



Research Article

Bringing Order to Chaos: The Case of Collision-Resistant Chameleon-Hashes

David Derler

DFINITY, Zurich, Switzerland
david@dfinity.org

Kai Samelin

Hamburg, Germany
kaispapers@gmail.com

Daniel Slamanig

Research Institute CODE, Universität der Bundeswehr München, Munich, Germany
daniel.slamanig@unibw.de

Communicated by Stefano Tessaro

Received 18 July 2022 / Revised 6 May 2024 / Accepted 3 June 2024

Abstract. Chameleon-hash functions, introduced by Krawczyk and Rabin (NDSS'00), are trapdoor collision-resistant hash functions parametrized by a public key. If the corresponding secret key is known, arbitrary collisions for the hash function can be found efficiently. Chameleon-hash functions have prominent applications in the design of cryptographic primitives, such as lifting non-adaptively secure signatures to adaptively secure ones. Recently, this primitive also received a lot of attention as a building block in more complex cryptographic applications, ranging from editable blockchains to advanced signature and encryption schemes. We observe that, in latter applications, various different notions of collision-resistance are used, and it is not always clear if the respective notion really covers what seems intuitively required by the application. Therefore, we revisit existing collision-resistance notions in the literature, study their relations, and by means of selected applications discuss which practical impact different notions of collision-resistance might have. Moreover, we provide a stronger, and arguably more desirable, notion of collision-resistance than what is known from the literature (which we call full collision-resistance). Finally, we present a surprisingly simple, and efficient, black-box construction of chameleon-hash functions achieving this strong notion of full collision-resistance.

1. Introduction

A chameleon-hash function (CH) is a trapdoor collision-resistant hash function parametrized by a public key. If the corresponding secret key is known, arbitrary collisions for the hash function, i.e., distinct messages $m \neq m'$ yielding the same hash value h , can be

efficiently found. Over the years, they have proven to be a very useful tool in theory, as well as in practice. Exemplary, CHs have been suggested by Shamir and Tauman [52] to construct online/offline signatures [21,33,34] (cf. also Sect. 6). Moreover, Shamir and Tauman in [52] showed that CHs can be used to generically lift non-adaptively secure signature schemes to adaptively secure ones, which has subsequently been used for instance by Hohenberger and Waters [43] to obtain short signatures under the RSA assumption in the standard model. If CHs are tightly secure, they can be used to generically construct tightly secure signatures [14]. Likewise, CHs are used to generically construct strong one-time signatures as shown by Mohassel [49], inspired by a concrete construction by Groth [38]. Zhang [57] shows how to construct IND-CCA secure public-key encryption from tag-based encryption (TBE) or identity-based encryption (IBE) and CHs. Bellare and Ristov [11,12] made the interesting discovery that chameleon-hashes in the sense of Krawczyk and Rabin [46] are equivalent to Σ -protocols, i.e., three round public-coin honest-verifier zero-knowledge proofs of knowledge. CHs are also used to construct sanitizable signatures [4,17,18], i.e., signatures where a designated entity can modify certain parts of a signed message without invalidating the respective signature under controlled conditions. Furthermore, CHs have been used by Steinfeld et al. [54] to extend Schnorr and RSA signatures to the universal designated-verifier setting [53]. Also, different flavors of chameleon-hashing such as (hierarchical) identity-based [6,8] or policy-based chameleon-hash functions [26,51] have been studied.

In a more applied setting, CHs have shown to be valuable to construct integrity measurement and remote attestation mechanisms (denoted chameleon attestation) [2], and are used in vehicular ad-hoc networks (VANETs) [41] or handover authentication in mobile networks [22]. More recently, CHs have been used as a means to rewrite blocks in blockchains by replacing the hash function to chain blocks and/or to hash transactions by chameleon-hashes [5,26], to which we come back in Sect. 6. This brief discussion already shows that chameleon-hashes are used in a wide spectrum of different applications requiring different strength of the respective chameleon-hash. Consequently, authors often introduce some ad-hoc notion of collision-resistance for their applications, or even ignore that applications might require a stronger notion. Subsequently, we briefly discuss the different notions which are most commonly found in the literature.

1.1. Formalizing Chameleon-Hashes

The concept of chameleon-hashing dates back to the notion of trapdoor commitments introduced by Brassard et al. [16] and was firstly coined chameleon-hashing by Krawczyk and Rabin [46] with an instantiation based on the well-known trapdoor-commitment scheme by Pedersen [50]. Later, Ateniese and de Medeiros [7] observed that the initial collision-resistance notion (which we denote *W-CollRes*) is rather weak (it does not give the adversary access to any collisions), and, more importantly, it is also satisfied by chameleon-hashes suffering from the key-exposure problem. Namely, when seeing a single collision for some hash h , it allows to publicly extract the secret trapdoor. Thus, any further guarantees are lost. While this is a desirable property for the initial use in chameleon signatures [46], and is also sufficient for the lifting compiler to adaptively secure signatures [52] (as no collision is ever revealed), it is too weak for many other applications. The key-exposure freeness definition by Ateniese and de Medeiros [7] is

for the specific case of public-coin chameleon-hashing (where verifying the chameleon-hash is essentially re-computing it). To address this, Ateniese et al. [5] introduced a related notion called enhanced collision-resistance (which we denote **E-CollRes**) for the generalized case of secret-coin chameleon-hashing (which is the setting that we also consider). The latter notion allows the adversary to see collisions, but it is not allowed to see any collision for the target hash, i.e., the hash corresponding to the collision it computes. Hence, once a single collision for a hash h is seen, an adversary can potentially find arbitrary collisions for that particular hash h . Recently, Khalili et al. [45] have pointed out issues regarding the practicality of the concrete random-oracle model instantiation,¹ proposed by Ateniese et al. in [5], and propose alternative constructions in the standard model. In another work, Camenisch et al. [18] proposed an alternative collision-resistance notion which allows the adversary to see arbitrary collisions also for the target hash, but not for the target message, i.e., the message used in the collision output by the adversary has never been queried. In other words, once a collision for a message m is seen, an adversary is allowed to find arbitrary other hashes h' with the queried messages. Arguably, this notion seems more realistic as it is better compatible with practical applications (e.g., one can often make the messages unique by appending a tag/nonce), and thus we denote it as standard collision-resistance (or **S-CollRes**).

1.2. Motivation and Contribution

The previous discussion already illustrates that there are many different collision-resistance notions. While this does not necessarily point to an issue, we observe that it is not always clear whether the respective notion does really cover what is required by the respective application. Moreover, it is not clear if the last notion discussed above (**S-CollRes**) is already the most desirable notion, or, if even stronger notions are achievable, and do have practical relevance. Motivated by these observations, we provide the following contributions:

1.2.1. Relations among Properties

We discuss the different security notions of chameleon-hashes, and rigorously study relations among them. Most importantly, we, for the first time, clarify the picture of existing collision-resistance notions by showing implications, and separations, (cf. Figure 1 for an overview). In the course of showing separations, we also provide a construction of a chameleon-hash satisfying the **E-CollRes** notion, but which clearly demonstrates weaknesses of this notion.

1.2.2. Stronger Notion

We find that the strongest existing collision-resistance notions, i.e., **E-CollRes** and **S-CollRes** (which are incomparable), might still be too weak for practical applications,

¹The requirement for an invertible encoding into the group introduces an enormous efficiency penalty, and thus their instantiation is incomplete. Moreover, it was only recently shown that the proposed chameleon-hash fulfills our stronger definition [23].




Fig. 1. Relations between CH collision-resistance properties .

see, e.g., Sect. 6. In particular, even if **S-CollRes** is satisfied, the hash values might still be malleable leaving space for potential real-world attacks. Consequently, we propose a stronger notion coined full collision-resistance (or **F-CollRes** for short), which enforces that the adversary cannot (except with negligible probability) output *any* new collisions and covers what one intuitively expects from collision-resistance.

1.2.3. Black-Box Construction

We present a simple, yet elegant, black-box construction of a chameleon-hash function satisfying this strong **F-CollRes** notion. Considering the complexity of existing constructions in [5,45], this is somewhat surprising. To recall, the construction from Ateniese et al. [5] starts from a public-coin chameleon-hash function that satisfies **W-CollRes**, uses an IND-CPA-secure encryption-scheme to encrypt the randomness of the chameleon-hash and then uses a true-simulation extractable (tSE) NIZK [32],² which is, in turn, based on a NIZK and an IND-CCA secure public-key encryption scheme, to prove that the ciphertext is an encryption of the randomness. The constructions by Khalili et al. [45], which avoid the aforementioned issues with [5], are based on another new public-coin chameleon-hash function that satisfies **W-CollRes** and then either uses Groth–Sahai NIZK proofs [40] and the IND-CCA secure Cramer–Shoup encryption scheme [25] or a succinct non-interactive argument of knowledge (SNARK). Both constructions by Khalili et al. [45] basically follow the generic template in [5]. In contrast, our black-box construction of a **F-CollRes** chameleon-hash is constructed from perfectly correct (multi-challenge) IND-CPA secure encryption, e.g., ElGamal encryption, and a simulation-sound extractable non-interactive zero-knowledge proof (SSE-NIZK), e.g., applying the compiler of Faust et al. [35] to a Fiat-Shamir transformed Σ -protocol. The basic idea is that the chameleon-hash is the encryption c of the message m and the randomness of the chameleon-hash is a NIZK proof s.t. either c correctly encrypts m under the pk of CH *or* one knows the secret key sk corresponding to pk . Interestingly, already a perfectly binding commitment (without any hiding) is sufficient to achieve the **F-CollRes** notion, but instead a multi-challenge IND-CPA secure encryption scheme as a perfectly binding commitment is used to additionally achieve the indistinguishability property of the CH, i.e., that fresh and adapted hashes are indistinguishable, a notion that is considered standard for chameleon-hashes.

²In true-simulation extractability the simulator can only be used for statements inside the language.

1.2.4. Applications

We discuss how our stronger notion allows to strengthen the security of existing applications. In particular, in Sect. 6 we discuss what problems may be caused by different notions of collision-resistance within recent applications to redactable blockchains [5, 26]. Here, either the hash function to chain blocks in a blockchain or the hash functions to aggregate transactions within single blocks (usually by means of a Merkle-tree) are replaced by a chameleon-hash function. Moreover, we take a second look at online/offline signatures and discuss how chameleon-hashes providing a stronger collision-resistance notion than the *W-CollRes* notion used by Shamir and Tauman in [52] allows to re-use offline signatures and add more robustness at the cost of a more expensive offline phase and a slightly more costly online phase.

