revised by Gong et al. [30]), which was later extended for multiple signers/sanitizers [31,32], unlinkability (which means a derived signatures cannot be linked to its original) [33,34,35,36], trapdoor SSs (where a signer can choose additional sanitizers after signature generation) [37,38], non-interactive public-accountability (an outsider can determine which party is accountable for a given valid message/signature pair) [39], limiting the sanitizer to signer-chosen values [40,41,42], invisibility (meaning that an outsider cannot derive which parts of a message are sanitizable) [43,44,45] and the case of strongly unforgeable signatures [46]. All these extensions allow for additional use-cases of this primitive [2].

Additional related work is given in some recent surveys [2,11,47]. We stress that a slightly altered SS can be used to “mimic” an RS by defining a special symbol to which the specific block is sanitized to, which then marks the block as “redacted”. However, as shown by de Meer et al. [21], this has a negative impact on the privacy guarantees of the resulting scheme because the special symbol remains visible. For example,  $m' = (\text{I, like, fish})$  is clearly different from  $m' = (\text{I, } \perp, \perp, \text{ like, fish})$ . We stress that our scheme supports both possibilities, i.e., visible and non-visible (transparent) redactions, adding additional freedom.

## 2 Preliminaries and Notation

The security parameter is denoted by  $\lambda \in \mathbb{N}$ . All algorithms implicitly take  $1^\lambda$  as an additional input. We write  $a \leftarrow A(x)$  if  $a$  is assigned to the output of the deterministic algorithm  $A$  with input  $x$ . If an algorithm  $A$  is probabilistic, we write  $a \leftarrow_r A(x)$ . If we want to make the random coins  $r$  used explicit, we write  $a \leftarrow_r A(x; r)$ . Otherwise, we assume that they are drawn internally. An algorithm is efficient if it runs in probabilistic polynomial time (PPT) in the length of its input. For the remainder of this paper, all algorithms are PPT if not explicitly mentioned otherwise. Most algorithms may return a special error symbol  $\perp \notin \{0, 1\}^*$ , denoting an exception. If  $S$  is a set, we write  $a \leftarrow_r S$  to denote that  $a$  is chosen uniformly at random from  $S$ . For a message  $m = (m^1, m^2, \dots, m^{\ell_m})$ ,  $m^i$  is called a block and  $\ell_m \in \mathbb{N}$  denotes the number of blocks in  $m$ . If  $m$  is clear from the context, it is dropped from  $\ell_m$ . To shorten notation, we use  $[a..b]$  (both  $a$  and  $b$ ,  $b \geq a$ , are always positive natural numbers) for the set  $\{a, a + 1, \dots, b\}$ . 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  $\nu(n) \leq n^{-k}, \forall n > n_0$ .

### 2.1 Building Blocks

We now present our required building blocks. These include labeled IND-CCA2 secure public-key encryption schemes ( $\Pi$ ), sanitizable signature (SS), PRFs, and aggregate signatures ( $\Sigma$ ).

**Labeled Public-Key Encryption Schemes.** A labeled public-key encryption scheme  $\Pi$  allows to encrypt a message  $m$  using a given public key  $\text{pk}_\Pi$  and label  $\vartheta \in \{0, 1\}^*$ . In a nutshell, the given ciphertext leaks no information about the contained message, except its length, if the corresponding secret key  $\text{sk}_\Pi$  is not known:

**Definition 1 (Labeled Public-Key Encryption).** *A labeled public-key encryption scheme  $\Pi$  consists of four algorithms  $\{\text{PPGen}^\Pi, \text{KGen}^\Pi, \text{Enc}^\Pi, \text{Dec}^\Pi\}$ , such that:*

$\mathbf{Exp}_{\mathcal{A}, \Pi}^{\text{IND-CCA2}}(\lambda)$ :  
 $\text{pp}_{\Pi} \leftarrow_r \text{PPGen}^{\Pi}(1^{\lambda})$   
 $(\text{sk}_{\Pi}, \text{pk}_{\Pi}) \leftarrow_r \text{KGen}^{\Pi}(\text{pp}_{\Pi})$   
 $b \leftarrow_r \{0, 1\}$   
 $(m_0^*, m_1^*, \vartheta^*, \text{state}_{\mathcal{A}}) \leftarrow_r \mathcal{A}^{\text{Dec}^{\Pi}(\text{sk}_{\Pi}, \cdot, \cdot)}(\text{pk}_{\Pi})$   
 If  $|m_0^*| \neq |m_1^*| \vee m_0^* \notin \mathcal{M} \vee m_1^* \notin \mathcal{M}$ :  
 $c^* \leftarrow \perp$   
 Else:  
 $c^* \leftarrow_r \text{Enc}^{\Pi}(\text{pk}_{\Pi}, m_b^*, \vartheta^*)$   
 $a \leftarrow_r \mathcal{A}^{\text{Dec}^{\Pi'}(\text{sk}_{\Pi}, \cdot, \cdot)}(\text{state}_{\mathcal{A}}, c^*)$   
 where  $\text{Dec}^{\Pi'}(\text{sk}_{\Pi}, \cdot, \cdot)$  behaves as  $\text{Dec}^{\Pi}(\text{sk}_{\Pi}, \cdot, \cdot)$ ,  
 but returns  $\perp$ , if  $(c^*, \vartheta^*)$  is queried.  
 return 1, if  $a = b$   
 return 0

Fig. 1:  $\Pi$  IND-CCA2-Security

$\text{PPGen}^{\Pi}$ . This algorithm outputs the public parameters of the scheme:

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

It is assumed that  $\text{pp}_{\Pi}$  is an implicit input to all other algorithms.

$\text{KGen}^{\Pi}$ . On input  $\text{pp}_{\Pi}$ , this algorithm outputs the key pair:

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

$\text{Enc}^{\Pi}$ . On input the public key  $\text{pk}_{\Pi}$ , a message  $m$ , and a label  $\vartheta \in \{0, 1\}^*$ , this algorithm outputs a ciphertext  $c$ :

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

$\text{Dec}^{\Pi}$ . On input  $\text{sk}_{\Pi}$ ,  $\vartheta$  and a ciphertext  $c$ , this deterministic algorithm outputs a message  $m$  or  $\perp$ :

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

$\Pi$  IND-CCA2-Security. We need IND-CCA2-security (and perfect correctness) for our construction to work. Note,  $\mathcal{M}$  is some message space implicitly defined by  $\text{pk}_{\Pi}$  and  $\text{pp}_{\Pi}$ .

**Definition 2 (IND-CCA2-Security).** A labeled encryption scheme  $\Pi$  is IND-CCA2-secure, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that:

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

The corresponding experiment is depicted in Figure 1.

A suitable instantiation is CS-encryption with labels [48].

**Aggregate Signatures.** Standard digital signatures allow the holder of a secret key  $\text{sk}_{\Sigma}$  to sign a message  $m$ , while with knowledge of the corresponding public key  $\text{pk}_{\Sigma}$ , everyone can verify whether a given signature was actually endorsed by the signer, i.e., the owner of  $\text{pk}_{\Sigma}$  [49]. An aggregate signature scheme ( $\Sigma$ ) allows to aggregate multiple signatures into a single (short) value [9].

As a side contribution, we introduce a new framework for aggregate signatures, where one can also de-aggregate signatures, and a novel *aggregate* uniqueness definition. This was, to the best of our knowledge, only considered in a white-box fashion in the context of BGLS-signatures [9], but was never formally defined in a black-box way, or only in some ad-hoc fashion [50].

**Definition 3 (Aggregate Signatures).** *An aggregate signature scheme  $\Sigma$  with explicit de-aggregation consists of six algorithms  $\{\text{PPGen}_\Sigma, \text{KGen}_\Sigma, \text{Sign}_\Sigma, \text{AVerf}_\Sigma, \text{Agg}_\Sigma, \text{DAgg}_\Sigma\}$  such that:*

$\text{PPGen}_\Sigma$ . *This algorithm outputs the public parameters:*

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

*We assume that  $\text{pp}_\Sigma$  contains  $1^\lambda$  and is implicit input to all other algorithms.*

$\text{KGen}_\Sigma$ . *On input  $\text{pp}_\Sigma$ , this algorithm outputs the key pair of a signer:*

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

$\text{Sign}_\Sigma$ . *On input the secret key  $\text{sk}_\Sigma$  and a message  $m$ , this algorithm outputs a signature:*

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

$\text{AVerf}_\Sigma$ . *On input a set  $\{(\text{pk}_{\Sigma,i}, m_i)\}$  of public keys/messages and an aggregate signature  $\sigma$ , this algorithm outputs a decision bit  $d \in \{0, 1\}$ :*

$$d \leftarrow \text{AVerf}_\Sigma(\{(\text{pk}_{\Sigma,i}, m_i)\}, \sigma)$$

$\text{Agg}_\Sigma$ . *On input a set of key/message tuples with corresponding aggregated signatures  $\mathcal{S}_{\text{agg}} = \{(\text{pk}_{\Sigma,i,j}, m_{i,j}), \sigma_i\}$ , where the sets  $\{(\text{pk}_{\Sigma,i,j}, m_{i,j})\}$  are pair-wise disjoint, this algorithm outputs a new aggregate signature  $\sigma'$  (or  $\perp$  if  $\mathcal{S}_{\text{agg}} = \emptyset$ ):*

$$\sigma'' \leftarrow \text{Agg}_\Sigma(\{(\text{pk}_{\Sigma,i,j}, m_{i,j}), \sigma_i\})$$

*We note that Boneh et al. require that each message appears at most once [9]. However, there are also mitigation strategies [51,9].*

$\text{DAgg}_\Sigma$ . *On input a set  $\{(\text{pk}_{\Sigma,i,j}, m_{i,j}), \sigma_i\}$  of public key/message tuples, along with an aggregate signature  $\sigma_i$  (protecting  $\{(\text{pk}_{\Sigma,i,j}, m_{i,j})\}$ , being pair-wise disjoint), an additional set  $\{(\text{pk}_{\Sigma,k}, m_k)\}$  with aggregate signature  $\sigma_k$ , this algorithm outputs a new aggregate signature  $\sigma'$  for  $\{(\text{pk}_{\Sigma,k}, m_k)\} \cup \bigcup_i \{(\text{pk}_{\Sigma,i,j}, m_{i,j})\}$ :*

$$\sigma' \leftarrow \text{DAgg}_\Sigma(\{(\text{pk}_{\Sigma,i,j}, m_{i,j}), \sigma_i\}, \{(\text{pk}_{\Sigma,k}, m_k)\}, \sigma_k)$$

Clearly, an aggregate signature may also protect a single message, i.e., degenerate to a “normal” signature.

Moreover, we require the usual correctness properties to hold. Namely, honestly generated signatures verify, which must also be true for honestly generated aggregates. Likewise, honestly generated signatures stemming from de-aggregation must also verify.

From a security perspective, we require existential unforgeability under chosen-message attacks (eUNF-CMA), a strict form of uniqueness, correctness, and that a third party cannot remove signatures from an aggregate, if it does not know the signatures for the messages to be removed. We now formally define each of those properties.

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.

```

Exp $\mathcal{A}, \Sigma$ eUNF-CMA( $\lambda$ )
   $\text{pp}_\Sigma \leftarrow_r \text{PPGen}_\Sigma(1^\lambda)$ 
   $(\text{sk}_\Sigma, \text{pk}_\Sigma) \leftarrow_r \text{KGen}_\Sigma(\text{pp}_\Sigma)$ 
   $\mathcal{Q} \leftarrow \emptyset$ 
   $(\{(\text{pk}_i^*, m_i^*)\}, m^*, \sigma^*) \leftarrow_r \mathcal{A}^{\text{Sign}'_\Sigma(\text{sk}_\Sigma, \cdot)}(\text{pk}_\Sigma)$ 
  where  $\text{Sign}'_\Sigma(\text{sk}_\Sigma, m)$ :
     $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{m\}$ 
    return  $\sigma \leftarrow \text{Sign}_\Sigma(\text{sk}_\Sigma, m)$ 
  return 1, if  $\text{AVerf}_\Sigma(\{(\text{pk}_i^*, m_i^*)\} \cup \{(\text{pk}_\Sigma, m^*)\}, \sigma^*) = 1 \wedge m^* \notin \mathcal{Q}$ 
  return 0

```

Fig. 2:  $\Sigma$  Unforgeability

```

Exp $\mathcal{A}, \Sigma$ Uniqueness( $\lambda$ )
   $\text{pp}_\Sigma \leftarrow_r \text{PPGen}_\Sigma(1^\lambda)$ 
   $(\{(\text{pk}_i^*, m_i^*)\}, \sigma^*, \sigma'^*) \leftarrow_r \mathcal{A}(\text{pp}_\Sigma)$ 
  return 1, if  $\text{AVerf}_\Sigma(\{(\text{pk}_i^*, m_i^*)\}, \sigma^*) = 1 \wedge$ 
     $\text{AVerf}_\Sigma(\{(\text{pk}_i^*, m_i^*)\}, \sigma'^*) = 1 \wedge \sigma^* \neq \sigma'^*$ 
  return 0

```

Fig. 3:  $\Sigma$  Uniqueness

**Definition 4** ( $\Sigma$  Unforgeability). *We say a  $\Sigma$  scheme is unforgeable, if for every PPT adversary  $\mathcal{A}$ , there exists a negligible function  $\nu$  such that:*

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

The corresponding experiment is depicted in Figure 2.

Uniqueness for aggregate signatures requires that for each set  $\{(\text{pk}_i, m_i)\}$  at most one signature can be found, even if all values can be adversarially generated. In contrast to Kuchta and Manulis [52], we require a slightly different uniqueness notion, i.e., the *complete* signature must be unique, and not only some part of it. Additionally, all values, but the public parameters, are explicitly generated by the adversary.

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

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

The corresponding experiment is depicted in Figure 3.

*Remark 1.* It should be obvious that uniqueness for aggregate signatures implies uniqueness for normal signatures. However, the converse is not true by the following argument: Append a random bit for every generated *aggregate*, and remove it before proceeding with the other algorithms, appending another random bit once they are finished, if the resulting signature is an aggregate.

Our final definition requires that an adversary cannot de-aggregate signatures from an aggregate signature, if the adversary never saw a signature (or aggregate) for the messages for which it tries to remove the signature from the aggregate. This resembles the  $k$ -element extraction assumption by Boneh et al. [9,53], but

for general aggregate signatures. To find out whether the adversary actually wins, we need to explicitly disregard aggregates which the adversary could find using non-avoidable transitivity relationships and adversarial signatures added “on top”. We therefore use the algorithm Closure which finds all “trivial” relations.

---

**Algorithm 1:** Algorithm Closure

---

**Input:**  $\mathcal{S} = \{(\text{pk}_i, m_i)\}$ ,  $\mathcal{D} = \{(\text{pk}_j, m_j)\}$   
**Output:**  $\mathcal{D}' = \{(\text{pk}_k, m_k)\}$

```

1  $l \leftarrow 0, l' \leftarrow 0$ 
2  $\mathcal{D}' \leftarrow \mathcal{D} \cup \mathcal{S}$ 
3 do
4    $\mathcal{T} \leftarrow \mathcal{D}'$ 
5    $l \leftarrow |\mathcal{T}|$ 
6   foreach  $\{(\text{pk}_l, m_l)\} \in \mathcal{T}$  do
7     foreach  $\{(\text{pk}_m, m_m)\} \neq \{(\text{pk}_l, m_l)\} \in \mathcal{T}$  do
8       if  $\{(\text{pk}_l, m_l)\} \subsetneq \{(\text{pk}_m, m_m)\}$  then
9          $\mathcal{D}' \leftarrow \mathcal{D}' \cup (\{(\text{pk}_m, m_m)\} \setminus \{(\text{pk}_l, m_l)\})$ 
10   $l' \leftarrow |\mathcal{D}'|$ 
11 while  $l \neq l'$ 
12 return  $\mathcal{D}'$ 

```

---

```

Exp $\mathcal{A}, \Sigma$ NoExtract( $\lambda$ )
   $\text{pp}_\Sigma \leftarrow_r \text{PPGen}_\Sigma(1^\lambda)$ 
   $\mathcal{M} \leftarrow \emptyset$ 
   $\mathcal{K} \leftarrow \emptyset$ 
   $(\{(\text{pk}_i^*, m_i^*)\}, \sigma^*) \leftarrow_r \mathcal{A}^{\text{KGen}'_\Sigma(\text{pp}_\Sigma), \text{Sign}'_\Sigma(\cdot)}(\text{pp}_\Sigma)$ 
  where  $\text{KGen}'_\Sigma(\text{pp}_\Sigma)$ :
     $(\text{sk}_i, \text{pk}_i) \leftarrow_r \text{KGen}_\Sigma(\text{pp}_\Sigma)$ 
     $\mathcal{K} \leftarrow \mathcal{K} \cup (\text{sk}_i, \text{pk}_i)$ 
    return  $\text{pk}_i$ 
  where  $\text{Sign}'_\Sigma(\{(\text{pk}_i, m_i)\})$ :
    return  $\perp$ , if  $\exists i : (\cdot, \text{pk}_i) \notin \mathcal{K}$ 
    for all  $(\text{pk}_i, m_i)$ , let  $\sigma_i \leftarrow \text{Sign}_\Sigma(\text{sk}_i, m_i)$ 
     $\sigma' \leftarrow \text{Agg}_\Sigma(\{(\text{pk}_i, m_i), \sigma_i\})$ 
     $\mathcal{M} \leftarrow \text{Closure}(\{(\text{pk}_i, m_i)\}, \mathcal{M})$ 
    return  $\sigma'$ 
   $\mathcal{C} \leftarrow \{(\text{pk}_i^*, m_i^*) \mid (\text{pk}_i^*, \cdot) \in \mathcal{K}\}$ 
  return 1, if  $\text{AVerf}_\Sigma(\{(\text{pk}_i^*, m_i^*)\}, \sigma^*) = 1 \wedge \mathcal{C} \notin \mathcal{M}$ 
  return 0

```

Fig. 4:  $\Sigma$  Extraction Secure

**Definition 6 ( $\Sigma$  No-Extraction).** We say a  $\Sigma$  scheme provides no-extraction, if for every PPT adversary  $\mathcal{A}$ , there exists a negligible function  $\nu$  such that:

$$\Pr \left[ \mathbf{Exp}_{\mathcal{A}, \Sigma}^{\text{NoExtract}}(\lambda) = 1 \right] \leq \nu(\lambda)$$

The corresponding experiment is depicted in Figure 4, where the algorithm Closure is given in Algorithm 1.

We additionally require that, regardless of the message  $m$ , each signature is of constant size. This can easily be achieved by hashing the message  $m$  using a collision-resistant hash-function prior to signing. A suitable instantiation is BGLS-signatures [9], if one can enforce distinct messages (what we do).

**Definition 7 (Pseudo-Random Functions).** A pseudo-random function PRF consists of two algorithms  $\{\text{KGen}_{\text{PRF}}, \text{Eval}_{\text{PRF}}\}$  such that:

$\text{KGen}_{\text{PRF}}$ . On input  $\lambda$ , this algorithm outputs a function key  $x \in \{0, 1\}^\lambda$ :

$$x \leftarrow_r \text{KGen}_{\text{PRF}}(1^\lambda)$$

$\text{Eval}_{\text{PRF}}$ . On input a function key  $x$  and  $p \in \{0, 1\}^\lambda$ , this deterministic algorithm outputs  $p' \in \{0, 1\}^\lambda$ :

$$p' \leftarrow \text{Eval}_{\text{PRF}}(x, p)$$

*Security.* For security, it is required that PRF is actually pseudo-random.

```

Exp $\mathcal{A}, \text{PRF}$ PR( $\lambda$ )
 $x \leftarrow_r \text{KGen}_{\text{PRF}}(1^\lambda)$ 
 $b \leftarrow_r \{0, 1\}$ 
 $f \leftarrow_r F_\lambda$ 
 $a \leftarrow_r \mathcal{A}^{\text{Eval}_{\text{PRF}}(x, \cdot)}(1^\lambda)$ 
  where oracle  $\text{Eval}_{\text{PRF}}(x, p)$ :
    return  $\perp$ , if  $p \notin \{0, 1\}^\lambda$ 
    if  $b = 0$ , return  $\text{Eval}_{\text{PRF}}(x, p)$ 
    return  $f(p)$ 
  return 1, if  $a = b$ 
  return 0

```

Fig. 5: PRF Pseudo-Randomness

**Definition 8 (PRF Pseudo-Randomness).** A pseudo-random function PRF is called *pseudo-random*, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that:

$$\left| \Pr[\mathbf{Exp}_{\mathcal{A}, \text{PRF}}^{\text{PR}}(\lambda) = 1] - \frac{1}{2} \right| \leq \nu(\lambda)$$

The corresponding experiment is depicted in Figure 5, where  $F_\lambda = \{f : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda\}$  is the set of all functions  $f$  mapping a value  $v \in \{0, 1\}^\lambda$  to another value  $v' \in \{0, 1\}^\lambda$ .

**Sanitizable Signatures.** Subsequently, we restate the definitions of (standard) SSs [43,29,46]. To recap, a SS allows a semi-trusted third party, named the sanitizer, to alter signer-chosen blocks to arbitrary bit-strings. The sanitizer holds its own key-pair and can be held accountable, if it sanitizes a message.

The following framework is essentially the one given by Camenisch et al. [44], which is itself based on existing work [29]. However, some additional notation is required beforehand. The variable  $\mathbb{A}^{\text{SS}}$  contains the set of indices of the modifiable blocks, as well as  $\ell$ , denoting the total number of blocks in the message  $m$ . For example, let  $\mathbb{A}^{\text{SS}} = (\{1, 2, 4\}, 4)$ . Then,  $m$  must contain four blocks ( $\ell = 4$ ) and all but the third are admissible. Note,  $\mathbb{A}^{\text{SS}}$  can be encoded in a length-invariant way by using a sequence of bits, e.g.,  $(1, 1, 0, 1)$



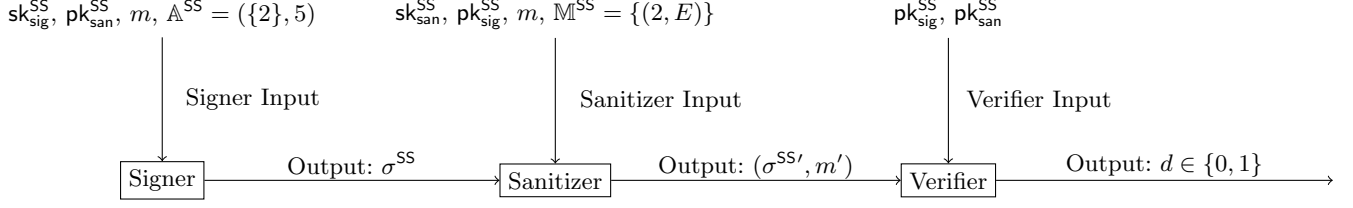


Fig. 6: Example workflow of an SS. The message  $m$  is set to  $(H, A, L, L, O)$  and is sanitized  $m' = (H, E, L, L, O)$

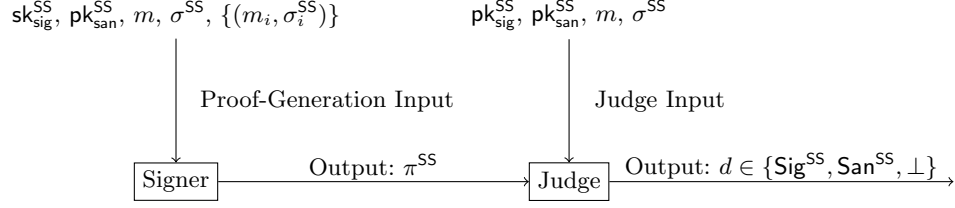


Fig. 7: Additional algorithms  $\text{Proof}^{\text{SS}}$  and  $\text{Judge}^{\text{SS}}$  for an accountable SS

for  $\mathbb{A}^{\text{SS}} = (\{1, 2, 4\}, 4)$ . The variable  $\mathbb{M}^{\text{SS}}$  is a set containing pairs  $(i, m^{i'})$  for those blocks that are modified, meaning that  $m^i$  is replaced by  $m^{i'}$ . We use the shorthand notation  $m' \leftarrow \mathbb{M}^{\text{SS}}(m)$  to denote the result of this replacement, while  $\mathbb{M}^{\text{SS}} \prec (m, \mathbb{A}^{\text{SS}})$  means that  $\mathbb{M}^{\text{SS}}$  is a valid modification instruction w.r.t.  $m$  and  $\mathbb{A}^{\text{SS}}$ . Likewise, we use  $\mathbb{A}^{\text{SS}} \prec m$  to denote that  $\mathbb{A}^{\text{SS}}$  is valid description of the admissible blocks w.r.t.  $m$ . An example workflow is depicted in Figure 6 and Figure 7. Both are derived from Bilzhaue et al. [2].

**Definition 9 (Sanitizable Signatures).** A sanitizable signature scheme  $\text{SS}$  consists of eight PPT algorithms  $\{\text{PPGen}^{\text{SS}}, \text{KGSig}^{\text{SS}}, \text{KGSan}^{\text{SS}}, \text{Sign}^{\text{SS}}, \text{Verify}^{\text{SS}}, \text{Sanit}^{\text{SS}}, \text{Proof}^{\text{SS}}, \text{Judge}^{\text{SS}}\}$  such that:

$\text{PPGen}^{\text{SS}}$ . This algorithm generates the public parameters:

$$\text{pp}_{\text{SS}} \leftarrow_r \text{PPGen}^{\text{SS}}(1^\lambda)$$

We assume that  $\text{pp}_{\text{SS}}$  is implicitly input to all other algorithms.

