

Chameleon-Hashes with Ephemeral Trapdoors And Applications to Invisible Sanitizable Signatures*

Jan Camenisch¹, David Derler², Stephan Krenn³,
Henrich C. Pöhls⁴, Kai Samelin^{1,5}, and Daniel Slamanig²

¹ IBM Research – Zurich, Rüschlikon, Switzerland
{jca|ksa}@zurich.ibm.com

² IAIK, Graz University of Technology, Graz, Austria
{david.derler|daniel.slamanig}@tugraz.at

³ AIT Austrian Institute of Technology GmbH, Vienna, Austria
stephan.krenn@ait.ac.at

⁴ ISL & Chair of IT-Security, University of Passau, Passau, Germany
hp@sec.uni-passau.de

⁵ TU Darmstadt, Darmstadt, Germany

Abstract. A chameleon-hash function is a hash function that involves a trapdoor the knowledge of which allows one to find arbitrary collisions in the domain of the function. In this paper, we introduce the notion of *chameleon-hash functions with ephemeral trapdoors*. Such hash functions feature additional, i.e., ephemeral, trapdoors which are chosen by the party computing a hash value. The holder of the main trapdoor is then unable to find a second pre-image of a hash value unless also provided with the ephemeral trapdoor used to compute the hash value. We present a formal security model for this new primitive as well as provably secure instantiations. The first instantiation is a generic black-box construction from any secure chameleon-hash function. We further provide three direct constructions based on standard assumptions. Our new primitive has some appealing use-cases, including a solution to the long-standing open problem of *invisible* sanitizable signatures, which we also present.

1 Introduction

Chameleon-hash functions, also called trapdoor-hash functions, are hash functions that feature a trapdoor that allows one to find arbitrary collisions in the domain of the functions. However, chameleon-hash functions are collision resistant as long as the corresponding trapdoor (or secret key) is not known. More precisely, a party who is privy of the trapdoor is able to find arbitrary collisions in the domain of the function. Example instantiations include trapdoor-commitment, and equivocal commitment schemes.

* The full version of this paper is available as IACR Cryptology ePrint Archive Report 2017/011. J. Camenisch and K. Samelin were supported by the EU ERC PERCY, grant agreement n°32131. D. Derler, S. Krenn, H. C. Pöhls and D. Slamanig were supported by EU H2020 project PRISMACLOUD, grant agreement n°644962.

One prominent application of this primitive are chameleon signatures [47]. Here, the intended recipient—who knows the trapdoor—of a signature σ for a message m can equivocate it to another message m' of his choice. This, in turn, means that a signature σ cannot be used to convince any other party of the authenticity of m , as the intended recipient could have “signed” arbitrary messages on its own. Many other applications appear in the literature, some of which we discuss in the related work section. However, all current constructions are “all-or-nothing” in that a party who computes a hash with respect to some public key cannot prevent the trapdoor holder from finding collisions. This can be too limiting for some use-cases.

Contribution. We introduce a new primitive dubbed chameleon-hash functions with ephemeral trapdoors. In a nutshell, this primitive requires that a collision in the hash function can be computed only when two secrets are known, i.e., the main trapdoor, and an ephemeral one. The main trapdoor is the secret key corresponding to the chameleon-hash function public key, while the second, ephemeral, trapdoor is generated by the party computing the hash value. The latter party can then decide whether the holder of the long-term secret key shall be able to equivocate the hash by providing or withholding the second trapdoor information. We present a formal security model for this new primitive. Furthermore, we present stronger definitions for existing chameleon-hash functions not considered before, including the new notion of uniqueness, and show how to construct chameleon-hash functions being secure in this stronger model. These new notions may also be useful in other scenarios.

Additionally, we provide four provably secure constructions for chameleon-hash functions with ephemeral trapdoors. The first is bootstrapped, while the three direct constructions are built on RSA-like and the DL assumption. Our new primitive has some interesting applications, including the first provably secure instantiation of *invisible* sanitizable signatures, which we also present. Additional applications of our new primitive may include revocable signatures [43], but also simulatable equivocable commitments [34]. However, in contrast to equivocable commitments, we want that parties can actually equivocate, not only a simulator. Therefore, we chose to call this primitive a chameleon-hash function rather than a commitment. Note, the primitive is different from “double-trapdoor chameleon-hash functions” [13, 25, 49], where knowing one out of two secrets is enough to produce collisions.

Related Work and State-of-the-Art. Chameleon-hash functions were introduced by Krawczyk and Rabin [47], and are based on some first ideas given by Brassard et al. [12]. Later, they have been ported to the identity-based setting (ID-based chameleon-hash functions), where the holder of a master secret key can extract new secret keys for each identity [6, 8, 26, 57, 60]. These were mainly used to tackle the key-exposure problem [7, 47]. Key exposure means that seeing a single collision in the hash allows to find further collisions by extracting the corresponding trapdoor. This problem was then directly solved by the intro-

duction of “key-exposure free” chameleon-hash functions [7, 36, 37, 57], which prohibit extracting the (master) secret key. This allows for the partial re-use of generated key material. Brzuska et al. then proposed a formal framework for tag-based chameleon-hashes secure under random-tagging attacks, i.e., random identities [15].

Beside this “plain” usage of the aforementioned primitive, chameleon-hash functions also proved useful in other areas such as on/offline signatures [27, 32, 58], (tightly) secure signature schemes [11, 44, 52], but also sanitizable signature schemes [4, 15, 41] and identity-based encryption schemes [61]. Moreover they are useful in context of trapdoor-commitments, direct anonymous attestation, Σ -protocols, and distributed hashing [3, 9, 12, 34].

Additional related work is discussed when presenting the application of our new primitive.

2 Preliminaries

Let us give our notation, the required assumptions, building blocks, and the extended framework for chameleon-hashes (without ephemeral trapdoors) first.

Notation. $\lambda \in \mathbb{N}$ denotes our security parameter. All algorithms implicitly take 1^λ as an additional input. We write $a \leftarrow A(x)$ if a is assigned the output of algorithm A with input x . An algorithm is efficient if it runs in probabilistic polynomial time (ppt) in the length of its input. 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 $a \leftarrow S$ to denote that a is chosen uniformly at random from S . For a message $m = (m[1], m[2], \dots, m[\ell])$, we call $m[i]$ a block, while $\ell \in \mathbb{N}$ denotes the number of blocks in a message m . For a list we require that we have an injective, and efficiently reversible encoding, mapping the list to $\{0, 1\}^*$. In the definitions we speak of a general message space \mathcal{M} to be as generic as possible. For our instantiations, however, we let the message space \mathcal{M} be $\{0, 1\}^*$ to reduce unhelpful boilerplate notation. 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 $\nu(n) \leq n^{-k}$, $\forall n > n_0$. For certain security properties we require that values only have one canonical representation, e.g., a “4” is not the same as a “04”, even if written as elements of \mathbb{N} for brevity. Finally, for a group G we use G^* to denote $G \setminus \{1_G\}$.

2.1 Assumptions

Discrete Logarithm Assumption. Let $(G, g, q) \leftarrow \text{GGen}(1^\lambda)$ be a group generator for a multiplicatively written group G of prime-order q with $\log_2 q = \lambda$, generated by g , i.e., $\langle g \rangle = G$. The discrete-logarithm problem (DLP) associated to GGen is to find x when given G, g, q , and g^x with $x \leftarrow \mathbb{Z}_q$. The DL assumption

now states that the DLP is hard, i.e., that for every ppt adversary \mathcal{A} , there exists a negligible function ν such that:

$$\Pr[(G, g, q) \leftarrow \text{GGen}(1^\lambda), x \leftarrow \mathbb{Z}_q, x' \leftarrow \mathcal{A}(G, g, q, g^x) : x = x'] \leq \nu(\lambda).$$

We sometimes sample from \mathbb{Z}_q^* instead of \mathbb{Z}_q . This changes the view of an adversary only negligibly, and is thus not made explicit.

2.2 Building Blocks

Collision-Resistant Hash Function Families. A family $\{\mathcal{H}_{\mathcal{R}}^k\}_{k \in \mathcal{K}}$ of hash-functions $\mathcal{H}_{\mathcal{R}}^k : \{0, 1\}^* \rightarrow \mathcal{R}$ indexed by key $k \in \mathcal{K}$ is collision-resistant if for any ppt adversary \mathcal{A} there exists a negligible function ν such that:

$$\Pr[k \leftarrow \mathcal{K}, (v, v') \leftarrow \mathcal{A}(k) : \mathcal{H}_{\mathcal{R}}^k(v) = \mathcal{H}_{\mathcal{R}}^k(v') \wedge v \neq v'] \leq \nu(\lambda).$$

Public-Key Encryption Schemes. Public-key encryption allows to encrypt a message m using a given public key pk so that the resulting ciphertext can be decrypted using the corresponding secret key sk . More formally:

Definition 1 (Public-Key Encryption Schemes). A public-key encryption scheme Π is a triple $(\text{KGen}_{\text{enc}}, \text{Enc}, \text{Dec})$ of ppt algorithms such that:

KGen_{enc}. The algorithm KGen_{enc} on input security parameter λ outputs the private and public keys of the scheme: $(\text{sk}_{\text{enc}}, \text{pk}_{\text{enc}}) \leftarrow \text{KGen}_{\text{enc}}(1^\lambda)$.

Enc. The algorithm Enc gets as input the public key pk_{enc} , and the message $m \in \mathcal{M}$ and outputs a ciphertext c : $c \leftarrow \text{Enc}(\text{pk}_{\text{enc}}, m)$.

Dec. The algorithm Dec on input a private key sk_{enc} and a ciphertext c outputs a message $m \in \mathcal{M} \cup \{\perp\}$: $m \leftarrow \text{Dec}(\text{sk}_{\text{enc}}, c)$.

Definition 2 (Secure Public-Key Encryption Schemes). We call a public-key encryption scheme Π IND- T secure, if it is correct, and IND- T -secure with $T \in \{\text{CPA}, \text{CCA2}\}$.

The formal security definitions are given in the full version of this paper.

Non-Interactive Proof Systems. Let L be an NP-language with associated witness relation R , i.e., $L = \{x \mid \exists w : R(x, w) = 1\}$. Throughout this paper, we use the Camenisch-Stadler notation [20] to express the statements proven in non-interactive, simulation-sound extractable, zero-knowledge (as defined below). In particular, we write $\pi \leftarrow \text{NIZKPoK}\{(w) : R(x, w) = 1\}$ to denote the computation of a non-interactive, simulation-sound extractable, zero-knowledge proof, where all values not in the parentheses are assumed to be public. For example, let L be defined by the following NP-relation for a group $(G, g, q) \leftarrow \text{GGen}(1^\lambda)$:

$$((g, h, y, z), (a, b)) \in R \iff y = g^a \wedge z = g^b h^a.$$

Then, we write $\pi \leftarrow \text{NIZKPoK}\{(a, b) : y = g^a \wedge z = g^b h^a\}$ to denote the corresponding proof of knowledge of witness $(a, b) \in \mathbb{Z}_q^2$ with respect to the statement $(g, h, y, z) \in G^4$. Additionally, we use $\{\mathbf{false}, \mathbf{true}\} \leftarrow \text{Verify}(x, \pi)$ to denote the corresponding verification algorithm and $\text{crs} \leftarrow \text{Gen}(1^\lambda)$ to denote the crs generation algorithm. We do not make the crs explicit and, for proof systems where a crs is required, we assume it to be an implicit input to all algorithms.

Definition 3. *We call a NIZKPoK secure, if it is complete, simulation-sound extractable, and zero-knowledge.*

The corresponding definitions can be found in the full version of this paper.

Chameleon-Hashes. Let us formally define a “standard” chameleon-hash. The framework is based upon the work done by Ateniese et al. and Brzuska et al. [5, 15], but adapted to fit our notation. Additionally, we provide some extended security definitions.

Definition 4. *A chameleon-hash CH consists of five algorithms (CParGen, CKGen, CHash, CHashCheck, Adapt), such that:*

CParGen. *The algorithm CParGen on input security parameter λ outputs public parameters of the scheme: $\text{pp}_{\text{ch}} \leftarrow \text{CParGen}(1^\lambda)$. For brevity, we assume that pp_{ch} is implicit input to all other algorithms.*

CKGen. *The algorithm CKGen given the public parameters pp_{ch} outputs the private and public keys of the scheme: $(\text{sk}_{\text{ch}}, \text{pk}_{\text{ch}}) \leftarrow \text{CKGen}(\text{pp}_{\text{ch}})$.*

CHash. *The algorithm CHash gets as input the public key pk_{ch} , and a message m to hash. It outputs a hash h , and some randomness r : $(h, r) \leftarrow \text{CHash}(\text{pk}_{\text{ch}}, m)$.⁶*

CHashCheck. *The deterministic algorithm CHashCheck gets as input the public key pk_{ch} , a message m , randomness r , and a hash h . It outputs a decision $d \in \{\mathbf{false}, \mathbf{true}\}$ indicating whether the hash h is valid: $d \leftarrow \text{CHashCheck}(\text{pk}_{\text{ch}}, m, r, h)$.*

Adapt. *The algorithm Adapt on input of secret key sk_{ch} , the old message m , the old randomness r , hash h , and a new message m' outputs new randomness r' : $r' \leftarrow \text{Adapt}(\text{sk}_{\text{ch}}, m, m', r, h)$.*

Correctness. For a CH we require the correctness property to hold. In particular, we require that for all $\lambda \in \mathbb{N}$, for all $\text{pp}_{\text{ch}} \leftarrow \text{CParGen}(1^\lambda)$, for all $(\text{sk}_{\text{ch}}, \text{pk}_{\text{ch}}) \leftarrow \text{CKGen}(\text{pp}_{\text{ch}})$, for all $m \in \mathcal{M}$, for all $(h, r) \leftarrow \text{CHash}(\text{pk}_{\text{ch}}, m)$, for all $m' \in \mathcal{M}$, we have for all for all $r' \leftarrow \text{Adapt}(\text{sk}_{\text{ch}}, m, m', r, h)$, that $\mathbf{true} = \text{CHashCheck}(\text{pk}_{\text{ch}}, m, r, h) = \text{CHashCheck}(\text{pk}_{\text{ch}}, m', r', h)$. This definition captures perfect correctness. The randomness is drawn by CHash, and not outside. This was done to capture “private-coin” constructions [5].