1.3. Differences to the Conference Version

Compared to the conference version published at IACR PKC 2020 [29], this version in Sect. 3.3 includes a more complete treatment of indistinguishability and in particular stronger indistinguishability notions and their relations. Moreover, in Sect. 4.1 it includes examples of existing chameleon-hashes providing the *W-CollRes* and *S-CollRes* notions, and, in Sect. 4.2 the full proofs of our construction providing *E-CollRes*. Finally, in Sect. 6.2 as an additional application we discuss the use of chameleon-hashes with stronger collision-resistance notions in online/offline signatures.

1.4. Follow-up Work

Derler et al. in SCN'20 [28] show how to remove the requirement to rely on public-key encryption from the approach presented in this paper. In particular, they show how to construct fully collision-resistant chameleon-hashes based on SSE NIZKs and non-interactive commitment schemes. They then present an instantiation from the discrete logarithm (DL) problem and a concrete construction from the learning parity with noise (LPN) problem. Latter yields the first chameleon-hash from post-quantum assumptions that provides a collision-resistance notion stronger than *W-CollRes* (as, e.g., the lattice-based chameleon-hash by Cash et al. from EC'10 [19]). In PKC'24, Li and Liu [47] introduce a lattice-based *F-CollRes* chameleon-hash without resorting to random oracles or NIZK proofs by relying on the new notion of tagged chameleon hashes. Very recently, Bellare, Riepel and Shea [13] initiated the formal study of backdoored hash functions, which are closely related to chameleon-hashes, and introduce a notion of *F-CollRes* for such hash functions.

2. Preliminaries

2.1. Notation

With $\lambda \in \mathbb{N}$ we denote our security parameter. All algorithms implicitly take 1^λ as an additional input. We write $a \leftarrow_{\S} A(x)$ if the output of a probabilistic algorithm A with input x is assigned to a and use $a \leftarrow A(x)$ if A is deterministic. An algorithm

is efficient, if it runs in probabilistic polynomial time (PPT) in the length of its input. All algorithms are PPT, if not explicitly mentioned otherwise. If we want to make the random coins used by an algorithm A explicit, we use the notation $a \leftarrow_{\$} A(x; \xi)$. We write $(a; \xi) \leftarrow_{\$} A(x)$, if we need to access the random coins ξ internally drawn by A . Most algorithms may return a special error symbol $\perp \notin \{0, 1\}^*$, denoting an exception. Returning output ends execution of an algorithm or an oracle. To make the presentation in the security proofs more compact, we occasionally use $(a, \perp) \leftarrow_{\$} A(x)$ to indicate that the second output is either ignored or not returned by A . If S is a finite set, we write $a \leftarrow_{\$} S$ to denote that a is chosen uniformly at random from S . \mathcal{M} denotes a message space of a scheme, and we generally assume that \mathcal{M} is derivable from the scheme's public parameters or its public key. For a list we require that there is an injective, and efficiently reversible, encoding, that maps the list to $\{0, 1\}^*$. A function $\nu : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ is negligible, if it vanishes faster than every inverse polynomial, i.e., $\forall k \in \mathbb{N}, \exists n_0 \in \mathbb{N}$ such that $\nu(n) \leq n^{-k}, \forall n > n_0$.

2.2. Building Blocks

We now present the building blocks we require. These include key-verifiable multi-challenge IND-CPA (mcIND-CPA) secure public-key encryption schemes Ω , digital signature schemes Σ , and non-interactive zero-knowledge proofs Π .

2.2.1. Public-Key Encryption Schemes

Subsequently, we define public-key encryption schemes.

Definition 1. (Public-Key Encryption Scheme) A public-key encryption scheme Ω consists of five algorithms $\{\text{PG}_{\Omega}, \text{KG}_{\Omega}, \text{Enc}, \text{Dec}, \text{KVf}_{\Omega}\}$, such that:

PG_{Ω} . The algorithm PG_{Ω} outputs the public parameters of the scheme:

$$\text{pp}_{\Omega} \leftarrow_{\$} \text{PG}_{\Omega}(1^{\lambda}).$$

It is assumed that pp_{Ω} is an implicit input to all other algorithms.

KG_{Ω} . The algorithm KG_{Ω} outputs the key pair, on input pp_{Ω} :

$$(\text{sk}_{\Omega}, \text{pk}_{\Omega}) \leftarrow_{\$} \text{KG}_{\Omega}(\text{pp}_{\Omega}).$$

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

$$c \leftarrow_{\$} \text{Enc}(\text{pk}_{\Omega}, m).$$

Dec . The deterministic algorithm Dec outputs a message $m \in \mathcal{M} \cup \{\perp\}$ on input sk_{Ω} , and a ciphertext c :

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

```

Exp $\mathcal{A}, \Omega$ mc-IND-CPA( $\lambda$ ):
   $pp_\Omega \leftarrow_{\$} PG_\Omega(1^\lambda)$ 
   $(sk_\Omega, pk_\Omega) \leftarrow_{\$} KG_\Omega(pp_\Omega)$ 
   $b \leftarrow_{\$} \{0, 1\}$ 
   $a \leftarrow_{\$} \mathcal{A}^{Enc'(pk_\Omega, \cdot, b)}(pk_\Omega)$ 
  where  $Enc'$  on input  $pk_\Omega, m_0, m_1, b$ :
    If  $m_0 \notin \mathcal{M} \vee m_1 \notin \mathcal{M} \vee |m_0| \neq |m_1|$ :
       $c \leftarrow \perp$ 
    Else:
       $c \leftarrow_{\$} Enc(pk_\Omega, m_b)$ 
  return  $c$ 
return 1, if  $a = b$ 
return 0

```

Fig. 2. Multi-Challenge IND-CPA Security.

KVf_Ω . The deterministic algorithm KVf_Ω decides whether a given public key pk_Ω corresponds to a given secret key sk_Ω :

$$d \leftarrow KVf_\Omega(pk_\Omega, sk_\Omega).$$

Definition 2. (Correctness) A public key encryption scheme Ω is called correct, if for all security parameters $\lambda \in \mathbb{N}$, for all $pp_\Omega \leftarrow_{\$} PG_\Omega(1^\lambda)$, for all $(sk_\Omega, pk_\Omega) \leftarrow_{\$} KG_\Omega(pp_\Omega)$, for all $m \in \mathcal{M}$, for all $c \leftarrow_{\$} Enc(pk_\Omega, m)$, we have that $m = Dec(sk_\Omega, c)$ and that for all sk'_Ω we have that $KVf_\Omega(pk_\Omega, sk'_\Omega) = 1 \implies m = Dec(sk'_\Omega, c)$.

Definition 3. (Multi-Challenge IND-CPA Security) A public-key encryption scheme Ω is multi-challenge IND-CPA secure (mcIND-CPA), if for any PPT adversary \mathcal{A} there exists a negligible function ν such that:

$$\left| \Pr \left[\mathbf{Exp}_{\mathcal{A}, \Omega}^{mcIND-CPA}(\lambda) = 1 \right] - 1/2 \right| \leq \nu(\lambda).$$

The corresponding experiment is depicted in Fig. 2.

Bellare et al. have shown, via a hybrid argument, that mcIND-CPA is equivalent to standard, i.e., “single-message”, IND-CPA [9]. We opted for using mcIND-CPA, because it allows writing our proofs down more compactly, improving readability.

2.2.2. Digital Signature Schemes

Subsequently, we define signature schemes.

Definition 4. (Digital Signatures) A digital signature scheme Σ consists of four algorithms $\{PG_\Sigma, KG_\Sigma, Sgn_\Sigma, Vrf_\Sigma\}$ such that:

```

Exp $\mathcal{A}, \Sigma$ eUNF-CMA( $\lambda$ )
   $\text{pp}_\Sigma \leftarrow_{\$} \text{PG}_\Sigma(1^\lambda)$ 
   $(\text{sk}_\Sigma, \text{pk}_\Sigma) \leftarrow_{\$} \text{KG}_\Sigma(\text{pp}_\Sigma)$ 
   $\mathcal{Q} \leftarrow \emptyset$ 
   $(m^*, \sigma^*) \leftarrow_{\$} \mathcal{A}^{\text{Sgn}'_\Sigma(\text{sk}_\Sigma, \cdot)}(\text{pk}_\Sigma)$ 
  where  $\text{Sgn}'_\Sigma$  on input  $\text{sk}_\Sigma$  and  $m$ :
     $\sigma \leftarrow_{\$} \text{Sgn}_\Sigma(\text{sk}_\Sigma, m)$ 
    set  $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{m\}$ 
    return  $\sigma$ 
  return 1, if  $\text{Vrf}_\Sigma(\text{pk}_\Sigma, m^*, \sigma^*) = 1 \wedge m^* \notin \mathcal{Q}$ 
  return 0

```

Fig. 3. Unforgeability .

PG_Σ . The algorithm PG_Σ outputs the public parameters

$$\text{pp}_\Sigma \leftarrow_{\$} \text{PG}_\Sigma(1^\lambda).$$

We assume that pp_Σ is implicit input to all other algorithms.

KG_Σ . The algorithm KG_Σ outputs the public and private key of the signer, where λ is the security parameter:

$$(\text{sk}_\Sigma, \text{pk}_\Sigma) \leftarrow_{\$} \text{KG}_\Sigma(\text{pp}_\Sigma).$$

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

$$\sigma \leftarrow_{\$} \text{Sgn}_\Sigma(\text{sk}_\Sigma, m).$$

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

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

Definition 5. (Correctness) A digital signature scheme Σ is called correct, if for all security parameters $\lambda \in \mathbb{N}$, for all $\text{pp}_\Sigma \leftarrow_{\$} \text{PG}_\Sigma(1^\lambda)$, for all $(\text{sk}_\Sigma, \text{pk}_\Sigma) \leftarrow_{\$} \text{KG}_\Sigma(\text{pp}_\Sigma)$, for all $m \in \mathcal{M}$, $\text{Vrf}_\Sigma(\text{pk}_\Sigma, m, \text{Sgn}_\Sigma(\text{sk}_\Sigma, m)) = 1$ is true.