$\text{KGSig}^{\text{SS}}$ . On input  $\text{pp}_{\text{SS}}$ , this algorithm outputs a key pair of the signer:

$$(\text{sk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{sig}}^{\text{SS}}) \leftarrow_r \text{KGSig}^{\text{SS}}(\text{pp}_{\text{SS}})$$

$\text{KGSan}^{\text{SS}}$ . On input  $\text{pp}_{\text{SS}}$ , this algorithm outputs a key pair of the sanitizer:

$$(\text{sk}_{\text{san}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}) \leftarrow_r \text{KGSan}^{\text{SS}}(\text{pp}_{\text{SS}})$$

$\text{Sign}^{\text{SS}}$ . On input of the public key  $\text{pk}_{\text{san}}^{\text{SS}}$ ,  $\mathbb{A}^{\text{SS}}$ , a message  $m$  and  $\text{sk}_{\text{sig}}^{\text{SS}}$ , this algorithm outputs a signature  $\sigma^{\text{SS}}$ :

$$\sigma^{\text{SS}} \leftarrow_r \text{Sign}^{\text{SS}}(\text{sk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}, m, \mathbb{A}^{\text{SS}})$$

$\text{Verify}^{\text{SS}}$ . On input  $\text{pk}_{\text{sig}}^{\text{SS}}$ ,  $\text{pk}_{\text{san}}^{\text{SS}}$ , message  $m$ , and a signature  $\sigma^{\text{SS}}$ , this deterministic algorithm outputs a decision  $d \in \{0, 1\}$ :

$$d \leftarrow \text{Verify}^{\text{SS}}(\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}, m, \sigma^{\text{SS}})$$

$\text{Sanit}^{\text{SS}}$ . On input  $\text{sk}_{\text{san}}^{\text{SS}}$ ,  $\text{pk}_{\text{sig}}^{\text{SS}}$ , message  $m$ , a signature  $\sigma^{\text{SS}}$ , and  $\mathbb{A}^{\text{SS}}$ , this algorithm outputs a sanitized signature  $\sigma^{\text{SS}'}$ :

$$(m', \sigma^{\text{SS}'}) \leftarrow_r \text{Sanit}^{\text{SS}}(\text{sk}_{\text{san}}^{\text{SS}}, \text{pk}_{\text{sig}}^{\text{SS}}, m, \sigma^{\text{SS}}, \mathbb{M}^{\text{SS}})$$

$\mathbf{Exp}_{\mathcal{A}, \text{SS}}^{\text{Unforgeability}}(\lambda)$   
 $\text{pp}_{\text{SS}} \leftarrow_r \text{PPGen}^{\text{SS}}(1^\lambda)$   
 $(\text{sk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{sig}}^{\text{SS}}) \leftarrow_r \text{KG}_{\text{sig}}^{\text{SS}}(\text{pp}_{\text{SS}})$   
 $(\text{sk}_{\text{san}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}) \leftarrow_r \text{KG}_{\text{san}}^{\text{SS}}(\text{pp}_{\text{SS}})$   
 $(m^*, \sigma^{\text{SS}*}) \leftarrow_r \mathcal{A}_{\text{Proof}^{\text{SS}}(\text{sk}_{\text{sig}}^{\text{SS}}, \dots, \text{Sanit}^{\text{SS}}(\text{sk}_{\text{san}}^{\text{SS}}, \dots, \cdot))}(\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}})$   
 $\forall i \in [1..q]$ , let  $(\text{pk}_{\text{san}, i}^{\text{SS}}, m_i, \mathbb{A}_i^{\text{SS}})$  and  $\sigma_i^{\text{SS}}$   
index the queries/answers to/from  $\text{Sign}^{\text{SS}}$   
 $\forall i \in [1..q']$ , let  $(\text{pk}_{\text{sig}, j}^{\text{SS}}, m_j, \sigma_j^{\text{SS}}, \mathbb{M}_j)$  and  $(m'_j, \sigma_j^{\text{SS}'})$   
index the queries/answers to/from  $\text{Sanit}^{\text{SS}}$   
return 1, if  $\text{Verify}^{\text{SS}}(\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}, m^*, \sigma^{\text{SS}*}) = 1 \wedge$   
 $\forall i \in [1..q] : (\text{pk}_{\text{san}}^{\text{SS}}, m^*, \sigma^{\text{SS}*}) \neq (\text{pk}_{\text{san}, i}^{\text{SS}}, m_i, \sigma_i^{\text{SS}}) \wedge$   
 $\forall j \in [1..q'] : (\text{pk}_{\text{sig}}^{\text{SS}}, m^*, \sigma^{\text{SS}*}) \neq (\text{pk}_{\text{sig}, j}^{\text{SS}}, m'_j, \sigma_j^{\text{SS}'})$   
return 0

Fig. 8: SS Unforgeability

$\text{Proof}^{\text{SS}}$ . On input  $\text{sk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}$ , a message  $m, \sigma^{\text{SS}}$ , and a set of polynomially many additional signature/message pairs  $\{(\sigma_i^{\text{SS}}, m_i)\}$ , this algorithm outputs a proof  $\pi^{\text{SS}}$  to pinpoint the accountable party for a given signature:

$$\pi^{\text{SS}} \leftarrow_r \text{Proof}^{\text{SS}}(\text{sk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}, m, \sigma^{\text{SS}}, \{(\sigma_i^{\text{SS}}, m_i)\})$$

$\text{Judge}^{\text{SS}}$ . On input  $\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}$ , message  $m, \sigma^{\text{SS}}$ , and a proof  $\pi^{\text{SS}}$ , this algorithm outputs a decision  $d \in \{\text{Sig}^{\text{PS}}, \text{San}^{\text{PS}}, \perp\}$ :

$$d \leftarrow \text{Judge}^{\text{SS}}(\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}, m, \sigma^{\text{SS}}, \pi^{\text{SS}})$$

Correctness was already specified by Brzuska et al. [29].

**SSs Security Definitions.** We now introduce the security properties required. These are the ones given by Beck et al. [43], but altered for the used notation, already incorporating the strong definitions by Krenn et al. [46], but a stronger notion of invisibility, where the adversary is now able to query arbitrary signatures to the sanitization oracle. Moreover, we consider neither unlinkability nor non-interactive public-accountability, as it depends on the context whether these properties are required [43,44]. However, non-interactive public-accountability is easy to achieve, e.g., by signing the resulting signature again [39].

*Unforgeability.* This definition requires that an adversary  $\mathcal{A}$  not having any secret keys is not able to produce any new valid signature  $\sigma^*$  on a message  $m^*$  which it has never seen, even if  $\mathcal{A}$  has full oracle access.

**Definition 10 (Unforgeability).** An SS is unforgeable, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that:

$$\Pr[\mathbf{Exp}_{\mathcal{A}, \text{SS}}^{\text{Unforgeability}}(\lambda) = 1] \leq \nu(\lambda),$$

where the corresponding experiment is defined in Figure 8.

*Immutability.* This definition prohibits that an adversary  $\mathcal{A}$  can generate a verifying signature  $\sigma^{\text{SS}*}$  for a message  $m^*$  not derivable from the signatures given by an honest signer, even if it can generate the sanitizer's key pair.

**Definition 11 (Immutability).** An SS is immutable, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that:

$$\Pr[\mathbf{Exp}_{\mathcal{A}, \text{SS}}^{\text{Immutability}}(\lambda) = 1] \leq \nu(\lambda),$$

where the corresponding experiment is defined in Figure 9.

$\mathbf{Exp}_{\mathcal{A},SS}^{\text{Immutability}}(\lambda)$   
 $\text{pp}_{SS} \leftarrow_r \text{PPGen}^{SS}(1^\lambda)$   
 $(\text{sk}_{\text{sig}}^{SS}, \text{pk}_{\text{sig}}^{SS}) \leftarrow_r \text{KG}_{\text{sig}}^{SS}(\text{pp}_{SS})$   
 $(m^*, \sigma^{SS*}, \text{pk}_{\text{san}}^{SS*}) \leftarrow_r \mathcal{A}^{\text{Sign}^{SS}(\text{sk}_{\text{sig}}^{SS}, \cdot, \cdot, \cdot), \text{Proof}^{SS}(\text{sk}_{\text{sig}}^{SS}, \cdot, \cdot, \cdot)}(\text{pk}_{\text{sig}}^{SS})$   
 $\forall i \in [1..q]$ , let  $(\text{pk}_{\text{san},i}^{SS}, m_i, \mathbb{A}_i^{SS})$   
 index the queries to  $\text{Sign}^{SS}$   
 return 1, if  $\text{Verify}^{SS}(\text{pk}_{\text{sig}}^{SS}, \text{pk}_{\text{san}}^{SS*}, m^*, \sigma^{SS*}) = 1 \wedge$   
 $\forall i \in [1..q] : (\text{pk}_{\text{san}}^{SS*} \neq \text{pk}_{\text{san},i}^{SS} \vee$   
 $m^* \notin \{\mathbb{M}(m_i) \mid \mathbb{M} \text{ with } \mathbb{M}^{SS}(\mathbb{A}_i^{SS}) = 1\})$   
 return 0

Fig. 9: SS Immutability

$\mathbf{Exp}_{\mathcal{A},SS}^{\text{Privacy}}(\lambda)$   
 $\text{pp}_{SS} \leftarrow_r \text{PPGen}^{SS}(1^\lambda)$   
 $(\text{sk}_{\text{sig}}^{SS}, \text{pk}_{\text{sig}}^{SS}) \leftarrow_r \text{KG}_{\text{sig}}^{SS}(\text{pp}_{SS})$   
 $(\text{sk}_{\text{san}}^{SS}, \text{pk}_{\text{san}}^{SS}) \leftarrow_r \text{KG}_{\text{san}}^{SS}(\text{pp}_{SS})$   
 $b \leftarrow_r \{0, 1\}$   
 $a \leftarrow_r \mathcal{A}^{\text{Sign}^{SS}(\text{sk}_{\text{sig}}^{SS}, \cdot, \cdot, \cdot), \text{Sanit}^{SS}(\text{sk}_{\text{san}}^{SS}, \cdot, \cdot, \cdot)}_{\text{Proof}^{SS}(\text{sk}_{\text{sig}}^{SS}, \cdot, \cdot, \cdot), \text{LoRSan}(\cdot, \cdot, \cdot, \cdot, \text{sk}_{\text{sig}}^{SS}, \text{sk}_{\text{san}}^{SS}, b)}(\text{pk}_{\text{sig}}^{SS}, \text{pk}_{\text{san}}^{SS})$   
 where  $\text{LoRSan}(m_0, m_1, \mathbb{M}_0^{SS}, \mathbb{M}_1^{SS}, \mathbb{A}^{SS}, \text{sk}_{\text{sig}}^{SS}, \text{sk}_{\text{san}}^{SS}, b)$ :  
 return  $\perp$ , if  $\mathbb{M}_0^{SS} \neq \mathbb{M}_1^{SS} \vee \mathbb{A}^{SS} \neq m_0 \vee \mathbb{A}^{SS} \neq m_1 \vee$   
 $\mathbb{M}_0^{SS}(m_0) \neq \mathbb{M}_1^{SS}(m_1) \vee \mathbb{A}^{SS} \neq m_0 \vee \mathbb{A}^{SS} \neq m_1$   
 $\sigma^{SS} \leftarrow_r \text{Sign}^{SS}(\text{sk}_{\text{sig}}^{SS}, \text{pk}_{\text{san}}^{SS}, m_b, \mathbb{A}^{SS})$   
 return  $(m', \sigma^{SS'}) \leftarrow_r \text{Sanit}^{SS}(\text{sk}_{\text{san}}^{SS}, \text{pk}_{\text{sig}}^{SS}, m_b, \sigma^{SS}, \mathbb{M}_b^{SS})$   
 return 1, if  $a = b$   
 return 0

Fig. 10: SS Privacy

*Privacy.* This definition prohibits that an adversary  $\mathcal{A}$  can learn anything about sanitized parts. This is similar to the definition of standard encryption schemes.

**Definition 12 (Privacy).** *An SS is private, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that:*

$$\left| \Pr[\mathbf{Exp}_{\mathcal{A},SS}^{\text{Privacy}}(\lambda)] - 1/2 \right| \leq \nu(\lambda),$$

where the corresponding experiment is defined in Figure 10.

*Transparency.* This definition prohibits that an adversary  $\mathcal{A}$  does not learn whether a signature  $\sigma^{SS}$  was generated through  $\text{Sign}^{SS}$  or  $\text{Sanit}^{SS}$ . We stress that the adversary cannot query signatures obtained by the Sign/Sanit-oracle to the  $\text{Proof}^{SS'}$ -oracle to avoid trivial attacks.

**Definition 13 (Transparency).** *An SS is transparent, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that:*

$$\left| \Pr[\mathbf{Exp}_{\mathcal{A},SS}^{\text{Transparency}}(\lambda)] - 1/2 \right| \leq \nu(\lambda),$$

where the corresponding experiment is defined in Figure 11.

**Exp**<sub>A,SS</sub><sup>Transparency</sup>( $\lambda$ )

$pp_{SS} \leftarrow_r \text{PPGen}^{SS}(1^\lambda)$   
 $(sk_{sig}^{SS}, pk_{sig}^{SS}) \leftarrow_r \text{KG}_{sig}^{SS}(pp_{SS})$   
 $(sk_{san}^{SS}, pk_{san}^{SS}) \leftarrow_r \text{KG}_{san}^{SS}(pp_{SS})$   
 $b \leftarrow_r \{0, 1\}$   
 $Q \leftarrow \emptyset$

$a \leftarrow_r \mathcal{A}_{\text{Proof}^{SS'}(sk_{sig}^{SS}, \cdot, \cdot, \cdot), \text{Sign/Sanit}(\cdot, \cdot, sk_{sig}^{SS}, sk_{san}^{SS}, b)}^{\text{Sign}^{SS}(sk_{sig}^{SS}, \cdot, \cdot, \cdot), \text{Sanit}^{SS}(sk_{san}^{SS}, \cdot, \cdot, \cdot)}(pk_{sig}^{SS}, pk_{san}^{SS})$   
 where  $\text{Proof}^{SS'}(sk_{sig}^{SS}, m, \sigma^{SS}, \{(m_i, \sigma_i^{SS}) \mid i \in \mathbb{N}\})$ :  
 return  $\perp$ , if  $pk_{san}^{SS'} = pk_{san}^{SS} \wedge$   
 $((m, \sigma^{SS}) \in Q \vee Q \cap \{(m_i, \sigma_i^{SS})\} \neq \emptyset)$   
 return  $\text{Proof}^{SS}(sk_{sig}^{SS}, pk_{san}^{SS'}, m, \sigma^{SS}, \{(m_i, \sigma_i^{SS})\})$   
 where  $\text{Sign/Sanit}(m, \mathbb{M}^{SS}, \mathbb{A}^{SS}, sk_{sig}^{SS}, sk_{san}^{SS})$ :  
 $\sigma^{SS} \leftarrow_r \text{Sign}^{SS}(sk_{sig}^{SS}, pk_{san}^{SS}, m, \mathbb{A}^{SS})$   
 $(m', \sigma^{SS'}) \leftarrow_r \text{Sanit}^{SS}(sk_{san}^{SS}, pk_{sig}^{SS}, m, \sigma^{SS}, \mathbb{M}^{SS})$   
 if  $b = 1$ :  
 $\sigma^{SS'} \leftarrow_r \text{Sign}^{SS}(sk_{sig}^{SS}, pk_{san}^{SS}, m', \mathbb{A}^{SS})$   
 If  $\sigma^{SS'} \neq \perp$ , set  $Q \leftarrow Q \cup \{(m', \sigma^{SS'})\}$   
 return  $(m', \sigma^{SS'})$   
 return 1, if  $a = b$   
 return 0

Fig. 11: SS Transparency

**Exp**<sub>A,SS</sub><sup>SigAccountability</sup>( $\lambda$ )

$pp_{SS} \leftarrow_r \text{PPGen}^{SS}(1^\lambda)$   
 $(sk_{san}^{SS}, pk_{san}^{SS}) \leftarrow_r \text{KG}_{san}^{SS}(pp_{SS})$   
 $(pk_{sig}^{SS*}, \pi^{SS*}, m^*, \sigma^{SS*}) \leftarrow_r \mathcal{A}_{\text{Sanit}^{SS}(sk_{san}^{SS}, \cdot, \cdot, \cdot)}(pk_{san}^{SS})$   
 $\forall i \in [1..q]$ , let  $(m'_i, \sigma_i^{SS'})$  and  $(m_i, \mathbb{M}_i^{SS}, \sigma_i^{SS}, pk_{sig}^{SS, i})$   
 index the answers/queries from/to  $\text{Sanit}^{SS}$   
 return 1, if  $\text{Verify}^{SS}(pk_{sig}^{SS*}, pk_{san}^{SS}, m^*, \sigma^{SS*}) = 1 \wedge$   
 $\forall i \in [1..q] : (pk_{sig}^{SS*}, m^*, \sigma^{SS*}) \neq (pk_{sig}^{SS, i}, m'_i, \sigma_i^{SS'}) \wedge$   
 $\text{Judge}^{SS}(pk_{sig}^{SS*}, pk_{san}^{SS}, m^*, \sigma^{SS*}, \pi^{SS*}) = \text{San}^{SS}$   
 return 0

Fig. 12: SS Signer-Accountability

*Signer-Accountability.* Signer-accountability prohibits that an adversary can generate a bogus proof that makes  $\text{Judge}^{SS}$  decide that the sanitizer is responsible for a given signature/message pair  $(m^*, \sigma^{SS*})$ , but the sanitizer has never generated this pair. This is even true, if the adversary can generate the signer's key pair.

**Definition 14 (Signer-Accountability).** An SS is signer-accountable, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that:

$$\Pr[\mathbf{Exp}_{\mathcal{A},SS}^{\text{SigAccountability}}(\lambda) = 1] \leq \nu(\lambda),$$

where the corresponding experiment is defined in Figure 12.

*Sanitizer-Accountability.* Sanitizer-accountability prohibits that an adversary can generate a bogus signature/message pair  $(m^*, \sigma^{SS*})$  that makes  $\text{Proof}^{SS}$  output a (honestly generated) generated proof  $\pi^{SS}$  which

```

Exp $\mathcal{A}, \text{SS}$ SanAccountability( $\lambda$ )
   $\text{pp}_{\text{SS}} \leftarrow_r \text{PPGen}^{\text{SS}}(1^\lambda)$ 
   $(\text{sk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{sig}}^{\text{SS}}) \leftarrow_r \text{KG}_{\text{sig}}^{\text{SS}}(\text{pp}_{\text{SS}})$ 
   $(m^*, \sigma^{\text{SS}*}, \text{pk}_{\text{san}}^{\text{SS}*}) \leftarrow_r \mathcal{A}^{\text{Sign}^{\text{SS}}(\text{sk}_{\text{sig}}^{\text{SS}}, \dots), \text{Proof}^{\text{SS}}(\text{sk}_{\text{sig}}^{\text{SS}}, \dots)}(\text{pk}_{\text{sig}}^{\text{SS}})$ 
   $\forall i \in [1..q]$ , let  $(\text{pk}_{\text{san}, i}^{\text{SS}}, m_i, \mathbb{A}_i^{\text{SS}})$  and  $\sigma_i^{\text{SS}}$ 
  index the queries/answers to/from  $\text{Sign}^{\text{SS}}$ 
   $\pi^{\text{SS}} \leftarrow_r \text{Proof}^{\text{SS}}(\text{sk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}*}, m^*, \sigma^{\text{SS}*}, \{(m_i, \sigma_i^{\text{SS}}) \mid 0 < i \leq q\})$ 
  return 1, if  $\text{Verify}^{\text{SS}}(\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}*}, m^*, \sigma^{\text{SS}*}) = 1 \wedge$ 
   $\forall i \in [1..q] : (\text{pk}_{\text{san}}^{\text{SS}*}, m^*, \sigma^{\text{SS}*}) \neq (\text{pk}_{\text{san}, i}^{\text{SS}}, m_i, \sigma_i^{\text{SS}}) \wedge$ 
   $\text{Judge}^{\text{SS}}(\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}*}, m^*, \sigma^{\text{SS}*}, \pi^{\text{SS}}) = \text{Sig}^{\text{SS}}$ 
  return 0

```

Fig. 13: SS Sanitizer-Accountability

points to the signer, but  $(m^*, \sigma^{\text{SS}*})$  has never been generated by the signer. This is even true, if the adversary can generate the sanitizer’s key pair.

**Definition 15 (Sanitizer-Accountability).** *An SS is sanitizer-accountable, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that:*

$$\Pr[\mathbf{Exp}_{\mathcal{A}, \text{SS}}^{\text{SanAccountability}}(\lambda) = 1] \leq \nu(\lambda),$$

where the corresponding experiment is defined in Figure 13.

*Invisibility.* Depending on the context, an additional privacy guarantee may be required. In particular, invisibility prohibits that an outsider holding no secret keys can decide which parts of a message  $m$  are sanitizable. Note, the signing oracle can be simulated using the LoRADM oracle and setting  $\mathbb{A}_0^{\text{SS}} = \mathbb{A}_1^{\text{SS}}$ . The notation  $\mathbb{A}_0^{\text{SS}} \cap \mathbb{A}_1^{\text{SS}}$  means that only those indices are admissible which are admissible in  $\mathbb{A}_0^{\text{SS}}$  and  $\mathbb{A}_1^{\text{SS}}$ .

We stress, however, that we introduce a slightly stronger variant than discussed in prior work [43]; Namely, we allow that the adversary can query *any* signature to  $\text{Sanit}^{\text{SS}'}$ , and not only the ones generated honestly. In particular, now only if the signature was created by one of the oracles, we enforce the restriction  $\mathbb{A}_0^{\text{SS}} \cap \mathbb{A}_1^{\text{SS}}$ . We stress, however, that all strongly invisible constructions proposed so far are also unforgeable, and thus such a signature can never be generated by the adversary.

**Definition 16 (Invisibility).** *An SS is invisible, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that:*

$$\left| \Pr[\mathbf{Exp}_{\mathcal{A}, \text{SS}}^{\text{Invisibility}}(\lambda)] - 1/2 \right| \leq \nu(\lambda),$$

where the corresponding experiment is defined in Figure 14.

### 3 Invisible Redactable Signatures

We now introduce the new notion of *invisible* redactable signatures. In contrast to standard RSs [28], such schemes hide which parts of a message can be redacted from outsiders. As in standard RSs anyone can redact — and one can thus trivially decide which blocks are redactable — we need to introduce a “designated redactor”, which is the only party able to decide this question. Thus, the designated redactor is the sanitizer in an RS. This is related to already existing definitions, but in a different context [54,20]. However, before we start defining and constructing our invisible RS, we need to settle some additional notation.

```

Exp $\mathcal{A}, SS$ Invisibility( $\lambda$ )
   $pp_{SS} \leftarrow_r \text{PPGen}^{SS}(1^\lambda)$ 
   $(sk_{sig}^{SS}, pk_{sig}^{SS}) \leftarrow_r \text{KG}_{sig}^{SS}(pp_{SS})$ 
   $(sk_{san}^{SS}, pk_{san}^{SS}) \leftarrow_r \text{KG}_{san}^{SS}(pp_{SS})$ 
   $b \leftarrow_r \{0, 1\}$ 
   $Q \leftarrow \emptyset$ 
   $a \leftarrow_r \mathcal{A}^{\text{Sanit}^{SS'}(sk_{san}^{SS}, \dots), \text{Proof}^{SS'}(sk_{sig}^{SS}, \dots), \text{LoRADM}(sk_{sig}^{SS}, \dots, b)}(pk_{sig}^{SS}, pk_{san}^{SS})$ 
  where  $\text{LoRADM}(sk_{sig}^{SS}, pk_{san}^{SS'}, m, \mathbb{A}_0^{SS}, \mathbb{A}_1^{SS}, b)$ :
    return  $\perp$ , if  $\mathbb{A}_0^{SS} \not\prec m \wedge \mathbb{A}_1^{SS} \not\prec m$ 
    return  $\perp$ , if  $pk_{san}^{SS} \neq pk_{san}^{SS'} \wedge \mathbb{A}_0^{SS} \neq \mathbb{A}_1^{SS}$ 
     $\sigma^{SS} \leftarrow_r \text{Sign}^{SS}(sk_{sig}^{SS}, pk_{san}^{SS'}, m, \mathbb{A}_b^{SS})$ 
    if  $pk_{san}^{SS'} = pk_{san}^{SS}$ , let  $Q \leftarrow Q \cup \{(m, \sigma^{SS}, \mathbb{A}_0^{SS} \cap \mathbb{A}_1^{SS})\}$ 
    return  $\sigma^{SS}$ 
  where  $\text{Sanit}^{SS'}(pk_{sig}^{SS'}, sk_{san}^{SS}, m, \mathbb{M}^{SS}, \sigma^{SS})$ :
    return  $\perp$ , if  $pk_{sig}^{SS'} = pk_{sig}^{SS} \wedge \exists (m, \sigma^{SS}, \mathbb{A}^{SS}) \in Q : \mathbb{M}^{SS} \not\prec (m, \mathbb{A}^{SS})$ 
     $(m', \sigma^{SS'}) \leftarrow_r \text{Sanit}^{SS'}(pk_{sig}^{SS'}, sk_{san}^{SS}, m, \mathbb{M}^{SS}, \sigma^{SS})$ 
    if  $pk_{sig}^{SS'} = pk_{sig}^{SS} \wedge \exists (m, \sigma^{SS}, \mathbb{A}^{SS'}) \in Q : \mathbb{M}^{SS} \prec (m, \mathbb{A}^{SS'})$ ,
       $Q \leftarrow Q \cup \{(m', \sigma^{SS'}, \mathbb{A}^{SS'})\}$ 
    return  $(m', \sigma^{SS'})$ 
  return 1, if  $a = b$ 
  return 0

```

Fig. 14: SS Invisibility

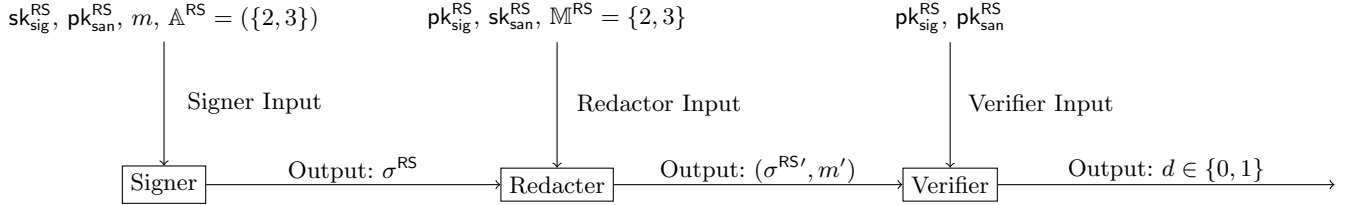


Fig. 15: Example workflow of a designated redactor RS. The message  $m$  is set to (I, do, not, like, crypto). After redacting,  $m'$  is (I, like, crypto). As we only consider private RSs, the redacted parts are not visible.