⁶ The randomness r is also sometimes called “check value” [5].

```

Experiment Indistinguishability $_{\mathcal{A}}^{\text{CH}}(\lambda)$ 
pp $_{\text{ch}} \leftarrow \text{CParGen}(1^\lambda)$ 
(sk $_{\text{ch}}, \text{pk}_{\text{ch}}) \leftarrow \text{CKGen}(\text{pp}_{\text{ch}})$ 
 $b \leftarrow \{0, 1\}$ 
 $a \leftarrow \mathcal{A}^{\text{HashOrAdapt}(\text{sk}_{\text{ch}}, \cdot, \cdot, b), \text{Adapt}(\text{sk}_{\text{ch}}, \cdot, \cdot, \cdot)}(\text{pk}_{\text{ch}})$ 
  where oracle HashOrAdapt on input sk $_{\text{ch}}, m, m', b$ :
    ( $h, r$ )  $\leftarrow \text{CHash}(\text{pk}_{\text{ch}}, m')$ 
    ( $h', r'$ )  $\leftarrow \text{CHash}(\text{pk}_{\text{ch}}, m)$ 
     $r'' \leftarrow \text{Adapt}(\text{sk}_{\text{ch}}, m, m', r', h')$ 
    If  $r = \perp \vee r'' = \perp$ , return  $\perp$ 
    if  $b = 0$ :
      return ( $h, r$ )
    if  $b = 1$ :
      return ( $h', r''$ )
return 1, if  $a = b$ 
return 0

```

Fig. 1. Indistinguishability

Indistinguishability. Indistinguishability requires that the randomnesses r does not reveal if it was obtained through CHash or Adapt. The messages are chosen by the adversary. We relax the perfect indistinguishability definition of Brzuska et al. [15] to a computational version, which is enough for most use-cases, including ours.

Note that we need to return \perp in the HashOrAdapt oracle, as the adversary may try to enter a message $m \notin \mathcal{M}$, even if $\mathcal{M} = \{0, 1\}^*$, which makes the algorithm output \perp . If we would not do this, the adversary could trivially decide indistinguishability. For similar reasons these checks are also included in other definitions.

Definition 5 (Indistinguishability). *A chameleon-hash CH is indistinguishable, if for any efficient adversary \mathcal{A} there exists a negligible function ν such that $\left| \Pr[\text{Indistinguishability}_{\mathcal{A}}^{\text{CH}}(\lambda) = 1] - \frac{1}{2} \right| \leq \nu(\lambda)$. The corresponding experiment is depicted in Fig. 1.*

Collision Resistance. Collision resistance says, that even if an adversary has access to an adapt oracle, it cannot find any collisions for messages other than the ones queried to the adapt oracle. Note, this is an even stronger definition than key-exposure freeness [7]: key-exposure freeness only requires that one cannot find a collision for some new “tag”, i.e., for some auxiliary value for which the adversary has never seen a collision.

Definition 6 (Collision-Resistance). *A chameleon-hash CH is collision-resistant, if for any efficient adversary \mathcal{A} there exists a negligible function ν such that $\Pr[\text{CollRes}_{\mathcal{A}}^{\text{CH}}(1^\lambda) = 1] \leq \nu(\lambda)$. The corresponding experiment is depicted in Fig. 2.*

Experiment $\text{CollRes}_{\mathcal{A}}^{\text{CH}}(\lambda)$
 $\text{pp}_{\text{ch}} \leftarrow \text{CParGen}(1^\lambda)$
 $(\text{sk}_{\text{ch}}, \text{pk}_{\text{ch}}) \leftarrow \text{CKGen}(\text{pp}_{\text{ch}})$
 $\mathcal{Q} \leftarrow \emptyset$
 $(m^*, r^*, m'^*, r'^*, h^*) \leftarrow \mathcal{A}^{\text{Adapt}'(\text{sk}_{\text{ch}}, \cdot, \cdot, \cdot)}(\text{pk}_{\text{ch}})$
 where oracle Adapt' on input $\text{sk}_{\text{ch}}, m, m', r, h$:
 Return \perp , if $\text{CHashCheck}(\text{pk}_{\text{ch}}, m, r, h) \neq \text{true}$
 $r' \leftarrow \text{Adapt}(\text{sk}_{\text{ch}}, m, m', r, h)$
 If $r' = \perp$, return \perp
 $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{m, m'\}$
 return r'
 return 1, if $\text{CHashCheck}(\text{pk}_{\text{ch}}, m^*, r^*, h^*) = \text{CHashCheck}(\text{pk}_{\text{ch}}, m'^*, r'^*, h^*) = \text{true} \wedge$
 $m'^* \notin \mathcal{Q} \wedge m^* \neq m'^*$
 return 0

Fig. 2. Collision Resistance

Experiment $\text{Uniqueness}_{\mathcal{A}}^{\text{CH}}(\lambda)$
 $\text{pp}_{\text{ch}} \leftarrow \text{CParGen}(1^\lambda)$
 $(\text{pk}^*, m^*, r^*, r'^*, h^*) \leftarrow \mathcal{A}(\text{pp}_{\text{ch}})$
 return 1, if $\text{CHashCheck}(\text{pk}^*, m^*, r^*, h^*) = \text{CHashCheck}(\text{pk}^*, m^*, r'^*, h^*) = \text{true}$
 $\wedge r^* \neq r'^*$
 return 0

Fig. 3. Uniqueness

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

Definition 7 (Uniqueness). *A chameleon-hash CH is unique, if for any efficient adversary \mathcal{A} there exists a negligible function ν such that $\Pr[\text{Uniqueness}_{\mathcal{A}}^{\text{CH}}(1^\lambda) = 1] \leq \nu(\lambda)$. The corresponding experiment is depicted in Fig. 3.*

Definition 8 (Secure Chameleon-Hashes). *We call a chameleon-hash CH secure, if it is correct, indistinguishable, and collision-resistant.*

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

In the full version of this paper, we show how to construct a unique chameleon-hash satisfying our strong notions, based on the ideas by Brzuska et al. [15].

3 Chameleon-Hashes with Ephemeral Trapdoors

As already mentioned, a chameleon-hash with ephemeral trapdoor (CHET) allows to prevent the holder of the trapdoor sk_{ch} from finding collisions, as long as no additional ephemeral trapdoor etd is known. This additional ephemeral trapdoor is chosen freshly for each new hash, and providing, or withholding, this trapdoor

thus allows to decide upon each hash computation if finding a collision is possible for the holder of the long-term trapdoor. Hence, we need to introduce a new framework given next, which is also accompanied by suitable security definitions.

Definition 9 (Chameleon-Hashes with Ephemeral Trapdoors). *A chameleon-hash with ephemeral trapdoors CHET is a tuple of five algorithms (CParGen, CKGen, CHash, CHashCheck, Adapt), such that:*

CParGen. *The algorithm CParGen on input security parameter λ outputs the public parameters: $\text{pp}_{\text{ch}} \leftarrow \text{CParGen}(1^\lambda)$. For simplicity, we assume that pp_{ch} is an implicit input to all other algorithms.*

CKGen. *The algorithm CKGen given the public parameters pp_{ch} outputs the long-term private and public keys of the scheme: $(\text{sk}_{\text{ch}}, \text{pk}_{\text{ch}}) \leftarrow \text{CKGen}(\text{pp}_{\text{ch}})$.*

CHash. *The algorithm CHash gets as input the public key pk_{ch} , and a message m to hash. It outputs a hash h , randomness r , and the trapdoor information: $(h, r, \text{etd}) \leftarrow \text{CHash}(\text{pk}_{\text{ch}}, m)$.*

CHashCheck. *The deterministic algorithm CHashCheck gets as input the public key pk_{ch} , a message m , a hash h , and randomness r . It outputs a decision bit $d \in \{\text{false}, \text{true}\}$, indicating whether the given hash is correct: $d \leftarrow \text{CHashCheck}(\text{pk}_{\text{ch}}, m, r, h)$.*

Adapt. *The algorithm Adapt gets as input sk_{ch} , the old message m , the old randomness r , the new message m' , the hash h , and the trapdoor information etd and outputs new randomness r' : $r' \leftarrow \text{Adapt}(\text{sk}_{\text{ch}}, m, m', r, h, \text{etd})$.*

Correctness. For each CHET we require the correctness properties to hold. In particular, we require that for all security parameters $\lambda \in \mathbb{N}$, for all $\text{pp}_{\text{ch}} \leftarrow \text{CParGen}(1^\lambda)$, for all $(\text{sk}_{\text{ch}}, \text{pk}_{\text{ch}}) \leftarrow \text{CKGen}(\text{pp}_{\text{ch}})$, for all $m \in \mathcal{M}$, for all $(h, r, \text{etd}) \leftarrow \text{CHash}(\text{pk}_{\text{ch}}, m)$, we have $\text{CHashCheck}(\text{pk}_{\text{ch}}, m, r, h) = \text{true}$, and additionally for all $m' \in \mathcal{M}$, for all $r' \leftarrow \text{Adapt}(\text{sk}_{\text{ch}}, m, m', r, h, \text{etd})$, we have $\text{CHashCheck}(\text{pk}_{\text{ch}}, m', r', h) = \text{true}$. This definition captures perfect correctness. We also require some security guarantees, which we introduce next.

Indistinguishability. Indistinguishability requires that the randomnesses r does not reveal if it was obtained through CHash or Adapt. In other words, an outsider cannot decide whether a message is the original one or not.

Definition 10 (Indistinguishability). *A chameleon-hash with ephemeral trapdoor CHET is indistinguishable, if for any efficient adversary \mathcal{A} there exists a negligible function ν such that $\left| \Pr[\text{Indistinguishability}_{\mathcal{A}}^{\text{CHET}}(\lambda) = 1] - \frac{1}{2} \right| \leq \nu(\lambda)$. The corresponding experiment is depicted in Fig. 4.*

Public Collision Resistance. Public collision resistance requires that, even if an adversary has access to an Adapt oracle, it cannot find any collisions by itself. Clearly, the collision must be fresh, i.e., must not be produced using the Adapt oracle.

Experiment $\text{Indistinguishability}_{\mathcal{A}}^{\text{CHET}}(\lambda)$

```

ppch ← CParGen(1λ)
(skch, pkch) ← CKGen(ppch)
b ← {0, 1}
a ←  $\mathcal{A}^{\text{HashOrAdapt}(\text{sk}_{\text{ch}}, \cdot, \cdot, b), \text{Adapt}(\text{sk}_{\text{ch}}, \cdot, \cdot, \cdot, \cdot)}$ (pkch)
where oracle HashOrAdapt on input skch, m, m', b:
  let (h, r, etd) ← CHash(pkch, m')
  let (h', r', etd') ← CHash(pkch, m)
  let r'' ← Adapt(skch, m, m', r', h', etd')
  if r'' = ⊥ ∨ r' = ⊥, return ⊥
  if b = 0:
    return (h, r, etd)
  if b = 1:
    return (h', r'', etd')
return 1, if a = b
return 0

```

Fig. 4. Indistinguishability

Experiment $\text{PublicCollRes}_{\mathcal{A}}^{\text{CHET}}(\lambda)$

```

ppch ← CParGen(1λ)
(skch, pkch) ← CKGen(ppch)
Q ← ∅
(m*, r*, m'*, r'*, h*) ←  $\mathcal{A}^{\text{Adapt}'(\text{sk}_{\text{ch}}, \cdot, \cdot, \cdot, \cdot)}$ (pkch)
where oracle Adapt' on input skch, m, m', r, etd, h:
  return ⊥, if CHashCheck(pkch, m, r, h) = false
  r' ← Adapt(skch, m, m', r, h, etd)
  If r' = ⊥, return ⊥
  Q ← Q ∪ {m, m'}
  return r'
return 1, if CHashCheck(pkch, m*, r*, h*) = true ∧
  CHashCheck(pkch, m'*, r'*, h*) = true ∧
  m'* ∉ Q ∧ m* ≠ m'*
return 0

```

Fig. 5. Public Collision-Resistance

Definition 11 (Public Collision-Resistance). A chameleon-hash with ephemeral trapdoor CHET is publicly collision-resistant, if for any efficient adversary \mathcal{A} there exists a negligible function ν such that $\Pr[\text{PublicCollRes}_{\mathcal{A}}^{\text{CHET}}(1^\lambda) = 1] \leq \nu(\lambda)$. The corresponding experiment is depicted in Fig. 5.