We require existential unforgeability under adaptively chosen message attacks (eUNF-CMA security). In a nutshell, unforgeability requires that an adversary \mathcal{A} cannot (except with negligible probability) come up with a signature for a message m^* for which the adversary did not see any signature before, even if the adversary \mathcal{A} is allowed to adaptively query for signatures on messages of its own choice.

Definition 6. (Unforgeability) We say a digital signature scheme Σ scheme is unforgeable, if for every PPT adversary \mathcal{A} , there exists a negligible function ν such that:

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

The corresponding experiment is depicted in Fig. 3.

For Construction 3, we require that the size of signatures is independent of the size of the signed messages.

2.2.3. Non-Interactive Proof Systems.

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

Definition 7. (Non-Interactive Proof System) A non-interactive proof system Π for language L consists of three algorithms $\{\text{PG}_{\Pi}, \text{Prf}_{\Pi}, \text{Vfy}_{\Pi}\}$, such that:

PG_{Π} . The algorithm PG_{Π} outputs public parameters of the scheme, where λ is the security parameter:

$$\text{crs}_{\Pi} \leftarrow_{\S} \text{PG}_{\Pi}(1^{\lambda}).$$

Prf_{Π} . The algorithm Prf_{Π} outputs the proof π , on input of the CRS crs_{Π} , statement x to be proven, and the corresponding witness w :

$$\pi \leftarrow_{\S} \text{Prf}_{\Pi}(\text{crs}_{\Pi}, x, w).$$

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

$$d \leftarrow \text{Vfy}_{\Pi}(\text{crs}_{\Pi}, x, \pi).$$

Definition 8. (Correctness) A non-interactive proof system is called correct, if for all $\lambda \in \mathbb{N}$, for all $\text{crs}_{\Pi} \leftarrow_{\S} \text{PG}_{\Pi}(1^{\lambda})$, for all $x \in L$, for all w such that $R(x, w) = 1$, for all $\pi \leftarrow_{\S} \text{Prf}_{\Pi}(\text{crs}_{\Pi}, x, w)$, it holds that $\text{Vfy}_{\Pi}(\text{crs}_{\Pi}, x, \pi) = 1$.

In the context of (zero-knowledge) proof-systems, correctness is sometimes also referred to as completeness. In addition, we require two standard security notions for zero-knowledge proofs of knowledge: zero-knowledge and simulation-sound extractability. We define them analogously to the definitions given in [27].

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

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

Fig. 4. Zero-Knowledge .

$\mathbf{Exp}_{\mathcal{A}, \Pi, \mathcal{E}}^{\text{SimSoundExt}}(\lambda)$
 $(\text{crs}_{\Pi}, \tau, \zeta) \leftarrow_{\S} \mathcal{E}_1(1^\lambda)$
 $\mathcal{Q} \leftarrow \emptyset$
 $(x^*, \pi^*) \leftarrow_{\S} \mathcal{A}^{\text{SIM}(\cdot)}(\text{crs}_{\Pi})$
 where SIM on input x :
 obtain $\pi \leftarrow_{\S} \text{SIM}_2(\text{crs}_{\Pi}, \tau, x)$
 $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{x, \pi\}$
 return π
 $w^* \leftarrow_{\S} \mathcal{E}_2(\text{crs}_{\Pi}, \zeta, x^*, \pi^*)$
 return 1, if $\forall \text{fy}_{\Pi}(x^*, \pi^*) = 1 \wedge R(x^*, w^*) = 0 \wedge (x^*, \pi^*) \notin \mathcal{Q}$
 return 0

Fig. 5. Simulation Sound Extractability .

Definition 9. (Zero-Knowledge) A non-interactive proof system Π for language L is zero-knowledge, if for any PPT adversary \mathcal{A} , there exists an PPT simulator $\text{SIM} = (\text{SIM}_1, \text{SIM}_2)$ such that there exist negligible functions ν_1 and ν_2 such that

$$\left| \Pr [\text{crs}_{\Pi} \leftarrow_{\S} \text{PG}_{\Pi}(1^\lambda) : \mathcal{A}(\text{crs}_{\Pi}) = 1] - \Pr [(\text{crs}_{\Pi}, \tau) \leftarrow_{\S} \text{SIM}_1(1^\lambda) : \mathcal{A}(\text{crs}_{\Pi}) = 1] \right| \leq \nu_1(\lambda),$$

and that

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

where the corresponding experiment is depicted in Fig. 4.

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

Note that the definition of simulation-sound extractability of [38] is stronger than ours in the sense that the adversary also gets the trapdoor ζ as input. However, in our context this weaker notion (previously also used, e.g., in [1, 32]) suffices.

Definition 10. (Simulation-Sound Extractability) A zero-knowledge non-interactive proof system Π for language L is said to be simulation-sound extractable, if for any PPT adversary \mathcal{A} , there exists a PPT extractor $\mathcal{E} = (\mathcal{E}_1, \mathcal{E}_2)$, such that

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

and that there exist a negligible function ν so that

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

where $\text{SIM} = (\text{SIM}_1, \text{SIM}_2)$ is as in Definition 9 and the corresponding experiment is depicted in Fig. 5.

3. Chameleon-Hashes, Revisited

In this section, we present the formal framework for chameleon-hashes, their security properties with a special focus on the collision-resistance notion, and then show relations and separations between the security properties.

3.1. Framework

We now present the framework for chameleon-hashes. We rely on the most recent comprehensive framework by Camenisch et al. [18], which is, in turn, based upon work done by Ateniese et al. and Brzuska et al. [5, 17].

Definition 11. A chameleon-hash CH is a tuple of five PPT algorithms (CHPG, CHKG, CHash, CHCheck, CHAdapt), such that:

CHPG. The algorithm CHPG, on input a security parameter λ outputs public parameters of the scheme:

$$\text{pp}_{\text{ch}} \leftarrow_{\$} \text{CHPG}(1^\lambda).$$

We assume that pp_{ch} is implicit input to all other algorithms.

CHKG. The algorithm CHKG, on input the public parameters pp_{ch} outputs the private and public keys of the scheme:

$$(\text{sk}_{\text{ch}}, \text{pk}_{\text{ch}}) \leftarrow_{\$} \text{CHKG}(\text{pp}_{\text{ch}}).$$

CHash. The algorithm **CHash** gets as input the public key \mathbf{pk}_{ch} , and a message m to hash. It outputs a hash h , and some randomness r ³:

$$(h, r) \leftarrow_{\S} \text{CHash}(\mathbf{pk}_{\text{ch}}, m).$$

CHCheck. The deterministic algorithm **CHCheck** gets as input the public key \mathbf{pk}_{ch} , a message m , randomness r , and a hash h . It outputs a bit $d \in \{0, 1\}$, indicating whether the hash h is valid:



$$d \leftarrow \text{CHCheck}(\mathbf{pk}_{\text{ch}}, m, r, h).$$


CHAdapt. The algorithm **CHAdapt** on input of a secret key \mathbf{sk}_{ch} , the message m , new message m' , randomness r , and hash h outputs new randomness r' :

$$r' \leftarrow_{\S} \text{CHAdapt}(\mathbf{sk}_{\text{ch}}, m, m', r, h).$$

Definition 12. (Correctness) A chameleon-hash is called correct, if for all security parameters $\lambda \in \mathbb{N}$, for all $\mathbf{pp}_{\text{ch}} \leftarrow_{\S} \text{CHPG}(1^\lambda)$, for all $(\mathbf{sk}_{\text{ch}}, \mathbf{pk}_{\text{ch}}) \leftarrow_{\S} \text{CHKG}(\mathbf{pp}_{\text{ch}})$, for all $m \in \mathcal{M}$, for all $(h, r) \leftarrow_{\S} \text{CHash}(\mathbf{pk}_{\text{ch}}, m)$, for all $m' \in \mathcal{M}$, we have for all $r' \leftarrow_{\S} \text{CHAdapt}(\mathbf{sk}_{\text{ch}}, m, m', r, h)$, that $1 = \text{CHCheck}(\mathbf{pk}_{\text{ch}}, m, r, h) = \text{CHCheck}(\mathbf{pk}_{\text{ch}}, m', r', h)$.

3.2. Collision-Resistance, Revisited

In this section we revisit existing collision-resistance notions, introduce a stronger and more desirable notion of collision-resistance dubbed *full collision-resistance* (or **F-CollRes** for short) and discuss how these notions differ. The main idea behind collision-resistance in general is to argue that an adversary that has no access to the secret key \mathbf{sk}_{ch} cannot find any collisions, i.e., pairs (m, r) and (m', r') and hash value h s.t. $\text{CHCheck}(\mathbf{pk}_{\text{ch}}, m, r, h) = \text{CHCheck}(\mathbf{pk}_{\text{ch}}, m', r', h) = 1$. In the weakest case, the adversary has no access to any other collisions, whereas in stronger notions the adversary is explicitly allowed to obtain collisions for arbitrary hashes via a **CHAdapt'** oracle (we indicate these by using  boxes). We present all the different notions in Fig. 6, where we indicate the differences in the winning conditions by using  boxes.

In all the experiments the challenger generates a key pair $(\mathbf{sk}_{\text{ch}}, \mathbf{pk}_{\text{ch}})$ honestly (along with some public parameters) and the adversary is then initialized with \mathbf{pk}_{ch} . We now discuss the differences of the single collision-resistance notions, where in the weakest case the adversary has no access to an **CHAdapt'** oracle (which allows the adversary to adaptively ask for collisions with messages and hashes of its own choice), but in all other cases the adversary does. To vertically align the experiments, we insert  boxes for lines which are missing in one experiment but are present in the other.

Weak Collision-Resistance (W-CollRes) [46] The adversary \mathcal{A} wins, if it can come up with a collision for the given public key.

³We note that the randomness r is also sometimes called “check value” [5].



Fig. 6. The $\text{Exp}_{A,CH}^X\text{-CollRes}$ experiment with $X \in \{W, E, S, F\}$.

Enhanced Collision-Resistance (E-CollRes) [5] The adversary gets access to a collision-finding oracle $\text{CHAdapt}'$, which outputs a collision for adversarially chosen hashes, but also keeps track of each queried *hash* h using the list \mathcal{Q} . The adversary wins, if it comes up with a collision for the given public key for an adversarially chosen hash h^* never input to $\text{CHAdapt}'$.