*Additional Notation.* In the following, let  $m = (m^1, m^2, \dots, m^\ell)$  be some message, while  $\mathbb{A}^{RS} \subseteq [1..\ell]$  denotes the admissible redactions, i.e., if  $i \in \mathbb{A}^{RS}$ , then  $m^i$  can be redacted by the designated redactor. The variable  $\mathbb{M}^{RS} \subseteq [1..\ell]$  denotes how a message  $m$  is to be modified, i.e., each block  $m^i$ ,  $i \in \mathbb{M}^{RS}$ , is removed from  $m$  to form the redacted message  $m'$ . In comparison to Derler et al. [28], however, we already define how those data-structures look like for preciseness. Additionally, as done for SSs, we use the shorthand notation  $m' \leftarrow \mathbb{M}^{RS}(m)$  to denote such a redaction. The notation  $\mathbb{M}^{RS} \prec (m, \mathbb{A}^{RS})$  means that  $\mathbb{M}^{RS}$  is a valid modification instruction w.r.t.  $m$  and  $\mathbb{A}^{RS}$ . Likewise, we use  $\mathbb{A}^{RS} \prec m$  to denote that  $\mathbb{A}^{RS}$  is valid description of the admissible blocks w.r.t.  $m$ .

An example workflow is depicted in Figure 15, derived from the work done by Bilzhaue et al. [2].

### 3.1 Framework

The following definitions for RSs are derived from Derler et al. [28], merged with the ideas given by Pöhls and Samelin [20], while also supporting parameter generation. Moreover, we do not consider additional “redaction information”  $\text{RED}^{RS}$ , as given by Derler et al. [28], as we have a designated redactor anyway.

**Definition 17 (Invisible Redactable Signatures).** An invisible redactable signature RS consists of the following six algorithms, i.e.,  $\{\text{PPGen}^{\text{RS}}, \text{KG}_{\text{sig}}^{\text{RS}}, \text{KG}_{\text{san}}^{\text{RS}}, \text{Sign}^{\text{RS}}, \text{Verify}^{\text{RS}}, \text{Red}^{\text{RS}}\}$ , such that:

$\text{PPGen}^{\text{RS}}$ . The algorithm  $\text{PPGen}^{\text{RS}}$  generates the public parameters:

$$\text{pp}_{\text{RS}} \leftarrow_r \text{PPGen}^{\text{RS}}(1^\lambda)$$

We assume that  $\text{pp}_{\text{RS}}$  is implicitly input to all other algorithms.

$\text{KG}_{\text{sig}}^{\text{RS}}$ . The algorithm  $\text{KG}_{\text{sig}}^{\text{RS}}$  generates a key pair:

$$(\text{sk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{sig}}^{\text{RS}}) \leftarrow_r \text{KG}_{\text{sig}}^{\text{RS}}(\text{pp}_{\text{RS}})$$

$\text{KG}_{\text{san}}^{\text{RS}}$ . The algorithm  $\text{KG}_{\text{san}}^{\text{RS}}$  generates a key pair:

$$(\text{sk}_{\text{san}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}) \leftarrow_r \text{KG}_{\text{san}}^{\text{RS}}(\text{pp}_{\text{RS}})$$

$\text{Sign}^{\text{RS}}$ . The algorithm  $\text{Sign}^{\text{RS}}$  outputs a signature  $\sigma^{\text{RS}}$  and some redaction information  $\text{RED}^{\text{RS}}$  on input of  $\text{sk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \mathbb{A}^{\text{RS}}$  and a message  $m$ :

$$\sigma^{\text{RS}} \leftarrow_r \text{Sign}^{\text{RS}}(\text{sk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, m, \mathbb{A}^{\text{RS}})$$

Note, it is assumed that  $\mathbb{A}^{\text{RS}}$  can always be derived.

$\text{Verify}^{\text{RS}}$ . The deterministic algorithm  $\text{Verify}^{\text{RS}}$  verifies a signature  $\sigma^{\text{RS}}$ , i.e., outputs a decision  $d \in \{0, 1\}$  w.r.t.  $\text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}$ , and a message  $m$ :

$$d \leftarrow \text{Verify}^{\text{RS}}(\text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, m, \sigma^{\text{RS}})$$

$\text{Red}^{\text{RS}}$ . The algorithm  $\text{Red}^{\text{RS}}$  outputs a derived signature  $\sigma^{\text{RS}'}$  and a derived message  $m'$ , along with the new admissible blocks  $\mathbb{A}^{\text{RS}'}$ , on input of  $\text{sk}_{\text{san}}^{\text{RS}}, \text{pk}_{\text{sig}}^{\text{RS}}$ , a signature  $\sigma^{\text{RS}}$ , and some modification instruction  $\mathbb{M}^{\text{RS}}$ :

$$(\sigma^{\text{RS}'}, m', \mathbb{A}^{\text{RS}'}) \leftarrow_r \text{Red}^{\text{RS}}(\text{sk}_{\text{san}}^{\text{RS}}, \text{pk}_{\text{sig}}^{\text{RS}}, m, \sigma^{\text{RS}}, \mathbb{M}^{\text{RS}})$$

*Correctness.* We also require that an RS is correct. We call an RS correct, if for all  $\lambda \in \mathbb{N}$ , for all  $\text{pp}_{\text{RS}} \leftarrow_r \text{PPGen}^{\text{RS}}(1^\lambda)$ , for all  $(\text{sk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{sig}}^{\text{RS}})$ , for all  $(\text{sk}_{\text{san}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}})$ , for all  $m$ , for all  $\mathbb{A}^{\text{RS}} \in \{\mathbb{A}^{\text{RS}'} \mid \mathbb{A}^{\text{RS}'} \prec m\}$ , for all  $\sigma^{\text{RS}} \leftarrow_r \text{Sign}^{\text{RS}}(\text{sk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, m, \mathbb{A}^{\text{RS}})$  we have that  $\text{Verify}^{\text{RS}}(\text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, m, \sigma^{\text{RS}}) = 1$  and also for all  $\mathbb{M}^{\text{RS}} \in \{\mathbb{M}^{\text{RS}'} \mid \mathbb{M}^{\text{RS}'} \prec (m, \mathbb{A}^{\text{RS}})\}$ , for all  $(\sigma^{\text{RS}'}, m', \mathbb{A}^{\text{RS}'}) \leftarrow_r \text{Red}^{\text{RS}}(\text{sk}_{\text{san}}^{\text{RS}}, \text{pk}_{\text{sig}}^{\text{RS}}, m, \sigma^{\text{RS}}, \mathbb{M}^{\text{RS}})$  we have that  $\text{Verify}^{\text{RS}}(\text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, m', \sigma^{\text{RS}'}) = 1$ .

### 3.2 Security Requirements

We now introduce our security model for RSs. This is an extended version derived from Derler et al. [28], which is, in turn, derived from Brzuska et al. [17]. Note, moreover, that we do not need accountability, as in our construction accountability is given by the SS, much like Pöhls and Samelin [20] and Bilzhouse et al. [55].

*Unforgeability.* This definition requires that an adversary  $\mathcal{A}$  cannot derive a message which is not derivable from any signed messages. We stress that, even though the set  $\bigcup_{i=1}^q \{\mathbb{M}^{\text{RS}}(m_i) \mid \mathbb{M}^{\text{RS}} \prec (m_i, \mathbb{A}_i^{\text{RS}})\}$  may grow exponentially, membership is trivially to decide, i.e., in polynomial time.

**Definition 18 (Unforgeability).** An RS is unforgeable, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that:

$$\Pr[\text{Exp}_{\mathcal{A}, \text{RS}}^{\text{Unforgeability}}(\lambda) = 1] \leq \nu(\lambda)$$

where the corresponding experiment is defined in Figure 16.

**Exp**<sub>A,RS</sub><sup>Unforgeability</sup>( $\lambda$ )

$\text{pp}_{\text{RS}} \leftarrow_r \text{PPGen}^{\text{RS}}(1^\lambda)$

$(\text{sk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{sig}}^{\text{RS}}) \leftarrow_r \text{KG}_{\text{sig}}^{\text{RS}}(\text{pp}_{\text{RS}})$

$(\text{sk}_{\text{san}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}) \leftarrow_r \text{KG}_{\text{san}}^{\text{RS}}(\text{pp}_{\text{RS}})$

$\mathcal{Q} \leftarrow \emptyset$

$(m^*, \sigma^{\text{RS}*}) \leftarrow_r \mathcal{A}^{\text{Sign}^{\text{RS}'(\text{sk}_{\text{sig}}^{\text{RS}}, \cdot, \cdot)}, \text{Red}^{\text{RS}'(\cdot, \text{sk}_{\text{san}}^{\text{RS}}, \cdot, \cdot)}}(\text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}})$

where  $\text{Sign}^{\text{RS}'(\text{sk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}'}, m, \mathbb{A}^{\text{RS}})$ :

$\sigma^{\text{RS}} \leftarrow_r \text{Sign}^{\text{RS}}(\text{sk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}'}, m, \mathbb{A}^{\text{RS}})$

if  $\text{pk}_{\text{san}}^{\text{RS}'} = \text{pk}_{\text{san}}^{\text{RS}}$ , let  $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(\sigma^{\text{RS}}, m)\}$

return  $\sigma^{\text{RS}}$

and  $\text{Red}^{\text{RS}'(\text{pk}_{\text{sig}}^{\text{RS}'}, \text{sk}_{\text{san}}^{\text{RS}}, m, \sigma^{\text{RS}}, \mathbb{M}^{\text{RS}})$ :

$(\sigma^{\text{RS}'}, m') \leftarrow_r \text{Red}^{\text{RS}}(\text{sk}_{\text{san}}^{\text{RS}}, \text{pk}_{\text{sig}}^{\text{RS}'}, m, \sigma^{\text{RS}}, \mathbb{M}^{\text{RS}})$

if  $\text{pk}_{\text{sig}}^{\text{RS}'} = \text{pk}_{\text{sig}}^{\text{RS}} \wedge \sigma^{\text{RS}'} \neq \perp$ , let  $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(\sigma^{\text{RS}'}, m')\}$

return  $\sigma^{\text{RS}'}$

return 1, if  $\text{Verify}^{\text{RS}}(\text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, m^*, \sigma^{\text{RS}*}) = 1 \wedge (\sigma^{\text{RS}*}, m^*) \notin \mathcal{Q}$

return 0

Fig. 16: RS Unforgeability

**Exp**<sub>A,RS</sub><sup>Immutability</sup>( $\lambda$ )

$\text{pp}_{\text{RS}} \leftarrow_r \text{PPGen}^{\text{RS}}(1^\lambda)$

$\mathcal{Q} \leftarrow \emptyset$

$(\text{sk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{sig}}^{\text{RS}}) \leftarrow_r \text{KG}_{\text{sig}}^{\text{RS}}(\text{pp}_{\text{RS}})$

$(m^*, \sigma^{\text{RS}*}, \text{pk}^*) \leftarrow_r \mathcal{A}^{\text{Sign}^{\text{RS}}(\text{sk}_{\text{sig}}^{\text{RS}}, \cdot, \cdot)}(\text{pk}_{\text{sig}}^{\text{RS}})$

where  $\text{Sign}^{\text{RS}}(\text{pk}_{\text{san}}^{\text{RS}}, m, \mathbb{A}^{\text{RS}})$

$\mathcal{Q}[\text{pk}_{\text{san}}^{\text{RS}}] \leftarrow \mathcal{Q}[\text{pk}_{\text{san}}^{\text{RS}}] \cup \{\mathbb{M}^{\text{RS}}(m) \mid \mathbb{M}^{\text{RS}} \prec (m, \mathbb{A}^{\text{RS}})\}$

return 1, if  $\text{Verify}^{\text{RS}}(\text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}^*, m^*, \sigma^{\text{RS}*}) = 1 \wedge m^* \notin \mathcal{Q}[\text{pk}^*]$

return 0

Fig. 17: RS Immutability

*Immutability.* This definition requires that an adversary  $\mathcal{A}$ , which even can generate  $\text{sk}_{\text{san}}^{\text{RS}}$ , cannot derive a message which is not derivable from any signed messages. We stress that, even though the set  $\bigcup_{i=1}^q \{\mathbb{M}^{\text{RS}}(m_i) \mid \mathbb{M}^{\text{RS}} \prec (m_i, \mathbb{A}_i^{\text{RS}})\}$  may grow exponentially, membership is trivially to decide, i.e., in polynomial time.

**Definition 19 (Immutability).** An RS is immutable, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that:

$$\Pr[\mathbf{Exp}_{\mathcal{A},\text{RS}}^{\text{Immutability}}(\lambda) = 1] \leq \nu(\lambda)$$

where the corresponding experiment is defined in Figure 17.

*Privacy.* This definition prohibits that an adversary  $\mathcal{A}$  can learn anything about redacted parts. This is similar to the definition for SSs.

**Definition 20 (Privacy).** An RS is private, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that:

$$\left| \Pr[\mathbf{Exp}_{\mathcal{A},\text{RS}}^{\text{Privacy}}(\lambda)] - 1/2 \right| \leq \nu(\lambda)$$

where the corresponding experiment is defined in Figure 18.







generating a signature on the overall tag and all public keys using the long-term key. Then, the signer has to generate a signature on each “left-of”-relation (See Figure 21 for a graphical explanation) of each block using the overall tag and the two block-specific tags, again using the long-term keys. The ephemeral secret key and each signature which should be redactable are then encrypted for the designated redactor. To avoid leaking how many blocks are redactable, the plaintext is padded accordingly. Namely, at most  $\ell + \frac{\ell(\ell-1)}{2}$  signatures are given to the adversary. Moreover, to prohibit tampering with the ciphertext, it is signed along with the overall tag and all public keys using the ephemeral secret signing key. Finally, all generated signatures are aggregated. Verification simply checks the aggregate signature on the values received. For redaction, the designated redactor decrypts the ciphertext, re-generates the signature on the ciphertext and de-aggregates it. For every block to be redacted, the designated redactor simply de-aggregates all related signatures. Finally, it generates a new ciphertext with the remaining signatures, once more padded accordingly, signs it and aggregates it.

<p><math>\text{PPGen}^{\text{RS}}(1^\lambda)</math>. Generate <math>\text{pp}_\Pi \leftarrow_r \text{PPGen}^\Pi(1^\lambda)</math> and <math>\text{pp}_\Sigma \leftarrow_r \text{PPGen}^\Sigma(1^\lambda)</math>. Return <math>(\text{pp}_\Pi, \text{pp}_\Sigma)</math>.</p> <p><math>\text{KGen}_{\text{sig}}^{\text{RS}}(\text{pp}_{\text{RS}})</math>. Generate <math>(\text{sk}_\Sigma, \text{pk}_\Sigma) \leftarrow_r \text{KGen}_\Sigma(\text{pp}_\Sigma)</math>. Return <math>(\text{sk}_\Sigma, \text{pk}_\Sigma)</math>.</p> <p><math>\text{KGen}_{\text{san}}^{\text{RS}}(\text{pp}_{\text{RS}})</math>. Generate <math>(\text{sk}_\Pi, \text{pk}_\Pi) \leftarrow_r \text{KGen}^\Pi(\text{pp}_\Pi)</math>. Return <math>(\text{sk}_\Pi, \text{pk}_\Pi)</math>.</p> <p><math>\text{Sign}^{\text{RS}}(\text{sk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, m, \mathbb{A}^{\text{RS}})</math>. The algorithm proceeds as follows, where <math>m = (m^1, \dots, m^\ell)</math>:</p> <ul style="list-style-type: none"> <li>– Draw <math>\ell + 1</math> tags <math>\tau_i \leftarrow_r \{0, 1\}^\lambda</math> and generate a new signature key pair <math>(\text{sk}_{\Sigma'}, \text{pk}_{\Sigma'}) \leftarrow_r \text{KGen}_\Sigma(\text{pp}_\Sigma)</math>.</li> <li>– For <math>i \in [1.. \ell]</math>, let <math>m'^i</math> be the augmented block <math>(m^i, \tau_0, \tau_i, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'})</math>.</li> <li>– Sign <math>\tau_0</math> and all public keys, i.e., let <math>\sigma_0 \leftarrow \text{Sign}_\Sigma(\text{sk}_\Sigma, (\tau_0, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'}))</math>, and sign each augmented block, i.e., let <math>\sigma_i \leftarrow \text{Sign}_\Sigma(\text{sk}_{\Sigma'}, m'^i)</math>. Finally, sign each “left-of” relation, i.e., for all <math>i \in [1.. \ell]</math>, for all <math>j &lt; i</math> let <math>\sigma_{i,j} \leftarrow \text{Sign}_\Sigma(\text{sk}_{\Sigma'}, (m'^i, \tau_j))</math>.</li> <li>– Encrypt the ephemeral secret key, all signatures related to the redactable blocks and some fake values to make the length of the encryption constant. In particular, generate <math>c \leftarrow_r \text{Enc}^\Pi(\text{pk}_\Pi, (\text{sk}_{\Sigma'}, \{\sigma_i\}_{i \in \mathbb{A}^{\text{RS}}}, \{\sigma_{i,j}\}_{i \in \mathbb{A}^{\text{RS}}, j \in [1.. \ell]}, t), (\text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'}))</math>, where <math>t</math> is a random string of length <math>\ell + \frac{\ell(\ell-1)}{2} -  (\{\sigma_i\}_{i \in \mathbb{A}^{\text{RS}}} \cup \{\sigma_{i,j}\}_{i \in \mathbb{A}^{\text{RS}}, j \in [1.. \ell]}) </math> times the size of a single signature. This essentially makes the ciphertext always the same size, regardless of <math>\mathbb{A}^{\text{RS}}</math>.</li> <li>– Sign <math>c</math> using the ephemeral signature key, i.e., let <math>\sigma_c \leftarrow \text{Sign}_\Sigma(\text{sk}_{\Sigma'}, (c, \tau_0, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'}))</math>.</li> <li>– Aggregate all signatures generated, i.e., let <math>\sigma_a \leftarrow \text{Agg}_\Sigma(\mathcal{S}_1 \cup \mathcal{S}_2 \cup \mathcal{S}_3 \cup \mathcal{S}_4)</math>, where <math>\mathcal{S}_1 = \{((\text{pk}_\Sigma, (\tau_0, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'})), \sigma_0)\}</math>, <math>\mathcal{S}_2 = \{((\text{pk}_{\Sigma'}, (c, \tau_0, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'})), \sigma_c)\}</math>, <math>\mathcal{S}_3 = \{((\text{pk}_\Sigma, m'^i), \sigma_i)\}_{i \in [1.. \ell]}</math>, <math>\mathcal{S}_4 = \{((\text{pk}_\Sigma, (m'^i, \tau_j)), \sigma_{i,j})\}_{i \in [1.. \ell], j \in [1.. \ell], i &lt; j}</math>.</li> <li>– Return <math>(\sigma_a, c, \text{pk}_{\Sigma'}, (\tau_i)_{i \in [0.. \ell]})</math>.</li> </ul> <p><math>\text{Verify}^{\text{RS}}(\text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, m, \sigma^{\text{RS}})</math>. Parse <math>\sigma^{\text{RS}}</math> as <math>(\sigma_a, c, \text{pk}_{\Sigma'}, (\tau_i)_{i \in [0.. \ell]})</math>. Let <math>m'^i</math> be the augmented block <math>(m^i, \tau_0, \tau_i, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'})</math>. Return <math>\text{AVerf}_\Sigma(\mathcal{S}_1 \cup \mathcal{S}_2 \cup \mathcal{S}_3 \cup \mathcal{S}_4, \sigma_a)</math>, where <math>\mathcal{S}_1 = \{(\text{pk}_\Sigma, (\tau_0, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'}))\}</math>, <math>\mathcal{S}_2 = \{(\text{pk}_{\Sigma'}, (c, \tau_0, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'}))\}</math>, <math>\mathcal{S}_3 = \{(\text{pk}_\Sigma, m'^i)\}_{i \in [1.. \ell]}</math>, <math>\mathcal{S}_4 = \{(\text{pk}_\Sigma, (m'^i, \tau_j))\}_{i \in [1.. \ell], j \in [1.. \ell], i &lt; j}</math>.</p> <p><math>\text{Red}^{\text{RS}}(\text{sk}_{\text{san}}^{\text{RS}}, \text{pk}_{\text{sig}}^{\text{RS}}, m, \sigma^{\text{RS}}, \mathbb{M}^{\text{RS}})</math>. Parse <math>\sigma^{\text{RS}}</math> as <math>(\sigma_a, c, \text{pk}_{\Sigma'}, (\tau_i)_{i \in [0.. \ell_m]})</math> and proceed as follows:</p> <ul style="list-style-type: none"> <li>– If <math>\text{Verify}^{\text{RS}}(\text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, m, \sigma^{\text{RS}}) \neq 1</math>, return false. Let <math>m'' \leftarrow \mathbb{M}^{\text{RS}}(m)</math>.</li> <li>– Let <math>(\text{sk}_{\Sigma'}, \{\sigma_i\}_{i \in \mathbb{A}^{\text{RS}}}, \{\sigma_{i,j}\}_{i \in \mathbb{A}^{\text{RS}}, j \in [1.. \ell_m]}) \leftarrow \text{Dec}^\Pi(\text{sk}_{\text{san}}^{\text{RS}}, c, (\text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'}))</math>.</li> <li>– For each <math>i</math>, let <math>m'^i</math> be the augmented block <math>(m^i, \tau_0, \tau_i, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'})</math>. Likewise, for each <math>i</math>, let <math>m''^i</math> be the augmented block <math>(m''^i, \tau_0, \tau_i, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'})</math>.</li> <li>– Let <math>\sigma_c \leftarrow \text{Sign}_\Sigma(\text{sk}_{\Sigma'}, (c, \tau_0, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'}))</math>.</li> <li>– Update <math>\mathbb{A}^{\text{RS}}</math> to <math>\mathbb{A}^{\text{RS}'}</math> by removing all indices in <math>\mathbb{M}_2^{\text{RS}}</math> and adjusting the remaining indices by reducing each <math>i</math> in <math>\mathbb{A}^{\text{RS}}</math> by <math> \{j \in \mathbb{M}^{\text{RS}} : j &lt; i\} </math>.</li> <li>– De-aggregate the signatures for the cipher and the messages (and relations) to be removed, i.e., compute <math>\sigma'_a \leftarrow \text{DAgg}_\Sigma((\mathcal{S}_1 \cup \mathcal{S}_2 \cup \mathcal{S}_3), (\mathcal{S}_4 \cup \mathcal{S}_5 \cup \mathcal{S}_6 \cup \mathcal{S}_7), \sigma_a)</math>, where <math>\mathcal{S}_1 = \bigcup\{(\{(\text{pk}_{\text{sig}}^{\text{RS}}, m'_i)\}, \sigma_i)\}_{i \in \mathbb{M}^{\text{RS}}}</math>, <math>\mathcal{S}_2 = \{(\{(\text{pk}_{\Sigma'}, (c, \tau_0, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'})), \sigma_c)\})\}</math>, <math>\mathcal{S}_3 = \bigcup\{(\{(\text{pk}_\Sigma, (m'^i, \tau_j)), \sigma_{i,j}\})\}_{i \in \mathbb{M}^{\text{RS}}, j \in \mathbb{M}^{\text{RS}}}</math>, <math>\mathcal{S}_4 = \{(\text{pk}_\Sigma, (\tau_0, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'}))\}</math>, <math>\mathcal{S}_5 = \{(\{(\text{pk}_{\Sigma'}, (c, \tau_0, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'}))\})\}</math>, <math>\mathcal{S}_6 = \bigcup\{(\{(\text{pk}_\Sigma, m'^i)\})\}_{i \in [1.. \ell_m]}</math>, <math>\mathcal{S}_7 = \bigcup\{(\{(\text{pk}_\Sigma, (m'^i, \tau_j))\})\}_{i \in [1.. \ell_m], j \in [1.. \ell_m], i &lt; j}</math>.</li> <li>– Generate <math>c' \leftarrow_r \text{Enc}^\Pi(\text{pk}_{\text{san}}^{\text{RS}}, (\text{sk}_{\Sigma'}, \{\sigma_i\}_{i \in \mathbb{A}^{\text{RS}'}, \sigma_{i,j}\}_{i \in \mathbb{A}^{\text{RS}'}, j \in [1.. \ell_m]}), (\text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'}))</math>, where <math>t'</math> is a random string of length <math>\ell_{m''} + \frac{\ell_{m''}(\ell_{m''}-1)}{2} -  (\{\sigma_i\}_{i \in \mathbb{A}^{\text{RS}'}} \cup \{\sigma_{i,j}\}_{i \in \mathbb{A}^{\text{RS}'}, j \in [1.. \ell_m]}) </math> times the size of a single signature.</li> <li>– Sign <math>c'</math>, i.e., let <math>\sigma'_c \leftarrow \text{Sign}_\Sigma(\text{sk}_{\Sigma'}, (c', \tau_0, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'}))</math>.</li> <li>– Aggregate <math>\sigma'_c</math> onto <math>\sigma'_a</math> by calculating <math>\sigma''_a \leftarrow \text{Agg}_\Sigma(\mathcal{S}_1 \cup \{(\mathcal{S}_2 \cup \mathcal{S}_3), \sigma'_a\})</math>, where <math>m'' = \mathbb{M}^{\text{RS}}(m)</math> is of length <math>\ell_{m''}</math>, <math>\mathcal{S}_1 = \{(\{(\text{pk}_{\Sigma'}, (c, \tau_0, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'})), \sigma'_c\})\}</math>, <math>\mathcal{S}_2 = \bigcup\{(\{(\text{pk}_{\text{sig}}^{\text{RS}}, m''^i)\})\}_{i \in [1.. \ell_{m''}]}</math>, and <math>\mathcal{S}_3 = \bigcup\{(\{(\text{pk}_{\text{sig}}^{\text{RS}}, (m''^i, \tau_j))\})\}_{i \in [1.. \ell_{m''}], j \in [1.. \ell_{m''}], i &lt; j}</math>.</li> <li>– Return <math>((\sigma''_a, c', \text{pk}_{\Sigma'}, (\tau_i)_{i \in [0.. \ell_{m''}]})</math>, <math>m''</math>, <math>\mathbb{A}^{\text{RS}'})</math>.</li> </ul>
---

Construction 1: Construction of an invisible RS

The proof of the following theorem is given in Appendix A.