Private Collision-Resistance. Private collision resistance requires that even the holder of the secret key sk_{ch} cannot find collisions as long as etd is unknown. This is formalized by a honest hashing oracle which does not return etd . Hence, \mathcal{A} 's goal is to return an actual collision on a non-adversarially generated hash h , for which it does not know etd .

Definition 12 (Private Collision-Resistance). A chameleon-hash with ephemeral trapdoor CHET is privately collision-resistant, if for any efficient adversary

```

Experiment PrivateCollResACHET( $\lambda$ )
ppch  $\leftarrow$  CParGen( $1^\lambda$ )
 $\mathcal{Q} \leftarrow \emptyset$ 
(pk*, state)  $\leftarrow$   $\mathcal{A}$ (ppch)
(m*, r*, m'*, r'*, h*)  $\leftarrow$   $\mathcal{A}^{\text{CHash}'(\text{pk}^*, \cdot)}$ (state)
  where oracle CHash' on input pk*, m:
    (h, r, etd)  $\leftarrow$  CHash(pk*, m)
    If h =  $\perp$ , return  $\perp$ 
     $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(h, m)\}$ 
    return (h, r)
return 1, if CHashCheck(pk*, m*, r*, h*) = true  $\wedge$ 
  CHashCheck(pk*, m'*, r'*, h*) = true  $\wedge$ 
  (h*, m*)  $\notin$   $\mathcal{Q} \wedge$  (h*,  $\cdot$ )  $\in$   $\mathcal{Q}$ 
return 0

```

Fig. 6. Private Collision-Resistance

```

Experiment UniquenessACHET( $\lambda$ )
ppch  $\leftarrow$  CParGen( $1^\lambda$ )
(pk*, m*, r*, r'*, h*)  $\leftarrow$   $\mathcal{A}$ (ppch)
return 1, if CHashCheck(pk*, m*, r*, h*) = CHashCheck(pk*, m*, r'*, h*) = true  $\wedge$ 
  r*  $\neq$  r'*
return 0

```

Fig. 7. Uniqueness

A there exists a negligible function ν such that $\Pr[\text{PrivateCollRes}_A^{\text{CHET}}(1^\lambda) = 1] \leq \nu(\lambda)$. The corresponding experiment is depicted in Fig. 6.

Uniqueness. Uniqueness requires that it is hard to come up with two different randomness values for the same message m^* and hash value h^* , where pk^* is adversarially chosen.

Definition 13 (Uniqueness). A chameleon-hash with ephemeral trapdoor CHET is unique, if for any efficient adversary \mathcal{A} there exists a negligible function ν such that $\Pr[\text{Uniqueness}_A^{\text{CHET}}(1^\lambda) = 1] \leq \nu(\lambda)$. The corresponding experiment is depicted in Fig. 7.

Definition 14 (Secure Chameleon-Hashes with Ephemeral Trapdoors). We call a chameleon-hash with ephemeral trapdoor CHET secure, if it is correct, indistinguishable, publicly collision-resistant, and privately collision-resistant.

Note, we do not require that a secure CHET is unique, as it depends on the use-case whether this strong security notion is required.

4 Constructions

Regarding constructions of CHET schemes, we first ask the natural question whether CHETs can be built from existing primitives in a black-box way. Interestingly, we can show how to elegantly “bootstrap” a CHET scheme in a black-box

fashion from *any* existing secure (and unique) chameleon-hash. Since, however, a secure chameleon-hash does not exist to date, we show how to construct it in the full version of this paper, based on the ideas by Brzuska et al. [15]. If one does not require uniqueness, one can, e.g., resort to the recent scheme given by Ateniese et al. [5].

We then proceed in presenting three direct constructions, two based on the DL assumption, and one based on an RSA-like assumption. While the DL-based constructions are not unique, the construction from RSA-like assumptions even achieves uniqueness. We however note that this strong security notion is not required in all use-cases. For example, in our application scenario (cf. Section 5), the CHETs do not need to be unique.

4.1 Black-Box Construction: Bootstrapping

We now present a black-box construction from any existing chameleon-hash. Namely, we show how one can achieve our desired goals by combining two instances of a secure chameleon-hash CH.

Construction 1 (Bootstrapped Construction) *We omit obvious checks for brevity. Let CHET be defined as:*

CParGen. *The algorithm CParGen does the following:*

1. Return $\text{pp}_{\text{ch}} \leftarrow \text{CH.CParGen}(1^\lambda)$.

CKGen. *The algorithm CKGen generates the key pair in the following way:*

1. Return $(\text{sk}_{\text{ch}}^1, \text{pk}_{\text{ch}}^1) \leftarrow \text{CH.CKGen}(\text{pp}_{\text{ch}})$.

CHash. *To hash a message m , w.r.t. public key pk_{ch}^1 do:*

1. Let $(\text{sk}_{\text{ch}}^2, \text{pk}_{\text{ch}}^2) \leftarrow \text{CH.CKGen}(\text{pp}_{\text{ch}})$.
2. Let $(h_1, r_1) \leftarrow \text{CH.CHash}(\text{pk}_{\text{ch}}^1, m)$.
3. Let $(h_2, r_2) \leftarrow \text{CH.CHash}(\text{pk}_{\text{ch}}^2, m)$.
4. Return $((h_1, h_2, \text{pk}_{\text{ch}}^2), (r_1, r_2), \text{sk}_{\text{ch}}^2)$.

CHashCheck. *To check whether a given hash $h = (h_1, h_2, \text{pk}_{\text{ch}}^2)$ is valid on input $\text{pk}_{\text{ch}} = \text{pk}_{\text{ch}}^1$, m , $r = (r_1, r_2)$, do:*

1. Let $b_1 \leftarrow \text{CH.CHashCheck}(\text{pk}_{\text{ch}}^1, m, r_1, h_1)$.
2. Let $b_2 \leftarrow \text{CH.CHashCheck}(\text{pk}_{\text{ch}}^2, m, r_2, h_2)$.
3. If $b_1 = \text{false} \vee b_2 = \text{false}$, return **false**.
4. Return **true**.

Adapt. *To find a collision w.r.t. m , m' , randomness $r = (r_1, r_2)$, hash $h = (h_1, h_2, \text{pk}_{\text{ch}}^2)$, $\text{etd} = \text{sk}_{\text{ch}}^2$, and $\text{sk}_{\text{ch}} = \text{sk}_{\text{ch}}^1$ do:*

1. If **false** = $\text{CHashCheck}(\text{pk}_{\text{ch}}^1, m, r, h)$, return \perp .
2. Compute $r'_1 \leftarrow \text{CH.Adapt}(\text{sk}_{\text{ch}}^1, m, m', r_1, h_1)$.
3. Compute $r'_2 \leftarrow \text{CH.Adapt}(\text{sk}_{\text{ch}}^2, m, m', r_2, h_2)$.
4. Return (r'_1, r'_2) .

The proof of the following theorem can be found in the full version of this paper.

Theorem 1. *If CH is secure and unique, then the chameleon-hash with ephemeral trapdoors CHET in Construction 1 is secure, and unique.*

This construction is easy to understand and only uses standard primitives. The question is now, if we can also directly construct CHET, which we answer to the affirmative subsequently.

4.2 A First Direct Construction

We now present a direct construction in groups where the DLP is hard using some ideas related to Pedersen commitments [53]. In a nutshell, the long-term secret is the discrete logarithm x between two elements g and h (i.e., $g^x = h$) of the long-term public key, while the ephemeral trapdoor is the randomness of the “commitment”. To prohibit that a seen collision allows to extract the long-term secret key x , both trapdoors are hidden in a NIZKPoK. To make the “commitment” equivocable, it is then again randomized. To avoid that the holder of sk_{ch} needs to store state, the randomness is encrypted to a public key of a IND-CCA2 secure encryption scheme contained in pk_{ch} . Security then directly follows from the DL assumption, IND-CCA2, the collision-resistance of the used hash function, and the extractability property of the NIZKPoK system. For brevity we assume that the NP-languages involved in the NIZKPoKs are implicitly defined by the scheme. Note, this construction is not unique.

Construction 2 (CHET in Known-Order Groups) *Let $\{\mathcal{H}_{\mathbb{Z}_q^*}^k\}_{k \in \mathcal{K}}$ denote a family of collision-resistant hash functions $\mathcal{H}_{\mathbb{Z}_q^*}^k : \{0, 1\}^* \rightarrow \mathbb{Z}_q^*$ indexed by a key $k \in \mathcal{K}$ and let CHET be as follows:*

CParGen. *The algorithm CParGen generates the public parameters in the following way:*

1. Let $(G, g, p) \leftarrow \text{GGen}(1^\lambda)$.
2. Let $k \leftarrow \mathcal{K}$ for the hash function.
3. Let $\text{crs} \leftarrow \text{Gen}(1^\lambda)$.⁷
4. Return $((G, g, q), k, \text{crs})$.

CKGen. *The algorithm CKGen generates the key pair in the following way:*

1. Draw random $x \leftarrow \mathbb{Z}_q^*$. Set $h \leftarrow g^x$.
2. Generate $\pi_{\text{pk}} \leftarrow \text{NIZKPoK}\{(x) : h = g^x\}$.
3. Let $(\text{sk}_{\text{enc}}, \text{pk}_{\text{enc}}) \leftarrow \text{II.KGen}_{\text{enc}}(1^\lambda)$.
4. Return $((x, \text{sk}_{\text{enc}}), (h, \pi_{\text{pk}}, \text{pk}_{\text{enc}}))$.

CHash. *To hash m w.r.t. $\text{pk}_{\text{ch}} = (h, \pi_{\text{pk}}, \text{pk}_{\text{enc}})$ do:*

1. Return \perp , if $h \notin G^*$.
2. If π_{pk} is not valid, return \perp .
3. Draw random $r \leftarrow \mathbb{Z}_q^*$.
4. Draw random $\text{etd} \leftarrow \mathbb{Z}_q^*$.
5. Let $h' \leftarrow g^{\text{etd}}$.
6. Generate $\pi_t \leftarrow \text{NIZKPoK}\{(\text{etd}) : h' = g^{\text{etd}}\}$.
7. Encrypt r , i.e., let $C \leftarrow \text{II.Enc}(\text{pk}_{\text{enc}}, r)$.
8. Let $a \leftarrow \mathcal{H}_{\mathbb{Z}_q^*}^k(m)$.
9. Let $p \leftarrow h^r$.
10. Generate $\pi_p \leftarrow \text{NIZKPoK}\{(r) : p = h^r\}$.
11. Let $b \leftarrow ph^{a}$.
12. Return $((b, h', \pi_t), (p, C, \pi_p), \text{etd})$.

⁷ Actually we need one crs per language, but we do not make this explicit here.

- CHashCheck.** To check whether a given hash (b, h', π_t) is valid on input $\text{pk}_{\text{ch}} = (h, \pi_{\text{pk}}, \text{pk}_{\text{enc}})$, $m, r = (p, C, \pi_p)$, do:
1. Return **false**, if $p \notin G^* \vee h' \notin G^*$.
 2. If either π_p, π_t , or π_{pk} are not valid, return \perp .
 3. Let $a \leftarrow \mathcal{H}_{\mathbb{Z}_q^*}^k(m)$.
 4. Return **true**, if $b = ph'^a$.
 5. Return **false**.
- Adapt.** To find a collision w.r.t. $m, m', (b, h', \pi_t)$, randomness (p, C, π_p) , and trapdoor information etd , and $\text{sk}_{\text{ch}} = (x, \text{sk}_{\text{enc}})$ do:
1. If **false** = CHashCheck($\text{pk}_{\text{ch}}, m, (p, C, \pi_p), (b, h', \pi_t)$), return \perp .
 2. Decrypt C , i.e., $r \leftarrow \Pi.\text{Dec}(\text{sk}_{\text{enc}}, C)$. If $r = \perp$, return \perp .
 3. If $h' \neq g^{\text{etd}}$, return \perp .
 4. Let $a \leftarrow \mathcal{H}_{\mathbb{Z}_q^*}^k(m)$.
 5. Let $a' \leftarrow \mathcal{H}_{\mathbb{Z}_q^*}^k(m')$.
 6. If $p \neq g^{xr}$, return \perp .
 7. If $a = a'$, return (p, C, π_p) .
 8. Let $r' \leftarrow \frac{rx + a \cdot \text{etd} - a' \cdot \text{etd}}{x}$.
 9. Let $p' \leftarrow h^{r'}$.
 10. Encrypt r' , i.e., let $C' \leftarrow \Pi.\text{Enc}(\text{pk}_{\text{enc}}, r')$.
 11. Generate $\pi'_p \leftarrow \text{NIZKPoK}\{r' : p' = h^{r'}\}$.
 12. Return (p', C', π'_p) .

Some of the checks can already be done in advance, e.g., at a PKI, which only generates certificates, if the restrictions on each public key are fulfilled.

The proof of the following Theorem is given in the full version of this paper.

Theorem 2. *If the DL assumption in G holds, $\mathcal{H}_{\mathbb{Z}_{|G|}^*}^k$ is collision-resistant, Π is IND-CCA2 secure, and NIZKPoK is secure, then the chameleon-hash with ephemeral trapdoors CHET in Construction 2 is secure.*

Two further constructions, one based on the DL assumption in gap-groups, and one based on RSA-like assumptions (in the random oracle model, which is also unique), are given in in the full version of this paper.