Standard Collision-Resistance (S-CollRes) [18] The adversary gets access to a collision-finding oracle $\text{CHAdapt}'$, which outputs a collision for the adversarially chosen hash, but also keeps track of each of the queried *messages* m and m' , using the list \mathcal{Q} . The adversary wins, if it comes up with a collision for the given public key for an adversarially chosen h^* for which the message m^* output by the adversary was never queried to the collision-finding oracle.

Full Collision-Resistance (F-CollRes). The adversary gets access to a collision-finding oracle $\text{CHAdapt}'$, which outputs a collision for the adversarially chosen hash, but also keeps track of each of the queried *hash/message pair* (h, m) and (h, m') , using the list \mathcal{Q} . The adversary wins, if it comes up with a hash/message

pair (h^*, m^*) , for the given public key, never queried to or output from the collision-finding oracle.⁴

Now, we formally define security with respect to all the collision-resistance notions.

Definition 13. (*X* Collision-Resistance) A chameleon-hash CH offers *X* collision-resistance with $X \in \{W, E, S, F\}$, if for any PPT adversary \mathcal{A} there exists a negligible function ν such that

$$\Pr[\mathbf{Exp}_{\mathcal{A}, \text{CH}}^{X\text{-CollRes}}(\lambda) = 1] \leq \nu(\lambda),$$



where the corresponding experiment is depicted in Fig. 6.

3.2.1. Discussion of the Notions

W-CollRes is the notion introduced in the first work on chameleon-hashes by Krawczyk and Rabin [46] and essentially represents the binding notion of a trapdoor-commitment scheme. Note that due to not giving access to a collision-finding oracle it gives no guarantees whatsoever if the adversary sees a single collision for any hash computed for the given public key.⁵ The **E-CollRes** notion has been introduced by Ateniese et al. [5] and we note that there exists a definition in the setting of public-coin chameleon-hashes, i.e., where the **CHCheck** algorithm simply re-runs the **CHash**, which is called key-exposure freeness [7, 20]. It captures requirements similar to the ones captured by **E – CollRes**, but it is not directly comparable as we are considering the more general secret-coin setting. We note that the **E-CollRes** notion allows the adversary to come up with arbitrary collisions for hashes it has seen a collision for. The **S – CollRes** notion has been introduced by Camenisch et al. [18], and it captures all of the intuitive requirements of real-world applications of chameleon-hashes. Yet, it still allows the hash itself to be malleable which might still be problematic in certain applications. Finally, our new **F-CollRes** notion enforces that the adversary cannot (except with negligible probability) output any new collisions and seems to be the most desirable notion for collision-resistance.

3.3. Indistinguishability, Revisited

In a nutshell, indistinguishability requires that an adversary cannot decide whether randomness was obtained through **CHash** or **CHAdapt**.

We present the respective formal security games in Fig. 7. We highlight differences by using  boxes, and missing parts using  boxes.

⁴In the case (h^*, m^*) is the new hash/message pair, simply switch names.

⁵A slightly stronger notion has been proposed by Zhang in [57] where the adversary sees a hash on a random message and is then given a single collision on a message of its choice. We do not cover this notion here as it seems to be tailored to the specific applications in [57] and all notions stronger than **W-CollRes** considered here cover more general cases.

3.3.1. (Normal) Indistinguishability (N-Ind)

Normal Indistinguishability (we sometimes refer to this notion simply as “Indistinguishability”, as this is the standard name in the literature) requires that the randomness r does not reveal if it was obtained through **CHash** or **CHAdapt**.

Upon setup, the challenger generates a key pair $(\text{sk}_{\text{ch}}, \text{pk}_{\text{ch}})$ for **CH** (along with some public parameters pp_{ch}), and draws a bit $b \leftarrow_{\$} \{0, 1\}$. The challenger initializes the adversary with the pk_{ch} and gives the adversary access to a **HashOrAdapt** oracle, which allows the adversary to submit two messages m, m' . Depending on the bit b , the challenger then either hashes m' directly ($b = 0$), or first hashes m , and then adapts m to m' ($b = 1$). The resulting hash/randomness pair (h, r) (or (h', r') resp.) is the oracle’s output to the adversary. The adversary’s objective is to guess the bit b . Note that all keys are generated honestly. The adversary gets access to a collision-finding oracle **CHAdapt** for arbitrary hashes, meaning that the adversary may also input hashes generated by the **HashOrAdapt**-oracle.

Samelin and Slamanig recently introduced *full* indistinguishability [51], which, in turn, generalizes the notion of *strong* indistinguishability by Derler et al. [26]. In their notion, the adversary is even allowed to generate the keys which are used for hashing and adapting (in the strong version, the adversary only knows all keys, but cannot generate them). See below for more information. Finally, we introduce an additional notion, dubbed enhanced indistinguishability, where the adversary not only receives the secret key generated, but the randomness r used for generation. This notion may be useful in context where randomness leaks to the adversary.

3.3.2. Strong Indistinguishability (S-Ind)

Strong indistinguishability requires that a randomness r does not reveal whether it was generated using **CHash** or **CHAdapt**, even if the adversary \mathcal{A} additionally receives the generated secret key. This also means that the collision-finding oracle can be dropped, as the adversary can find collisions on its own.

3.3.3. Enhanced Indistinguishability (E-Ind)

Enhanced indistinguishability requires that a randomness r does not reveal whether it was generated using **CHash** or **CHAdapt**, even if the adversary \mathcal{A} knows the randomness ξ used to generate the secret key. Again, this also means that the collision-finding oracle can be dropped, as the adversary can find collisions on its own.

3.3.4. Full Indistinguishability (F-Ind)

Full indistinguishability requires that a randomness r does not reveal whether it was generated using **CHash** or **CHAdapt**, even if the adversary \mathcal{A} controls all values, but the public parameters.⁶ Once more, this also means that the collision-finding oracle can be dropped, as the adversary can find collisions on its own.

⁶Lifting this definition to also cover those parameters is straightforward.

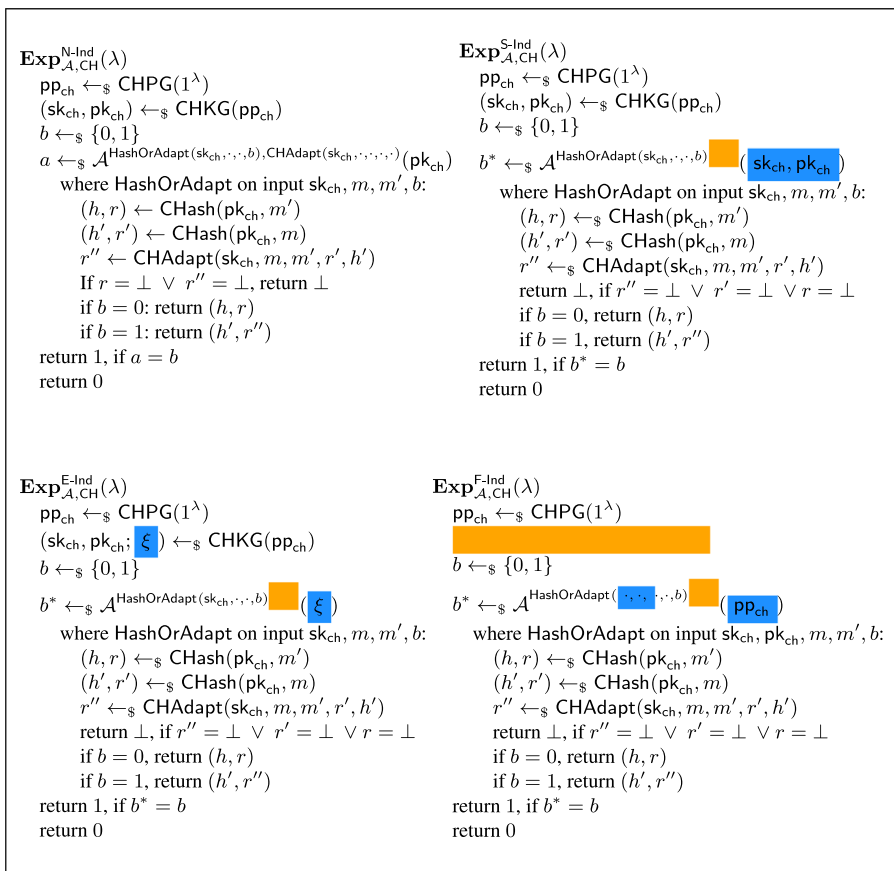


Fig. 7. The $\text{Exp}_{A,CH}^{X\text{-Ind}}$ experiment with $X \in \{N, S, E, F\}$ (Color Figure online).

Definition 14. (X Indistinguishability) A chameleon-hash CH offers X indistinguishability with $X \in \{N, S, E, F\}$, if for any PPT adversary \mathcal{A} there exists a negligible function ν such that

$$\left| \Pr[\text{Exp}_{A,CH}^{X\text{-Ind}}(\lambda) = 1] - 1/2 \right| \leq \nu(\lambda).$$

The corresponding experiments are depicted in Fig. 7.

We only consider normal indistinguishability as fundamental for chameleon-hashes, but examine stronger notions to achieve a more complete picture of the relations. We also stress that there may be scenarios where some sort of indistinguishability is not required or even hindering.


```

Exp $\mathcal{A}, \text{CH}$ Uniqueness( $\lambda$ )
   $\text{pp}_{\text{ch}} \leftarrow_{\$} \text{CHPG}(1^\lambda)$ 
   $(\text{pk}^*, m^*, r^*, r'^*, h^*) \leftarrow_{\$} \mathcal{A}(\text{pp}_{\text{ch}})$ 
  return 1, if  $\text{CHCheck}(\text{pk}^*, m^*, r^*, h^*) = \text{CHCheck}(\text{pk}^*, m^*, r'^*, h^*) = 1 \wedge r^* \neq r'^*$ 
  return 0

```

Fig. 8. Uniqueness .

3.4. Uniqueness

Camenisch et al. [18] defined a property called uniqueness. Uniqueness requires that for each hash/message pair, exactly one randomness can be found, even if the adversary \mathcal{A} controls all values, but the public parameters.⁷

Definition 15. (Uniqueness) A chameleon-hash CH is unique, if for any PPT adversary \mathcal{A} there exists a negligible function ν such that

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

The corresponding experiment is depicted in Fig. 8.

We do not consider uniqueness as a fundamental property, as there are only very few applications requiring this notion [18,51]. However, to obtain a more complete picture with respect to the relations of the security properties, we also investigate uniqueness.