**Theorem 1.** *If  $\Pi$  is correct and IND-CCA2 secure, while  $\Sigma$  is correct, unforgeable, unique and non-extractable, then the construction of an RS, given in Construction 1, is correct, unforgeable, immutable, private, transparent, and invisible.*

## 4 Non-Accountable Invisible SS

In our construction of fully invisible PS (See Section 5), we use non-accountable, yet invisible *and transparent*, SS. As accountability is one of the main concerns of SSs [3], this notion has, for obvious reasons, not been considered before [4]. However, as we show, in certain contexts such a notion has its merits besides proving implications [45].

*Background.* In a nutshell, a non-accountable invisible SS, from now on denoted by  $SS'$ , behaves as a standard SS, but the algorithms  $\text{Proof}^{SS}$  and  $\text{Judge}^{SS}$  are simply set to  $\perp$ , i.e., effectively all algorithms related to accountability are dropped, clearly also affecting the correctness definition [29]. This also means that an  $SS'$  may still achieve all security notions, but sanitizer-accountability and signer-accountability. This is exactly what our construction, given in Construction 2, achieves.

The reason for doing so is that accountability of the signatures can result from a different mechanism, what is exactly what we do in our final construction using an accountable SS.

*Construction.* Our construction is given in Construction 2. The basic idea is that each block is signed using a fresh ephemeral signature key. If a block is admissible, the corresponding secret key is given to the sanitizer. This paradigm follows already existing ideas [45]. However, their scheme does not achieve transparency, while ours is not accountable.

In more detail, the signer holds a long-term key-pair for a  $\Sigma$ , while the sanitizer holds a key-pair for  $\Pi$ . At signing, the signer generates a fresh ephemeral key-pair for each block in the message  $m$  to sign. If a block is admissible, the randomness used to generate those key-pairs is derived from a PRF to achieve a smaller signature size. If a block is not admissible, fresh random coins are drawn. The secret key  $x$  for the PRF is encrypted to the sanitizer. All ephemeral public keys and the resulting ciphertext are signed using the long-term keys. For sanitizing, the sanitizer reconstructs the secret key  $x$  for the PRF and with it the signing keys for each admissible block, and then simply signs the blocks to be sanitized.

It is easy to see that this construction is invisible and does not provide any form of accountability, while we stress that we cannot avoid the encryption due to recent results [45]. Moreover, strictly speaking, our construction is even transparent in the sense of Brzuska et al. [29], i.e., the proof-restriction is not necessary.

We stress that we could also aggregate all signatures generated. However, to keep the description short, we opted for not doing this.

The proof of the following theorem is found in Appendix B.

**Theorem 2.** *If  $\Pi$  is correct and IND-CCA2 secure, PRF pseudorandom and correct, while  $\Sigma$  is correct and unforgeable, the construction of a  $SS'$ , given in Construction 2, is correct, unforgeable, immutable, private, transparent, and invisible.*

## 5 Fully Invisible Protean Signatures

We now present our framework for PSs, which is taken from Krenn et al. [4].

To recap, a PS allows to remove and alter signer-chosen parts of a signed message by a semi-trusted third party, i.e., the sanitizer. The sanitizer can also be held accountable, if it chose to edit a signed message. For the framework, we settle some additional notation, which is derived from the ones used for RSs and SSs, to ease understanding.

$\text{PPGen}^{\text{SS}}(1^\lambda)$ . Generate  $\text{pp}_\Pi \leftarrow_r \text{PPGen}^\Pi(1^\lambda)$  and  $\text{pp}_\Sigma \leftarrow_r \text{PPGen}^\Sigma(1^\lambda)$ . Return  $(\text{pp}_\Pi, \text{pp}_\Sigma)$ .  
 $\text{KG}_{\text{sig}}^{\text{SS}}(\text{pp}_{\text{SS}})$ . Generate  $(\text{sk}_\Sigma, \text{pk}_\Sigma) \leftarrow_r \text{KGen}_\Sigma(\text{pp}_\Sigma)$ . Return  $(\text{sk}_\Sigma, \text{pk}_\Sigma)$ .  
 $\text{KG}_{\text{san}}^{\text{SS}}(\text{pp}_{\text{SS}})$ . Generate  $(\text{sk}_\Pi, \text{pk}_\Pi) \leftarrow_r \text{KGen}^\Pi(\text{pp}_\Pi)$ . Return  $(\text{sk}_\Pi, \text{pk}_\Pi)$ .  
 $\text{Sign}^{\text{SS}}(\text{sk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}, m, \mathbb{A}^{\text{SS}})$ . The algorithm proceeds as follows, where  $m = (m^1, \dots, m^\ell)$ :

- If  $\mathbb{A}^{\text{SS}} \prec m$  does not hold, return  $\perp$ .
- Draw  $x \leftarrow_r \text{KGen}_{\text{PRF}}(1^\lambda)$ .
- For  $i \in [1..\ell]$ , if  $i \in \mathbb{A}^{\text{SS}}.1$ , let  $r_i \leftarrow \text{Eval}_{\text{PRF}}(x, i)$  and  $r_i \leftarrow_r \{0, 1\}^\lambda$  otherwise. Set  $(\text{sk}_i, \text{pk}_i) \leftarrow_r \text{KGen}_\Sigma(\text{pp}_\Sigma; r_i)$ .
- Encrypt the seed and  $\mathbb{A}^{\text{SS}}$ , i.e., let  $c \leftarrow_r \text{Enc}^\Pi(\text{pk}_{\text{san}}^{\text{SS}}, (x, \mathbb{A}^{\text{SS}}), (\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}, (\text{pk}_i)_{[1..\ell]}))$ .
- Sign the public keys, and the ciphertext, i.e., let  $\sigma_s \leftarrow \text{Sign}_\Sigma(\text{sk}_\Sigma, (\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}, (\text{pk}_i)_{[1..\ell]}, c))$ .
- Sign each randomized block using each  $\text{sk}_i$ , i.e., let  $\sigma_i \leftarrow \text{Sign}_\Sigma(\text{sk}_i, ((\text{pk}_i, m^i)_{[1..\ell]}, c, \sigma_s, \text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}))$ .
- Return  $(c, \sigma_s, (\text{pk}_i, \sigma_i)_{i \in [1..\ell]})$ .

 $\text{Verify}^{\text{SS}}(\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}, m, \sigma^{\text{SS}})$ . Parse  $\sigma^{\text{SS}}$  as  $(c, \sigma_s, (\text{pk}_i, \sigma_i)_{i \in [1..\ell]})$ . If  $\text{Verf}_\Sigma(\text{pk}_{\text{sig}}^{\text{SS}}, (\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}, (\text{pk}_i)_{[1..\ell]}, c), \sigma_s) = 0$ , return 0. If, for all  $i \in [1..\ell]$ , we have that  $\text{Verf}_\Sigma(\text{pk}_i, ((\text{pk}_i, m^i), c, \sigma_s, \text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}), \sigma_i) = 1$ , return 1. Otherwise, return 0.  
 $\text{Sanit}^{\text{SS}}(\text{sk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{sig}}^{\text{SS}}, m, \sigma^{\text{SS}}, \mathbb{M}^{\text{SS}})$ . The algorithm parses  $\sigma^{\text{SS}}$  as  $(c, \sigma_s, (\text{pk}_i, \sigma_i)_{i \in [1..\ell]})$  and proceeds as follows:

- If  $\text{Verify}^{\text{SS}}(\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}, m, \sigma^{\text{SS}}) \neq 1$ , return  $\perp$ .
- Let  $(x, \mathbb{A}^{\text{SS}}) \leftarrow \text{Dec}^\Pi(\text{sk}_{\text{san}}^{\text{SS}}, c, (\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}, (\text{pk}_i)_{[1..\ell]}))$ . If decryption fails, return  $\perp$ .
- For each  $i \in \mathbb{A}^{\text{SS}}.1$ , let  $(\text{sk}'_i, \text{pk}'_i) \leftarrow_r \text{KGen}_\Sigma(\text{pp}_\Sigma; \text{Eval}_{\text{PRF}}(x, i))$ .
- If  $\mathbb{M}^{\text{SS}} \setminus \mathbb{A}^{\text{SS}}.1 \neq \emptyset$ , return  $\perp$ .
- For each  $(i, m^{i'}) \in \mathbb{M}^{\text{SS}}$ , let  $\sigma'_i \leftarrow \text{Sign}_\Sigma(\text{sk}'_i, ((\text{pk}_i, m^{i'})_{[1..\ell]}, c, \sigma_s, \text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}))$ .
- Return  $(\mathbb{M}^{\text{SS}}(m), (c, \sigma_s, (\text{pk}_i, \sigma'_i)_{i \in [1..\ell]}))$ .

Construction 2: Construction of a non-accountable invisible  $\text{SS}'$

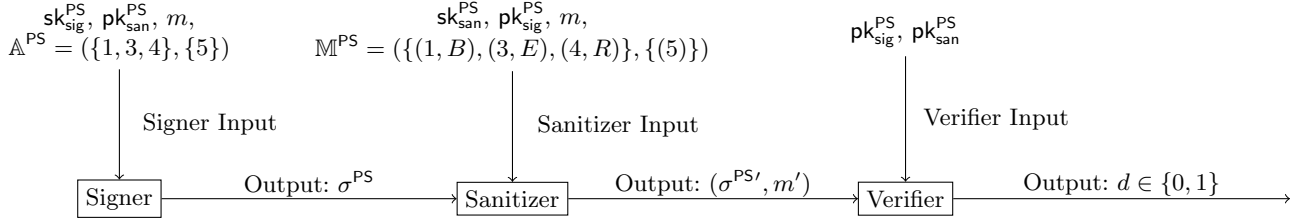


Fig. 22: Example workflow of a PS. The message  $m$  is set to  $(H, E, L, L, O)$  and is modified to  $(B, E, E, R)$ .

## 5.1 Framework

For the framework, we use the following notation. The variable  $\mathbb{A}^{\text{PS}}$  is a list containing the set of indices of the editable blocks, as well as the blocks which can be redacted. For example, let  $\mathbb{A}^{\text{PS}} = (\{1, 2\}, \{4\})$ . Then, the first and second block are editable, while only the fourth block can be redacted. The variable  $\mathbb{M}^{\text{PS}}$  is a list containing a set of pairs  $(i, m^{i'})$  for those blocks that are modified, meaning that  $m^i$  is replaced by  $m^{i'}$  and a set of indices to be redacted. In more detail, if  $\mathbb{M}^{\text{PS}} = (\{(1, b), (2, b)\}, \{3\})$  means that the first two blocks are altered to contain a  $b$ , while the third block is redacted.

We use the shorthand notation  $m' \leftarrow \mathbb{M}^{\text{PS}}(m)$  to denote the result of this replacement, while  $\mathbb{M}^{\text{PS}} \prec (m, \mathbb{A}^{\text{PS}})$  means that  $\mathbb{M}^{\text{PS}}$  is a valid modification instruction w.r.t.  $m$  and  $\mathbb{A}^{\text{PS}}$ . Likewise, we use  $\mathbb{A}^{\text{PS}} \prec m$  to denote that  $\mathbb{A}^{\text{PS}}$  is valid description of the admissible blocks w.r.t.  $m$ .

An example workflow is depicted in Figure 22 and Figure 23. To ease understanding and the description of our construction, we define that the replacements are done first and the redactions afterwards.

**Definition 23 (Protean Signatures).** A protean signature scheme PS consists of the following eight PPT algorithms  $(\text{PPGen}^{\text{PS}}, \text{KG}_{\text{sig}}^{\text{PS}}, \text{KG}_{\text{san}}^{\text{PS}}, \text{Sign}^{\text{PS}}, \text{Verify}^{\text{PS}}, \text{Edit}^{\text{PS}}, \text{Proof}^{\text{PS}}, \text{Judge}^{\text{PS}})$  such that:

$\text{PPGen}^{\text{PS}}$ . The algorithm  $\text{PPGen}^{\text{PS}}$  generates the public parameters:

$$\text{pp}_{\text{PS}} \leftarrow_r \text{PPGen}^{\text{PS}}(1^\lambda)$$

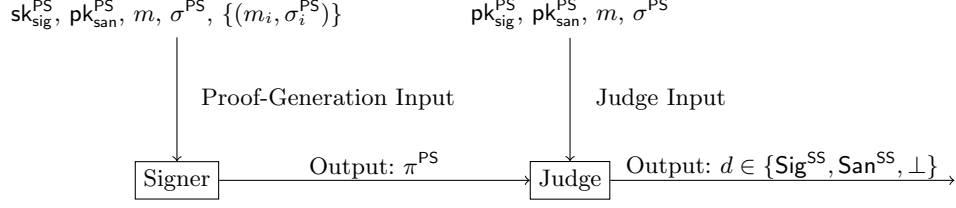


Fig. 23: Proof-generation and Judge<sup>PS</sup>

We assume that  $\text{pp}_{\text{PS}}$  is implicitly input to all other algorithms.

$\text{KG}_{\text{sig}}^{\text{PS}}$ . The algorithm  $\text{KG}_{\text{sig}}^{\text{PS}}$  generates the key pair of the signer:

$$(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{sig}}^{\text{PS}}) \leftarrow_r \text{KG}_{\text{sig}}^{\text{PS}}(\text{pp}_{\text{PS}})$$

$\text{KG}_{\text{san}}^{\text{PS}}$ . The algorithm  $\text{KG}_{\text{san}}^{\text{PS}}$  generates the key pair of the sanitizer:

$$(\text{sk}_{\text{san}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}) \leftarrow_r \text{KG}_{\text{san}}^{\text{PS}}(\text{pp}_{\text{PS}})$$

$\text{Sign}^{\text{PS}}$ . The algorithm  $\text{Sign}^{\text{PS}}$  generates a signature  $\sigma^{\text{PS}}$  on input of the public key  $\text{pk}_{\text{san}}^{\text{PS}}$ ,  $\mathbb{A}^{\text{PS}}$ , a message  $m$ , and  $\text{sk}_{\text{sig}}^{\text{PS}}$ :

$$\sigma^{\text{PS}} \leftarrow_r \text{Sign}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}, m, \mathbb{A}^{\text{PS}})$$

It is assumed that  $\mathbb{A}^{\text{PS}}$  can be derived from any verifying signature  $\sigma^{\text{PS}}$ , if  $\text{sk}_{\text{san}}^{\text{PS}}$  is known.

$\text{Verify}^{\text{PS}}$ . The algorithm  $\text{Verify}^{\text{PS}}$  verifies a signature  $\sigma^{\text{PS}}$ , i.e., outputs a decision  $d \in \{0, 1\}$  w.r.t.  $\text{pk}_{\text{san}}^{\text{PS}}$ ,  $\text{pk}_{\text{sig}}^{\text{PS}}$ , and a message  $m$ :

$$d \leftarrow \text{Verify}^{\text{PS}}(\text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}, m, \sigma^{\text{PS}})$$

$\text{Edit}^{\text{PS}}$ . The algorithm  $\text{Edit}^{\text{PS}}$  generates a sanitized signature  $\sigma^{\text{PS}'}$  and updated  $\mathbb{A}^{\text{PS}'}$ , given inputs  $\text{sk}_{\text{san}}^{\text{PS}}$ ,  $\mathbb{A}^{\text{PS}}$ , a message  $m$ , a signature  $\sigma$ , and  $\text{pk}_{\text{sig}}^{\text{PS}}$ :

$$(m', \sigma^{\text{PS}'}, \mathbb{A}^{\text{PS}'}) \leftarrow_r \text{Edit}^{\text{PS}}(\text{sk}_{\text{san}}^{\text{PS}}, \text{pk}_{\text{sig}}^{\text{PS}}, \sigma, m, \mathbb{M}^{\text{PS}})$$

$\text{Proof}^{\text{PS}}$ . The algorithm  $\text{Proof}^{\text{PS}}$  outputs a proof  $\pi^{\text{PS}}$  on input  $m$ ,  $\sigma^{\text{PS}}$ ,  $\text{sk}_{\text{sig}}^{\text{PS}}$ ,  $\text{pk}_{\text{san}}^{\text{PS}}$ , and a set of polynomially many additional signature/message pairs  $\{(\sigma_i^{\text{PS}}, m^i)\}$ . The proof  $\pi^{\text{PS}}$  is used by the next algorithm to pinpoint the accountable party for a given signature:

$$\pi^{\text{PS}} \leftarrow_r \text{Proof}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}, m, \sigma^{\text{PS}}, \{(\sigma_i^{\text{PS}}, m^i)\})$$

$\text{Judge}^{\text{PS}}$ . The algorithm  $\text{Judge}^{\text{PS}}$  outputs a decision  $d \in \{\text{Sig}^{\text{PS}}, \text{San}^{\text{PS}}, \perp\}$  indicating whether the message/signature pair has been created by the signer, or the sanitizer:

$$d \leftarrow \text{Judge}^{\text{PS}}(\text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}, m, \sigma^{\text{PS}}, \pi^{\text{PS}})$$

## 5.2 PSs Security Definitions

We now introduce the security properties for PSs. Clearly, the goals are similar to the ones for SSs and RSs. However, due to the extended capabilities, the semantic is quite different, while we need to take extra care for changed indices after redactions.

$\mathbf{Exp}_{\mathcal{A}, \text{PS}}^{\text{Unforgeability}}(\lambda)$   
 $\text{pp}_{\text{PS}} \leftarrow_r \text{PPGen}^{\text{PS}}(1^\lambda)$   
 $(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{sig}}^{\text{PS}}) \leftarrow_r \text{KG}_{\text{sig}}^{\text{PS}}(\text{pp}_{\text{PS}})$   
 $(\text{sk}_{\text{san}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}) \leftarrow_r \text{KG}_{\text{san}}^{\text{PS}}(\text{pp}_{\text{PS}})$   
 $(m^*, \sigma^{\text{PS}*}) \leftarrow_r \mathcal{A}^{\text{Sign}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \cdot, \cdot), \text{Edit}^{\text{PS}}(\text{sk}_{\text{san}}^{\text{PS}}, \cdot, \cdot, \cdot), \text{Proof}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \cdot, \cdot, \cdot)}(\text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$   
 $\forall i \in [1..q]$ , let  $(\text{pk}_{\text{san}, i}^{\text{PS}}, m_i, \mathbb{A}_i^{\text{PS}})$  and  $\sigma_i^{\text{PS}}$   
index the queries/answers to/from  $\text{Sign}^{\text{PS}}$   
 $\forall i \in [1..q']$ , let  $(\text{pk}_{\text{sig}, j}^{\text{PS}}, m_j, \sigma_j^{\text{PS}}, \mathbb{M}_j)$  and  $(m'_j, \sigma_j^{\text{PS}'}, \mathbb{A}'_j)$   
index the queries/answers to/from  $\text{Edit}^{\text{PS}}$   
return 1, if  $\text{Verify}^{\text{PS}}(\text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}, m^*, \sigma^{\text{PS}*}) = 1 \wedge$   
 $\forall i \in [1..q] : (\text{pk}_{\text{san}}^{\text{PS}}, m^*, \sigma^{\text{PS}*}) \neq (\text{pk}_{\text{san}, i}^{\text{PS}}, m_i, \sigma_i^{\text{PS}}) \wedge$   
 $\forall j \in [1..q'] : (\text{pk}_{\text{sig}}^{\text{PS}}, m^*, \sigma^{\text{PS}*}) \neq (\text{pk}_{\text{sig}, j}^{\text{PS}}, m_j, \sigma_j^{\text{PS}'})$   
return 0

Fig. 24: PS Unforgeability

$\mathbf{Exp}_{\mathcal{A}, \text{PS}}^{\text{Immutability}}(\lambda)$   
 $\text{pp}_{\text{PS}} \leftarrow_r \text{PPGen}^{\text{PS}}(1^\lambda)$   
 $(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{sig}}^{\text{PS}}) \leftarrow_r \text{KG}_{\text{sig}}^{\text{PS}}(\text{pp}_{\text{PS}})$   
 $(m^*, \sigma^{\text{PS}*}, \text{pk}_{\text{san}}^{\text{PS}*}) \leftarrow_r \mathcal{A}^{\text{Sign}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \cdot, \cdot), \text{Proof}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \cdot, \cdot, \cdot)}(\text{pk}_{\text{sig}}^{\text{PS}})$   
 $\forall i \in [1..q]$ , let  $(\text{pk}_{\text{san}, i}^{\text{PS}}, m_i, \mathbb{A}_i^{\text{PS}})$   
index the queries to  $\text{Sign}^{\text{PS}}$   
return 1, if  $\text{Verify}^{\text{PS}}(\text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}*}, m^*, \sigma^{\text{PS}*}) = 1 \wedge$   
 $\forall i \in [1..q] : (\text{pk}_{\text{san}}^{\text{PS}*} \neq \text{pk}_{\text{san}, i}^{\text{PS}} \vee$   
 $m^* \notin \{\mathbb{M}(m_i) \mid \mathbb{M} \text{ with } \mathbb{M} \prec (m_i, \mathbb{A}_i^{\text{PS}})\})$   
return 0

Fig. 25: PS Immutability

**Unforgeability.** This definition requires that an adversary  $\mathcal{A}$  not having any secret keys is not able to produce any valid signature  $\sigma^{\text{PS}*}$  on a message  $m^*$  which it has never not seen, even if  $\mathcal{A}$  has full oracle access, i.e., this captures “strong unforgeability” [46].

**Definition 24 (Unforgeability).** A PS is unforgeable, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that

$$\Pr[\mathbf{Exp}_{\mathcal{A}, \text{PS}}^{\text{Unforgeability}}(\lambda) = 1] \leq \nu(\lambda),$$

where the corresponding experiment is defined in Figure 24.

**Immutability.** This definition prohibits that an adversary  $\mathcal{A}$  can generate a verifying signature  $\sigma^{\text{PS}*}$  for a message  $m^*$  not derivable from the signatures given by an honest signer, even if it can generate the sanitizer’s key pair.

**Definition 25 (Immutability).** A PS is immutable, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that

$$\Pr[\mathbf{Exp}_{\mathcal{A}, \text{PS}}^{\text{Immutability}}(\lambda) = 1] \leq \nu(\lambda),$$

where the corresponding experiment is defined in Figure 25.

```

Exp $\mathcal{A}, \text{PS}$ Privacy( $\lambda$ )
   $\text{pp}_{\text{PS}} \leftarrow_r \text{PPGen}^{\text{PS}}(1^\lambda)$ 
   $(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{sig}}^{\text{PS}}) \leftarrow_r \text{KG}_{\text{sig}}^{\text{PS}}(\text{pp}_{\text{PS}})$ 
   $(\text{sk}_{\text{san}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}) \leftarrow_r \text{KG}_{\text{san}}^{\text{PS}}(\text{pp}_{\text{PS}})$ 
   $b \leftarrow_r \{0, 1\}$ 
   $a \leftarrow_r \mathcal{A}_{\text{ProofPS}}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{sig}}^{\text{PS}}, \text{sk}_{\text{san}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}, \text{LoREdit}(\cdot, \cdot, \cdot, \cdot, \text{sk}_{\text{sig}}^{\text{PS}}, \text{sk}_{\text{san}}^{\text{PS}}, b))(\text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ 
  where  $\text{LoREdit}(m_0, m_1, \mathbb{M}_0^{\text{PS}}, \mathbb{M}_1^{\text{PS}}, \mathbb{A}_0^{\text{PS}}, \mathbb{A}_1^{\text{PS}}, \text{sk}_{\text{sig}}^{\text{PS}}, \text{sk}_{\text{san}}^{\text{PS}}, b)$ 
   $\sigma_i^{\text{PS}} \leftarrow_r \text{Sign}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}, m_i, \mathbb{A}_i^{\text{PS}})$  for  $i \in \{0, 1\}$ 
   $(m'_i, \sigma_i^{\text{PS}'}, \mathbb{A}_i^{\text{PS}'}) \leftarrow_r \text{Edit}^{\text{PS}}(\text{sk}_{\text{san}}^{\text{PS}}, \text{pk}_{\text{sig}}^{\text{PS}}, m_i, \sigma_i^{\text{PS}}, \mathbb{M}_i^{\text{PS}})$  for  $i \in \{0, 1\}$ 
  return  $\perp$ , if  $m'_0 \neq m'_1 \vee \mathbb{A}_0^{\text{PS}'} \neq \mathbb{A}_1^{\text{PS}'}$ 
  return  $(m'_b, \sigma_b^{\text{PS}'}, \mathbb{A}_b^{\text{PS}'})$ 
return 1, if  $a = b$ 
return 0

```

Fig. 26: PS Privacy

**Privacy.** This definition prohibits that an adversary  $\mathcal{A}$  can learn anything about edited (redacted or sanitized) parts.

**Definition 26 (Privacy).** A PS is private, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that

$$\left| \Pr[\mathbf{Exp}_{\mathcal{A}, \text{PS}}^{\text{Immutability}}(\lambda)] - 1/2 \right| \leq \nu(\lambda),$$

where the corresponding experiment is defined in Figure 26.

**Transparency.** This definition requires that an adversary  $\mathcal{A}$  does not learn whether a signature  $\sigma^{\text{PS}}$  was generated through  $\text{Sign}^{\text{PS}}$  or  $\text{Edit}^{\text{PS}}$ .

**Definition 27 (Transparency).** A PS is transparent, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that

$$\left| \Pr[\mathbf{Exp}_{\mathcal{A}, \text{PS}}^{\text{Transparency}}(\lambda)] - 1/2 \right| \leq \nu(\lambda),$$

where the corresponding experiment is defined in Figure 27.

**Signer-Accountability.** Signer-accountability prohibits that an adversary can generate a bogus proof that makes  $\text{Judge}^{\text{PS}}$  decide that the sanitizer is responsible for a given signature/message pair  $(m^*, \sigma^{\text{PS}*})$ , but the sanitizer has never generated this pair. This is even true, if the adversary can generate the signer's key pair.

**Definition 28 (Signer-Accountability).** A PS is signer-accountable, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that

$$\Pr[\mathbf{Exp}_{\mathcal{A}, \text{PS}}^{\text{SigAccountability}}(\lambda) = 1] \leq \nu(\lambda),$$

where the corresponding experiment is defined in Figure 28.