5 Application: Invisible Sanitizable Signatures

Informally, security of digital signatures requires that a signature σ on a message m becomes invalid as soon as a single bit of m is altered [40]. However, there are many real-life use-cases in which a subsequent change to signed data by a semi-trusted party without invalidating the signature is desired. As a simplified example, consider a patient record which is signed by a medical doctor. The accountant, which charges the insurance company, only requires knowledge of the treatments and the patient's insurance number. This protects the patient's privacy. In this constellation, having the data re-signed by the M.D. whenever subsets of the record need to be forwarded to some party induces too much

overhead to be practical in real scenarios or may even be impossible due to availability constraints.

Sanitizable signature schemes (SSS) [4] address these shortcomings. They allow the signer to determine which blocks $m[i]$ of a given message $m = (m[1], m[2], \dots, m[i], \dots, m[\ell])$ are admissible. Any such admissible block can be changed to a different bitstring $m[i]' \in \{0, 1\}^*$, where $i \in \{1, 2, \dots, \ell\}$, by a semi-trusted party named the sanitizer. This party is identified by a private/public key pair and the sanitization process described before requires the private key. In a nutshell, sanitization of a message m results in an altered message $m' = (m[1]', m[2]', \dots, m[i]', \dots, m[\ell]')$, where $m[i] = m[i]'$ for every non-admissible block, and also a signature σ' , which verifies under the original public key. Thus, authenticity of the message is still ensured. In the prior example, for the server storing the data it is possible to already black-out the sensitive parts of a signed document without any additional communication with the M.D. and in particular without access to the signing key of the M.D.

Real-world applications of SSSs include the already mentioned privacy-preserving handling of patient data, secure routing, privacy-preserving document disclosure, credentials, and blank signatures [4, 17, 18, 19, 24, 30, 42].

Our Contribution. We introduce the notion of *invisible* SSSs. This strong privacy notion requires that a third party not holding any secret keys cannot decide whether a specific block is admissible, i.e., can be sanitized. This has already been discussed by Ateniese et al. [4] in the first work on sanitizable signatures, but they neither provide a formal framework nor a provably secure construction. However, we identify some use-cases where such a notion is important, and we close this gap by introducing a new framework for SSSs, along with an extended security model. Moreover, we propose a construction being provably secure in our framework. Our construction paradigm is based on IND-CPA secure encryption schemes, standard, yet unique, chameleon-hashes, and strongly unforgeable signature schemes. These can be considered standard tools nowadays. We pair those with a chameleon-hash with ephemeral trapdoors.

Motivation. At PKC '09, Brzuska et al. formalized the most common security model of SSSs [15]. For our work, the most important property they are addressing is “weak transparency”. It means that although a third party sees which blocks of a message are admissible, it cannot decide whether some block has already been sanitized by a sanitizer. More precisely, their formalization explicitly requires that the third party is always able to decide whether a given block in a message is admissible. However, as this may invade privacy, having a construction which hides this additional information is useful as well. To address this problem the notion of “strong transparency” has been informally proposed in the original work by Ateniese et al. [4].

Examples. To make the usefulness of such a stronger privacy property more visible, consider the following two application scenarios.

In the first scenario, we consider that a document is the output of a workflow that requires several—potentially heavy—computations to become ready. We assume that the output of each workflow step could be produced by one party alone, but could also be outsourced. However, if the party decides to outsource the production of certain parts of the document it wants the potential involvement of other parties to stay hidden, e.g., the potential and actual outsourcing might be considered a trade secret. In order to regain some control that all tasks are done only by authorized subordinates, the document—containing template parts—is signed with a sanitizable signature. Such an approach, i.e., to use SSS for workflow control, was proposed in [29].

The second one is motivated by an ongoing legal debate in Germany.⁸ Consider a school class where a pupil suffers from dyslexia⁹ and thus can apply for additional help to compensate the illness. One way to compensate this is to consider spelling mistakes less when giving grades. Assume that only the school’s principal shall decide to what extent a certain grade shall be improved. Of course, this shall only be possible for pupils who are actually handicapped. For the pupil with dyslexia, e.g., known to the teacher of the class in question, the grade is marked as sanitizable by the principal. The legal debate in Germany is about an outsider, e.g., future employer, who should not be able to decide that grades had the potential to be altered and of course also not see for which pupils the grades have been altered to preserve their privacy. To achieve this, standard sanitizable signature schemes are clearly not enough, as they do not guarantee that an outsider cannot derive which blocks are potentially sanitizable, i.e., which pupil is actually handicapped. We offer a solution to this problem, where an outsider cannot decide which block is admissible, i.e., can be altered.

State-of-the-Art. SSSs have been introduced by Ateniese et al. [4]. Brzuska et al. formalized most of the current security properties [15]. These have been later extended for (strong) unlinkability [17, 19, 35] and non-interactive public accountability [18, 19]. Some properties discussed by Brzuska et al. [15] have then been refined by Gong et al. [41]. Namely, they also consider the admissible blocks in the security games, while still requiring that these are visible to everyone. Recently, Krenn et al. further refined the security properties to also account for the signatures, not only the message [48].¹⁰ We use the aforementioned results as our starting point for the extended definitions.

Also, several extensions such as limiting the sanitizer to signer-chosen values [21, 31, 46, 56], trapdoor SSSs (which allow to add new sanitizers after signature generation by the signer) [23, 59], multi-sanitizer and -signer environments for SSSs [16, 19, 22], and sanitization of signed and encrypted data [33]

⁸ See for example the ruling from the German Federal Administrative Court (BVerwG) 29.07.2015, Az.: 6 C 33.14, 6 C 35.14.

⁹ A disorder involving difficulty in learning to read or interpret words, letters and other symbols.

¹⁰ We want to stress that Krenn et al. [48] also introduce “strong transparency”, which is not related to the definition given by Ateniese et al. [4].

have been considered. SSSs have also been used as a tool to make other primitives accountable [55], and to build other primitives [10, 51]. Also, SSSs and data-structures being more complex than lists have been considered [56]. Our results carry over to the aforementioned extended settings with only minor additional adjustments. Implementations of SSSs have also been presented [18, 19, 50, 54].

Of course, computing on signed messages is a broad field. We can therefore only give a small overview. Decent and comprehensive overviews of other related primitives, however, have already been published [2, 14, 28, 38, 39].

5.1 Additional Building Blocks

We assume that the reader is familiar with digital signatures, PRGs, and PRFs, and only introduce the notation used in the following. A PRF consists of a key generation algorithm KGen_{prf} and an evaluation algorithm Eval_{prf} ; similarly, a PRG consists of an evaluation algorithm Eval_{prg} . Finally, a digital signature scheme Σ consists of a key generation algorithm KGen_{sig} , a signing algorithm Sign , and a verification algorithm Verify . The formal definition and security notions are given in the full version of this paper.

5.2 Our Framework for Sanitizable Signature Schemes

Subsequently, we introduce our framework for SSSs. Our definitions are based on existing work [15, 18, 19, 41, 48]. However, due to our goals, we need to modify the current framework to account for the fact that the admissible blocks are only visible to the sanitizer. We do not consider “non-interactive public accountability” [18, 19, 45], which allows a third party to decide which party is accountable, as transparency is mutually exclusive to this property, but is very easy to achieve, e.g., by signing the sanitizable signature again [18].

Before we present the formal definition, we settle some notation. The variable ADM contains the set of indices of the modifiable blocks, as well as the number ℓ of blocks in a message m . We write $\text{ADM}(m) = \text{true}$, if ADM is valid w.r.t. m , i.e., ADM contains the correct ℓ and all indices are in m . For example, let $\text{ADM} = (\{1, 2, 4\}, 4)$. Then, m must contain four blocks, while all but the third will be admissible. If we write $m_i \in \text{ADM}$, we mean that m_i is admissible. MOD is a set containing pairs $(i, m[i]')$ for those blocks that shall be modified, meaning that $m[i]$ is replaced with $m[i]'$. We write $\text{MOD}(\text{ADM}) = \text{true}$, if MOD is valid w.r.t. ADM , meaning that the indices to be modified are contained in ADM . To allow a compact presentation of our construction we write $\tilde{X}_{n,m}$ with $n \leq m$ for the vector $(X_n, X_{n+1}, X_{n+2}, \dots, X_{m-1}, X_m)$.

Definition 15 (Sanitizable Signatures). *A sanitizable signature scheme SSS consists of eight ppt algorithms ($\text{SSSParGen}, \text{KGen}_{\text{sig}}, \text{KGen}_{\text{san}}, \text{Sign}, \text{Sanit}, \text{Verify}, \text{Proof}, \text{Judge}$) such that*

SSSParGen. *The algorithm SSSParGen , on input security parameter λ , generates the public parameters: $\text{pp}_{\text{SSS}} \leftarrow \text{SSSParGen}(1^\lambda)$. We assume that pp_{SSS} is implicitly input to all other algorithms.*

KGen_{sig}. The algorithm KGen_{sig} takes the public parameters pp_{SSS} and returns the signer’s private key and the corresponding public key: $(\text{pk}_{\text{sig}}, \text{sk}_{\text{sig}}) \leftarrow \text{KGen}_{\text{sig}}(\text{pp}_{\text{SSS}})$.

KGen_{san}. The algorithm KGen_{san} takes the public parameters pp_{SSS} and returns the sanitizer’s private key and the corresponding public key: $(\text{pk}_{\text{san}}, \text{sk}_{\text{san}}) \leftarrow \text{KGen}_{\text{san}}(\text{pp}_{\text{SSS}})$.

Sign. The algorithm Sign takes as input a message m , sk_{sig} , pk_{san} , as well as a description ADM of the admissible blocks. If $\text{ADM}(m) = \text{false}$, this algorithm returns \perp . It outputs a signature $\sigma \leftarrow \text{Sign}(m, \text{sk}_{\text{sig}}, \text{pk}_{\text{san}}, \text{ADM})$.

Sanit. The algorithm Sanit takes a message m , modification instruction MOD , a signature σ , pk_{sig} , and sk_{san} . It outputs m' together with σ' : $(m', \sigma') \leftarrow \text{Sanit}(m, \text{MOD}, \sigma, \text{pk}_{\text{sig}}, \text{sk}_{\text{san}})$ where $m' \leftarrow \text{MOD}(m)$ is message m modified according to the modification instruction MOD .

Verify. The algorithm Verify takes as input the signature σ for a message m w.r.t. the public keys pk_{sig} and pk_{san} and outputs a decision $d \in \{\text{true}, \text{false}\}$: $d \leftarrow \text{Verify}(m, \sigma, \text{pk}_{\text{sig}}, \text{pk}_{\text{san}})$.

Proof. The algorithm Proof takes as input sk_{sig} , a message m , a signature σ , a set of polynomially many additional message/signature pairs $\{(m_i, \sigma_i)\}$ and pk_{san} . It outputs a string $\pi \in \{0, 1\}^*$ which can be used by the Judge to decide which party is accountable given a message/signature pair (m, σ) : $\pi \leftarrow \text{Proof}(\text{sk}_{\text{sig}}, m, \sigma, \{(m_i, \sigma_i) \mid i \in \mathbb{N}\}, \text{pk}_{\text{san}})$.

Judge. The algorithm Judge takes as input a message m , a signature σ , pk_{sig} , pk_{san} , as well as a proof π . Note, this means that once a proof π is generated, the accountable party can be derived by anyone for that message/signature pair (m, σ) . It outputs a decision $d \in \{\text{Sig}, \text{San}\}$, indicating whether the message/signature pair has been created by the signer, or the sanitizer: $d \leftarrow \text{Judge}(m, \sigma, \text{pk}_{\text{sig}}, \text{pk}_{\text{san}}, \pi)$.

Correctness of Sanitizable Signature Schemes. We require the usual correctness requirements to hold. In a nutshell, every signed and sanitized message/signature pair should verify, while a honestly generated proof on a honestly generated message/signature pair should point to the correct accountable party. We refer to [15] for a formal definition, which straightforwardly extends to our framework.

5.3 Security of Sanitizable Signature Schemes

Next, we introduce our security model, where our definitions already incorporate newer insights [15, 19, 41, 48]. In particular, we mostly consider the “strong” definitions by Krenn et al. [48] as the new state-of-the-art. Due to our goals, we also see the data-structure corresponding to the admissible blocks, i.e., ADM , as an asset which needs protection, which addresses the work done by Gong et al. [41]. All formal definitions can be found in the full version of this paper.

Unforgeability. No one should be able to generate any new signature not seen before without having access to any private keys.

Immutability. Sanitizers must only be able to perform allowed modifications.

In particular, a sanitizer must not be able to modify non-admissible blocks.

Privacy. Similar to semantic security for encryption schemes, privacy captures the inability of an attacker to derive any knowledge about sanitized parts.

Transparency. An attacker cannot tell whether a specific message/signature pair has been sanitized or not.

Accountability. For signer-accountability, a signer should not be able to accuse a sanitizer if the sanitizer is actually not responsible for a given message, and vice versa for sanitizer-accountability.

5.4 Invisibility of SSSs

Next, we introduce the new property of invisibility. Basically, invisibility requires that an outsider cannot decide which blocks of a given message are admissible. With $\text{ADM}_0 \cap \text{ADM}_1$, we denote the intersection of the admissible blocks, ignoring the length of the messages.