4. Relationships between Properties of Chameleon-Hashes

Below we show relations and separations between the security properties of chameleon-hashes. Before doing so, we recall in Sect. 4.1 examples of chameleon-hashes providing the W -CollRes and S -CollRes notions, respectively.

4.1. Existing Constructions of Chameleon-Hashes

4.1.1. Instantiation of a Weakly Collision-Resistant CH

We recall the initial CH construction by Krawczyk and Rabin [46] in Construction 1.

Note that a collision-resistant hash function is applied to the message prior to chameleon-hashing to extend the domain, which is a standard technique. Seeing a collision (if not resulting from the collision-resistant hash function) allows to extract the sk_{ch} by computing $x \leftarrow (H(m) - H(m')) / (r' - r) \bmod q$.

4.1.2. Instantiation of a Standard Collision-Resistant CH

We recall a construction by Camenisch et al. from [18] in Construction 2. Before we do so, we recall some background on the setup the scheme requires: Let $(N, p, q, e, d) \leftarrow_{\$}$

⁷Lifting this definition to also cover those parameters is straightforward.

CHPG(1^λ) : Outputs the public parameters (\mathbb{G}, g, q, H) , where $(\mathbb{G}, g, q) \leftarrow \mathbf{GGen}(1^\lambda)$ is a group \mathbb{G} of prime order q generated by g and $H : \{0, 1\}^* \rightarrow \mathbb{Z}_q^*$ is a hash function chosen uniformly at random from a family of collision-resistant hash functions.

CHKG(pp_{ch}) : Parse pp_{ch} as (\mathbb{G}, g, q, H) and return $(\text{sk}_{\text{ch}}, \text{pk}_{\text{ch}}) \leftarrow (x, g^x)$, where

$$x \leftarrow_{\S} \mathbb{Z}_q^*.$$

CHash(pk_{ch}, m) : Return (h, r) , where

$$r \leftarrow_{\S} \mathbb{Z}_q^*, \text{ and } h \leftarrow g^{H(m)} \text{pk}_{\text{ch}}^r.$$

CHCheck($\text{pk}_{\text{ch}}, m, h, r$) : Return 1 if the following holds, and 0 otherwise:

$$h = g^{H(m)} \text{pk}_{\text{ch}}^r.$$

CHAdapt($\text{sk}_{\text{ch}}, m, m', h, r$) : Output \perp , if **CHCheck**($\text{pk}_{\text{ch}}, m, h, r$) $\neq 1$. Otherwise return r' , where

$$r' \leftarrow \frac{H(m) + xr - H(m')}{x}.$$

Construction 1: DL-based chameleon-hash

RSAG(1^λ) be an instance generator which returns an RSA modulus $N = pq$, where p and q are distinct primes, $e > 1$ is an integer co-prime to $\varphi(n)$, and $de \equiv 1 \pmod{\varphi(n)}$. The scheme requires that **RSAG** always outputs moduli of the same bit-length, based on λ , and that the one-more RSA assumption holds [10].

CHPG(1^λ) : Output the public parameters $\text{pp}_{\text{ch}} \leftarrow (1^\lambda, e)$, where e is prime and $e > N'$ with $N' = \max_{\xi} \{N \in \mathbb{N} : (N, \cdot, \cdot, \cdot, \cdot) \leftarrow_{\S} \mathbf{RSAG}(1^\lambda; \xi)\}$.

CHKG(pp_{ch}) : Run $(N, p, q, \cdot, \cdot) \leftarrow_{\S} \mathbf{RSAG}(1^\lambda)$, choose a hash function $H : \{0, 1\}^* \rightarrow \mathbb{Z}_N^*$ (modeled as a random-oracle), compute d s.t. $ed \equiv 1 \pmod{\varphi(N)}$, set $\text{sk}_{\text{ch}} \leftarrow d$, $\text{pk}_{\text{ch}} \leftarrow_{\S} (N, H)$, and return $(\text{sk}_{\text{ch}}, \text{pk}_{\text{ch}})$.

CHash(pk_{ch}, m) : Parse $\text{pk}_{\text{ch}} = (N, H)$ and a message m , choose $r \leftarrow_{\S} \mathbb{Z}_N^*$, compute $h \leftarrow H(m)r^e \pmod{N}$, and output (h, r) .

CHCheck($\text{pk}_{\text{ch}}, m, h, r$) : Parse $\text{pk}_{\text{ch}} = (N, H)$, compute $h' \leftarrow H(m)r^e \pmod{N}$, and output 1 if $h' = h$ and 0 otherwise.

CHAdapt($\text{sk}_{\text{ch}}, m, m', h, r$) : Output \perp , if **CHCheck**($\text{pk}_{\text{ch}}, m, h, r$) $\neq 1$. Otherwise, let $x \leftarrow H(m)$, $x' \leftarrow H(m')$, $y \leftarrow xr^e \pmod{N}$ and return $r' \leftarrow (y(x'^{-1}))^d \pmod{N}$.

Construction 2: RSA-based Chameleon-Hash

4.2. Collision-Resistance Properties

We start by analyzing how the various collision-resistance notions are related.

Theorem 1. *Standard collision-resistance is strictly stronger than weak collision-resistance.*

Proof. We first prove that standard collision-resistance implies weak collision-resistance. Then we give a counterexample showing that the other direction of the implication does not hold.

S – CollRes \implies **W – CollRes**: Assume \mathcal{A} to be an adversary who breaks weak collision-resistance. We now construct an adversary \mathcal{B} which breaks standard collision-resistance. In particular, \mathcal{B} proceeds as follows. It receives pp_{ch} and pk_{ch} from its own challenger, and uses both to initialize \mathcal{A} . Whenever \mathcal{A} outputs a winning tuple $(m^*, r^*, m'^*, r'^*, h^*)$, \mathcal{B} returns that tuple to its own challenger. As the collision-finding oracle was never queried, that tuple also makes \mathcal{B} win the standard collision-resistance game with the same probability \mathcal{A} wins the weak collision-resistance game.

W – CollRes $\not\Rightarrow$ **S – CollRes**: The CH by Krawczyk and Rabin [46] provides a counterexample: it is weakly collision-resistant, but does not offer standard collision-resistance. Observe that it is possible to trivially extract the secret key from a collision. That collision is obtained from the collision-finding oracle in the standard collision-resistance game (cf. Section 4.1 for more details). \square

Theorem 2. *Enhanced collision-resistance is strictly stronger than weak collision-resistance.*

Proof. The proof is identical to the one of Theorem 1. \square

Theorem 3. *Full collision-resistance is strictly stronger than standard collision-resistance.*

Proof. We first prove that full collision-resistance implies standard collision-resistance and then give a counterexample showing that the other direction of the implication does not hold.

F – CollRes \implies **S – CollRes**: Assume \mathcal{A} to be an adversary who breaks standard collision-resistance. Now we construct an adversary \mathcal{B} which breaks full collision-resistance. In particular, \mathcal{B} proceeds as follows. It receives pp_{ch} and pk_{ch} from its own challenger, and uses both to initialize \mathcal{A} . All queries to the collision-finding oracle are relayed to \mathcal{B} 's own oracle. Whenever \mathcal{A} outputs a winning tuple $(m^*, r^*, m'^*, r'^*, h^*)$, \mathcal{B} returns that tuple to its own challenger. As $m^* \neq m'^*$ must be true, and m^* was never queried to \mathcal{A} 's collision-finding oracle, this also means that (h^*, m^*) was never queried to \mathcal{B} 's oracle, thus meeting the winning condition.

S – CollRes $\not\Rightarrow$ **F – CollRes** : The scheme by Camenisch et al. [18] (See Construction 2) provides a counterexample: it offers standard collision-resistance, but does not offer full collision-resistance. In particular, their construction is re-randomizable (cf. Section 4.1 for more details).

In more detail, to show that this construction is not fully collision-resistant, consider the following strategy: Receive $\text{pk}_{\text{ch}} = (N, H)$ and $\text{pp}_{\text{ch}} = e$. Compute $(h, r) \leftarrow_{\S} \text{CHash}(\text{pk}_{\text{ch}}, m)$, with m random. Then, ask for an adaption (h, r, m) to (h, r', m') , for

some random $m' \neq m$. Then, compute $h^* \leftarrow h2^e \pmod N$, $r_1^* \leftarrow 2r \pmod N$, and $r_2^* \leftarrow 2r' \pmod N$. Because no collision for h^* was computed, this construction cannot be fully collision-resistant. Note, this works, as $H(m)(2r)^e \equiv h2^e \pmod N$ for any input. Also note that the attack above also breaks enhanced collision-resistance (we will later use this to derive a corollary). \square

Theorem 4. *Full collision-resistance is strictly stronger than enhanced collision-resistance.*

Before we provide the proof of Theorem 4, we provide a novel construction of a chameleon-hash satisfying the E-CollRes notion that is used to separate the notions F-CollRes and E-CollRes.

4.2.1. Construction

Our CH presented below provides E-CollRes, but allows to efficiently find arbitrary collisions for a given hash, once a single collision was seen. However, it is not possible to find collisions for any other hash. The main idea is to encrypt a message m using a mcIND-CPA secure encryption scheme Ω and use the ciphertext as the hash. The randomness r of the chameleon-hash is the public key $\text{pk}_{\Omega'}$ of a freshly sampled key-pair $(\text{sk}_{\Omega'}, \text{pk}_{\Omega'})$ of Ω , the encryption c' of a signature σ under $\text{pk}_{\Omega'}$ and a SSE NIZK π for the following language:

$$L := \{(\text{pk}_{\Omega}, \text{pk}_{\Sigma}, h, m) \mid \exists (\sigma, \xi) : \\ h = \text{Enc}(\text{pk}_{\Omega}, m; \xi) \vee \text{Vrf}_{\Sigma}(\text{pk}_{\Sigma}, h, \sigma) = 1\}. \quad (1)$$

Informally, this language requires the prover to show that it either knows the randomness ξ attesting that h is a well-formed encryption of m , or a valid signature σ for h . The basic idea of the construction is that when computing a hash, the witness ξ is used. The randomness includes an encryption of the signature (initially one on 0) under the public key $\text{pk}_{\Omega'}$. Note that the trick is that for adaption one computes a signature σ for h , uses σ as a witness, and includes an encryption of σ under $\text{pk}_{\Omega'}$ in the randomness. Clearly, now seeing a single collision allows to compute arbitrary collisions for the hash h .