**Exp** <sub>$\mathcal{A}, \text{PS}$</sub> <sup>Transparency</sup>( $\lambda$ )

$\text{pp}_{\text{PS}} \leftarrow_r \text{PPGen}^{\text{PS}}(1^\lambda)$   
 $(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{sig}}^{\text{PS}}) \leftarrow_r \text{KG}_{\text{sig}}^{\text{PS}}(\text{pp}_{\text{PS}})$   
 $(\text{sk}_{\text{san}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}) \leftarrow_r \text{KG}_{\text{san}}^{\text{PS}}(\text{pp}_{\text{PS}})$   
 $b \leftarrow_r \{0, 1\}$   
 $\mathcal{Q} \leftarrow \emptyset$

$a \leftarrow_r \mathcal{A}_{\text{Proof}^{\text{PS}'}}^{\text{Sign}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \cdot, \cdot), \text{Edit}^{\text{PS}}(\text{sk}_{\text{san}}^{\text{PS}}, \cdot, \cdot, \cdot)}(\text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$   
where  $\text{Proof}^{\text{PS}'}$ ( $\text{sk}_{\text{sig}}^{\text{PS}}, m, \sigma^{\text{PS}}, \{(m_i, \sigma_i^{\text{PS}}) \mid i \in \mathbb{N}\}$ ):  
return  $\perp$ , if  $\text{pk}_{\text{san}}^{\text{PS}} = \text{pk}_{\text{san}}^{\text{PS}} \wedge ((m, \sigma^{\text{PS}}) \in \mathcal{Q} \vee \mathcal{Q} \cap \{(m_i, \sigma_i^{\text{PS}})\} \neq \emptyset)$   
return  $\text{Proof}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}'}, m, \sigma^{\text{PS}}, \{(m_i, \sigma_i^{\text{PS}})\})$   
where  $\text{Sign/Edit}(m, \mathbb{M}^{\text{PS}}, \mathbb{A}^{\text{PS}}, \text{sk}_{\text{sig}}^{\text{PS}}, \text{sk}_{\text{san}}^{\text{PS}}, b)$ :  
 $\sigma^{\text{PS}} \leftarrow_r \text{Sign}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}, m, \mathbb{A}^{\text{PS}})$   
 $(m', \sigma^{\text{PS}'}, \mathbb{A}^{\text{PS}'}) \leftarrow_r \text{Edit}^{\text{PS}}(\text{sk}_{\text{san}}^{\text{PS}}, \text{pk}_{\text{sig}}^{\text{PS}}, m, \sigma^{\text{PS}}, \mathbb{M}^{\text{PS}})$   
if  $b = 1$ :  
 $\sigma^{\text{PS}'} \leftarrow_r \text{Sign}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}, m', \mathbb{A}^{\text{PS}'})$   
if  $\sigma^{\text{PS}'} \neq \perp$ , set  $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(m', \sigma^{\text{PS}'})\}$   
return  $(m', \sigma^{\text{PS}'})$   
return 1, if  $a = b$   
return 0

Fig. 27: PS Transparency

**Exp** <sub>$\mathcal{A}, \text{PS}$</sub> <sup>SigAccountability</sup>( $\lambda$ )

$\text{pp}_{\text{PS}} \leftarrow_r \text{PPGen}^{\text{PS}}(1^\lambda)$   
 $(\text{sk}_{\text{san}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}) \leftarrow_r \text{KG}_{\text{san}}^{\text{PS}}(\text{pp}_{\text{PS}})$   
 $(\text{pk}_{\text{sig}}^{\text{PS}*}, \pi^{\text{PS}*}, m^*, \sigma^{\text{PS}*}) \leftarrow_r \mathcal{A}_{\text{Edit}^{\text{PS}}(\text{sk}_{\text{san}}^{\text{PS}}, \cdot, \cdot, \cdot)}(\text{pk}_{\text{san}}^{\text{PS}})$   
 $\forall i \in [1..q]$ , let  $(m'_i, \sigma_i^{\text{PS}'}, \mathbb{A}^{\text{PS}'}_j)$  and  $(m_i, \mathbb{M}_i^{\text{PS}}, \sigma_i^{\text{PS}}, \text{pk}_{\text{sig}, i}^{\text{PS}})$   
index the answers/queries from/to  $\text{Edit}^{\text{PS}}$   
return 1, if  $\text{Verify}^{\text{PS}}(\text{pk}_{\text{sig}}^{\text{PS}*}, \text{pk}_{\text{san}}^{\text{PS}}, m^*, \sigma^{\text{PS}*}) = 1 \wedge$   
 $\forall i \in [1..q] : (\text{pk}_{\text{sig}}^{\text{PS}*}, m^*, \sigma^{\text{PS}*}) \neq (\text{pk}_{\text{sig}, i}^{\text{PS}}, m'_i, \sigma_i^{\text{PS}'}) \wedge$   
 $\text{Judge}^{\text{PS}}(\text{pk}_{\text{sig}}^{\text{PS}*}, \text{pk}_{\text{san}}^{\text{PS}}, m^*, \sigma^{\text{PS}*}, \pi^{\text{PS}*}) = \text{San}^{\text{PS}}$   
return 0

Fig. 28: PS Signer-Accountability

**Sanitizer-Accountability.** Sanitizer-accountability prohibits that an adversary can generate a bogus signature/message pair  $(m^*, \sigma^{\text{PS}*})$  that makes  $\text{Proof}^{\text{PS}}$  output an honestly generated proof  $\pi^{\text{PS}}$  which points to the signer, but  $(m^*, \sigma^{\text{PS}*})$  has never been generated by the signer. This is even true, if the adversary can generate the sanitizer's key pair.

**Definition 29 (Sanitizer-Accountability).** A PS is sanitizer-accountable, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that

$$\Pr[\mathbf{Exp}_{\mathcal{A}, \text{PS}}^{\text{SanAccountability}}(\lambda) = 1] \leq \nu(\lambda),$$

where the corresponding experiment is defined in Figure 29.

**Exp** <sub>$\mathcal{A}, \text{PS}$</sub> <sup>SanAccountability</sup>( $\lambda$ )

$\text{pp}_{\text{PS}} \leftarrow_r \text{PPGen}^{\text{PS}}(1^\lambda)$   
 $(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{sig}}^{\text{PS}}) \leftarrow_r \text{KG}_{\text{sig}}^{\text{PS}}(\text{pp}_{\text{PS}})$   
 $(m^*, \sigma^{\text{PS}*}, \text{pk}_{\text{san}}^{\text{PS}*}) \leftarrow_r \mathcal{A}^{\text{Sign}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \dots), \text{Proof}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \dots)}(\text{pk}_{\text{sig}}^{\text{PS}})$   
 $\forall i \in [1..q]$ , let  $(\text{pk}_{\text{san}, i}^{\text{PS}}, m_i, \mathbb{A}_i^{\text{PS}})$  and  $\sigma_i^{\text{PS}}$   
index the queries/answers to/from  $\text{Sign}^{\text{PS}}$   
 $\pi^{\text{PS}} \leftarrow_r \text{Proof}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}*}, m^*, \sigma^{\text{PS}*}, \{(m_i, \sigma_i^{\text{PS}}) \mid 0 < i \leq q\})$   
return 1, if  $\text{Verify}^{\text{PS}}(\text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}*}, m^*, \sigma^{\text{PS}*}) = 1 \wedge$   
 $\forall i \in [1..q] : (\text{pk}_{\text{san}}^{\text{PS}*}, m^*, \sigma^{\text{PS}*}) \neq (\text{pk}_{\text{san}, i}^{\text{PS}}, m_i, \sigma_i^{\text{PS}}) \wedge$   
 $\text{Judge}^{\text{PS}}(\text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}*}, m^*, \sigma^{\text{PS}*}, \pi^{\text{PS}}) = \text{Sig}^{\text{PS}}$   
return 0

Fig. 29: PS Sanitizer-Accountability

**Exp** <sub>$\mathcal{A}, \text{PS}$</sub> <sup>Invisibility</sup>( $\lambda$ )

$\text{pp}_{\text{PS}} \leftarrow_r \text{PPGen}^{\text{PS}}(1^\lambda)$   
 $(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{sig}}^{\text{PS}}) \leftarrow_r \text{KG}_{\text{sig}}^{\text{PS}}(\text{pp}_{\text{PS}})$   
 $(\text{sk}_{\text{san}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}) \leftarrow_r \text{KG}_{\text{san}}^{\text{PS}}(\text{pp}_{\text{PS}})$   
 $b \leftarrow_r \{0, 1\}$   
 $\mathcal{Q} \leftarrow \emptyset$   
 $a \leftarrow_r \mathcal{A}^{\text{Edit}^{\text{PS}' }(\text{sk}_{\text{san}}^{\text{PS}}, \dots), \text{Proof}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \dots), \text{LoRADM}(\text{sk}_{\text{sig}}^{\text{PS}}, \dots, b)}(\text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$   
where  $\text{LoRADM}(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}'}, m, \mathbb{A}_0^{\text{PS}}, \mathbb{A}_1^{\text{PS}}, b)$ :  
return  $\perp$ , if  $\mathbb{A}_0^{\text{PS}} \not\prec m \wedge \mathbb{A}_1^{\text{PS}} \not\prec m$   
return  $\perp$ , if  $\text{pk}_{\text{san}}^{\text{PS}} \neq \text{pk}_{\text{san}}^{\text{PS}' } \wedge \mathbb{A}_0^{\text{PS}} \neq \mathbb{A}_1^{\text{PS}}$   
 $\sigma^{\text{PS}} \leftarrow_r \text{Sign}^{\text{PS}}(\text{sk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}'}, m, \mathbb{A}_b^{\text{PS}})$   
if  $\text{pk}_{\text{san}}^{\text{PS}' } = \text{pk}_{\text{san}}^{\text{PS}}$   
 $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(m, \sigma^{\text{PS}}, (\mathbb{A}_0^{\text{PS}}.1 \cap \mathbb{A}_1^{\text{PS}}.1, \mathbb{A}_0^{\text{PS}}.2 \cap \mathbb{A}_1^{\text{PS}}.2))\}$   
return  $\sigma^{\text{PS}}$   
where  $\text{Edit}^{\text{PS}' }(\text{pk}_{\text{sig}}^{\text{PS}'}, \text{sk}_{\text{san}}^{\text{PS}}, \sigma^{\text{PS}}, m, \mathbb{M}^{\text{PS}})$ :  
return  $\perp$ , if  $\text{pk}_{\text{sig}}^{\text{PS}' } = \text{pk}_{\text{sig}}^{\text{PS}} \wedge \exists (m, \sigma^{\text{PS}}, \mathbb{A}) \in \mathcal{Q} : \mathbb{M}^{\text{PS}} \not\prec (m, \mathbb{A})$   
 $(m', \sigma^{\text{PS}' }, \mathbb{A}^{\text{PS}' }) \leftarrow_r \text{Edit}^{\text{PS}}(\text{pk}_{\text{sig}}^{\text{PS}}, \text{sk}_{\text{san}}^{\text{PS}}, m, \mathbb{M}^{\text{PS}}, \sigma^{\text{PS}})$   
if  $\text{pk}_{\text{sig}}^{\text{PS}' } = \text{pk}_{\text{sig}}^{\text{PS}} \wedge \exists (m, \sigma^{\text{PS}}, \mathbb{A}^{\text{PS}' }) \in \mathcal{Q} : \mathbb{M}^{\text{PS}} \prec (m, \mathbb{A}^{\text{PS}' })$ ,  
 $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(m', \sigma^{\text{PS}' }, \mathbb{A}^{\text{PS}' })\}$   
return  $(m', \sigma^{\text{PS}' })$   
return 1, if  $a = b$   
return 0

Fig. 30: PS Invisibility

### 5.3 (Full) Invisibility

Invisibility prohibits that an outsider can decide which blocks can be edited and also which blocks can be redacted. Note, the signing oracle can be simulated by using the same  $\mathbb{A}^{\text{PS}}$  in the LoRADM oracle. Moreover, as done for SSs (See Section 2), we define a slightly stronger variant than defined by Krenn et al. [4]: the adversary  $\mathcal{A}$  can now query arbitrary signature to the LoRADM-oracle.

**Definition 30 (Invisibility).** A PS is (fully) invisible, if for any PPT adversary  $\mathcal{A}$  there exists a negligible function  $\nu$  such that

$$\left| \Pr[\mathbf{Exp}_{\mathcal{A}, \text{PS}}^{\text{Invisibility}}(\lambda)] - 1/2 \right| \leq \nu(\lambda),$$

where the corresponding experiment is defined in Figure 30.

On a formal level, the main difference to the invisibility definition of Krenn et al. [4] is that in their work the LoRADM-oracle aborts if the indices of the redactable blocks specified by  $\mathbb{A}_0^{\text{PS}}$  and  $\mathbb{A}_1^{\text{PS}}$  are not the same, i.e., if  $\mathbb{A}_0^{\text{PS}}.2 \neq \mathbb{A}_1^{\text{PS}}.2$ , in order to avoid trivial distinguishers based on the redactable blocks; by removing this condition in our definition, it is now guaranteed that also the redactable blocks remain hidden from an outsider, thereby formalizing our ambition that an outsider should neither be able to decide which parts of a message are editable *nor which parts are redactable*.

## 5.4 Construction

We now present our construction of a PS. It is essentially the same as given by Krenn et al. [4], but with some minor, yet very important, quirks.

The basic idea is to combine RSs and SSs by bridging them using unique tags. In more detail, each block  $m^i \in m$  is signed using an invisible non-accountable SS', while an additional (non-admissible) tag  $\tau$  is used to identify the “overall” message  $m$  the block  $m^i$  belongs to. Moreover, each block  $m^i$  is also assigned a (non-admissible) additional tag  $\tau_i$ , along with all public keys, used by an invisible RS to allow for redactions. Thus, there are  $\ell_m \sigma_i^{\text{SS}}$ , where each signature protects  $(m^i, \tau, \tau_i, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ . If a block  $m_i$  is sanitizable, it is marked as admissible within  $\mathbb{A}_i^{\text{SS}}$ . This allows to sanitize the block  $m^i$ . Then, each tag  $\tau_i$  is put into an RS to allow for transparent redactions, additionally bound to the non-redactable “overall” tag  $\tau$  and all (non-redactable) public keys. If a block  $m^i$  is non-redactable, this is marked in  $\mathbb{A}^{\text{RS}}$ . Thus,  $\sigma^{\text{RS}}$  protects  $(\tau_1, \dots, \tau_{\ell_m}, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ . Finally, to achieve accountability, all tags, all signatures generated so far, the resulting values are signed again using an additional, non-invisible but accountable, SS, while in this outer SS everything, but the public keys and the tag  $\tau$  are admissible. To maintain transparency, the overall message  $m$  is a *single* block in the outer SS.

In more detail, the outer signature  $\sigma_0^{\text{SS}}$  protects  $(m, \sigma^{\text{RS}}, (\tau_i, \sigma_i^{\text{SS}})_i, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ . Thus, changing the message or any signature requires changing  $\sigma_0^{\text{SS}}$ , implying accountability. Upon sanitization of a block  $m^i$ ,  $\sigma_i^{\text{SS}}$  is sanitized, while the outer signature  $\sigma_0^{\text{SS}}$  needs to be adjusted as well. For redaction of a block  $m^i$ ,  $\sigma^{\text{RS}}$  is adjusted and the corresponding signature is no longer given out. This also means that  $\sigma_0^{\text{SS}}$  must be adjusted.

The resulting construction is depicted in Construction 3. To give a graphical overview of the construction idea, see Figure 31 (before editing) and Figure 32 (after editing). Moreover, we do not consider unlinkability [33], as it seems to be very hard to achieve with the underlying construction paradigm, especially considering that at the time of writing this paper no SSs which are unlinkable and invisible at the same time were known yet. Note that in a parallel work, Bultel et al. [57] presented such a scheme; however, it seems hard to combine their scheme with our construction paradigm.

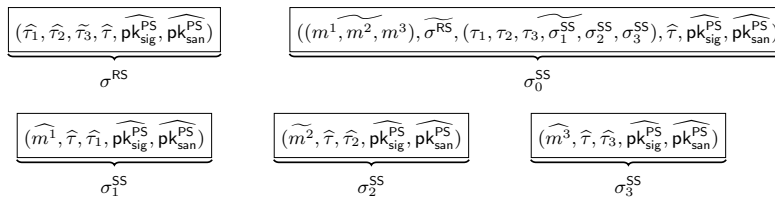


Fig. 31: Our main construction idea. Let  $\mathbb{A}^{\text{PS}} = (\{2\}, \{3\})$  and  $m = (m^1, m^2, m^3)$  for preciseness, i.e., only the second block of  $m$  is sanitizable, while only the last block of  $m$  is redactable. Redactable elements for the RS (or sanitizable for the SS) are marked with a tilde, i.e.,  $\tilde{\cdot}$ . Blocks which are not redactable (or sanitizable resp.) are marked with a hat, i.e.,  $\hat{\cdot}$ .

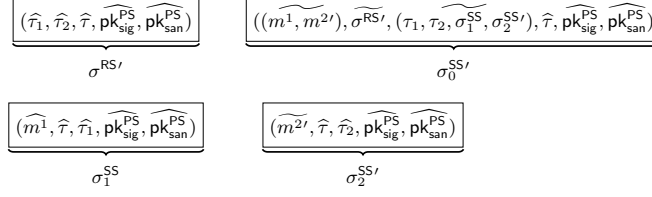


Fig. 32: State after sanitization. Here, block  $m^3$  was redacted and  $m^2$  was changed to  $m^{2'}$ . Block  $m^1$  must stay the same.

<p><b>PPGen<sup>PS</sup>(1<sup>λ</sup>).</b> Let <math>pp_{SS} \leftarrow_r SS.PPGen^{SS}(1^\lambda)</math>, <math>pp_{SS'} \leftarrow_r SS'.PPGen^{SS}(1^\lambda)</math>, and <math>pp_{RS} \leftarrow_r PPGen^{RS}(1^\lambda)</math>.  Return <math>pp_{PS} = (pp_{SS}, pp_{SS'}, pp_{RS})</math>.</p> <p><b>KG<sub>sig</sub><sup>PS</sup>(pp<sub>PS</sub>).</b> Let <math>(sk_{sig}^{SS}, pk_{sig}^{SS}) \leftarrow_r SS.KG_{sig}^{SS}(pp_{SS})</math>, <math>(sk_{sig}^{SS'}, pk_{sig}^{SS'}) \leftarrow_r SS'.KG_{sig}^{SS}(pp_{SS'})</math>, and <math>(sk_{sig}^{RS}, pk_{sig}^{RS}) \leftarrow_r KG_{sig}^{RS}(pp_{RS})</math>.  Return <math>(sk_{sig}^{PS}, pk_{sig}^{PS}) = ((sk_{sig}^{SS}, sk_{sig}^{SS'}, sk_{sig}^{RS}), (pk_{sig}^{SS}, pk_{sig}^{SS'}, pk_{sig}^{RS}))</math>.</p> <p><b>KG<sub>san</sub><sup>PS</sup>(pp<sub>PS</sub>).</b> Let <math>(sk_{san}^{SS}, pk_{san}^{SS}) \leftarrow_r SS.KG_{san}^{SS}(pp_{SS})</math>, <math>(sk_{san}^{SS'}, pk_{san}^{SS'}) \leftarrow_r SS'.KG_{san}^{SS}(pp_{SS'})</math>, and <math>(sk_{san}^{RS}, pk_{san}^{RS}) \leftarrow_r KG_{san}^{RS}(pp_{RS})</math>.  Return <math>(sk_{san}^{PS}, pk_{san}^{PS}) = ((sk_{san}^{SS}, sk_{san}^{SS'}, sk_{san}^{RS}), (pk_{san}^{SS}, pk_{san}^{SS'}, pk_{san}^{RS}))</math>.</p>
<p><b>Sign<sup>PS</sup>(sk<sub>sig</sub><sup>PS</sup>, pk<sub>sig</sub><sup>PS</sup>, m, A<sup>PS</sup>).</b> The algorithm proceeds as follows:</p> <ul style="list-style-type: none"> <li>– If <math>\mathbb{A}^{PS} \prec m = (m^1, m^2, \dots, m^\ell)</math> does not hold, return <math>\perp</math>, otherwise parse <math>\mathbb{A}^{PS} = (\mathbb{A}_1^{PS}, \mathbb{A}_2^{PS})</math>.</li> <li>– Draw <math>\tau \leftarrow_r \{0, 1\}^\lambda</math>.</li> <li>– For all <math>i \in [1..\ell_m]</math>, let <math>\sigma_i^{SS} \leftarrow_r SS'.Sign^{SS}(sk_{sig}^{SS'}, pk_{sig}^{SS'}, (m^i, \tau, \tau_i, pk_{sig}^{PS}, pk_{san}^{PS}), \mathbb{A}_i^{SS})</math>, where each <math>\tau_i \leftarrow_r \{0, 1\}^\lambda</math>. Furthermore, if <math>i \in \mathbb{A}_1^{PS}</math>, let <math>\mathbb{A}_i^{SS} = (\{1\}, 5)</math> and <math>\mathbb{A}_i^{SS} = (\emptyset, 5)</math> otherwise.</li> <li>– Let <math>\sigma^{RS} \leftarrow_r Sign^{RS}(sk_{sig}^{RS}, pk_{sig}^{RS}, m', \mathbb{A}^{RS})</math>, where <math>\mathbb{A}^{RS} = \mathbb{A}_2^{PS}</math> and the message <math>m' = (\tau_1, \dots, \tau_\ell, \tau, pk_{sig}^{PS}, pk_{san}^{PS})</math>.</li> <li>– Generate <math>\sigma_0^{SS} \leftarrow_r SS.Sign^{SS}(sk_{sig}^{SS}, pk_{sig}^{SS}, (m, \sigma^{RS}, (\tau_i, \sigma_i^{SS})_{1 \leq i \leq \ell_m}, \tau, pk_{sig}^{PS}, pk_{san}^{PS}), \mathbb{A}_0^{SS})</math>, where <math>\mathbb{A}_0^{SS} = (\{1, 2, 3\}, 6)</math>.</li> <li>– Return <math>((\sigma_i^{SS})_{0 \leq i \leq \ell_m}, \sigma^{RS}, \tau, (\tau_i)_{1 \leq i \leq \ell_m})</math>.</li> </ul> <p><b>Verify<sup>PS</sup>(pk<sub>sig</sub><sup>PS</sup>, pk<sub>san</sub><sup>PS</sup>, m, σ<sup>PS</sup>).</b> If <math>Verify^{SS}(pk_{sig}^{SS}, pk_{san}^{SS}, (m, \sigma^{RS}, \tau, (\tau_i, \sigma_i^{SS})_{1 \leq i \leq \ell_m}, pk_{sig}^{PS}, pk_{san}^{PS}), \sigma_0^{SS}) = 0</math>, return 0. If <math>Verify^{RS}(pk_{sig}^{RS}, (\tau_1, \dots, \tau_\ell, \tau, pk_{sig}^{PS}, pk_{san}^{PS}), \sigma^{RS}) = 0</math>, return 0. If for any <math>i \in [1..\ell_m] : SS'.Verify^{SS}(pk_{sig}^{SS'}, pk_{san}^{SS'}, (m^i, \tau, \tau_i, pk_{sig}^{PS}, pk_{san}^{PS}), \sigma_i^{SS}) = 0</math>, return 0. Return 1.</p> <p><b>Edit<sup>PS</sup>(sk<sub>san</sub><sup>PS</sup>, pk<sub>sig</sub><sup>PS</sup>, σ<sup>PS</sup>, m, M<sup>PS</sup>).</b> The algorithm proceeds as follows:</p> <ul style="list-style-type: none"> <li>– If <math>Verify^{PS}(pk_{sig}^{PS}, pk_{san}^{PS}, m, \sigma^{PS}) = 0</math>, return <math>\perp</math>, otherwise parse <math>M^{PS} = (M_1^{PS}, M_2^{PS})</math> and continue.</li> <li>– For all <math>(i, m^{i'}) \in M_1^{PS}</math>, let <math>(m^{i'}, \sigma_i^{SS'}) \leftarrow_r SS'.Sanit^{SS}(sk_{san}^{SS'}, pk_{sig}^{SS'}, (m^i, \tau, \tau_i, pk_{sig}^{PS}, pk_{san}^{PS}), \sigma_i^{SS}, \{(0, m^{i'})\})</math>. Return <math>\perp</math> if <math>\sigma_i^{SS'} = \perp</math>, otherwise set <math>\sigma_i^{SS} \leftarrow \sigma_i^{SS'}</math>.</li> <li>– Generate <math>(\sigma^{RS'}, \cdot, \cdot) \leftarrow_r Red^{RS}(pk_{sig}^{RS}, m'', \sigma^{RS}, M_2^{PS}, RED^{RS})</math>, where <math>m'' = (\tau_1, \dots, \tau_\ell, \tau, pk_{sig}^{PS}, pk_{san}^{PS})</math>. If <math>\sigma^{RS'} = \perp</math>, return <math>\perp</math>.</li> <li>– Let <math>(m'_0, \sigma_0^{SS'}) \leftarrow_r SS.Sanit^{SS}(sk_{san}^{SS}, pk_{sig}^{SS}, (m, \sigma^{RS}, (\tau_i, \sigma_i^{SS})_{i \in [1..\ell]}, \tau, pk_{sig}^{PS}, pk_{san}^{PS}), \sigma_0^{SS}, \{(1, m'), (2, \sigma^{RS'}), (3, (\tau_i, \sigma_i^{SS'})_{i \in [1..\ell] \setminus M_2^{PS}})\})</math>. If <math>\sigma_0^{SS'} = \perp</math>, return <math>\perp</math>.</li> <li>– Update <math>\mathbb{A}^{PS}</math> to <math>\mathbb{A}^{PS'}</math> by removing all indices in <math>M_2^{PS}</math> and adjusting the remaining indices by reducing each <math>i</math> in <math>\mathbb{A}^{PS}</math> by <math> \{j \in M_2^{PS} : j &lt; i\} </math>.</li> <li>– Return <math>(M^{PS}(m), ((\sigma_i^{SS})_{i \in [1..\ell] \setminus M_2^{PS}}, \sigma^{RS'}, \tau, (\tau_i)_{i \in [1..\ell] \setminus M_2^{PS}}), \mathbb{A}^{PS'})</math>.</li> </ul> <p><b>Proof<sup>PS</sup>(sk<sub>sig</sub><sup>PS</sup>, pk<sub>san</sub><sup>PS</sup>, m, σ<sup>PS</sup>, {(σ<sub>i</sub><sup>PS</sup>, m<sub>i</sub>)}).</b> If for any <math>(\sigma_i^{PS}, m_i)</math>, <math>Verify^{PS}(pk_{sig}^{PS}, pk_{san}^{PS}, m_i, \sigma_i^{PS}) = 0</math>, return <math>\perp</math>. If <math>Verify^{PS}(pk_{sig}^{PS}, pk_{san}^{PS}, m, \sigma^{PS}) = 0</math>, return <math>\perp</math>. Return <math>SS.Proof^{SS}(sk_{sig}^{SS}, pk_{san}^{SS}, m', \sigma^{SS}, \{(\sigma_i^{SS}, m_i')\})</math>, where <math>m' = (m, \sigma^{RS}, (\tau_i)_{1 \leq i \leq \ell_m}, \tau, pk_{sig}^{PS}, pk_{san}^{PS})</math> and each <math>m_i' = (m_i, \sigma_i^{RS}, (\tau_{i,j})_{1 \leq j \leq \ell_{m_i}}, \tau_i, pk_{sig}^{PS}, pk_{san}^{PS})</math>.</p> <p><b>Judge<sup>PS</sup>(pk<sub>sig</sub><sup>PS</sup>, pk<sub>san</sub><sup>PS</sup>, m, σ<sup>PS</sup>, π<sup>PS</sup>).</b> If <math>Verify^{PS}(pk_{sig}^{PS}, pk_{san}^{PS}, m, \sigma^{PS}) = 0</math>, return <math>\perp</math>. Return <math>SS.Judge^{SS}(pk_{sig}^{SS}, pk_{san}^{SS}, m', \sigma^{SS}, \pi^{PS})</math>, where <math>m' = (m, \sigma^{RS}, (\tau_i)_{1 \leq i \leq \ell_m}, \tau, pk_{sig}^{PS}, pk_{san}^{PS})</math>.</p>

### Construction 3: Our fully invisible PS scheme

The proof of the following Theorem is found in Appendix C.