In a nutshell, the adversary can query an LoRADM oracle which either makes ADM_0 or ADM_1 admissible in the final signature. Of course, the adversary has to be restricted to $\text{ADM}_0 \cap \text{ADM}_1$ for sanitization requests for signatures originating from those created by LoRADM and their derivatives to avoid trivial attacks. The sign oracle can be simulated by querying the LoRADM oracle with $\text{ADM}_0 = \text{ADM}_1$. We stress that our invisibility definition is very strong, as it also takes the signatures into account, much like the definitions given by Krenn et al. [48]. One can easily alter our definition to only account for the messages in question, e.g., if one wants to avoid strongly unforgeable signatures, or even allow re-randomizable signatures. An adjustment is straightforward.

Definition 16 (Invisibility). *An SSS is invisible, if for any efficient adversary \mathcal{A} there exists a negligible function ν such that $\left| \Pr[\text{Invisibility}_{\mathcal{A}}^{\text{SSS}}(\lambda) = 1] - \frac{1}{2} \right| \leq \nu(\lambda)$, where the corresponding experiment is defined in Fig. 8.*

It is obvious that invisibility is not implied by any other property. In a nutshell, taking any secure SSS, it is sufficient to non-malleably append ADM to each block $m[i]$ to prevent invisibility. Clearly, all other properties of such a construction are still preserved.

Definition 17 (Secure SSS). *We call an SSS secure, if it is correct, private, unforgeable, immutable, sanitizer-accountable, signer-accountable, and invisible.*

We do neither consider non-interactive public accountability nor unlinkability nor transparency as essential security requirements, as it depends on the concrete use-case whether these properties are required.

5.5 Construction

We now introduce our construction and use the construction paradigm of Ateniese et al. [4], enriching it with several ideas of prior work [15, 41, 50]. The main idea

Experiment $\text{Invisibility}_{\mathcal{A}}^{\text{SSS}}(\lambda)$

$\text{pp}_{\text{SSS}} \leftarrow \text{SSSParGen}(1^\lambda)$
 $(\text{pk}_{\text{sig}}, \text{sk}_{\text{sig}}) \leftarrow \text{KGen}_{\text{sig}}(\text{pp}_{\text{SSS}})$
 $(\text{pk}_{\text{san}}, \text{sk}_{\text{san}}) \leftarrow \text{KGen}_{\text{san}}(\text{pp}_{\text{SSS}})$
 $b \leftarrow \{0, 1\}$
 $\mathcal{Q} \leftarrow \emptyset$
 $a \leftarrow \mathcal{A}^{\text{Sanit}'(\cdot, \cdot, \cdot, \text{sk}_{\text{san}}), \text{Proof}(\text{sk}_{\text{sig}}, \cdot, \cdot, \cdot), \text{LoRADM}(\cdot, \cdot, \cdot, \text{sk}_{\text{sig}}, b)}(\text{pk}_{\text{sig}}, \text{pk}_{\text{san}})$
where oracle LoRADM on input of $m, \text{ADM}_0, \text{ADM}_1, \text{sk}_{\text{sig}}, b$:
return \perp , if $\text{ADM}_0(m) \neq \text{ADM}_1(m)$
let $\sigma \leftarrow \text{Sign}(m, \text{sk}_{\text{sig}}, \text{pk}_{\text{san}}, \text{ADM}_b)$
let $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(m, \sigma, \text{ADM}_0 \cap \text{ADM}_1)\}$
return σ
where oracle Sanit' on input of $m, \text{MOD}, \sigma, \text{pk}'_{\text{sig}}, \text{sk}_{\text{san}}$:
return \perp , if $\text{pk}'_{\text{sig}} = \text{pk}_{\text{sig}} \wedge \nexists (m, \sigma, \text{ADM}) \in \mathcal{Q} : \text{MOD}(\text{ADM}) = \text{true}$
let $(m', \sigma') \leftarrow \text{Sanit}(m, \text{MOD}, \sigma, \text{pk}'_{\text{sig}}, \text{sk}_{\text{san}})$
if $\text{pk}'_{\text{sig}} = \text{pk}_{\text{sig}} \wedge \exists (m, \sigma, \text{ADM}') \in \mathcal{Q} : \text{MOD}(\text{ADM}') = \text{true}$,
let $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(m', \sigma', \text{ADM}')\}$
return (m', σ')
return 1, if $a = b$
return 0

Fig. 8. Invisibility

is to hash each block using a chameleon-hash with ephemeral trapdoors, and then sign the hashes. The main trick we introduce to limit the sanitizer is that only those etd_i are given to the sanitizer, for which the respective block $m[i]$ should be sanitizable. To hide whether a given block is sanitizable, each etd_i is encrypted; a sanitizable block contains the real etd_i , while a non-admissible block encrypts a 0, where 0 is assumed to be an invalid etd . For simplicity, we require that the IND-CPA secure encryption scheme Π allows that each possible etd , as well as 0, is in the message space \mathcal{M} of Π , which can be achieved using standard embedding and padding techniques, or using KEM/DEM combinations [1]. To achieve accountability, we generate additional “tags” for a “standard” chameleon-hash (which binds everything together) in a special way, namely we use PRFs and PRGs, which borrows ideas from the construction given by Brzuska et al. [15].

Construction 3 (Secure and Transparent SSS) *The secure and transparent SSS construction is as follows:*

SSSParGen. *To generate the public parameters, do the following steps:*

1. Let $\text{pp}_{\text{ch}} \leftarrow \text{CHET.CParGen}(1^\lambda)$.
2. Let $\text{pp}'_{\text{ch}} \leftarrow \text{CH.CParGen}(1^\lambda)$.
3. Return $\text{pp}_{\text{SSS}} = (\text{pp}_{\text{ch}}, \text{pp}'_{\text{ch}})$.

KGen_{sig}. *To generate the key pair for the signer, do the following steps:*

1. Let $(\text{pk}_s, \text{sk}_s) \leftarrow \Sigma.\text{KGen}_{\text{sig}}(1^\lambda)$.
2. Pick a key for a PRF, i.e., $\kappa \leftarrow \text{PRF.KGen}_{\text{prf}}(1^\lambda)$.
3. Return $(\text{pk}_s, (\kappa, \text{sk}_s))$.

KGen_{san}. *To generate the key pair for the sanitizer, do the following steps:*

1. Let $(\text{pk}_{\text{ch}}, \text{sk}_{\text{ch}}) \leftarrow \text{CHET.CKGen}(\text{pp}_{\text{ch}})$.

2. Let $(\text{pk}'_{\text{ch}}, \text{sk}'_{\text{ch}}) \leftarrow \text{CH.CKGen}(\text{pp}'_{\text{ch}})$.
 3. Let $(\text{pk}_{\text{enc}}, \text{sk}_{\text{enc}}) \leftarrow \text{II.KGen}_{\text{enc}}(1^\lambda)$.
 4. Return $((\text{pk}_{\text{ch}}, \text{pk}'_{\text{ch}}, \text{pk}_{\text{enc}}), (\text{sk}_{\text{ch}}, \text{sk}'_{\text{ch}}, \text{sk}_{\text{enc}}))$.
- Sign.** To generate a signature σ , on input of $m = (m[1], m[2], \dots, m[\ell])$, $\text{sk}_{\text{sig}} = (\kappa, \text{sk}_s)$, $\text{pk}_{\text{san}} = (\text{pk}_{\text{ch}}, \text{pk}'_{\text{ch}}, \text{pk}_{\text{enc}})$, and ADM do the following steps:
1. If $\text{ADM}(m) \neq \text{true}$, return \perp .
 2. Draw $x_0 \leftarrow \{0, 1\}^\lambda$.
 3. Let $x'_0 \leftarrow \text{PRF.Eval}_{\text{prf}}(\kappa, x_0)$.
 4. Let $\tau \leftarrow \text{PRG.Eval}_{\text{prg}}(x'_0)$.
 5. For each $i \in \{1, 2, \dots, \ell\}$ do:
 - (a) Set $(h_i, r_i, \text{etd}_i) \leftarrow \text{CHET.CHash}(\text{pk}_{\text{ch}}, (i, m[i], \text{pk}_{\text{sig}}))$.
 - (b) If block i is not admissible, let $\text{etd}_i \leftarrow 0$.
 - (c) Compute $c_i \leftarrow \text{II.Enc}(\text{pk}_{\text{enc}}, \text{etd}_i)$.
 6. Set $(h_0, r_0) \leftarrow \text{CH.CHash}(\text{pk}'_{\text{ch}}, (0, m, \tau, \ell, \tilde{h}_{1,\ell}, \tilde{c}_{1,\ell}, \tilde{r}_{1,\ell}, \text{pk}_{\text{sig}}))$.
 7. Set $\sigma' \leftarrow \Sigma.\text{Sign}(\text{sk}_s, (x_0, \tilde{h}_{0,\ell}, \tilde{c}_{1,\ell}, \text{pk}_{\text{san}}, \text{pk}_{\text{sig}}, \ell))$.
 8. Return $\sigma = (\sigma', x_0, \tilde{r}_{0,\ell}, \tau, \tilde{c}_{1,\ell}, \tilde{h}_{0,\ell})$.
- Verify.** To verify a signature $\sigma = (\sigma', x_0, \tilde{r}_{0,\ell}, \tau, \tilde{c}_{1,\ell}, \text{etd}_0, \tilde{h}_{0,\ell})$, on input of $m = (m[1], m[2], \dots, m[\ell])$, w.r.t. to $\text{pk}_{\text{sig}} = \text{pk}_s$ and $\text{pk}_{\text{san}} = (\text{pk}_{\text{ch}}, \text{pk}'_{\text{ch}}, \text{pk}_{\text{enc}})$, do:
1. For each $i \in \{1, 2, \dots, \ell\}$ do:
 - (a) Set $b_i \leftarrow \text{CHET.CHashCheck}(\text{pk}_{\text{ch}}, (i, m[i], \text{pk}_{\text{sig}}), r_i, h_i)$. If any $b_i = \text{false}$, return **false**.
 2. Let $b_0 \leftarrow \text{CH.CHashCheck}(\text{pk}'_{\text{ch}}, (0, m, \tau, \ell, \tilde{h}_{1,\ell}, \tilde{c}_{1,\ell}, \tilde{r}_{1,\ell}, \text{pk}_{\text{sig}}), r_0, h_0)$.
 3. If $b_0 = \text{false}$, return **false**.
 4. Return $d \leftarrow \Sigma.\text{Verify}(\text{pk}_s, (x_0, \tilde{h}_{0,\ell}, \tilde{c}_{1,\ell}, \text{pk}_{\text{san}}, \text{pk}_{\text{sig}}, \ell), \sigma')$.
- Sanit.** To sanitize a signature $\sigma = (\sigma', x_0, \tilde{r}_{0,\ell}, \tau, \tilde{c}_{1,\ell}, h_{0,\ell})$, on input of $m = (m[1], m[2], \dots, m[\ell])$, w.r.t. to $\text{pk}_{\text{sig}} = \text{pk}_s$, $\text{sk}_{\text{san}} = (\text{sk}_{\text{ch}}, \text{sk}'_{\text{ch}}, \text{sk}_{\text{enc}})$, and MOD do:
1. Verify the signature, i.e., run $d \leftarrow \text{SSS.Verify}(m, \sigma, \text{pk}_{\text{sig}}, \text{pk}_{\text{san}})$. If $d = \text{false}$, return \perp .
 2. Decrypt each c_i for $i \in \{1, 2, \dots, \ell\}$, i.e., let $\text{etd}_i \leftarrow \text{II.Dec}(\text{sk}_{\text{enc}}, c_i)$. If any decryption fails, return \perp .
 3. For each index $i \in \text{MOD}$ check that $\text{etd}_i \neq 0$. If not, return \perp .
 4. For each block $m[i]' \in \text{MOD}$ do:
 - (a) Let $r'_i \leftarrow \text{CHET.Adapt}(\text{sk}_{\text{ch}}, (i, m[i], \text{pk}_{\text{sig}}), (i, m[i]', \text{pk}_{\text{sig}}), r_i, \text{etd}_i)$.
 - (b) If $r'_i = \perp$, return \perp .
 5. For each block $m[i]' \notin \text{MOD}$ do:
 - (a) Let $r'_i \leftarrow r_i$.
 6. Let $m' \leftarrow \text{MOD}(m)$.
 7. Draw $\tau' \leftarrow \{0, 1\}^{2\lambda}$.
 8. Let $r'_0 \leftarrow \text{CH.Adapt}(\text{sk}'_{\text{ch}}, (0, m, \tau, \ell, \tilde{h}_{1,\ell}, \tilde{c}_{1,\ell}, \tilde{r}_{1,\ell}, \text{pk}_{\text{sig}}), (0, m', \tau', \ell, \tilde{h}_{1,\ell}, \tilde{c}_{1,\ell}, \tilde{r}'_{1,\ell}, \text{pk}_{\text{sig}}), r_0, h_0)$.
 9. Return $(m', (\sigma', x_0, \tilde{r}'_{0,\ell}, \tau', \tilde{c}_{1,\ell}, \tilde{h}_{0,\ell}))$.
- Proof.** To create a proof π , on input of $m = (m[1], m[2], \dots, m[\ell])$, a signature σ , w.r.t. to pk_{san} and sk_{sig} , and $\{(m_i, \sigma_i) \mid i \in \mathbb{N}\}$ do:
1. Return \perp , if **false** = $\text{SSS.Verify}(m, \sigma, \text{pk}_{\text{sig}}, \text{pk}_{\text{san}})$.