This CH can be instantiated by instantiating Σ as structure-preserving signatures (SPS) in type-III bilinear groups (assuming SXDH), e.g., Groth's SPS [39]. Thus, Ω can be ElGamal [37] in one of the base-groups. The algorithm KVf_{Ω} is simply checking whether $g^{\text{sk}_{\Omega}} = g^x = \text{pk}_{\Omega}$, while for Π , a suitable instantiation is a Fiat-Shamir transformed Σ -protocol in the random-oracle model [36], which also works very well with ElGamal encryption and Groth's signature scheme.

Subsequently, we use $\boxed{\text{frameboxes}}$ and \rightsquigarrow to highlight the changes we make in the algorithms throughout a sequence of games (and we only show the changes).

CHPG(1^λ) : Fix a public-key encryption scheme Ω , a signature scheme Σ , and a compatible NIZK proof system for language L in (1). Return $\text{pp}_{\text{ch}} = (\text{pp}_\Omega, \text{pp}_\Sigma, \text{crs}_\Pi)$, where

$$\text{pp}_\Omega \leftarrow_{\S} \text{PG}_\Omega(1^\lambda), \text{pp}_\Sigma \leftarrow_{\S} \text{PG}_\Sigma(1^\lambda), \text{ and } \text{crs}_\Pi \leftarrow_{\S} \text{PG}_\Pi(1^\lambda).$$

CHKG(pp_{ch}) : Return $(\text{sk}_{\text{ch}}, \text{pk}_{\text{ch}}) = ((\text{sk}_\Omega, \text{sk}_\Sigma), (\text{pp}_{\text{ch}}, \text{pk}_\Omega, \text{pk}_\Sigma, \sigma_0))$, where

$$(\text{sk}_\Omega, \text{pk}_\Omega) \leftarrow_{\S} \text{KG}_\Omega(\text{pp}_\Omega), (\text{sk}_\Sigma, \text{pk}_\Sigma) \leftarrow_{\S} \text{KG}_\Sigma(\text{pp}_\Sigma), \text{ and } \sigma_0 \leftarrow_{\S} \text{Sgn}_\Sigma(\text{sk}_\Sigma, 0).$$

0 is considered some special invalid hash value for CH.

CHash(pk_{ch}, m) : Parse pk_{ch} as $((\text{pp}_\Omega, \text{crs}_\Pi), \text{pk}_\Omega)$, and return $(h, r) = (c, (\pi, c', \text{pk}_\Omega'))$, where

$$(c; \xi) \leftarrow_{\S} \text{Enc}(\text{pk}_\Omega, m), (\text{sk}_\Omega', \text{pk}_\Omega') \leftarrow_{\S} \text{KG}_\Omega(\text{pp}_\Omega), c' \leftarrow_{\S} \text{Enc}(\text{pk}_\Omega', \sigma_0), \text{ and } \\ \pi \leftarrow_{\S} \text{Prf}_\Pi(\text{crs}_\Pi, (\text{pk}_\Omega, \text{pk}_\Sigma, c, m), (\perp, \xi))$$

CHCheck($\text{pk}_{\text{ch}}, m, r, h$) : Parse pk_{ch} as $((\text{pp}_\Omega, \text{crs}_\Pi), \text{pk}_\Omega)$ and r as $(\pi, c', \text{pk}_\Omega')$, and return 1 if the following holds, and 0 otherwise:

$$m \in \mathcal{M} \wedge \text{Vfy}_\Pi(\text{crs}_\Pi, (\text{pk}_\Omega, \text{pk}_\Sigma, h, m), \pi) = 1.$$

CHAdapt($\text{sk}_{\text{ch}}, m, m', r, h$) : Parse sk_{ch} as sk_Ω . Verify that $m' \in \mathcal{M}$, **CHCheck**($\text{pk}_{\text{ch}}, m, r, h$) = 1, and return \perp if not. Otherwise, return $r' = (\pi', c'', \text{pk}_\Omega')$, where

$$\sigma \leftarrow_{\S} \text{Sgn}_\Sigma(\text{sk}_\Sigma, h), c'' \leftarrow_{\S} \text{Enc}(\text{pk}_\Omega', \sigma), \text{ and } \\ \pi' \leftarrow_{\S} \text{Prf}_\Pi(\text{crs}_\Pi, (\text{pk}_\Omega, \text{pk}_\Sigma, h, m'), (\sigma, \perp)).$$

Construction 3: Enhanced Collision-Resistant Chameleon-Hash

Theorem 5. *If Ω , Σ , and Π are correct, then Construction 3 is correct.*

Correctness follows from inspection and the (perfect) correctness of the used primitives.

While indistinguishability is technically not needed for proving the separation we are after in this section, we nevertheless prove it here for completeness.

Theorem 6. *If Ω is mcIND-CPA secure and Π is zero-knowledge, then Construction 3 is indistinguishable (N-Ind).*

Proof. To prove indistinguishability, we use a sequence of games:

Game 0: The original indistinguishability game.

Game 1: As Game 0, but we modify the algorithms **CHPG**, **CHash**, and **CHAdapt** used within the game as follows:

CHPG'(1^λ) :

$$\text{crs}_\Pi \leftarrow_{\S} \text{PG}_\Pi(1^\lambda) \rightsquigarrow (\text{crs}_\Pi, \tau) \leftarrow_{\S} \text{SIM}_1(1^\lambda).$$

CHash'(pk_{ch}, m) :

$$\pi \leftarrow_{\S} \text{Prf}_\Pi(\text{crs}_\Pi, (\text{pk}_\Omega, \text{pk}_\Sigma, h, m), (\perp, \xi)) \rightsquigarrow \pi \leftarrow_{\S} \text{SIM}_2(\text{crs}_\Pi, \tau, (\text{pk}_\Omega, \text{pk}_\Sigma, h, m))$$

CHAdapt'(sk_{ch}, m, m', r, h) :

$$\pi' \leftarrow_{\S} \text{Prf}_\Pi(\text{crs}_\Pi, (\text{pk}_\Omega, \text{pk}_\Sigma, h, m'), (\sigma, \perp)) \rightsquigarrow \pi' \leftarrow_{\S} \text{SIM}_2(\text{crs}_\Pi, \tau, (\text{pk}_\Omega, \text{pk}_\Sigma, h, m'))$$

Transition – Game 0 → Game 1 : We bound the probability for an adversary to detect this game change by presenting a hybrid game, which, depending on a zero-knowledge challenger \mathcal{C}^{ZK} , either produces the distribution in Game 0 or Game 1, respectively. In particular, assume the following changes:

CHPG''(1^λ) :

$$(\text{crs}_\Pi, \tau) \leftarrow_{\S} \text{SIM}_1(1^\lambda) \rightsquigarrow \text{crs}_\Pi \leftarrow_{\S} \mathcal{C}^{\text{ZK}}.$$

CHash''(pk_{ch}, m) :

$$\pi \leftarrow_{\S} \text{SIM}_2(\text{crs}_\Pi, \tau, (\text{pk}_\Omega, \text{pk}_\Sigma, h, m)) \rightsquigarrow \pi \leftarrow_{\S} \mathcal{C}^{\text{ZK}}.P_b((\text{pk}_\Omega, \text{pk}_\Sigma, h, m), (\perp, \xi)).$$

CHAdapt''(sk_{ch}, m, m', r, h) :

$$\pi' \leftarrow_{\S} \text{SIM}_2(\text{crs}_\Pi, \tau, (\text{pk}_\Omega, \text{pk}_\Sigma, h, m')) \rightsquigarrow \pi' \leftarrow_{\S} \mathcal{C}^{\text{ZK}}.P_b((\text{pk}_\Omega, \text{pk}_\Sigma, h, m'), (\sigma, \perp)).$$

Clearly, if the challenger's internal bit is 0 we simulate the distribution in Game 0, whereas we simulate the distribution in Game 1 otherwise. We have that $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{zk}}(\lambda)$.

Game 2: As Game 1, but we further modify the CHash algorithm as follows:

CHash'''(pk_{ch}, m) :

$$(c; \xi) \leftarrow_{\S} \text{Enc}(\text{pk}_\Omega, m) \rightsquigarrow (c; \xi) \leftarrow_{\S} \text{Enc}(\text{pk}_\Omega, 0).$$

Transition – Game 1 → Game 2 : We bound the probability for an adversary to distinguish between two consecutive games by introducing a hybrid game which uses a mcIND-CPA challenger to interpolate between two consecutive games:

$\text{CHKG}(\text{pp}_{\text{ch}})'$: Return $(\text{sk}_{\text{ch}}, \text{pk}_{\text{ch}}) = ((\perp, \text{sk}_{\Sigma}), (\text{pp}_{\text{ch}}, \text{pk}_{\Omega}, \text{pk}_{\Sigma}, \sigma_0))$, where

$$\begin{aligned} (\text{sk}_{\Omega}, \text{pk}_{\Omega}) &\leftarrow_{\S} \text{KG}_{\Omega}(\text{pp}_{\Omega}) \rightsquigarrow \boxed{\text{pk}_{\Omega} \leftarrow_{\S} \mathcal{C}^{\text{mc-cpa}}}, \\ (\text{sk}_{\Sigma}, \text{pk}_{\Sigma}) &\leftarrow_{\S} \text{KG}_{\Sigma}(\text{pp}_{\Sigma}), \text{ and } \sigma_0 \leftarrow_{\S} \text{Sgn}_{\Sigma}(\text{sk}_{\Sigma}, 0). \end{aligned}$$

0 is considered some special invalid hash value for CH.

$\text{CHash}''''(\text{pk}_{\text{ch}}, m)$:

$$(c; \xi) \leftarrow_{\S} \text{Enc}(\text{pk}_{\Omega}, 0) \rightsquigarrow \boxed{(c; \perp) \leftarrow_{\S} \mathcal{C}^{\text{mc-cpa}}.\text{Enc}'(m, 0)}.$$

Now, depending on the challenger's bit, we either simulate Game 1 or Game 2. Thus we have that $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{mc-cpa}}(\lambda)$

$\text{Game3}_i (1 \leq i \leq q)$: As Game 3_{i-1} (resp. Game 2 if $i = 0$) but we modify the **HashOrAdapt** as follows. We let q be an upper bound on the queries to the **HashOrAdapt** oracle. Up to query number i , we do the following:

$\text{HashOrAdapt}''''(\text{sk}_{\text{ch}}, m, m', b)$: In CHash

$$c' \leftarrow_{\S} \text{Enc}(\text{pk}_{\Omega}', \sigma_0) \rightsquigarrow \boxed{c' \leftarrow_{\S} \text{Enc}(\text{pk}_{\Omega}', 0)}.$$

and in CHAdapt

$$c' \leftarrow_{\S} \text{Enc}(\text{pk}_{\Omega}', \sigma) \rightsquigarrow \boxed{c' \leftarrow_{\S} \text{Enc}(\text{pk}_{\Omega}', 0)}.$$

For every query after query i we simulate **HashOrAdapt** as in Game 2.