**Theorem 3.** *If SS is unforgeable, immutable, private, transparent, signer-accountable, and sanitizer-accountable, RS is correct, unforgeable, immutable, private, transparent, and invisible, while SS' is unforgeable, immutable, private, transparent, and invisible, then the construction of a PS, given in Construction 3, is correct, unforgeable, private, transparent, immutable, signer-accountable, sanitizer-accountable, and (fully) invisible.*

**Implementation.** To show that the construction really is practical, we provide an evaluation of our implementation. To maintain comparability with existing measurements [4], we have chosen to use the same

parameters (as far as possible). Namely, our implementation is done in Java 10 and measured on an Intel i5-2400@3.10GHz with 16GiB of RAM. As the (aggregate) signature scheme, we implemented BGLS [9], while for  $\mathcal{H}$  we chose CS-encryption [48]. As the underlying groups we chose IAIK’s ECCelerate pairings library. In particular, the underlying pairing curves are “SNARK\_2” ( $\mathcal{H}$  uses only the first group). As the outer  $\mathcal{S}$ , we use the construction given by Gong et al. [30], but altered using the results given by Krenn et al. [58] and Beck et al. [43] to meet the stronger security definitions. Namely, we use the unique chameleon-hash [59] by Krenn et al. [58], with 2’048 Bit moduli, also hash the nonce, encrypt the original signature to the signer, and use unique signatures as done by Beck et al. [43], to achieve the stronger unforgeability, privacy and accountability definitions (but invisibility).

We did not implement any particular optimizations with the following exceptions: BGLS signatures allow a significant performance gain at verification, if the aggregate contains many signatures signed under the same public key [9]. As this is the case for the used RS, we chose to implement this optimization. An additional optimization we have implemented is that the signatures  $\sigma_i^{\mathcal{S}\mathcal{S}}$  for the invisible  $\mathcal{S}\mathcal{S}'$  are signed again using the outer  $\mathcal{S}$ , and thus are not required to be unique in this context — and can be replaced by standard unforgeable signatures.

As Krenn et al. [4], we evaluated our implementation with 32 blocks, whereas 25% were marked as admissible, and an additional 25% as redactable. For editing, 50% of the admissible blocks were sanitized and redacted. We omit proof generation and the judge, as they are simple database look-ups, and parameter generation as it is a one-time setup. The overall results are depicted in Figure 33a, Figure 33b, and Table 33c. They are based on 1’000 runs, while verification was measured after sanitization. Note, however, that we have measured  $\text{Edit}^{\text{PS}}$  without verifying each signature to see the actual runtime of the editing part, and not the verification one.

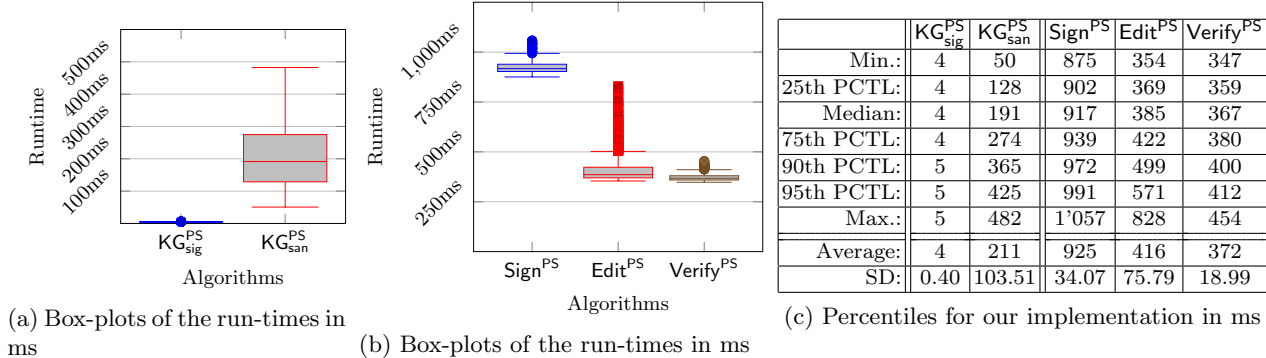


Fig. 33: Performance Evaluation Results

As it can easily be seen, even our not entirely optimized implementation is significantly faster than the original construction, and even offers stronger security guarantees; for comparison, in [4], signing, editing, and verification take about 28s, 8s, and 5s on average, respectively. Interestingly enough, our construction is even faster than the invisible  $\mathcal{S}\mathcal{S}$  introduced by Beck et al. [43], while, as already clarified by Krenn et al. [4], a PS can simply mimic a  $\mathcal{S}\mathcal{S}$  by prohibiting redactions.

## 6 Conclusion

We have strengthened the state-of-the-art definition of invisibility for Protean Signatures (PS) to also account for the redactable parts. In particular, using our new notion an outsider can neither decide which parts are

redactable nor which parts are editable. To achieve this, we introduced the new notions of invisible redactable signatures (RS), non-accountable invisible sanitizable signature schemes (SS), and a novel framework for aggregate signatures which explicitly allow for de-aggregation of signatures. Using those primitives, our resulting provably secure construction becomes practically efficient, proven by our prototypical implementation.

**Acknowledgements.** The projects leading to this work have received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 783119 (SECRETAS), No 780315 (SEMIOTICS), and No 830929 (CyberSec4Europe) respectively.

## References

1. Shafi Goldwasser, Silvio Micali, and Ronald L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM J. Comput.*, 17(2):281–308, 1988.
2. Arne Bilzhaue, 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.
3. 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.
4. 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.
5. Daniel Slamanig and Stefan Rass. Generalizations and extensions of redactable signatures with applications to electronic healthcare. In Bart De Decker and Ingrid Schaumüller-Bichl, editors, *Communications and Multimedia Security, 11th IFIP TC 6/TC 11 International Conference, CMS 2010, Linz, Austria, May 31 - June 2, 2010. Proceedings*, volume 6109 of *Lecture Notes in Computer Science*, pages 201–213. Springer, 2010.
6. Zhen Yu Wu, Chih-Wen Hsueh, Cheng-Yu Tsai, Feipei Lai, Hung-Chang Lee, and Yu-Fang Chung. Redactable signatures for signed CDA documents. *J. Medical Systems*, 36(3):1795–1808, 2012.
7. 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.
8. 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.
9. Dan Boneh, Craig Gentry, Ben Lynn, and Hovav Shacham. Aggregate and verifiably encrypted signatures from bilinear maps. In Eli Biham, editor, *Advances in Cryptology - EUROCRYPT 2003, International Conference on the Theory and Applications of Cryptographic Techniques, Warsaw, Poland, May 4-8, 2003, Proceedings*, volume 2656 of *Lecture Notes in Computer Science*, pages 416–432. Springer, 2003.
10. 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.
11. 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](http://www.prismacloud.eu), 2015.
12. 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.
13. Dan Boneh and David Mandell Freeman. Homomorphic signatures for polynomial functions. In Kenneth G. Paterson, editor, *Advances in Cryptology - EUROCRYPT 2011 - 30th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tallinn, Estonia, May 15-19, 2011. Proceedings*, volume 6632 of *Lecture Notes in Computer Science*, pages 149–168. Springer, 2011.
14. Giulia Traverso, Denise Demirel, and Johannes A. Buchmann. *Homomorphic Signature Schemes - A Survey*. Springer Briefs in Computer Science. Springer, 2016.

15. Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In Hugo Krawczyk, editor, *Public-Key Cryptography - PKC 2014 - 17th International Conference on Practice and Theory in Public-Key Cryptography, Buenos Aires, Argentina, March 26-28, 2014. Proceedings*, volume 8383 of *Lecture Notes in Computer Science*, pages 501–519. Springer, 2014.
16. Rotem Tsabary. An equivalence between attribute-based signatures and homomorphic signatures, and new constructions for both. In Yael Kalai and Leonid Reyzin, editors, *Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part II*, volume 10678 of *Lecture Notes in Computer Science*, pages 489–518. Springer, 2017.
17. Christina Brzuska, Heike Busch, Özgür Dagdelen, Marc Fischlin, Martin Franz, Stefan Katzenbeisser, Mark Manulis, Cristina Onete, Andreas Peter, Bertram Poettering, and Dominique Schröder. Redactable signatures for tree-structured data: Definitions and constructions. In Zhou and Yung [60], pages 87–104.
18. Ashish Kundu and Elisa Bertino. Privacy-preserving authentication of trees and graphs. *Int. J. Inf. Sec.*, 12(6):467–494, 2013.
19. Kai Samelin, Henrich Christopher Pöhls, Arne Bilzhause, Joachim Posegga, and Hermann de Meer. On structural signatures for tree data structures. In Feng Bao, Pierangela Samarati, and Jianying Zhou, editors, *Applied Cryptography and Network Security - 10th International Conference, ACNS 2012, Singapore, June 26-29, 2012. Proceedings*, volume 7341 of *Lecture Notes in Computer Science*, pages 171–187. Springer, 2012.
20. Henrich Christopher Pöhls and Kai Samelin. Accountable redactable signatures. In *10th International Conference on Availability, Reliability and Security, ARES 2015, Toulouse, France, August 24-27, 2015*, pages 60–69. IEEE Computer Society, 2015.
21. 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.
22. Kai Samelin, Henrich Christopher Pöhls, Arne Bilzhause, Joachim Posegga, and Hermann de Meer. Redactable signatures for independent removal of structure and content. In Mark Dermot Ryan, Ben Smyth, and Guilin Wang, editors, *Information Security Practice and Experience - 8th International Conference, ISPEC 2012, Hangzhou, China, April 9-12, 2012. Proceedings*, volume 7232 of *Lecture Notes in Computer Science*, pages 17–33. Springer, 2012.
23. Stuart Haber, Yasuo Hatano, Yoshinori Honda, William G. Horne, Kunihiko Miyazaki, Tomas Sander, Satoru Tezoku, and Danfeng Yao. Efficient signature schemes supporting redaction, pseudonymization, and data deidentification. In Masayuki Abe and Virgil D. Gligor, editors, *Proceedings of the 2008 ACM Symposium on Information, Computer and Communications Security, ASIACCS 2008, Tokyo, Japan, March 18-20, 2008*, pages 353–362. ACM, 2008.
24. Tetsuya Izu, Noboru Kunihiro, Kazuo Ohta, Makoto Sano, and Masahiko Takenaka. Sanitizable and deletable signature. In Kyo-II Chung, Kiwook Sohn, and Moti Yung, editors, *Information Security Applications, 9th International Workshop, WISA 2008, Jeju Island, Korea, September 23-25, 2008, Revised Selected Papers*, volume 5379 of *Lecture Notes in Computer Science*, pages 130–144. Springer, 2008.
25. Tetsuya Izu, Noboru Kunihiro, Kazuo Ohta, Makoto Sano, and Masahiko Takenaka. Yet another sanitizable signature from bilinear maps. In *Proceedings of the The Forth International Conference on Availability, Reliability and Security, ARES 2009, March 16-19, 2009, Fukuoka, Japan*, pages 941–946. IEEE Computer Society, 2009.
26. Kunihiko Miyazaki, Mitsuru Iwamura, Tsutomu Matsumoto, Ryōichi Sasaki, Hiroshi Yoshiura, Satoru Tezuka, and Hideki Imai. Digitally signed document sanitizing scheme with disclosure condition control. *IEICE Transactions*, 88-A(1):239–246, 2005.
27. Henrich Christopher Pöhls, Kai Samelin, and Joachim Posegga. Sanitizable signatures in XML signature - performance, mixing properties, and revisiting the property of transparency. In Javier López and Gene Tsudik, editors, *Applied Cryptography and Network Security - 9th International Conference, ACNS 2011, Nerja, Spain, June 7-10, 2011. Proceedings*, volume 6715 of *Lecture Notes in Computer Science*, pages 166–182, 2011.
28. 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.
29. Christina Brzuska, Marc Fischlin, Tobias Freudenreich, Anja Lehmann, Marcus Page, Jakob Schelbert, Dominique Schröder, and Florian Volk. Security of sanitizable signatures revisited. In 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*, pages 317–336. Springer, 2009.
30. 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.
31. 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 155 of *LNI*, pages 117–128. GI, 2009.
32. 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.
  33. 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.
  34. 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.
  35. 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.
  36. Russell W. F. Lai, Tao Zhang, Sherman S. M. Chow, and Dominique Schröder. Efficient sanitizable signatures without random oracles. In Ioannis G. Askoxylakis, Sotiris Ioannidis, Sokratis K. Katsikas, and Catherine A. Meadows, editors, *Computer Security - ESORICS 2016 - 21st European Symposium on Research in Computer Security, Heraklion, Greece, September 26-30, 2016, Proceedings, Part I*, volume 9878 of *Lecture Notes in Computer Science*, pages 363–380. Springer, 2016.
  37. 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.
  38. Dae Hyun Yum, Jae Woo Seo, and Pil Joong Lee. Trapdoor sanitizable signatures made easy. In Zhou and Yung [60], pages 53–68.
  39. 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.
  40. 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.
  41. 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.
  42. Marek Klonowski and Anna Lauks. Extended sanitizable signatures. In Min Surp Rhee and Byoungcheon Lee, editors, *Information Security and Cryptology - ICISC 2006, 9th International Conference, Busan, Korea, November 30 - December 1, 2006, Proceedings*, volume 4296 of *Lecture Notes in Computer Science*, pages 343–355. Springer, 2006.
  43. 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.
  44. 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.
  45. 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.
  46. 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.
47. Esha Ghosh, Olga Ohrimenko, and Roberto Tamassia. Zero-knowledge authenticated order queries and order statistics on a list. In Tal Malkin, Vladimir Kolesnikov, Allison Bishop Lewko, and Michalis Polychronakis, editors, *Applied Cryptography and Network Security - 13th International Conference, ACNS 2015, New York, NY, USA, June 2-5, 2015, Revised Selected Papers*, volume 9092 of *Lecture Notes in Computer Science*, pages 149–171. Springer, 2015.
  48. 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.
  49. Ronald L. Rivest, Adi Shamir, and Leonard M. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126, 1978.
  50. Einar Mykletun, Maithili Narasimha, and Gene Tsudik. Signature bouquets: Immutability for aggregated/condensed signatures. In Pierangela Samarati, Peter Y. A. Ryan, Dieter Gollmann, and Refik Molva, editors, *Computer Security - ESORICS 2004, 9th European Symposium on Research Computer Security, Sophia Antipolis, France, September 13-15, 2004, Proceedings*, volume 3193 of *Lecture Notes in Computer Science*, pages 160–176. Springer, 2004.
  51. Mihir Bellare, Chanathip Namprempre, and Gregory Neven. Unrestricted aggregate signatures. In Lars Arge, Christian Cachin, Tomasz Jurdzinski, and Andrzej Tarlecki, editors, *Automata, Languages and Programming, 34th International Colloquium, ICALP 2007, Wroclaw, Poland, July 9-13, 2007, Proceedings*, volume 4596 of *Lecture Notes in Computer Science*, pages 411–422. Springer, 2007.
  52. Veronika Kuchta and Mark Manulis. Unique aggregate signatures with applications to distributed verifiable random functions. In Michel Abdalla, Cristina Nita-Rotaru, and Ricardo Dahab, editors, *Cryptology and Network Security - 12th International Conference, CANS 2013, Paraty, Brazil, November 20-22, 2013, Proceedings*, volume 8257 of *Lecture Notes in Computer Science*, pages 251–270. Springer, 2013.
  53. Jean-Sébastien Coron and David Naccache. Boneh et al.’s k-element aggregate extraction assumption is equivalent to the diffie-hellman assumption. In Chi-Sung Lai, editor, *Advances in Cryptology - ASIACRYPT 2003, 9th International Conference on the Theory and Application of Cryptology and Information Security, Taipei, Taiwan, November 30 - December 4, 2003, Proceedings*, volume 2894 of *Lecture Notes in Computer Science*, pages 392–397. Springer, 2003.
  54. David Derler, Stephan Krenn, and Daniel Slamanig. Signer-anonymous designated-verifier redactable signatures for cloud-based data sharing. In Sara Foresti and Giuseppe Persiano, editors, *Cryptology and Network Security - 15th International Conference, CANS 2016, Milan, Italy, November 14-16, 2016, Proceedings*, volume 10052 of *Lecture Notes in Computer Science*, pages 211–227, 2016.
  55. Arne Bilzhausen, Manuel Huber, Henrich C. Pöhls, and Kai Samelin. Cryptographically enforced four-eyes principle. In *11th International Conference on Availability, Reliability and Security, ARES 2016, Salzburg, Austria, August 31 - September 2, 2016*, pages 760–767. IEEE Computer Society, 2016.
  56. Kunihiko Miyazaki, Goichiro Hanaoka, and Hideki Imai. Digitally signed document sanitizing scheme based on bilinear maps. In Ferng-Ching Lin, Der-Tsai Lee, Bao-Shuh Paul Lin, Shihuhpyng Shieh, and Sushil Jajodia, editors, *Proceedings of the 2006 ACM Symposium on Information, Computer and Communications Security, ASIACCS 2006, Taipei, Taiwan, March 21-24, 2006*, pages 343–354. ACM, 2006.
  57. 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.
  58. Stephan Krenn, Henrich C. Pöhls, Kai Samelin, and Daniel Slamanig. Chameleon-hashes with dual long-term trapdoors and their applications. In Antoine Joux, Abderrahmane Nitaaj, 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.
  59. 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.
  60. 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*, 2010.

## A Proof of Theorem 1

We now provide the proof of security for Construction 1.

*Proof.* Correctness follows from inspection. Each security property is proven on its own. However, we already keep all queries and answers to and from the oracle. This does not change the view of the adversary.

*Unforgeability.* To prove that our scheme is unforgeable, we use a sequence of games:

**Game 0:** The original unforgeability game.

**Game 1:** We now abort, if we draw a tag twice.

*Transition - Game 0  $\rightarrow$  Game 1:* Due to the birthday paradox, this can only happen with negligible probability.  $|\Pr[S_0] - \Pr[S_1]| \leq \frac{q_s^2}{2^\lambda}$  follows, where  $q_s$  is the number of tags drawn. Note, this also means that no message under  $\text{pk}_{\text{sig}}^{\text{RS}}$  is signed twice.

**Game 2:** We now abort, if the adversary was able to generate a signature for  $\text{pk}_{\text{sig}}^{\text{RS}}$  which protects a message not signed by the signer.

*Transition - Game 1  $\rightarrow$  Game 2:* In this case,  $\mathcal{A}$  returns  $(m^*, (\sigma_a, c, \text{pk}_{\Sigma'}, (\tau_i)_{i \in [0..l]}))$  for which  $\sigma_a$  contains a signature not explicitly generated by the signer. We can use this forgery to construct an adversary  $\mathcal{B}$  which breaks the unforgeability of the underlying  $\Sigma$ . The reduction works as follows. It receives the parameters and the public  $\text{pk}$  to forge and directly embeds them into the values given to the adversary  $\mathcal{A}$ . Queries to the signing oracle are answered honestly; all inner signatures to be generated are delegated to  $\mathcal{B}$ 's own oracle (with the exception of the ephemeral signature, which can be generated honestly). Then, as by assumption  $\sigma_a$  protects at least one message which was not signed by the signing oracle,  $\mathcal{B}$  can return  $(\{\mathcal{S} \cup (\{\text{pk}_{\Sigma'}, (c, \tau_0, \text{pk}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'})\})\}, m^*, \sigma_a)$ , where  $\mathcal{S}$  is the set of all messages checked for the underlying aggregate w.r.t. to  $\text{pk}$  as derivable from the construction, but a single forged message  $m^*$  (which can easily be spotted) which can be arbitrarily chosen from the set of forged messages.  $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{aggsig-unf}}(\lambda)$  follows, as, by assumption, at least  $m^*$  must be fresh. Note, this also covers the case of changes public keys, tags and “mix-and-match” attacks by merging multiple signatures, as we already ruled out tag-collisions, while all public keys are always signed as well, and all signed values are bound to  $\tau_0$ .

**Game 3:** We now abort, if the adversary was able to generate a new signature for  $\text{pk}_{\Sigma'}$  for the challenge  $\text{pk}_{\text{san}}^{\text{RS}}$ , which was never generated by the challenger.

*Transition - Game 2  $\rightarrow$  Game 3:* In this case,  $\mathcal{A}$  returns  $(m^*, (\sigma_a, c, \text{pk}_{\Sigma'}, (\tau_i)_{i \in [0..l]}))$  for which  $\sigma_a$  contains a signature under  $\text{pk}_{\Sigma'}$  not explicitly generated by the signer. We can use this forgery to construct an adversary  $\mathcal{B}$  which breaks the unforgeability of the underlying  $\Sigma$ . The reduction works as follows. It receives the parameters and the public  $\text{pk}$  to forge and directly embeds them into the values given to the adversary  $\mathcal{A}$ . Queries to the signing oracle are answered honestly; all inner signatures to be generated are delegated to  $\mathcal{B}$ 's own oracle (with the exception of the ephemeral signature, which can be generated honestly). Then, as by assumption  $\sigma_a$  protects at least one message which was not signed by the signing oracle,  $\mathcal{B}$  can return  $(\{\mathcal{S} \cup (\{\text{pk}_{\Sigma'}, (c, \tau_0, \text{pk}, \text{pk}^*, \text{pk}_{\Sigma'})\})\}, m^*, \sigma_a)$ , where  $\mathcal{S}$  is the set of all messages checked for the underlying aggregate w.r.t. to  $\text{pk}$  as derivable from the construction, but a single forged message  $m^*$  (which can easily be spotted) which can be arbitrarily chosen from the set of forged messages.  $|\Pr[S_2] - \Pr[S_3]| \leq q_s \nu_{\text{aggsig-unf}}(\lambda)$ , where  $q_s$  is the number of signatures generated, follows, as the reduction  $\mathcal{B}$  has to guess where the adversary  $\mathcal{A}$  forges a signature. Note, this also covers the case of changes public keys, tags and “mix-and-match” attacks by merging multiple signatures, as we already ruled out tag-collisions, while all public keys are always signed as well, and all signed values are bound to  $\tau_0$ .

**Game 3:** As Game 2, but we abort, if the adversary was able to redact a non-redactable block.

*Transition - Game 2  $\rightarrow$  Game 3:* In this case, the adversary  $\mathcal{A}$  was able to remove a signature from the aggregate, which should not be possible. Thus, this means that the adversary  $\mathcal{A}$  was able to break the signature scheme. We show this by construction of an adversary  $\mathcal{B}$  which uses  $\mathcal{A}$  to break the no-extraction notion of the used  $\Sigma$ . The reduction works as follows. It receives the parameters and the public  $\text{pk}$  to forge. The public parameters are embedded honestly; all other values are generated as in the prior game. Next,  $\mathcal{B}$  draws a random index  $i \leftarrow_r [1..q_s]$ , where  $q_s$  is an upper bound on the number of queries to the signing oracle. Then, every  $j$ th query to the signing oracle, where  $i \neq j$ , is answered honestly. On the  $i$ th query, however,  $\mathcal{B}$  embeds the challenge  $\text{pk}$  and uses its own signing oracle to generate the signature  $\sigma_c$ . The other signatures can be generated honestly. Then, as by assumption  $\sigma_a$  contains a signature on the

string  $(c, \tau_0, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'}^{\text{RS}})$ ,  $\mathcal{B}$  can use to de-aggregate all other signatures from  $\sigma_a$  to obtain  $\sigma_c$  using its honestly generated  $\text{sk}_{\text{sig}}^{\text{RS}}$  and return  $(\emptyset, \emptyset, (c, \tau_0, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'}^{\text{RS}}), \sigma_c)$  as its own forgery. In the case that  $\text{pk}_{\Sigma'}^{\text{RS}} \neq \text{pk}$ ,  $\mathcal{B}$  aborts. Note, we have already ruled out forgeries of never signed messages, while each message signed is fresh due to the tags.  $|\Pr[S_2] - \Pr[S_3]| \leq \nu_{\text{aggsig-noExt}}(\lambda)$  follows.

**Game 4:** As Game 3, but we abort, if the adversary was able to generate a new aggregate  $\sigma_a$  which protects a set of messages returned by the challenger.

*Transition - Game 3  $\rightarrow$  Game 4:* In this case, the adversary  $\mathcal{A}$  was able to break the uniqueness of the underlying signature scheme. The reduction works as follows.  $\mathcal{B}$  receives the public parameters from its own challenger and embeds them accordingly. All other values are generated honestly and given to the adversary  $\mathcal{A}$ . Then, once the adversary  $\mathcal{A}$  outputs  $(m^*, \sigma^{\text{RS}^*})$ , and by assumption,  $\sigma_a^*$  was never seen, but the messages protected by an honestly generated signature  $\sigma_a$ ,  $\mathcal{B}$  can directly return  $(\mathcal{S}, \sigma_a, \sigma_a^*)$ , where  $\mathcal{S}$  is the set of public key/messages protected in the aggregate which can be derived as in the construction.  $|\Pr[S_3] - \Pr[S_4]| \leq \nu_{\text{aggsig-unique}}(\lambda)$  follows.