2. Verify each signature in the list, i.e., run $d_i \leftarrow \text{SSS.Verify}(m_i, \sigma_i, \text{pk}_{\text{sig}}, \text{pk}_{\text{san}})$. If for any $d_i = \text{false}$, return \perp .
 3. Go through the list of (m_i, σ_i) and find a (non-trivial) colliding tuple of the chameleon-hash with (m, σ) , i.e., $h_0 = h'_0$, where also $\text{true} = \text{CH.CHashCheck}(\text{pk}'_{\text{ch}}, (0, m, \tau, \ell, \tilde{h}_{1,\ell}, \tilde{c}_{1,\ell}, \tilde{r}_{1,\ell}, \text{pk}_{\text{sig}}), r_0, h_0)$, and $\text{true} = \text{CH.CHashCheck}(\text{pk}'_{\text{ch}}, (0, m', \tau', \ell, \tilde{h}'_{1,\ell}, \tilde{c}'_{1,\ell}, \tilde{r}'_{1,\ell}, \text{pk}_{\text{sig}}), r'_0, h'_0)$ for some different tag τ' or message m' . Let this signature/message pair be $(\sigma', m') \in \{(m_i, \sigma_i) \mid i \in \mathbb{N}\}$.
 4. Return $\pi = ((\sigma', m'), \text{PRF.Eval}_{\text{prf}}(\kappa, x_0))$, where x_0 is contained in (σ, m) .
- Judge.** To find the accountable party on input of $m = (m[1], m[2], \dots, m[\ell])$, a valid signature σ , w.r.t. to $\text{pk}_{\text{san}}, \text{pk}_{\text{sig}}$, and a proof π do:
1. Check if π is of the form $((\sigma', m'), v)$ with $v \in \{0, 1\}^\lambda$. If not, return **Sig**.
 2. Also return \perp , if $\text{false} = \text{SSS.Verify}(m', \sigma', \text{pk}_{\text{sig}}, \text{pk}_{\text{san}})$, or $\text{false} = \text{SSS.Verify}(m, \sigma, \text{pk}_{\text{sig}}, \text{pk}_{\text{san}})$.
 3. Let $\tau'' \leftarrow \text{PRG.Eval}_{\text{prg}}(v)$.
 4. If $\tau' \neq \tau''$, return **Sig**.
 5. If we have $h_0 = h'_0$, $\text{true} = \text{CH.CHashCheck}(\text{pk}_{\text{ch}}, (0, m, \tau, \ell, \tilde{h}_{1,\ell}, \tilde{c}_{1,\ell}, \text{pk}_{\text{sig}}), r_0, \text{pk}_{\text{sig}}, h_0) = \text{CH.CHashCheck}(\text{pk}'_{\text{ch}}, (0, m', \tau', \ell', \tilde{h}'_{1,\ell}, \tilde{c}'_{1,\ell}, \text{pk}_{\text{sig}}), r'_0, \text{pk}_{\text{sig}}, h'_0)$, $\tilde{c}_{1,\ell} = \tilde{c}'_{1,\ell}$, $x_0 = x'_0$, $\ell = \ell'$, and $\tilde{h}_{0,\ell} = \tilde{h}'_{0,\ell}$, return **San**.
 6. Return **Sig**.

Theorem 3. If Π is IND-CPA secure, Σ , PRF, PRG, CHET are secure, CH is secure and unique, Construction 3 is a secure and transparent SSS.

Note, CHET is not required to be unique. We prove each property on its own.

Proof. Correctness follows by inspection.

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

Game 0: The original unforgeability game.

Game 1: As Game 0, but we abort if the adversary outputs a forgery (m^*, σ^*) with $\sigma^* = (\sigma'^*, x_0^*, \tilde{r}_{0,\ell}^*, \tilde{\tau}^*, \tilde{c}_{1,\ell}^*, \tilde{h}_{0,\ell}^*)$, where $(\sigma'^*, (x_0, \tilde{h}_{0,\ell}, \tilde{c}_{1,\ell}, \text{pk}_{\text{san}}, \text{pk}_{\text{sig}}, \ell))$ was never obtained from the sign or sanitizing oracle. Let this event be E_1 .

Transition - Game 0 \rightarrow Game 1: Clearly, if $(\sigma'^*, (x_0, \tilde{h}_{0,\ell}, \tilde{c}_{1,\ell}, \text{pk}_{\text{san}}, \text{pk}_{\text{sig}}, \ell))$ was never obtained by the challenger, this tuple breaks the strong unforgeability of the underlying signature scheme. The reduction works as follows. We obtain a challenge public key pk_c from a strong unforgeability challenger and embed it as pk_{sig} . For every required “inner” signature σ' , we use the signing oracle provided by the challenger. Now, whenever E_1 happens, we can output σ'^* together with the message protected by σ'^* as a forgery to the challenger. That is, E_1 happens with exactly the same probability as a forgery. Further, both games proceed identically, unless E_1 happens. Taking everything together yields $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{unf-cma}}(\lambda)$.

Game 2: Among others, we now have established that the adversary can no longer win by modifying \mathbf{pk}_{sig} , and \mathbf{pk}_{san} . We proceed as in Game 1, but abort if the adversary outputs a forgery (m^*, σ^*) , where message m^* or any of the other values protected by the outer chameleon-hash were never returned by the signer or the sanitizer oracle. Let this event be E_2 .

Transition - Game 1 \rightarrow Game 2: The probability of the abort event E_2 to happen is exactly the probability of the adversary breaking collision freeness for the outer chameleon-hash. Namely, we already established that the adversary cannot tamper with the inner signature and therefore the hash value h_0^* must be from a previous oracle query. Now, assume that we obtain \mathbf{pk}'_{ch} from a collision freeness challenger. If E_2 happens, there must be a previous oracle query with associated values $(0, m, \tau, \ell, \tilde{h}_{1,\ell}, \tilde{c}_{1,\ell}, \tilde{r}_{1,\ell}, \mathbf{pk}_{\text{sig}})$ and r_0 so that h_0^* is a valid hash with respect to some those values and r_0 . Further, we also have that $(0, m, \tau, \ell, \tilde{h}_{1,\ell}, \tilde{c}_{1,\ell}, \tilde{r}_{1,\ell}, \mathbf{pk}_{\text{sig}}) \neq (0, m^*, \tau^*, \ell^*, \tilde{h}_{1,\ell}^*, \tilde{c}_{1,\ell}^*, \tilde{r}_{1,\ell}^*, \mathbf{pk}_{\text{sig}})$, and can thus output $((0, m^*, \tau^*, \ell^*, \tilde{h}_{1,\ell}^*, \tilde{c}_{1,\ell}^*, \tilde{r}_{1,\ell}^*, \mathbf{pk}_{\text{sig}}), r_0^*, (0, m, \tau, \ell, \tilde{h}_{1,\ell}, \tilde{c}_{1,\ell}, \tilde{r}_{1,\ell}, \mathbf{pk}_{\text{sig}}), r_0, h_0^*)$ as the collision. Thus, the probability that E_2 happens is exactly the probability of a collision for the chameleon-hash. Both games proceed identically, unless E_2 happens. $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{ch-coll-res}}(\lambda)$ follows.

Game 3: As Game 2, but we abort if the adversary outputs a forgery where only the randomness r_0 changed, i.e., we have previously generated a signature with respect to r_0 so that $r_0 \neq r_0^*$. Let this be event be E_3 .

Transition - Game 2 \rightarrow Game 3: If the abort event E_3 happens, the adversary breaks uniqueness of the chameleon-hash. In particular we have values $(0, m^*, \tau^*, \ell^*, \tilde{h}_{1,\ell}^*, \tilde{c}_{1,\ell}^*, \tilde{r}_{1,\ell}^*, \mathbf{pk}_{\text{sig}})$ in the forgery which also correspond to some previous query, but r_0 from the previous query is different from r_0^* . Obtaining \mathbf{pp}'_{ch} from a uniqueness challenger thus shows that E_3 happens with exactly the same probability as the adversary breaks uniqueness of the chameleon hash. Thus, we have that $|\Pr[S_2] - \Pr[S_3]| \leq \nu_{\text{ch-unique}}(\lambda)$.

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

Immutability. We prove immutability using a sequence of games.

Game 0: The immutability game.

Game 1: As Game 0, but we abort if the adversary outputs a forgery (m^*, σ^*) with $\sigma^* = (\sigma'^*, x_0^*, \tilde{r}_{0,\ell}^*, \tilde{\tau}^*, \tilde{c}_{1,\ell}^*, \tilde{h}_{0,\ell}^*)$ where $(\sigma'^*, (x_0, \tilde{h}_{0,\ell}, \tilde{c}_{1,\ell}, \mathbf{pk}_{\text{san}}, \mathbf{pk}_{\text{sig}}, \ell))$ was never obtained from the sign oracle.

Transition - Game 0 \rightarrow Game 1: Let us use E_1 to refer to the abort event. Clearly, if $(\sigma'^*, (x_0, \tilde{h}_{0,\ell}, \tilde{c}_{1,\ell}, \mathbf{pk}_{\text{san}}, \mathbf{pk}_{\text{sig}}, \ell))$ was never obtained by the challenger, this tuple breaks the strong unforgeability of the underlying signature scheme. The reduction works as follows. We obtain a challenge public key \mathbf{pk}_c from a strong unforgeability challenger and embed it as \mathbf{pk}_{sig} . For every required “inner” signature σ' , we use the signing oracle provided by the

challenger. Now, whenever E_1 happens, we can output σ'^* together with the message protected by σ'^* as a forgery to the challenger. That is, E_1 happens with exactly the same probability as a forgery of the underlying signature scheme. Further, both games proceed identically, unless E_1 happens. Taking everything together yields $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{unf-cma}}(\lambda)$.

Game 2: As Game 1, but the challenger aborts, if the message m^* is not derivable from any returned signature. Note, we already know that tampering with the signatures is not possible, and thus pk_{sig} , and pk_{san} , are fixed. The same is true for deleting or appending blocks, as ℓ is signed in every case. Let this event be denoted E_2 .

Transition - Game 1 \rightarrow Game 2: Now assume that E_2 is non-negligible. We can then construct an adversary \mathcal{B} which breaks the private collision-resistance of the underlying chameleon-hash with ephemeral trapdoors. Let the signature returned be $\sigma^* = (\sigma'^*, x_0^*, \tilde{r}_{0,\ell^*}^*, \tilde{\tau}^*, \tilde{c}_{1,\ell^*}^*, \tilde{h}_{0,\ell^*}^*)$, while \mathcal{A} 's public key is pk^* . Due to prior game hops, we know that \mathcal{A} cannot tamper with the ‘‘inner’’ signatures. Thus, there must exist another signature $\sigma = (\sigma'^*, x_0^*, \tilde{r}'_{0,\ell^*}, \tilde{\tau}'^*, \tilde{c}_{1,\ell^*}^*, \tilde{h}_{0,\ell^*}^*)$ returned by the signing oracle. This, however, also implies that there must exist an index $i \in \{1, 2, \dots, \ell^*\}$, for which we have $\text{CHET.CHashCheck}(\text{pk}_{\text{ch}}, (i, m^*[i], \text{pk}_{\text{sig}}, r_i^*, h_i^*)) = \text{CHET.CHashCheck}(\text{pk}_{\text{ch}}, (i, m'^*[i], \text{pk}_{\text{sig}}, r_i'^*, h_i^*)) = \text{true}$, where $m^*[i] \neq m'^*[i]$ by assumption. \mathcal{B} proceeds as follows. Let q_h be the number of ‘‘inner hashes’’ created. Draw an index $i \leftarrow \{1, 2, \dots, q_h\}$. For a query $i \neq j$, proceed as in the algorithms. If $i = j$, however, \mathcal{B} returns the current public key pk_c for the chameleon-hash with ephemeral trapdoors. This key is contained in pk_{san}^* . \mathcal{B} then receives back control, and queries its CHash oracle with $(i, m[i], \text{pk}_{\text{sig}})$, where i is the current index of the message m to be signed. Then, if $((i, m^*[i], \text{pk}_{\text{sig}}), r_i^*, (i, m'^*[i], \text{pk}_{\text{sig}}), r_i'^*, h_i^*)$ is the collision w.r.t. pk_c , it can directly return it. $|\Pr[S_1] - \Pr[S_2]| \leq q_h \nu_{\text{priv-coll}}(\lambda)$ follows, as \mathcal{B} has to guess where the collision will take place.

As each hop changes the view of the adversary only negligibly, immutability is proven, as the adversary has no other way to break immutability in Game 2.

Privacy. We prove privacy; we use a sequence of games.

Game 0: The original privacy game.

Game 1: As Game 0, but we abort if the adversary queries a verifying message-signature pair (m^*, σ^*) which was never returned by the signer or the sanitizer oracle, and queries it to the sanitization or proof generation oracle.