Transition - $\text{Game3}_i \rightarrow \text{Game3}_{i+1}$ (resp. $\text{Game 2} \rightarrow 3_1$) : We bound the probability for an adversary to distinguish between two consecutive games by introducing a hybrid game which interpolates between to subsequent games. Then, up to query number $i - 1$, we do the following:

$\text{HashOrAdapt}''''(\text{sk}_{\text{ch}}, m, m', b)$: In CHash

$$c' \leftarrow_{\S} \text{Enc}(\text{pk}_{\Omega}', \sigma_0) \rightsquigarrow \boxed{c' \leftarrow_{\S} \text{Enc}(\text{pk}_{\Omega}', 0)}.$$

and in CHAdapt

$$c' \leftarrow_{\S} \text{Enc}(\text{pk}_{\Omega}', \sigma) \rightsquigarrow \boxed{c' \leftarrow_{\S} \text{Enc}(\text{pk}_{\Omega}', 0)}.$$

In query number i we do the following:

HashOrAdapt^{''''}(sk_{ch}, m, m', b) :

$$(sk_{\Omega'}, pk_{\Omega'}) \leftarrow_{\S} KG_{\Omega}(pp_{\Omega}) \rightsquigarrow \boxed{(\perp, pk_{\Omega'}) \leftarrow_{\S} C^{mc-cpa}}$$

In CHash

$$c' \leftarrow_{\S} Enc(pk_{\Omega'}, 0) \rightsquigarrow \boxed{c' \leftarrow_{\S} C^{mc-cpa}.Enc'(\sigma_0, 0)}$$

and in CHAdapt

$$c' \leftarrow_{\S} Enc(pk_{\Omega'}, 0) \rightsquigarrow \boxed{c' \leftarrow_{\S} C^{mc-cpa}.Enc'(\sigma, 0)}$$

For every query after query i we simulate **HashOrAdapt** as in Game 2. Now, depending on the challenger's bit, we either simulate Game i or Game $i + 1$. Thus, we have that $|\Pr[S_2] - \Pr[S_{3_q}]| \leq q \cdot v_{mc-cpa}(\lambda)$, where q is the overall number of queries to **HashOrAdapt**.⁸

Now, the indistinguishability game is independent of the bit b , proving indistinguishability. \square

Theorem 7. *If Ω is perfectly correct, Σ is unforgeable, and Π is zero-knowledge as well as simulation-sound extractable, then Construction 3 provides enhanced collision-resistance.*

Proof. To prove enhanced collision-resistance, we use a sequence of games.

Game 0: The original enhanced collision-resistance game.

Game 1: As Game 0, but we modify the **CHPG** and the **CHAdapt** as follows:

CHPG'(1^λ) :

$$crs_{\Pi} \leftarrow_{\S} PG_{\Pi}(1^\lambda) \rightsquigarrow \boxed{crs_{\Pi}, \tau \leftarrow_{\S} SIM_1(1^\lambda)}$$

CHAdapt'(sk_{ch}, m, m', r, h) :

$$\pi' \leftarrow_{\S} Prf_{\Pi}(crs_{\Pi}, (pk_{\Omega}, pk_{\Sigma}, h, m'), (\sigma, \perp)) \rightsquigarrow \boxed{\pi' \leftarrow_{\S} SIM_2(crs_{\Pi}, \tau, (pk_{\Omega}, pk_{\Sigma}, h, m'))}$$

Transition - Game0 \rightarrow Game 1 : We bound the probability for an adversary to detect this game change by presenting a hybrid game, which, depending on a zero-knowledge challenger C^{zk} , either produces the distribution in Game 0 or Game 1, respectively.

⁸Note, if unrolled, using the bounds of Bellare et al. [9], $|\Pr[S_2] - \Pr[S_{3_q}]| \leq 2q \cdot v_{cpa}(\lambda)$ follows.

CHPG''(1^λ) :

$$(\text{crs}_\Pi, \tau) \leftarrow_{\S} \text{SIM}_1(1^\lambda) \rightsquigarrow \boxed{\text{crs}_\Pi \leftarrow_{\S} \mathcal{C}^{\text{zk}}}$$

CHAdapt''(sk_{ch}, m, m', r, h) :

$$\pi' \leftarrow_{\S} \text{SIM}_2(\text{crs}_\Pi, \tau, (\text{pk}_\Omega, \text{pk}_\Sigma, h, m')) \rightsquigarrow \boxed{\pi' \leftarrow_{\S} \mathcal{C}^{\text{zk}}.P_b((\text{pk}_\Omega, \text{pk}_\Sigma, h, m'), \sigma)}$$

Clearly, if the challenger's internal bit is 0, we simulate the distribution in Game 0, whereas we simulate the distribution in Game 1 otherwise. We have that $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{zk}}(\lambda)$.

Game 2: As Game 1, but we further modify the CHPG algorithm as follows:

CHPG'''(1^λ) :

$$(\text{crs}_\Pi, \tau) \leftarrow_{\S} \text{SIM}_1(1^\lambda) \rightsquigarrow \boxed{(\text{crs}_\Pi, \tau, \zeta) \leftarrow_{\S} \mathcal{E}_1(1^\lambda)}$$

Transition – Game 1 → Game 2 : Under simulation-sound extractability, Game 1 and Game 2 are indistinguishable. That is, $|\Pr[S_1] - \Pr[S_2]| = 0$.

Game 3: As Game 2, but we keep a list \mathcal{Q} of all hashes h previously submitted to the collision-finding oracle which are accepted by the CHCheck algorithm.

Transition – Game 2 → Game 3 : This change is conceptual, and thus, we have $|\Pr[S_2] - \Pr[S_3]| = 0$.

Game 4: As Game 3, but for every valid collision $(m^*, r^*, m'^*, r'^*, h^*)$ output by the adversary we observe that either (h^*, m^*, r^*) or (h^*, m'^*, r'^*) must be a “fresh” collision, i.e., $h^* \notin \mathcal{Q}$. We assume, without loss of generality, that (m'^*, r'^*) is the “fresh” collision. We run $(\text{sk}', \sigma') \leftarrow_{\S} \mathcal{E}_2(\text{crs}_\Pi, \zeta, (\text{pk}_\Omega, h^*, m'^*), r'^*)$ and abort if the extraction fails. We call this event E_1 .

Transition – Game 3 → Game 4 : Game 3 and Game 4 proceed identically, unless E_1 occurs. Assume, toward contradiction, that event E_1 occurs with non-negligible probability. We now construct an adversary \mathcal{B} which breaks the simulation-sound extractability property of the NIZK proof system with non-negligible probability. We engage with a simulation-sound extractability challenger \mathcal{C}^{sse} and modify the algorithms as follows:

CHPG''''(1^λ) :

$$(\text{crs}_\Pi, \tau, \zeta) \leftarrow_{\S} \mathcal{E}_1(1^\lambda) \rightsquigarrow \boxed{\text{crs}_\Pi \leftarrow_{\S} \mathcal{C}^{\text{sse}}}$$

CHAdapt''''(sk_{ch}, m, m', r, h) :

$$\pi' \leftarrow_{\S} \text{SIM}_2(\text{crs}_\Pi, \tau, (\text{pk}_\Omega, \text{pk}_\Sigma, h, m')) \rightsquigarrow \boxed{\pi' \leftarrow_{\S} \mathcal{C}^{\text{sse}}.\text{SIM}((\text{pk}_\Omega, \text{pk}_\Sigma, h, m'))}$$

In the end, we output $((\text{pk}_\Sigma, h^*, m'^*), r'^*)$ to the challenger. This shows that we have $|\Pr[S_3] - \Pr[S_4]| \leq \nu_{\text{sse}}(\lambda)$.

Reduction to eUNF-CMA: We are now ready to construct an adversary \mathcal{B} which breaks the unforgeability of the underlying Σ . Our adversary \mathcal{B} proceeds as follows. It receives pp_Σ and pk_Σ from its own challenger. To generate σ_0 , \mathcal{B} simply queries its signature oracle to obtain it on the message 0. It embeds them straightforwardly inside pp_{ch} and pk_{ch} to initialize \mathcal{A} . For adaption, a new signature σ' must be generated and encrypted. Those signatures are also obtained by querying the signature oracle. Now we know that we have extracted two witnesses (sk, σ) as well as (sk'', σ'') where one attests membership of $(\text{pk}_\Sigma, h^*, m'^*)$ in L , and one attests membership of $(\text{pk}_\Sigma, h^*, m'')$ for some fresh h^* in L . By the perfect correctness of the signature scheme, we know that at most one of them must be signature for h^* . However, as the signature was never queried, (h^*, σ) (or (h^*, σ'') resp.) must be a validating signature, breaking the unforgeability of the used Σ . Now, we have that $\Pr[S_4] \leq \nu_{\text{eunf-cma}}(\lambda)$. This concludes the proof. \square

We are now ready to present the proof of Theorem 4.

Proof. We first prove that full collision-resistance implies enhanced collision-resistance and then give a counterexample showing that the other direction of the implication does not hold.

F – CollRes \implies E – CollRes: Assume \mathcal{A} to be an adversary who breaks the enhanced collision-resistance. We can then construct an adversary \mathcal{B} which breaks the full collision-resistance. In particular, \mathcal{B} proceeds as follows. It receives pp_{ch} and pk_{ch} from its own challenger, and uses both to initialize \mathcal{A} . All queries to the collision-finding oracle are relayed to \mathcal{B} 's own oracle. Whenever \mathcal{A} outputs a winning tuple $(m^*, r^*, m'^*, r'^*, h^*)$, \mathcal{B} returns that tuple to its own challenger. As $m^* \neq m'^*$ must be true, and h^* was never queried to \mathcal{A} 's collision-finding oracle, this also means that (h^*, m^*) was never queried to \mathcal{B} 's oracle, thus meeting the winning condition.