**Game 5:** As Game 4, but we abort, if the adversary was able to exchange a signature on  $(c, \tau_0, \text{pk}_{\text{sig}}^{\text{RS}}, \text{pk}_{\text{san}}^{\text{RS}}, \text{pk}_{\Sigma'}^{\text{RS}})$  on the aggregate from some already seen aggregate.

*Transition - Game 4  $\rightarrow$  Game 5:* If an adversary outputs  $(m^*, \sigma^{\text{RS}^*})$ , where  $\sigma^{\text{RS}^*} = (\sigma_a, c, \text{pk}_{\Sigma'}^{\text{RS}}, (\tau_i)_{i \in [0..l]})$ , meeting the above winning conditions, we can construct an adversary  $\mathcal{B}$  which breaks the No-Extraction property of the underlying  $\Sigma$ . It proceeds as follows. It first receives the public parameters. It then queries its own challenge oracle to obtain a long-term public-key  $\text{pk}$ . Both are embedded honestly; other values are generated honestly. For every signing query,  $\mathcal{B}$  first requests an additional key  $\text{pk}'$  if  $\text{pk}_{\text{san}}^{\text{RS}}$  is the challenge one. If this is not the case,  $\mathcal{B}$  generates one honestly. It then proceeds as in the signing algorithm, but requests a full aggregate on all signatures generated under the keys generated. If  $\text{pk}_{\text{san}}^{\text{RS}}$  is not the challenge one, all signatures which are related to redacting are queried to the challenge oracle and embedded for the adversary. If  $\text{pk}_{\text{san}}^{\text{RS}}$  is the challenge one, it also gets all signatures, but the one for  $\sigma_0$  and  $\sigma_c$ . For redaction, if  $\text{pk}_{\text{san}}^{\text{RS}}$  and  $\text{pk}_{\text{sig}}^{\text{RS}}$  are the challenge ones,  $\mathcal{B}$  requests a complete new aggregate signature on the redacted message. Then, whenever  $\mathcal{A}$  outputs  $\sigma^{\text{RS}^*}$  meeting the winning requirements,  $\mathcal{B}$  can simply output  $(\mathcal{S}, \sigma_a)$ , where  $\mathcal{S}$  is as in the verification algorithm.  $|\Pr[S_4] - \Pr[S_5]| \leq \nu_{\text{aggsig-noExt}}(\lambda)$  follows.

Now, the adversary can no longer win the unforgeability game. Moreover, each hop changes the view of the adversary only negligibly, concluding the proof.

*Immutability.* To prove that our scheme is immutable, we use a sequence of games:

**Game 0:** The original immutability game.

**Game 1:** We now abort, if we draw a tag twice.

*Transition - Game 0  $\rightarrow$  Game 1:* Due to the birthday paradox, this can only happen with negligible probability.  $|\Pr[S_0] - \Pr[S_1]| \leq \frac{q_s^2}{2\lambda}$  follows, where  $q_s$  is the number of key pairs generated. Note, this also means that no message under  $\text{pk}_{\text{sig}}^{\text{RS}}$  is signed twice.

**Game 2:** We now abort, if the adversary was able to generate a signature for  $\text{pk}_{\text{sig}}^{\text{RS}}$  which protects a message not signed by the signer.

*Transition - Game 1  $\rightarrow$  Game 2:* In this case,  $\mathcal{A}$  returns  $(m^*, (\sigma_a, c, \text{pk}_{\Sigma'}^{\text{RS}}, (\tau_i)_{i \in [0..l]}), \text{pk}^*)$  for which  $\sigma_a$  contains a signature not explicitly generated by the signer. We can use this forgery to construct an adversary  $\mathcal{B}$  which breaks the unforgeability of the underlying  $\Sigma$ . The reduction works as follows. It receives the parameters and the public  $\text{pk}$  to forge and directly embeds them into the values given to the adversary  $\mathcal{A}$ . Queries to the signing oracle are answered honestly; all inner signatures to be generated are delegated to  $\mathcal{B}$ 's own oracle (with the exception of the ephemeral signature, which can be generated honestly). Then, as by assumption  $\sigma_a$  protects at least one message which was not signed by the signing oracle,  $\mathcal{B}$  can return  $(\{\mathcal{S} \cup \{\text{pk}_{\Sigma'}^{\text{RS}}, (c, \tau_0, \text{pk}, \text{pk}^*, \text{pk}_{\Sigma'}^{\text{RS}})\}\}, m^*, \sigma_a)$ , where  $\mathcal{S}$  is the set of all messages checked for the underlying aggregate w.r.t. to  $\text{pk}$  as derivable from the construction, but a single forged message  $m^*$  (which can easily be

spotted) which can be arbitrarily chosen from the set of forged messages.  $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{aggsig-uf}}(\lambda)$  follows, as, by assumption, at least  $m'^*$  must be fresh. Note, this also covers the case of changed public keys, tags and “mix-and-match” attacks by merging multiple signatures, as we already ruled out tag-collisions, while all public keys are always signed as well, and all signed values are bound to  $\tau_0$ .

**Game 3:** As Game 2, but we abort, if the adversary was able to redact a non-redactable block.

*Transition - Game 2  $\rightarrow$  Game 3:* In this case, the adversary  $\mathcal{A}$  was able to remove a signature from the aggregate, which should not be possible. Thus, this means that the adversary  $\mathcal{A}$  was able to break the signature scheme. We show this by construction of an adversary  $\mathcal{B}$  which uses  $\mathcal{A}$  to break the no-extraction notion of the used  $\Sigma$ . The reduction works as follows. It receives the parameters and the public  $\text{pk}$  to forge and directly embeds them into the values given to the adversary  $\mathcal{A}$ . Queries to the signing oracle are answered honestly; all inner signatures to be generated are delegated to  $\mathcal{B}$ 's own oracle as a bulk (with the exception of the ephemeral signature, which can be generated honestly). Then, as by assumption  $\sigma_a$  protects less messages as given as aggregate by the signing oracle,  $\mathcal{B}$  can return  $(\mathcal{S}, \sigma'_a)$ , where  $\mathcal{S}$  is the set of all messages checked for the underlying aggregate w.r.t. to  $\text{pk}$  as derivable from the construction. The signature  $\sigma_c$  can be removed using  $\text{sk}_{\Sigma'}$ , generating  $\sigma'_a$ , as we have already ruled out forgeries of never signed messages, while each message signed is fresh due to the tags.  $|\Pr[S_2] - \Pr[S_3]| \leq \nu_{\text{aggsig-noExt}}(\lambda)$  follows.

Now, the adversary can no longer win the immutability game. Moreover, each hop changes the view of the adversary only negligibly, concluding the proof.

*Privacy.* To prove that our scheme is private, we use a sequence of games:

**Game 0:** The original privacy game in the case  $b = 0$ .

**Game 1:** We now switch to  $b = 1$ .

*Transition - Game 0  $\rightarrow$  Game 1:* As the signatures are distributed exactly the same, our scheme is private in an information theoretical sense. Thus,  $|\Pr[S_0] - \Pr[S_1]| = 0$  follows.

*Transparency.* To prove that our scheme is transparent, we use a sequence of games:

**Game 0:** The original transparency game in the case  $b = 0$ .

**Game 1:** We now switch to  $b = 1$ .

*Transition - Game 0  $\rightarrow$  Game 1:* As the signatures are distributed exactly the same, our scheme is transparent in an information theoretical sense. Thus,  $|\Pr[S_0] - \Pr[S_1]| = 0$  follows.

*Invisibility.* To prove that our scheme is invisible, we use a sequence of games:

**Game 0:** The original invisibility game.

**Game 1:** We now abort, if the adversary queries some  $(\sigma^{\text{RS}}, m)$  for the challenge public keys which verifies, but was never returned by either LoRADM or  $\text{Red}^{\text{RS}'}$ .

*Transition - Game 0  $\rightarrow$  Game 1:* Given this adversary  $\mathcal{A}$ , we can construct an adversary  $\mathcal{B}$  which breaks the unforgeability of the RS. The reduction works as follows. First, it draws a random bit  $b \leftarrow_r \{0, 1\}$ . Then, it receives  $\text{pk}_{\text{sig}}^{\text{RS}}$  and  $\text{pk}_{\text{san}}^{\text{RS}}$  (along with the parameters) and then passes those keys to the adversary. Every redaction query is done using the  $\text{Red}^{\text{RS}'}$  provided (imposing the limitation the invisibility game gives). Likewise, queries to LoRADM are answered by using the  $\text{Sign}^{\text{RS}'}$  provided, but using  $\mathbb{A}_b^{\text{RS}}$  in the case the challenge  $\text{pk}_{\text{san}}^{\text{RS}}$  is queried. Then, whenever  $\mathcal{A}$  queries  $(\sigma^{\text{RS}}, m)$  for the challenge  $\text{pk}_{\text{san}}^{\text{RS}}$ ,  $\mathcal{B}$  can simply return  $(\sigma^{\text{RS}}, m)$  as its own forgery, as, by assumption,  $(\sigma^{\text{RS}}, m)$  was not seen before. Thus,  $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{rss-uf}}(\lambda)$  follows.

**Game 2:** We now start replace *each* ciphertext generated for the challenge  $\text{pk}_{\text{san}}^{\text{RS}}$  with an encryption of 0 (with the appropriate length).

*Transition - Game 1  $\rightarrow$  Game 2:* Assume that the adversary can distinguish this replacement with a non-negligible probability. We can then construct a reduction  $\mathcal{B}$  which uses  $\mathcal{A}$  to break the IND-CCA2 security of the underlying encryption scheme. The reduction works via a series of hybrids. Our reduction  $\mathcal{B}$  proceeds as follows. It receives  $\text{pk}$  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 2.i is also output by  $\mathcal{B}$ . Note, decryption queries for ciphertexts generated by the adversary can be queried to decryption oracle provided; the content for all other ciphertexts are known and thus the ciphertexts do not need to be decrypted at all. Thus,  $|\Pr[S_1] - \Pr[S_2]| \leq q\nu_{\text{ind-cca2}}(\lambda)$  follows, where  $q$  is the number of ciphertexts generated.

As now the game is independent of the bit  $b$ , invisibility is proven.

## B Proof of Theorem 2

We now provide the proof of security for Construction 2.

*Proof.* Correctness follows from inspection. Each security property is proven on its own. However, we already keep all queries and answers to and from the oracle. This does not change the view of the adversary.

*Unforgeability.* To prove that our scheme is unforgeable, we use a sequence of games:

**Game 0:** The original unforgeability game.

**Game 1:** We now abort, if we draw an  $x$  twice.

*Transition - Game 0  $\rightarrow$  Game 1:* Due to the birthday paradox, this can only happen with negligible probability.  $|\Pr[S_0] - \Pr[S_1]| \leq \frac{q_s^2}{2\lambda}$  follows, where  $q_s$  is the number of key pairs generated.

**Game 2:** We now start replace *each* ciphertext generated for the challenge  $\text{pk}_{\text{san}}^{\text{SS}}$  with an encryption of 0 (with the appropriate length).

*Transition - Game 1  $\rightarrow$  Game 2:* Assume that the adversary can distinguish this replacement with a non-negligible probability. We can then construct a reduction  $\mathcal{B}$  which uses  $\mathcal{A}$  to break the IND-CCA2 security of the underlying encryption scheme. The reduction works via a series of hybrids. Our reduction  $\mathcal{B}$  proceeds as follows. It receives  $\text{pk}$  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 2.i is also output by  $\mathcal{B}$ . Note, decryption queries for ciphertexts generated by the adversary can be queried to decryption oracle provided; the content for all other ciphertexts are known and thus the ciphertexts do not need to be decrypted at all. Thus,  $|\Pr[S_1] - \Pr[S_2]| \leq q\nu_{\text{ind-cca2}}(\lambda)$  follows, where  $q$  is the number of ciphertexts generated.

**Game 3:** We now replace all  $r_i$  with a purely random value  $r_i \leftarrow_r \{0, 1\}^\lambda$ .

*Transition - Game 2  $\rightarrow$  Game 3:* An adversary distinguishing this replacement can be turned into an adversary against the pseudo-randomness of the PRF. We prove this via a series of hybrids. Let Game 3.0 the same as Game 2. In Game 3.i  $\mathcal{B}$  uses its  $\text{Eval}'_{\text{PRF}}$  oracle to generate the random coins. All other values are generated as in prior game. Then, whatever  $\mathcal{A}$  outputs, is also output by  $\mathcal{B}$ .  $|\Pr[S_2] - \Pr[S_3]| \leq q_s\nu_{\text{prf-pr}}(\lambda)$  follows, where  $q_s$  is the number of calls to the signature-generation oracle.

**Game 4:** We now abort, if we draw some random coins twice.

*Transition - Game 3  $\rightarrow$  Game 4:* Due to the birthday paradox, this can only happen with negligible probability.  $|\Pr[S_3] - \Pr[S_4]| \leq \frac{q_s^2}{2\lambda}$  follows, where  $q_s$  is the number of key pairs generated.

**Game 5:** We now abort, if the adversary was able to generate a signature on a string of public keys and ciphertexts not signed.

*Transition - Game 4  $\rightarrow$  Game 5:* In this case, the adversary  $\mathcal{A}$  was able to generate a signature  $\sigma_s$  (contained in  $\sigma^{\text{SS}^*}$ ) on  $(\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}, (\text{pk}_i)_{[1..\ell]}, c)$  which was never generated by the signer. We can use this to construct a reduction  $\mathcal{B}$  which forges a signature of the underlying  $\Sigma$ . The reduction works as follows. It receives the  $\text{pk}$  (and the corresponding parameters) to forge, and embeds it accordingly into the parameters/public key. Everything else is generated honestly. For every signature generated,  $\mathcal{B}$  queries its signature-generation oracle; this signature is then embedded in the response. Finally, once  $\mathcal{A}$  outputs its forgery,  $\mathcal{B}$  can return  $(\emptyset, \emptyset, (\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}, (\text{pk}_i)_{[1..\ell]}, c), \sigma_s)$  as its own forgery.  $|\Pr[S_4] - \Pr[S_5]| \leq \nu_{\text{aggsig-uf}}(\lambda)$  follows.

**Game 6:** We now abort, if the adversary was able to generate a new signature  $\sigma'_s$  on a string of public keys already signed.

*Transition - Game 5  $\rightarrow$  Game 6:* In this case, the adversary  $\mathcal{A}$  was able to generate a new signature  $\sigma'_s$  (contained in  $\sigma^{\text{SS}^*}$ ) on  $(\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}, (\text{pk}_i)_{[1..\ell]}, c)$  which was never generated by the signer, but  $\sigma_s$  was. We can use this to construct a reduction  $\mathcal{B}$  which breaks the uniqueness of the underlying  $\Sigma$ . The reduction works as follows. It receives the corresponding parameters of the  $\Sigma$  to forge, and embeds it accordingly into the parameters. Everything else is generated honestly. Finally, once  $\mathcal{A}$  outputs its forgery,  $\mathcal{B}$  can return  $(\{(\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}, (\text{pk}_i)_{[1..\ell]}, c)\}, \sigma_s, \sigma'_s)$  as its own forgery.  $|\Pr[S_5] - \Pr[S_6]| \leq \nu_{\text{aggsig-unique}}(\lambda)$  follows.

**Game 7:** We now abort, if the adversary was able to generate a new inner signature  $\sigma'_i$  on a string signed before.

*Transition - Game 6  $\rightarrow$  Game 7:* In this case, the adversary  $\mathcal{A}$  was able to generate a new signature  $\sigma'_i$  (contained in  $\sigma^{\text{SS}^*}$ ) on on some string  $((\text{pk}_i, m^i)_{[1..\ell]}, c, \sigma_s, \text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}})$  which was signed before, where  $\sigma_i \neq \sigma'_i$  holds, but  $\sigma_s$  was signed. We can use this to construct a reduction  $\mathcal{B}$  which breaks the uniqueness of the underlying  $\Sigma$ . The reduction works as follows. It receives the corresponding parameters of the  $\Sigma$  to forge, and embeds it accordingly into the parameters. Everything else is generated honestly. Finally, once  $\mathcal{A}$  outputs its forgery,  $\mathcal{B}$  can return  $(\{(\text{pk}_i, ((\text{pk}_i, m^i)_{[1..\ell]}, c, \sigma_s, \text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}}))\}, \sigma_i, \sigma'_i)$  as its own forgery, where  $\text{pk}_i$  is the corresponding public key.  $|\Pr[S_6] - \Pr[S_7]| \leq \nu_{\text{aggsig-unique}}(\lambda)$  follows.

**Game 8:** We now abort, if the adversary outputs a validating  $(m^*, \sigma^{\text{SS}^*})$ , where  $\sigma^{\text{SS}^*} = (c, \sigma_s, (\text{pk}_i, \sigma_i)_{i \in [1..\ell]})$ , where  $m^*$  was never returned from any query to the signing or sanitization oracle.

*Transition - Game 7  $\rightarrow$  Game 8:* In this case, there must be a block  $m^{i^*}$  which was changed, but the adversary never saw a signature for that block. We can use this to break the unforgeability of the underlying signature scheme. Our reduction  $\mathcal{B}$  works as follows. Let  $q_s$  be an upper bound on the number of signature key pairs created. First,  $\mathcal{B}$  randomly selects an index  $i \leftarrow_r [1..q_s]$ . It receives the  $\text{pk}$  (and the corresponding parameters) to forge. The parameters are embedded honestly. Everything else is generated honestly. Then, once the  $i$ th inner signature is generated,  $\mathcal{B}$  embeds  $\text{pk}$  and uses its own signature-generation oracle to receive the corresponding signature. The result is embedded honestly. The same is true for sanitization: for every change,  $\mathcal{B}$  queries the signature-generation oracle to obtain a new signature. As, by assumption, at least one block must be fresh, but  $\mathcal{B}$  needs to guess where this happens,  $|\Pr[S_7] - \Pr[S_8]| \leq q_s \nu_{\text{aggsig-uf}}(\lambda)$  follows.

Now, the adversary can no longer win the unforgeability game. Moreover, each hop changes the view of the adversary only negligibly, concluding the proof.

*Immutability.* To prove that our scheme is immutable, we use a sequence of games:

**Game 0:** The original immutability game.

**Game 1:** We now abort, if an ephemeral public key (for which the corresponding secret keys are not given to the sanitizer) was drawn twice.

*Transition - Game 0  $\rightarrow$  Game 1:* Due to the birthday paradox, this can only happen with negligible probability.  $|\Pr[S_0] - \Pr[S_1]| \leq \frac{q_s^2}{2\lambda}$  follows, where  $q_s$  is the number of key pairs generated.

**Game 2:** We now abort, if the adversary outputs a validating  $(m^*, \sigma^{\text{SS}*}, \text{pk}_{\text{san}}^{\text{SS}*})$ , for which the ephemeral public keys or  $c$  have not been signed in that particular order.

*Transition - Game 1  $\rightarrow$  Game 2:* In this case, the adversary  $\mathcal{A}$  was able to generate a signature  $\sigma_s$  (contained in  $\sigma^{\text{SS}*}$ ) on  $(\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}^*, (\text{pk}_i)_{[1..l]}, c)$  which was never generated by the signer. We can use this to construct a reduction  $\mathcal{B}$  which forges a signature of the underlying  $\Sigma$ . The reduction works as follows. It receives the  $\text{pk}$  (and the corresponding parameters) to forge, and embeds it accordingly into the parameters/public key. Everything else is generated honestly. For every signature generated,  $\mathcal{B}$  queries its signature-generation oracle; this signature is then embedded in the response. Finally, once  $\mathcal{A}$  outputs its forgery,  $\mathcal{B}$  can return  $(\emptyset, \emptyset, (\text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}^*, (\text{pk}_i)_{[1..l]}, c), \sigma_s)$  as its own forgery.  $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{aggsig-unf}}(\lambda)$  follows.

**Game 3:** We now abort, if the adversary outputs a validating  $(m^*, \sigma^{\text{SS}*}, \text{pk}_{\text{san}}^{\text{SS}*})$ , where  $\sigma^{\text{SS}*} = (c, \sigma_s, (\text{pk}_i, \sigma_i)_{i \in [1..l]})$ , where  $m^*$  was never derivable from any query to the signing oracle.

*Transition - Game 2  $\rightarrow$  Game 3:* In this case, there must be a block  $m^{i*}$  which was changed, but the adversary does not have the corresponding secret key. Our reduction  $\mathcal{B}$  works as follows. Let  $q_s$  be an upper bound on the number of signature key pairs created for which the sanitizer does not receive the secret key, breaking the unforgeability of the underlying  $\Sigma$ . First,  $\mathcal{B}$  randomly selects an index  $i \leftarrow_r [1..q_s]$ . It receives the  $\text{pk}$  (and the corresponding parameters) to forge. The parameters are embedded honestly. Everything else is generated honestly. Then, once the  $i$ th inner signature, for which the sanitizer does not receive the secret key, is generated,  $\mathcal{B}$  embeds  $\text{pk}$  and uses its own signature-generation oracle to receive the corresponding signature. The result is embedded honestly. As, by assumption, at least one block must be fresh, but  $\mathcal{B}$  needs to guess where this happens,  $|\Pr[S_2] - \Pr[S_3]| \leq q_s \nu_{\text{aggsig-unf}}(\lambda)$  follows.

Now, the adversary can no longer win the immutability game. Moreover, each hop changes the view of the adversary only negligibly, concluding the proof.

*Privacy.* To prove that our scheme is private, we use a sequence of games:

**Game 0:** The original privacy game in the case  $b = 0$ .

**Game 1:** We now switch to  $b = 1$ .

*Transition - Game 0  $\rightarrow$  Game 1:* As the signatures are distributed exactly the same, our scheme is private in an information theoretical sense. Thus,  $|\Pr[S_0] - \Pr[S_1]| = 0$  follows.

*Transparency.* To prove that our scheme is transparent, we use a sequence of games:

**Game 0:** The original transparency game in the case  $b = 0$ .

**Game 1:** We now switch to  $b = 1$ .

*Transition - Game 0  $\rightarrow$  Game 1:* As the signatures are distributed exactly the same, our scheme is transparent in an information theoretical sense. Thus,  $|\Pr[S_0] - \Pr[S_1]| = 0$  follows.

*Invisibility.* To prove that our scheme is invisible, we use a sequence of games:

**Game 0:** The original invisibility game.

**Game 1:** We now abort, if the adversary queries some  $(\sigma^{\text{SS}}, m)$  for the challenge public keys which verifies, but was never returned by either LoRADM or  $\text{Sanit}^{\text{SS}'}$ .

*Transition - Game 0  $\rightarrow$  Game 1:* Given this adversary  $\mathcal{A}$ , we can construct an adversary  $\mathcal{B}$  which breaks the unforgeability of the SS. The reduction works as follows. First, it draws a random bit  $b \leftarrow_r \{0, 1\}$ . Then, it receives  $\text{pk}_{\text{sig}}^{\text{SS}}$  and  $\text{pk}_{\text{san}}^{\text{SS}}$  and then passes those keys to the adversary. Every redaction query is done using the  $\text{Sanit}^{\text{SS}'}$  provided (imposing the limitation the invisibility game gives). Likewise, queries to LoRADM are answered by using the  $\text{Sign}^{\text{SS}'}$  provided, but using  $\text{A}_b^{\text{SS}}$  in the case the challenge  $\text{pk}_{\text{san}}^{\text{SS}}$  is queried. Then, whenever  $\mathcal{A}$  queries  $(\sigma^{\text{SS}}, m)$  for the challenge  $\text{pk}_{\text{san}}^{\text{SS}}$ ,  $\mathcal{B}$  can simply return  $(\sigma^{\text{SS}}, m)$  as its own forgery, as, by assumption,  $(\sigma^{\text{SS}}, m)$  was not seen before. Thus,  $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{sss-unf}}(\lambda)$  follows.

**Game 2:** We now abort, if we draw an  $x$  twice.

*Transition - Game 0  $\rightarrow$  Game 1:* Due to the birthday paradox, this can only happen with negligible probability.  $|\Pr[S_1] - \Pr[S_2]| \leq \frac{q_s^2}{2\lambda}$  follows, where  $q_s$  is the number of key pairs generated.

**Game 3:** We now start replace *each* ciphertext generated for the challenge  $\text{pk}_{\text{san}}^{\text{SS}}$  with an encryption of 0 (with the appropriate length).

*Transition - Game 2  $\rightarrow$  Game 3:* Assume that the adversary can distinguish this replacement with a non-negligible probability. We can then construct a reduction  $\mathcal{B}$  which uses  $\mathcal{A}$  to break the IND-CCA2 security of the underlying encryption scheme. The reduction works via a series of hybrids. Our reduction  $\mathcal{B}$  proceeds as follows. It receives  $\text{pk}$  and (and the corresponding parameters) from its own challenger and embeds them correctly. All other values are generated as in Game 2. 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 3.0 is the same as Game 2 while in Game 3.1., however, we make the first replacement. Then, whatever  $\mathcal{A}$  outputs in Game 3.i is also output by  $\mathcal{B}$ . Note, decryption queries for ciphertexts generated by the adversary can be queried to decryption oracle provided; the content for all other ciphertexts are known and thus the ciphertexts do not need to be decrypted at all. Thus,  $|\Pr[S_2] - \Pr[S_3]| \leq q\nu_{\text{ind-cca2}}(\lambda)$  follows, where  $q$  is the number of ciphertexts generated.

**Game 4:** We now replace all  $r_i$  with a purely random value  $r_i \leftarrow_r \{0, 1\}^\lambda$ .

*Transition - Game 3  $\rightarrow$  Game 4:* An adversary distinguishing this replacement can be turned into an adversary against the pseudorandomness of the PRF. We prove this via a series of hybrids. Let Game 4.0 the same as Game 3. In Game 4.i  $\mathcal{B}$  uses its  $\text{Eval}'_{\text{PRF}}$  oracle to generate the random coins. All other values are generated as in prior game. Then, whatever  $\mathcal{A}$  outputs, is also output by  $\mathcal{B}$ .  $|\Pr[S_3] - \Pr[S_4]| \leq q_s\nu_{\text{prf-pr}}(\lambda)$  follows, where  $q_s$  is the number of calls to the signature-generation oracle.