Transition - Game 0 \rightarrow Game 1: Let us use E_1 to refer to the abort event. Clearly, whenever the adversary queries such a new pair, we can output it to break the unforgeability of our scheme, as this tuple is fresh. However, we have already proven that this can only happen with negligible probability. $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{sss-unf}}(\lambda)$ follows.

Game 2: As Game 1, but instead of hashing the blocks $(i, m_b[i], \text{pk}_{\text{sig}})$ for the inner chameleon-hashes using CHash, and then *Adapt* to $(i, m[i], \text{pk}_{\text{sig}})$, we directly apply CHash to $(i, m[i], \text{pk}_{\text{sig}})$.

Transition - Game 1 \rightarrow Game 2: Assume that the adversary can distinguish this hop. We can then construct an \mathcal{B} which wins the indistinguishability game. \mathcal{B} receives pk_c as it's own challenge, \mathcal{B} embeds pk_c as pk_{ch} , and proceeds honestly with the exception that it uses the `HashOrAdapt` oracle to generate the inner hashes. Then, whatever \mathcal{A} outputs, is also output by \mathcal{B} . $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{chet-ind}}(\lambda)$ follows.

Game 3: As Game 2, but instead of adapting $(0, m, \tau, \ell, \tilde{h}_{1,\ell}, \tilde{c}_{1,\ell}, \tilde{r}_{1,\ell}, \text{pk}_{\text{sig}})$ to the new values, directly use `CHash`.

Transition - Game 2 \rightarrow Game 3: Assume that the adversary can distinguish this hop. We can then construct an \mathcal{B} which wins the indistinguishability game. \mathcal{B} receives pk'_c as it's own challenge, \mathcal{B} embeds pk'_c as pk'_{ch} , and proceeds honestly with the exception that it uses the `HashOrAdapt` oracle to generate the outer hashes. Then, whatever \mathcal{A} outputs, is also output by \mathcal{B} . $|\Pr[S_2] - \Pr[S_3]| \leq \nu_{\text{ch-ind}}(\lambda)$ follows.

Clearly, we are now independent of the bit b . As each hop changes the view of the adversary only negligibly, privacy is proven.

Transparency. We prove transparency by showing that the distributions of sanitized and fresh signatures are indistinguishable. Note, the adversary is not allowed to query `Proof` for values generated by `Sanit/Sign`.

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

Game 1: As Game 0, but we abort if the adversary queries a valid message-signature pair (m^*, σ^*) which was never returned by any of the calls to the sanitization or signature generation oracle. Let us use E_1 to refer to the abort event.

Transition - Game 0 \rightarrow Game 1: Clearly, whenever the adversary queries such a new pair, we can output it to break the unforgeability of our scheme, as this tuple is fresh. A reduction is straightforward. Thus, we have $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{sss-unf}}(\lambda)$.

Game 2: As Game 1, but instead of computing $x'_0 \leftarrow \text{PRF.Eval}_{\text{prf}}(\lambda, x_0)$, we set $x'_0 \leftarrow \{0, 1\}^\lambda$ within every call to `Sign` in the `Sanit/Sign` oracle.

Transition - Game 1 \rightarrow Game 2: A distinguisher between these two games straightforwardly yields a distinguisher for the PRF. Thus, we have $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{ind-prf}}(\lambda)$.

Game 3: As Game 2, but instead of computing $\tau \leftarrow \text{PRG.Eval}_{\text{prg}}(x'_0)$, we set $\tau \leftarrow \{0, 1\}^{2\lambda}$ for every call to `Sign` within the `Sanit/Sign` oracle.

Transition - Game 2 \rightarrow Game 3: A distinguisher between these two games yields a distinguisher for the PRG using a standard hybrid argument. Thus, we have $|\Pr[S_2] - \Pr[S_3]| \leq q_s \nu_{\text{ind-prg}}(\lambda)$, where q_s is the number of calls to the PRG.

Game 4: As Game 3, but we abort if a tag τ was drawn twice. Let this event be E_4 .

Transition - Game 3 \rightarrow Game 4: As the tags τ are drawn completely random, event E_4 only happens with probability $\frac{q_t^2}{2^{2\lambda}}$, where q_t is the number of drawn tags. $|\Pr[S_3] - \Pr[S_4]| \leq \frac{q_t^2}{2^{2\lambda}}$ follows.

Game 5: As Game 4, but instead of hash and then adapting the inner chameleon-hashes, directly hash $(i, m[i], \mathbf{pk}_{\text{sig}})$.

Transition - Game 4 \rightarrow Game 5: Assume that the adversary can distinguish this hop. We can then construct an \mathcal{B} which wins the indistinguishability game. In particular, the reduction works as follows. \mathcal{B} receives \mathbf{pk}_c as it's own challenge, \mathcal{B} embeds \mathbf{pk}_c as \mathbf{pk}_{ch} , and proceeds honestly except that it uses the `HashOrAdapt` oracle to generate the inner hashes. Then, whatever \mathcal{A} outputs, is also output by \mathcal{B} . $|\Pr[S_4] - \Pr[S_5]| \leq \nu_{\text{ind-chet}}(\lambda)$ follows.

Game 6: As Game 5, but instead of hashing and then adapting the outer hash, we directly hash the message, i.e., $(0, m, \tau, \ell, \tilde{h}_{1,\ell}, \tilde{c}_{1,\ell}, \tilde{r}_{1,\ell}, \mathbf{pk}_{\text{sig}})$.

Transition - Game 5 \rightarrow Game 6: Assume that the adversary can distinguish this hop. We can then construct an \mathcal{B} which wins the indistinguishability game. In particular, the reduction works as follows. \mathcal{B} receives \mathbf{pk}'_c as it's own challenge, embeds \mathbf{pk}'_c as \mathbf{pk}'_{ch} , and proceeds honestly with the exception that it uses the `HashOrAdapt` oracle to generate the outer hashes. Then, whatever \mathcal{A} outputs, is also output by \mathcal{B} . $|\Pr[S_5] - \Pr[S_6]| \leq \nu_{\text{ind-ch}}(\lambda)$ follows.

We are now in the case $b = 1$, while each hop changes the view of the adversary only negligibly. This concludes the proof.

Signer-Accountability. We prove that our construction is signer-accountable by a sequence of games.

Game 0: The original signer-accountability game.

Game 1: As Game 0, but we abort if the sanitization oracle draws a tag τ' which is in the range of the PRG. Let this event be E_1 .

Transition - Game 0 \rightarrow Game 1: This hop is indistinguishable by a standard statistical argument: at most 2^λ values lie in the range of the PRG. $|\Pr[S_0] - \Pr[S_1]| \leq \frac{q_s 2^\lambda}{2^{2\lambda}} = \frac{q_s}{2^\lambda}$ follows, where q_s is the number of sanitizing requests. Note, this also means, that there exists no valid pre-image x_0 .

Game 2: As Game 1, but we now abort, if a tag was drawn twice by the sanitization oracles. Let this event be E_2 .

Transition - Game 1 \rightarrow Game 2: As the tags are drawn uniformly from $\{0, 1\}^{2\lambda}$, this case only happens with negligible probability. $|\Pr[S_1] - \Pr[S_2]| \leq \frac{q_s^2}{2^{2\lambda}}$ follows, where q_s is the number of sanitization oracle queries.

Game 3: As Game 2, but we now abort, if the adversary was able to find $(\mathbf{pk}^*, \pi^*, m^*, \sigma^*)$ for some message m^* with a τ^* which was never returned by the sanitization oracle. Let this event be E_3 .

Transition - Game 2 \rightarrow Game 3: In the previous games we have already established that the sanitizer oracle will never return a signature with respect to a tag τ in the range of the PRG. Thus, if event E_3 happens, we know by the condition checked in step 4 of `Judge` that at least one of the tags (either τ^* in σ^* , or τ^π in π^*) was chosen by the adversary, which, in further consequence, implies a collision for `CH`. Namely, assume that E_3 happens with non-negligible probability. Then we embed the challenge public key \mathbf{pk}_c

in pk'_{ch} , and use the provided adaption oracle to simulate the sanitizer oracle. If E_3 happens we can output $((0, m^*, \tau^*, \ell^*, \tilde{h}_{1,\ell^*}^*, \tilde{c}_{1,\ell^*}^*, \tilde{r}_{1,\ell^*}^*, \text{pk}^*), r_0^*, (0, m'^*, \tau'^*, \ell^*, \tilde{h}_{1,\ell^*}^*, \tilde{c}_{1,\ell^*}^*, \tilde{r}_{1,\ell^*}^*, \text{pk}^*), r_0^*, h_0^*)$, as a valid collision. These values can simply be compiled using π^* , m^* , and σ^* . $|\Pr[S_2] - \Pr[S_3]| \leq \nu_{\text{ch-coll-res}}(\lambda)$ follows.

Game 4: As Game 3, but we now abort, if the adversary was able to find $(\text{pk}^*, \pi^*, m^*, \sigma^*)$ for a new message m^* which was never returned by the sanitization oracle. Let this event be E_4 .

Transition - Game 3 \rightarrow Game 4: Assume that E_4 happens with non-negligible probability. In the previous games we have already established that the only remaining possibility for the adversary is to re-use tags τ^*, τ^π corresponding to some query/response to the sanitizer oracle. Then, m^* must be fresh, as it was never returned by the sanitization oracle by assumption. Thus, $((0, m^*, \tau^*, \ell^*, \tilde{h}_{1,\ell^*}^*, \tilde{c}_{1,\ell^*}^*, \tilde{r}_{1,\ell^*}^*, \text{pk}^*), r_0^*, (0, m'^*, \tau'^*, \ell^*, \tilde{h}_{1,\ell^*}^*, \tilde{c}_{1,\ell^*}^*, \tilde{r}_{1,\ell^*}^*, \text{pk}^*), r_0^*, h_0^*)$, is a valid collision. These values can simply be compiled using π^* , m^* , and σ^* . $|\Pr[S_3] - \Pr[S_4]| \leq \nu_{\text{ch-coll-res}}(\lambda)$ follows.

In the last game the adversary can no longer win; each hop only changes the view negligibly. This concludes the proof.

Sanitizer-Accountability. We prove that our construction is sanitizer-accountable by a sequence of games.

Game 0: The original sanitizer-accountability definition.

Game 1: As Game 0, but we abort if the adversary outputs a forgery $(m^*, \sigma^*, \text{pk}^*)$ with $\sigma^* = (\sigma'^*, x_0^*, \tilde{r}_{0,\ell^*}^*, \tilde{\tau}^*, \tilde{c}_{1,\ell^*}^*, \tilde{h}_{0,\ell^*}^*)$ where $(\sigma'^*, (x_0, \tilde{h}_{0,\ell}, \tilde{c}_{1,\ell}, \text{pk}^*, \text{pk}_{\text{sig}}, \ell))$ was never obtained from the signing oracle.

Transition - Game 0 \rightarrow Game 1: Let us use E_1 to refer to the abort event. Clearly, if $(\sigma'^*, (x_0, \tilde{h}_{0,\ell}, \tilde{c}_{1,\ell}, \text{pk}^*, \text{pk}_{\text{sig}}, \ell))$ was never obtained by the challenger, this tuple breaks the strong unforgeability of the underlying signature scheme. The reduction works as follows. We obtain a challenge public key pk_c from a strong unforgeability challenger and embed it as pk_{sig} . For every required “inner” signature σ' , we use the signing oracle provided by the challenger. Now, whenever E_1 happens, we can output σ'^* together with the message protected by σ'^* as a forgery to the challenger. That is, E_1 happens with exactly the same probability as a forgery. Further, both games proceed identically, unless E_1 happens. Taking everything together yields $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{unf-cma}}(\lambda)$.

Game 2: As Game 1, but we abort if the adversary outputs a forgery where only the randomness r_0 changed, i.e., we have previously generated a signature with respect to r_0 so that $r_0 \neq r_0^*$. Let this event be E_2 .

Transition - Game 1 \rightarrow Game 2: If the abort event E_2 happens, the adversary breaks uniqueness of the chameleon-hash. In particular we have values $(0, m^*, \tau^*, \ell^*, \tilde{h}_{1,\ell^*}^*, \tilde{c}_{1,\ell^*}^*, \tilde{r}_{1,\ell^*}^*, \text{pk}_{\text{sig}}^*)$ in the forgery which also correspond to some previous query, but r_0 from the previous query is different from r_0^* . Obtaining pp'_{ch} from a uniqueness challenger thus shows that E_2 happens with exactly

the same probability as the adversary breaks uniqueness of the chameleon hash and we have that $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{ch-unique}}(\lambda)$.

In Game 2 the forgery is different from any query/answer tuple obtained using `Sign` by definition. Due to the previous hops, the only remaining possibility is a collision in the outer chameleon-hash, i.e., for $h_0^* = h_0'^*$ we have $\text{CH.CHashCheck}(\text{pk}'^*, (0, m^*, \tau^*, \ell^*, \tilde{h}_{1,\ell^*}^*, \tilde{c}_{1,\ell^*}^*, \tilde{r}_{1,\ell^*}^*, \text{pk}_{\text{sig}}), r_0^*, h_0^*) = \text{CH.CHashCheck}(\text{pk}'^*, (0, m'^*, \tau'^*, \ell'^*, \tilde{h}_{1,\ell'^*}^*, \tilde{c}_{1,\ell'^*}^*, \tilde{r}_{1,\ell'^*}^*, \text{pk}_{\text{sig}}), r_0'^*, h_0'^*) = \text{true}$. In this case the `Judge` algorithm returns `San` and $\Pr[S_2] = 0$ which concludes the proof.