**Game 5:** We now abort, if we draw an  $r_i$  twice.

*Transition - Game 4  $\rightarrow$  Game 5:* Due to the birthday paradox, this can only happen with negligible probability.  $|\Pr[S_4] - \Pr[S_5]| \leq \frac{q_s^2}{2\lambda}$  follows, where  $q_s$  is the number of random coins drawn.

As now the game is independent of the bit  $b$ , invisibility is proven.

## C Proof of Theorem 3

We now provide the proof of security for Construction 3.

*Proof.* Correctness follows from inspection. Each security property is proven on its own. However, we already keep all queries and answers to and from the oracle. This does not change the view of the adversary. We also directly generate any keys required, but not received by the reduction's own challenger, honestly, also embedding them, without mentioning it, to shorten the proof.

*Unforgeability.* To prove that our scheme is unforgeable, we use a sequence of games:

**Game 0:** The original unforgeability game.

**Game 1:** We now abort, if the adversary outputs  $(m, \sigma^{\text{PS}})$ , where  $\sigma^{\text{PS}} = ((\sigma_i^{\text{SS}})_{0 \leq i \leq \ell_m}, \sigma^{\text{RS}}, c, \tau, (\tau_i)_{1 \leq i \leq \ell_m})$ , where any  $(m, \sigma^{\text{RS}}, c, (\tau_i, \sigma_i^{\text{SS}})_{1 \leq i \leq \ell_m}, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ , protected by  $\sigma_0^{\text{SS}}$ , has never been returned by the challenger.

*Transition - Game 0  $\rightarrow$  Game 1:* In this case,  $((m, \sigma^{\text{RS}}, c, (\tau_i, \sigma_i^{\text{SS}})_{1 \leq i \leq \ell_m}, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}), \sigma_0^{\text{SS}})$  is a valid forgery of the outer SS. A reduction is simple. Namely, the reduction  $\mathcal{B}$  receives the challenge keys  $\text{pk}_{\text{sig}}^{\text{SS}'}$  and  $\text{pk}_{\text{san}}^{\text{SS}'}$  from its own challenger, and embeds them into  $\text{pk}_{\text{sig}}^{\text{PS}}$  and  $\text{pk}_{\text{san}}^{\text{PS}}$ . Every underlying signing and sanitization request for the SSs is performed by the reduction's oracles. As, by assumption, the message protected by  $\sigma_0^{\text{SS}}$  must be fresh, it thus breaks the unforgeability of the underlying SS in any case. Thus,  $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{sss-unf}}(\lambda)$  follows.



**Game 2:** We now abort, if the adversary outputs  $(m, \sigma^{\text{PS}})$ , where  $\sigma^{\text{PS}} = ((\sigma_i^{\text{SS}})_{0 \leq i \leq \ell_m}, \sigma^{\text{RS}}, \tau, (\tau_i)_{1 \leq i \leq \ell_m})$ , where any  $((\sigma_i^{\text{SS}})_{0 \leq i \leq \ell_m}, \sigma^{\text{RS}}, \tau, (\tau_i)_{1 \leq i \leq \ell_m})$  protected by  $\sigma_0^{\text{SS}}$  was returned by the challenger, but  $\sigma_0^{\text{SS}}$  was never created by the challenger.

*Transition - Game 1  $\rightarrow$  Game 2:* In this case,  $((\sigma_i^{\text{SS}})_{0 \leq i \leq \ell_m}, \sigma^{\text{RS}}, \tau, (\tau_i)_{1 \leq i \leq \ell_m}, \sigma_0^{\text{SS}})$  is a valid forgery of the outer SS. The reduction works as in the prior hop. As, by assumption, the message protected by  $\sigma_0^{\text{SS}}$  must be fresh, it thus breaks the unforgeability of the underlying SS in any case. Thus,  $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{SS-unf}}(\lambda)$  follows.

Now, the adversary can no longer win the unforgeability game, as also each public key is bound to a tag. Moreover, each hop changes the view of the adversary only negligibly, concluding the proof.

*Immutability.* To prove that our scheme is immutable, we use a sequence of games:

**Game 0:** The original immutability game.

**Game 1:** We now abort, if the challenger draws a tag twice.

*Transition - Game 0  $\rightarrow$  Game 1:* The probability that this event happens is bounded by the birthday paradox.  $|\Pr[S_0] - \Pr[S_1]| \leq q_t^2/2^\lambda$  follows, where  $q_t$  is the number of drawn tags.

**Game 2:** We now abort, if the adversary outputs  $(m, \sigma^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ , where  $\sigma^{\text{PS}} = ((\sigma_i^{\text{SS}})_{0 \leq i \leq \ell_m}, \sigma^{\text{RS}}, \tau, (\tau_i)_{1 \leq i \leq \ell_m})$ , but  $\text{pk}_{\text{san}}^{\text{PS}}$  was never signed by the signing oracle w.r.t. to  $\tau$ .

*Transition - Game 1  $\rightarrow$  Game 2:* This breaks the immutability property of the outer SS. The reduction proceeds as follows. It receives  $\text{pk}_{\text{sig}}^{\text{SS}'}$  from its own challenger and embeds it into  $\text{pk}_{\text{sig}}^{\text{PS}}$ . Then, every signing query is performed by the reduction's own oracles. Then, after  $\mathcal{A}$  returned  $(m, \sigma^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ ,  $(m', \sigma_0^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}})$  with  $m' = (m, \sigma^{\text{RS}}, (\tau_i, \sigma_i^{\text{SS}})_{1 \leq i \leq \ell_m}, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$  is a valid forgery.  $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{SS-imm}}(\lambda)$  follows.

**Game 3:** We now abort, if the adversary outputs  $(m, \sigma^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ , where  $\sigma^{\text{PS}} = ((\sigma_i^{\text{SS}})_{0 \leq i \leq \ell_m}, \sigma^{\text{RS}}, \tau, (\tau_i)_{1 \leq i \leq \ell_m})$ , but  $\tau$  was never drawn by the challenger.

*Transition - Game 2  $\rightarrow$  Game 3:* As  $\tau$  is non-admissible, the adversary was able to generate a signature not derivable, breaking the immutability of the outer SS. The reduction works exactly as in the prior game.  $|\Pr[S_2] - \Pr[S_3]| \leq \nu_{\text{SS-imm}}(\lambda)$  follows.

**Game 4:** We now abort, if the adversary outputs  $(m, \sigma^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ , where  $\sigma^{\text{PS}} = ((\sigma_i^{\text{SS}})_{0 \leq i \leq \ell_m}, \sigma^{\text{RS}}, \tau, (\tau_i)_{1 \leq i \leq \ell_m})$ , but some  $\tau_i$  was never signed by the challenger w.r.t.  $\tau$  or the ordering is inconsistent.

*Transition - Game 3  $\rightarrow$  Game 4:* As each  $\tau_i$  is signed by the RS, the adversary was able to generate a forgery of the RS. It receives  $\text{pk}_{\text{sig}}^{\text{RS}'}$  from its own challenger and embeds it into  $\text{pk}_{\text{sig}}^{\text{PS}}$ . Then, every signing query is performed by the reduction's own oracles. Then,  $((\tau_1, \tau_2, \dots, \tau_\ell, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}), \sigma^{\text{RS}})$  is a valid forgery.  $|\Pr[S_4] - \Pr[S_5]| \leq \nu_{\text{RS-unf}}(\lambda)$  follows.

**Game 5:** We now abort, if the adversary outputs  $(m, \sigma^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ , where  $\sigma^{\text{PS}} = ((\sigma_i^{\text{SS}})_{0 \leq i \leq \ell_m}, \sigma^{\text{RS}}, \tau, (\tau_i)_{1 \leq i \leq \ell_m})$ , but it was able to redact a block not marked as redactable.

*Transition - Game 4  $\rightarrow$  Game 5:* Note, we already ruled out tag-collisions, and thus the messages are uniquely identifiable. The reduction is same as in the prior hop.  $|\Pr[S_4] - \Pr[S_5]| \leq \nu_{\text{RS-unf}}(\lambda)$  follows.

**Game 6:** We now abort, if the adversary outputs  $(m, \sigma^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ , where  $\sigma^{\text{PS}} = ((\sigma_i^{\text{SS}})_{0 \leq i \leq \ell_m}, \sigma^{\text{RS}}, \tau, (\tau_i)_{1 \leq i \leq \ell_m})$ , but it was able to sanitize a block with tag  $\tau_i$  which was not marked as sanitizable.

*Transition - Game 5  $\rightarrow$  Game 6:* Note, we already ruled out tag-collisions and thus the messages are uniquely identifiable. The reduction is same as in Game 3, but the reduction returns  $((m^i, \tau, \tau_i), \sigma_i^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}})$ .  $|\Pr[S_5] - \Pr[S_6]| \leq \nu_{\text{SS-imm}}(\lambda)$  follows.

Now, the adversary can no longer win the immutability game. Moreover, each hop changes the view of the adversary only negligibly, concluding the proof.

*Privacy.* To prove that our scheme is private, we use a sequence of games:

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

**Game 1:** Instead of signing  $(m_0^i, \tau, \tau_i, \text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}})$  in the inner SSs and adjusting them to  $(m^i, \tau, \tau_i, \text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}})$ , sign  $(m_1^i, \tau, \tau_i, \text{pk}_{\text{sig}}^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}})$  and adjust accordingly.

*Transition - Game 1  $\rightarrow$  Game 2:* Assume that the adversary can distinguish these two games. We can then construct a reduction  $\mathcal{B}$  which uses the adversary  $\mathcal{A}$  to break the privacy of the underlying SS'. Namely,  $\mathcal{B}$  proceeds as follows. It receives  $\text{pk}_{\text{sig}}^{\text{SS}'}$  and  $\text{pk}_{\text{san}}^{\text{SS}'}$ , and embeds them into  $\text{pk}_{\text{sig}}^{\text{PS}}$  and  $\text{pk}_{\text{san}}^{\text{PS}}$ . Then, every signing, editing and proof oracle queries are answered by  $\mathcal{B}$ 's own oracles. However, for the calls to the LoREdit oracle, the calls for the SSs are redirected to the LoRSan oracle and the result embedded to the answer. Clearly, the simulation is perfect. Then, whatever  $\mathcal{A}$  outputs is also output by  $\mathcal{B}$ .  $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{SS}'\text{-priv}}$  follows.

**Game 2:** Instead of signing  $(m_0, \sigma^{\text{RS}}, (\tau_i, \sigma_i^{\text{SS}})_{1 \leq i \leq \ell_m}, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$  in the outer SSs and adjusting them to  $(m', \sigma^{\text{RS}'}, (\tau_i, \sigma_i^{\text{SS}})_{1 \leq i \leq \ell_m}, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ , sign  $(m_1, \sigma^{\text{RS}}, (\tau_i, \sigma_i^{\text{SS}})_{1 \leq i \leq \ell_m}, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$  and adjust. Note, the distribution of  $\sigma^{\text{RS}}$  and the tags are still exactly the same, even if reused, as the redactions are still performed as in the case  $b = 0$ .

*Transition - Game 1  $\rightarrow$  Game 2:* Assume that the adversary can distinguish these two games. We can then construct a reduction  $\mathcal{B}$  which uses the adversary  $\mathcal{A}$  to break the privacy of the underlying SS. Namely,  $\mathcal{B}$  proceeds as follows. It receives  $\text{pk}_{\text{sig}}^{\text{SS}'}$  and  $\text{pk}_{\text{san}}^{\text{SS}'}$ , and embeds them into  $\text{pk}_{\text{sig}}^{\text{PS}}$  and  $\text{pk}_{\text{san}}^{\text{PS}}$ . Then, every signing, editing and proof oracle queries are answered by  $\mathcal{B}$ 's own oracles. However, for the calls to the LoREdit oracle, the calls for the SSs are redirected to the LoRSan oracle and the result embedded to the answer. Clearly, the simulation is perfect. Then, whatever  $\mathcal{A}$  outputs is also output by  $\mathcal{B}$ .  $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{SS}\text{-priv}}$  follows.

**Game 3:** Instead of signing  $(\tau_1, \tau_2, \dots, \tau_\ell, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$  in the RSs from the first message, use the second message and then redact as required. Note, the distribution of the tags are still exactly the same due to the uniform distribution.

*Transition - Game 2  $\rightarrow$  Game 3:* Assume that the adversary can distinguish these two games. We can then construct a reduction  $\mathcal{B}$  which uses the adversary  $\mathcal{A}$  to break the privacy of the underlying RS. Namely,  $\mathcal{B}$  proceeds as follows. It receives  $\text{pk}_{\text{sig}}^{\text{RS}'}$  and embeds them into  $\text{pk}_{\text{sig}}^{\text{PS}}$ . Then, every signing query is answered by  $\mathcal{B}$ 's own signing oracle. However, for the calls to the LoREdit oracle, the calls for the RSs are redirected to the LoRRedact oracle and the result embedded to the answer. Clearly, the simulation is perfect. Then, whatever  $\mathcal{A}$  outputs is also output by  $\mathcal{B}$ . Note, here we no longer need  $\text{RED}^{\text{RS}}$ , as this done via the oracles.  $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{RS}\text{-priv}}$  follows.

Now, we are in the case  $b = 1$ . As each hop only changes the view of the adversary negligibly, privacy is proven.

*Transparency.* To prove that our scheme is transparent, we use a sequence of games:

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

**Game 1:** Instead of signing  $(m_i, \tau, \tau_i, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$  in the inner SS's and adjusting them to  $(m', \tau, \tau_i, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ , directly sign  $(m'_i, \tau, \tau_i, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ . Again, the distribution of  $\sigma^{\text{RS}}$  and the tags are still exactly the same, even if reused, as the redactions are still performed as in the case  $b = 0$ . Note, the restrictions on the proof-oracle are still implicitly enforced.

*Transition - Game 1  $\rightarrow$  Game 2:* Assume that the adversary can distinguish these two games. We can then construct a reduction  $\mathcal{B}$  which uses the adversary  $\mathcal{A}$  to break the transparency of the underlying SS. Namely,  $\mathcal{B}$  proceeds as follows. It receives  $\text{pk}_{\text{sig}}^{\text{SS}'}$  and  $\text{pk}_{\text{san}}^{\text{SS}'}$ , and embeds them into  $\text{pk}_{\text{sig}}^{\text{PS}}$  and  $\text{pk}_{\text{san}}^{\text{PS}}$ . Then, every signing, editing and proof oracle queries are answered by  $\mathcal{B}$ 's own oracles. However, for the calls to the Sign/Edit oracle, the calls for the SSs are redirected to the Sign/Sanit oracle and the result embedded to the answer. Clearly, the simulation is perfect. Then, whatever  $\mathcal{A}$  outputs is also output by  $\mathcal{B}$ .  $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{SS}'\text{-tran}}$  follows.

**Game 2:** Instead of signing  $(m_i, \tau, \tau_i, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$  in the inner SS's and adjusting them to  $(m', \tau, \tau_i, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ , directly sign  $(m'_i, \tau, \tau_i, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ . Again, the distribution of  $\sigma^{\text{RS}}$  and the tags are still exactly the same, even if reused, as the redactions are still performed as in the case  $b = 0$ . Note, the restrictions on the proof-oracle are still implicitly enforced.

*Transition - Game 1  $\rightarrow$  Game 2:* Assume that the adversary can distinguish these two games. We can then construct a reduction  $\mathcal{B}$  which uses the adversary  $\mathcal{A}$  to break the transparency of the underlying SS. Namely,  $\mathcal{B}$  proceeds as follows. It receives  $\text{pk}_{\text{sig}}^{\text{SS}'}$  and  $\text{pk}_{\text{san}}^{\text{SS}'}$ , and embeds them into  $\text{pk}_{\text{sig}}^{\text{PS}}$  and  $\text{pk}_{\text{san}}^{\text{PS}}$ . Then, every signing, editing and proof oracle queries are answered by  $\mathcal{B}$ 's own oracles. However, for the calls to the Sign/Edit oracle, the calls for the SSs are redirected to the Sign/Sanit oracle and the result embedded to the answer. Clearly, the simulation is perfect. Then, whatever  $\mathcal{A}$  outputs is also output by  $\mathcal{B}$ .  $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{SS}'\text{-tran}}$  follows.

**Game 3:** Instead of signing  $(\tau_1, \tau_2, \dots, \tau_\ell, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$  in the RSs from the first message and then redacting it, directly sign the redacted messages. Note, the distribution of the tags are still exactly the same due to the uniform distribution.

*Transition - Game 2  $\rightarrow$  Game 3:* Assume that the adversary can distinguish these two games. We can then construct a reduction  $\mathcal{B}$  which uses the adversary  $\mathcal{A}$  to break the transparency of the underlying RS. Namely,  $\mathcal{B}$  proceeds as follows. It receives  $\text{pk}_{\text{sig}}^{\text{RS}'}$  and embeds them into  $\text{pk}_{\text{sig}}^{\text{PS}}$ . Then, every signing query is answered by  $\mathcal{B}$ 's own signing oracle. However, for the calls to the Sign/Edit oracle, the calls for the RSs are redirected to the Sign/Redact oracle and the result embedded to the answer. Clearly, the simulation is perfect. Then, whatever  $\mathcal{A}$  outputs is also output by  $\mathcal{B}$ . Note, here we no longer need  $\text{RED}^{\text{RS}}$ , as this done via the oracles and are already replaced with a 0.  $|\Pr[S_2] - \Pr[S_3]| \leq \nu_{\text{RS}\text{-tran}}$  follows.

Now, we are in the case  $b = 1$ . As each hop only changes the view of the adversary negligibly, transparency is proven.

*Signer-Accountability.* To prove that our scheme is signer-accountable, we use a sequence of games:

**Game 0:** The original signer-accountability game.

**Game 1:** We now abort, if the adversary outputs  $(\text{pk}_{\text{sig}}^{\text{PS}}, \pi^{\text{PS}}, m, \sigma^{\text{PS}})$ , where  $\sigma^{\text{PS}} = ((\sigma_i^{\text{SS}})_{0 \leq i \leq \ell_m}, \sigma^{\text{RS}}, \tau, (\tau_i)_{1 \leq i \leq \ell_m})$ , where any  $(m, \sigma^{\text{RS}}, (\tau_i, \sigma_i^{\text{SS}})_{1 \leq i \leq \ell_m}, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ , protected by  $\sigma_0^{\text{SS}}$ , has never been returned by the challenger.

*Transition - Game 0  $\rightarrow$  Game 1:* In this case,  $(\text{pk}_{\text{sig}}^{\text{SS}}, \pi^{\text{PS}}, (m, \sigma^{\text{RS}}, (\tau_i, \sigma_i^{\text{SS}})_{1 \leq i \leq \ell_m}, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}, \sigma_0^{\text{SS}})$  is a valid forgery of the outer SS. For the reduction,  $\mathcal{B}$  receives the challenge keys  $\text{pk}_{\text{san}}^{\text{SS}'}$  from its own challenger, and embeds them into  $\text{pk}_{\text{san}}^{\text{PS}}$ . Every underlying sanitization request for the SSs is performed by the reduction's oracles. As, by assumption, the proof is wrong for  $\sigma_0^{\text{SS}}$ , it breaks the signer-accountability of the underlying SS in any case. Thus,  $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{SS}\text{-sigacc}}(\lambda)$  follows.

**Game 2:**  $(\text{pk}_{\text{sig}}^{\text{PS}}, \pi^{\text{PS}}, m, \sigma^{\text{PS}})$ , where  $\sigma^{\text{PS}} = ((\sigma_i^{\text{SS}})_{0 \leq i \leq \ell_m}, \sigma^{\text{RS}}, \tau, (\tau_i)_{1 \leq i \leq \ell_m})$ , where  $(m, \sigma^{\text{RS}}, (\tau_i, \sigma_i^{\text{SS}})_{1 \leq i \leq \ell_m}, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$  is not new, but  $\sigma_0^{\text{SS}}$  has never been returned by the challenger.

*Transition - Game 1  $\rightarrow$  Game 2:* In this case,  $(\text{pk}_{\text{sig}}^{\text{SS}}, \pi^{\text{PS}}, (m, \sigma^{\text{RS}}, c, (\tau_i, \sigma_i^{\text{SS}})_{1 \leq i \leq \ell_m}, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}, \sigma_0^{\text{SS}})$  is a valid forgery of the outer SS. The reduction works as in the prior hop. As, by assumption, the proof is wrong for  $\sigma_0^{\text{SS}}$ , it breaks the signer-accountability of the underlying SS in any case. Thus,  $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{SS}\text{-sigacc}}(\lambda)$  follows.

Now, the adversary can no longer win the signer-accountability game. Moreover, each hop changes the view of the adversary only negligibly, concluding the proof.

*Sanitizer-Accountability.* To prove that our scheme is sanitizer-accountable, we use a sequence of games:

**Game 0:** The original sanitizer-accountability game.

**Game 1:** We now abort, if the adversary outputs  $(m, \sigma^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ , where  $\sigma^{\text{PS}} = ((\sigma_i^{\text{SS}})_{0 \leq i \leq \ell_m}, \sigma^{\text{RS}}, \tau, (\tau_i)_{1 \leq i \leq \ell_m})$ , where any  $(m, \sigma^{\text{RS}}, (\tau_i, \sigma_i^{\text{SS}})_{1 \leq i \leq \ell_m}, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ , protected by  $\sigma_0^{\text{SS}}$ , has never been returned by the challenger.

*Transition - Game 0  $\rightarrow$  Game 1:* Here,  $((m, \sigma^{\text{RS}}, (\tau_i, \sigma_i^{\text{SS}})_{1 \leq i \leq \ell_m}, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}), \sigma_0^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}})$  is a valid forgery of the outer SS. For the reduction,  $\mathcal{B}$  receives the challenge keys  $\text{pk}_{\text{sig}}^{\text{SS}'}$  from its own challenger, and embeds them into  $\text{pk}_{\text{sig}}^{\text{PS}}$ . Every underlying signing and proof-generation request for the SSs is performed by the reduction's oracles. As, by assumption, the signer outputs a wrong proof for  $\sigma_0^{\text{SS}}$ , it breaks the sanitizer-accountability of the underlying SS in any case. Thus,  $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{SSs-sanacc}}(\lambda)$  follows.

**Game 2:** We now abort, if the adversary outputs  $(m, \sigma^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$ , where  $\sigma^{\text{PS}} = ((\sigma_i^{\text{SS}})_{0 \leq i \leq \ell_m}, \sigma^{\text{RS}}, \tau, (\tau_i)_{1 \leq i \leq \ell_m})$ , where any  $(m, \sigma^{\text{RS}}, (\tau_i, \sigma_i^{\text{SS}})_{1 \leq i \leq \ell_m}, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}})$  is not new, but  $\sigma_0^{\text{SS}}$  has never been returned by the challenger.

*Transition - Game 1  $\rightarrow$  Game 2:* Here,  $((m, \sigma^{\text{RS}}, (\tau_i, \sigma_i^{\text{SS}})_{1 \leq i \leq \ell_m}, \tau, \text{pk}_{\text{sig}}^{\text{PS}}, \text{pk}_{\text{san}}^{\text{PS}}), \sigma_0^{\text{SS}}, \text{pk}_{\text{san}}^{\text{SS}})$  is a valid forgery of the outer SS. The reduction works as in the prior hop. As, by assumption, the signer outputs a wrong proof for  $\sigma_0^{\text{SS}}$ , it breaks the sanitizer-accountability of the underlying SS in any case. Thus,  $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{SSs-sanacc}}(\lambda)$  follows.

Now, the adversary can no longer win the sanitizer-accountability game. Moreover, each hop changes the view of the adversary only negligibly, concluding the proof.

*Invisibility.* To prove that our scheme is invisible, we use a sequence of games:

**Game 0:** The original invisibility game where  $b = 0$ .

**Game 1:** Instead of using  $\mathbb{A}_0^{\text{PS}.1}$  use  $\mathbb{A}_1^{\text{PS}.1}$  as  $\mathbb{A}^{\text{SS}}$  in the SS'.

*Transition - Game 0  $\rightarrow$  Game 1:* This does changes the view of the adversary only negligibly due to the invisibility of the underlying SS'. Namely, assume that an adversary  $\mathcal{A}$  can distinguish these games with non-negligible probability. We can then construct an adversary  $\mathcal{B}$  which breaks the invisibility guarantees of the used SS'. In particular,  $\mathcal{B}$  receives  $\text{pk}_{\text{sig}}^{\text{SS}'}$  and  $\text{pk}_{\text{san}}^{\text{SS}'}$ , and embeds them into  $\text{pk}_{\text{sig}}^{\text{PS}}$  and  $\text{pk}_{\text{san}}^{\text{PS}}$ . For all oracle queries,  $\mathcal{B}$  uses its own oracles to answer correctly, but makes block 1 in each underlying SS admissible or not using its own challenge oracle. Then, whatever  $\mathcal{A}$  outputs, is also output by  $\mathcal{B}$ .  $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{SSs'-invis}}$  follows.

**Game 2:** Instead of using  $\mathbb{A}_0^{\text{PS}.2}$  use  $\mathbb{A}_1^{\text{PS}.2}$  as  $\mathbb{A}^{\text{RS}}$  in the RS.

*Transition - Game 1  $\rightarrow$  Game 2:* This does changes the view of the adversary only negligibly due to the invisibility of the underlying RS. Namely, assume that an adversary  $\mathcal{A}$  can distinguish these games with non-negligible probability. We can then construct an adversary  $\mathcal{B}$  which breaks the invisibility guarantees of the used RS. In particular,  $\mathcal{B}$  receives  $\text{pk}_{\text{sig}}^{\text{RS}'}$  and  $\text{pk}_{\text{san}}^{\text{RS}'}$ , and embeds them into  $\text{pk}_{\text{sig}}^{\text{PS}}$  and  $\text{pk}_{\text{san}}^{\text{PS}}$ . For all oracle queries,  $\mathcal{B}$  uses its own oracles to answer correctly using its own challenge oracle (the last three blocks are never redactable). Then, whatever  $\mathcal{A}$  outputs, is also output by  $\mathcal{B}$ .  $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{RSs-invis}}$  follows.

Now, we are in the case  $b = 1$ . As each hop only changes the view of the adversary negligibly, invisibility is proven.