Invisibility. We prove that our construction is invisible by a sequence of games. The idea is to show that we can simulate the view of the adversary without giving out any useful information at all.

Game 0: The original invisibility game, i.e., the challenger runs the experiment as defined.

Game 1: As Game 0, but we abort if the adversary queries a valid message-signature pair (m^*, σ^*) which was never returned by the signer or the sanitizer oracle to the sanitization or proof generation oracle.

Transition - Game 0 \rightarrow Game 1: Let us use E_1 to refer to the abort event. Clearly, whenever the adversary outputs such a new pair, we can output it to break unforgeability of our scheme, as this tuple is fresh. However, we have already proven that this can only happen with negligible probability. $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{SSS-unf}}(\lambda)$ follows.

Game 2: As Game 1, but we internally keep all `etdi`.

Transition - Game 1 \rightarrow Game 2: This is only a conceptual change. $|\Pr[S_1] - \Pr[S_2]| = 0$ follows.

Game 3: As Game 2, but we encrypt only zeroes instead of the real `etdi` in LoRADM independent of whether block are admissible or not. Note, the challenger still knows all `etdi`, and can thus still sanitize correctly.

Transition - Game 2 \rightarrow Game 3: A standard reduction, using hybrids, shows that this hop is indistinguishable by the IND-CPA security of the encryption scheme used. $|\Pr[S_2] - \Pr[S_3]| \leq q_h \nu_{\text{ind-cpa}}(\lambda)$ follows, where q_h is the number of generated ciphertexts by LoRADM.¹¹

At this point, the distribution is independent of the LoRADM oracle. Note, the sanitization, and proof oracles, can be still be simulated without any restrictions, as each `etdi` is known to the challenger. Thus, the view the adversary receives is now completely independent of the bit b used in the invisibility definition. As each hop only changes the view of the adversary negligibly, our construction is thus proven to be invisible. \square

¹¹ We note that IND-CPA security of the encryption scheme Π is sufficient, as the abort in Game 1 ensures that the adversary can only submit queries with respect to ciphertexts which were previously generated in the reduction, i.e., where we can simply look up the respective values `etdi` instead of decryption.

Acknowledgements. We are grateful to the anonymous reviewers of PKC 2017 for providing valuable comments and suggestions that helped to significantly improve the presentation of the paper.

References

1. Abe, M., Gennaro, R., Kurosawa, K.: Tag-kem/dem: A new framework for hybrid encryption. *J. Cryptology* 21(1), 97–130 (2008)
2. Ahn, J.H., Boneh, D., Camenisch, J., Hohenberger, S., a. shelat, Waters, B.: Computing on authenticated data. In: TCC. pp. 1–20 (2012)
3. Alsouri, S., Dagdelen, Ö., Katzenbeisser, S.: Group-based attestation: Enhancing privacy and management in remote attestation. In: Trust. pp. 63–77 (2010)
4. Ateniese, G., Chou, D.H., de Medeiros, B., Tsudik, G.: Sanitizable signatures. In: ESORICS. pp. 159–177 (2005)
5. Ateniese, G., Magri, B., Venturi, D., Andrade, E.R.: Redactable blockchain - or - rewriting history in bitcoin and friends. IACR Cryptology ePrint Archive 2016, 757 (2016)
6. Ateniese, G., de Medeiros, B.: Identity-Based Chameleon Hash and Applications. In: Financial Cryptography. pp. 164–180 (2004)
7. Ateniese, G., de Medeiros, B.: On the key exposure problem in chameleon hashes. In: SCN. pp. 165–179 (2004)
8. Bao, F., Deng, R.H., Ding, X., Lai, J., Zhao, Y.: Hierarchical identity-based chameleon hash and its applications. In: ACNS. pp. 201–219 (2011)
9. Bellare, M., Ristov, T.: A characterization of chameleon hash functions and new, efficient designs. *J. Cryptology* 27(4), 799–823 (2014)
10. Bilzhausen, A., Huber, M., Pöhls, H.C., Samelin, K.: Cryptographically Enforced Four-Eyes Principle. In: ARES. pp. 760–767 (2016)
11. Blazy, O., Kakvi, S.A., Kiltz, E., Pan, J.: Tightly-secure signatures from chameleon hash functions. In: PKC. pp. 256–279 (2015)
12. Brassard, G., Chaum, D., Crépeau, C.: Minimum disclosure proofs of knowledge. *J. Comput. Syst. Sci.* 37(2), 156–189 (1988)
13. Bresson, E., Catalano, D., Gennaro, R.: Improved on-line/off-line threshold signatures. In: PKC. pp. 217–232 (2007)
14. Brzuska, C., Busch, H., Dagdelen, Ö., Fischlin, M., Franz, M., Katzenbeisser, S., Manulis, M., Onete, C., Peter, A., Poettering, B., Schröder, D.: Redactable Signatures for Tree-Structured Data: Definitions and Constructions. In: ACNS. pp. 87–104 (2010)
15. Brzuska, C., Fischlin, M., Freudenreich, T., Lehmann, A., Page, M., Schelbert, J., Schröder, D., Volk, F.: Security of Sanitizable Signatures Revisited. In: PKC. pp. 317–336 (2009)
16. Brzuska, C., Fischlin, M., Lehmann, A., Schröder, D.: Sanitizable signatures: How to partially delegate control for authenticated data. In: BIOSIG. pp. 117–128 (2009)
17. Brzuska, C., Fischlin, M., Lehmann, A., Schröder, D.: Unlinkability of Sanitizable Signatures. In: PKC. pp. 444–461 (2010)
18. Brzuska, C., Pöhls, H.C., Samelin, K.: Non-Interactive Public Accountability for Sanitizable Signatures. In: EuroPKI. pp. 178–193 (2012)
19. Brzuska, C., Pöhls, H.C., Samelin, K.: Efficient and Perfectly Unlinkable Sanitizable Signatures without Group Signatures. In: EuroPKI. pp. 12–30 (2013)

20. Camenisch, J., Stadler, M.: Efficient group signature schemes for large groups (extended abstract). In: CRYPTO. pp. 410–424 (1997)
21. Canard, S., Jambert, A.: On extended sanitizable signature schemes. In: CT-RSA. pp. 179–194 (2010)
22. Canard, S., Jambert, A., Lescuyer, R.: Sanitizable signatures with several signers and sanitizers. In: AFRICACRYPT. pp. 35–52 (2012)
23. Canard, S., Laguillaumie, F., Milhau, M.: Trapdoor sanitizable signatures and their application to content protection. In: ACNS. pp. 258–276 (2008)
24. Canard, S., Lescuyer, R.: Protecting privacy by sanitizing personal data: a new approach to anonymous credentials. In: ASIACCS. pp. 381–392 (2013)
25. Catalano, D., Raimondo, M.D., Fiore, D., Gennaro, R.: Off-line/on-line signatures: Theoretical aspects and experimental results. In: PKC. pp. 101–120 (2008)
26. Chen, X., Tian, H., Zhang, F., Ding, Y.: Comments and improvements on key-exposure free chameleon hashing based on factoring. In: Inscrypt. pp. 415–426 (2010)
27. Chen, X., Zhang, F., Susilo, W., Mu, Y.: Efficient generic on-line/off-line signatures without key exposure. In: ACNS. pp. 18–30 (2007)
28. Demirel, D., Derler, D., Hanser, C., Pöhls, H.C., Slamanig, D., Traverso, G.: PRISMACLOUD D4.4: Overview of Functional and Malleable Signature Schemes. Tech. rep., H2020 Prismacloud, www.prismacloud.eu (2015)
29. Derler, D., Hanser, C., Pöhls, H.C., Slamanig, D.: Towards authenticity and privacy preserving accountable workflows. In: Privacy and Identity Management. pp. 170–186 (2015)
30. Derler, D., Hanser, C., Slamanig, D.: Blank digital signatures: Optimization and practical experiences. In: Privacy and Identity Management for the Future Internet in the Age of Globalisation, vol. 457, pp. 201–215. Springer (2014)
31. Derler, D., Slamanig, D.: Rethinking privacy for extended sanitizable signatures and a black-box construction of strongly private schemes. In: ProvSec. pp. 455–474 (2015)
32. Even, S., Goldreich, O., Micali, S.: On-line/off-line digital signatures. *J. Cryptology* 9(1), 35–67 (1996)
33. Fehr, V., Fischlin, M.: Sanitizable signcryption: Sanitization over encrypted data (full version). IACR Cryptology ePrint Archive, Report 2015/765 (2015)
34. Fischlin, M.: Trapdoor Commitment Schemes and Their Applications. Ph.D. thesis, University of Frankfurt (2001)
35. Fleischhacker, N., Krupp, J., Malavolta, G., Schneider, J., Schröder, D., Simkin, M.: Efficient unlinkable sanitizable signatures from signatures with rerandomizable keys. In: PKC-1. pp. 301–330 (2016)
36. Gao, W., Li, F., Wang, X.: Chameleon hash without key exposure based on schnorr signature. *Computer Standards & Interfaces* 31(2), 282–285 (2009)
37. Gao, W., Wang, X., Xie, D.: Chameleon Hashes Without Key Exposure Based on Factoring. *J. Comput. Sci. Technol.* 22(1), 109–113 (2007)
38. Ghosh, E., Goodrich, M.T., Ohrimenko, O., Tamassia, R.: Fully-dynamic verifiable zero-knowledge order queries for network data. ePrint 2015, 283 (2015)
39. Ghosh, E., Ohrimenko, O., Tamassia, R.: Zero-knowledge authenticated order queries and order statistics on a list. In: ACNS. vol. 2015, pp. 149–171 (2015)
40. Goldwasser, S., Micali, S., Rivest, R.L.: A Digital Signature Scheme Secure Against Adaptive Chosen-Message Attacks. *SIAM Journal on Computing* 17, 281–308 (1988)
41. Gong, J., Qian, H., Zhou, Y.: Fully-secure and practical sanitizable signatures. In: Inscrypt. vol. 6584, pp. 300–317 (2011)

42. Hanser, C., Slamanig, D.: Blank digital signatures. In: ASIACCS. pp. 95 – 106 (2013)
43. Hanzlik, L., Kutylowski, M., Yung, M.: Hard invalidation of electronic signatures. In: ISPEC. pp. 421–436 (2015)
44. Hohenberger, S., Waters, B.: Short and stateless signatures from the RSA assumption. In: CRYPTO. pp. 654–670 (2009)
45. Höhne, F., Pöhls, H.C., Samelin, K.: Rechtsfolgen editierbarer signaturen. *Datenschutz und Datensicherheit* 36(7), 485–491 (2012)
46. Klonowski, M., Lauks, A.: Extended Sanitizable Signatures. In: ICISC. pp. 343–355 (2006)
47. Krawczyk, H., Rabin, T.: Chameleon Hashing and Signatures. In: NDSS. pp. 143–154 (2000)
48. Krenn, S., Samelin, K., Sommer, D.: Stronger security for sanitizable signatures. In: DPM. pp. 100–117 (2015)
49. Lai, R.W.F., Zhang, T., Chow, S.S.M., Schröder, D.: Efficient sanitizable signatures without random oracles. In: ESORICS-1. pp. 363–380 (2016)
50. de Meer, H., Pöhls, H.C., Posegga, J., Samelin, K.: Scope of security properties of sanitizable signatures revisited. In: ARES. pp. 188–197 (2013)
51. de Meer, H., Pöhls, H.C., Posegga, J., Samelin, K.: On the relation between redactable and sanitizable signature schemes. In: ESSoS. pp. 113–130 (2014)
52. Mohassel, P.: One-time signatures and chameleon hash functions. In: SAC. pp. 302–319 (2010)
53. Pedersen, T.P.: Non-interactive and information-theoretic secure verifiable secret sharing. In: CRYPTO. pp. 129–140 (1991)
54. Pöhls, H.C., Peters, S., Samelin, K., Posegga, J., de Meer, H.: Malleable signatures for resource constrained platforms. In: WISTP. pp. 18–33 (2013)
55. Pöhls, H.C., Samelin, K.: Accountable redactable signatures. In: ARES. pp. 60–69 (2015)
56. Pöhls, H.C., Samelin, K., Posegga, J.: Sanitizable Signatures in XML Signature - Performance, Mixing Properties, and Revisiting the Property of Transparency. In: ACNS. LNCS, vol. 6715, pp. 166–182. Springer (2011)
57. Ren, Q., Mu, Y., Susilo, W.: Mitigating Phishing by a New ID-based Chameleon Hash without Key Exposure. In: AusCERT. pp. 1–13 (2007)
58. Shamir, A., Tauman, Y.: Improved online/offline signature schemes. In: CRYPTO. pp. 355–367 (2001)
59. Yum, D.H., Seo, J.W., Lee, P.J.: Trapdoor sanitizable signatures made easy. In: ACNS. pp. 53–68 (2010)
60. Zhang, F., Safavi-naini, R., Susilo, W.: Id-based chameleon hashes from bilinear pairings. *IACR Cryptology ePrint Archive* 2003, 208 (2003)
61. Zhang, R.: Tweaking TBE/IBE to PKE transforms with Chameleon Hash Functions. In: ACNS. pp. 323–339 (2007)