E – CollRes $\not\implies$ F – CollRes : The scheme presented in Construction 3 gives a counterexample: it allows finding arbitrarily many collisions for a given hash h , if it sees a single one, but for no other $h' \neq h$. In more detail, to show that this construction is not fully collision-resistant, consider the following strategy. Receive $\text{pk}_{\text{ch}} = (\text{pk}_\Omega, \text{pk}_\Sigma)$ and $\text{pp}_{\text{ch}} = (\text{pp}_\Omega, \text{crs}_\Pi, \text{pp}_\Sigma)$. Compute $(h, r) \leftarrow_{\S} \text{CHash}(\text{pk}_{\text{ch}}, m)$, with m random. Also store the secret key sk_Ω' . Then, ask for an adaption (h, r, m) to (h, r', m') , where $r' = (\pi, c'', \text{pk}_\Omega')$, for some random m' . Then, compute $\sigma \leftarrow \text{Dec}(\text{sk}_\Omega', c'')$. Then arbitrary collisions for h are generated by executing CHAdapt in a similar way the owner of pk_{ch} does for finding collisions, due to the knowledge of σ for h . Because such collisions can only be generated for already seen collisions w.r.t. h , enhanced collision-resistance holds, but full collision-resistance does not. Also note that standard collision-resistance does not hold for Construction 3 for the same reason (we will later use this to derive a corollary). \square

Theorem 8. *Enhanced collision-resistance and standard collision-resistance together imply full collision-resistance.*

Proof. The theorem above is proven using a sequence of games.

Game 0: The original full collision-resistance game.

Game 1: As Game 0, we abort, if the adversary \mathcal{A} outputs $(m^*, r^*, m'^*, r'^*, h^*)$ such that the winning conditions are met, but h^* was never queried to the collision-finding oracle.

Transition – Game 0 \rightarrow Game 1 : If this is the case, we build an adversary \mathcal{B} which breaks the enhanced collision-resistance of the underlying scheme. Namely, \mathcal{B} receives pk_{Ch} and uses it to initialize \mathcal{A} . Every adaption query by \mathcal{A} is answered by \mathcal{B} using its own oracle. Once \mathcal{A} outputs $(m^*, r^*, m'^*, r'^*, h^*)$, \mathcal{B} returns $(m^*, r^*, m'^*, r'^*, h^*)$ to its own challenger. As h^* was never seen, \mathcal{B} wins its own game. $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{enh-collres}}(\lambda)$ follows.

Game 2: As Game 1, we abort, if the adversary \mathcal{A} outputs $(m^*, r^*, m'^*, r'^*, h^*)$ such that the winning conditions are met, but m^* was never queried to the collision-finding oracle.

Transition – Game 1 \rightarrow Game 2 : If this is the case, we build an adversary \mathcal{B} which breaks the standard collision-resistance of the underlying scheme. Namely, \mathcal{B} receives pk_{Ch} and uses it to initialize \mathcal{A} . Every adaption query by \mathcal{A} is answered by \mathcal{B} using its own oracle. Once \mathcal{A} outputs $(m^*, r^*, m'^*, r'^*, h^*)$, \mathcal{B} returns $(m^*, r^*, m'^*, r'^*, h^*)$ to its own challenger. As m^* was never seen, \mathcal{B} wins its own game. $|\Pr[S_1] - \Pr[S_2]| \leq \nu_{\text{st-collres}}(\lambda)$ follows.

In Game 2, the adversary can no longer win the full collision-resistance game. This proves the theorem. \square

The corollary below follows from the constructions used in the proofs of Theorem 3 and Theorem 4, which provide standard collision-resistance but not enhanced collision-resistance, and vice versa.

Corollary 1. *Standard collision-resistance and enhanced collision-resistance are independent.*

4.3. Relations Between Indistinguishability Notions.

We formally prove that full indistinguishability is strictly stronger than enhanced indistinguishability. Enhanced indistinguishability is strictly stronger than strong indistinguishability, which, in turn, is strictly stronger than indistinguishability (cf. Figure 9 for an overview).

Theorem 9. *Full Indistinguishability is strictly stronger than Enhanced Indistinguishability.*




Fig. 9. Relations between CH indistinguishability properties .

Proof. We first prove that full indistinguishability implies enhanced indistinguishability and then give a counterexample showing that the other direction of the implication does not hold.

F – Ind \implies E – Ind: Assume \mathcal{A} to be an adversary who wins the full indistinguishability game with some probability (non-negligibly) larger than $1/2$. Now, we construct an adversary \mathcal{B} which wins the enhanced indistinguishability game with the same probability. In particular, \mathcal{B} proceeds as follows. It receives pp_{ch} from its own challenger, generates $(\text{sk}_{\text{ch}}, \text{pk}_{\text{ch}}; \xi)$ honestly, and uses pp_{ch} and ξ to initialize \mathcal{A} . All queries to the collision-finding oracle are answered by querying \mathcal{B} 's own oracle (with the honestly generated keys). Whenever \mathcal{A} outputs a bit a , \mathcal{B} returns that bit to its own challenger. As the simulation is perfect, \mathcal{B} 's winning probability equals the one of \mathcal{A} .

E – Ind $\not\Rightarrow$ F – Ind : Let $\text{CH} := (\text{CHPG}, \text{CHKG}, \text{CHash}, \text{CHCheck}, \text{CHAdapt})$ be a fully indistinguishable chameleon-hash. We define a chameleon-hash $\text{CH}' := (\text{CHPG}', \text{CHKG}', \text{CHash}', \text{CHCheck}', \text{CHAdapt}')$, which internally uses CH as presented in Construction 4.

The basic idea is that in case particular random coins ξ are drawn, at each adaption, the message in question is augmented with a bit indicating that an adaption happened. As this particular randomness ($\xi = 0$) is never drawn with overwhelming probability, knowing the randomness does not help – being able to choose it, however, makes creating a distinguisher trivial.

$\text{CHPG}'(1^\lambda)$: Return $\text{pp}_{\text{ch}} \leftarrow_{\mathcal{S}} \text{CHPG}(1^\lambda)$.
 $\text{CHKG}'(\text{pp}_{\text{ch}})$: Draw $\xi \leftarrow_{\mathcal{S}} \{0, 1\}^\lambda$. Let $(\text{sk}_{\text{ch}}, \text{pk}'_{\text{ch}}) \leftarrow_{\mathcal{S}} \text{CHKG}(\text{pp}_{\text{ch}})$. Return $(\text{sk}_{\text{ch}}, (\text{pk}'_{\text{ch}}, 1))$, if $\xi = 0$, and $(\text{sk}_{\text{ch}}, (\text{pk}'_{\text{ch}}, 0))$ otherwise.
 $\text{CHash}(\text{pk}_{\text{ch}}, m)$: Parse pk_{ch} as $(\text{pk}'_{\text{ch}}, x)$. Return $(h, (r, 0))$, where $(h, r) \leftarrow_{\mathcal{S}} \text{CHash}(\text{pk}'_{\text{ch}}, (m, 0))$.
 $\text{CHCheck}(\text{pk}_{\text{ch}}, m, r, h)$: Parse r as (r', x) and pk_{ch} as $(\text{pk}'_{\text{ch}}, y)$. Return $\text{CHCheck}(\text{pk}'_{\text{ch}}, (m, x), r', h)$.
 $\text{CHAdapt}(\text{sk}_{\text{ch}}, m, m', r, h)$: Parse r as (r', x) and pk_{ch} as $(\text{pk}'_{\text{ch}}, y)$. If $y = 1$, let $r'' \leftarrow_{\mathcal{S}} \text{CHAdapt}(\text{sk}_{\text{ch}}, (m, x), (m', 1), r', h)$. Return $(r'', 1)$. Otherwise, let $r'' \leftarrow_{\mathcal{S}} \text{CHAdapt}(\text{sk}_{\text{ch}}, (m, 0), (m', 0), r', h)$. Return $(r'', 0)$.

Construction 4: E – Ind $\not\Rightarrow$ F – Ind

Clearly, if all parties generate their keys honestly (and thus a 0 is appended to the public key with overwhelming probability), the last bit appended to the randomness is always 0 after adaption, is never appended at hashing, and is independent of the message hashed. If, however, the adversary can choose randomness ξ , it can generate a pk_{ch} with an appended 1, thus making adaption append a 1, while hashing still appends a 0. This trivially breaks full indistinguishability. \square

Theorem 10. *Enhanced Indistinguishability is strictly stronger than Strong Indistinguishability.*

Proof. We first prove that enhanced indistinguishability implies strong indistinguishability and then give a counterexample showing that the other direction of the implication does not hold.

$E - \text{Ind} \implies S - \text{Ind}$: Assume \mathcal{A} to be an adversary who wins the strong indistinguishability game with non-negligible probability. Using \mathcal{A} we construct an adversary \mathcal{B} which wins the enhanced indistinguishability game with the same probability: \mathcal{B} receives pp_{ch} and r from its own challenger, generating $(\text{sk}_{\text{ch}}, \text{pk}_{\text{ch}}) \leftarrow_{\$} \text{CHKG}(\text{pp}_{\text{ch}}, \xi)$. It uses $(\text{sk}_{\text{ch}}, \text{pk}_{\text{ch}})$ to initialize \mathcal{A} . All queries to the collision-finding oracle are answered by querying \mathcal{B} 's own oracle. Whenever \mathcal{A} outputs a bit a , \mathcal{B} returns that bit to its own challenger. As the simulation is perfect, \mathcal{B} 's winning probability equals the one of \mathcal{A} .

$S - \text{Ind} \not\Rightarrow E - \text{Ind}$: Let $\text{CH} := (\text{CHPG}, \text{CHKG}, \text{CHash}, \text{CHCheck}, \text{CHAdapt})$ be chameleon-hash with enhanced indistinguishability. We define a chameleon-hash $\text{CH}' := (\text{CHPG}', \text{CHKG}', \text{CHash}', \text{CHCheck}', \text{CHAdapt}')$, which internally uses CH as presented in Construction 5.

The basic idea is that at key generation a key pair $(\text{sk}_{\Omega}, \text{pk}_{\Omega})$ for an encryption scheme is generated. The secret key sk_{Ω} is discarded and thus not part of sk_{ch} . At each hashing, the message is also encrypted using the public key pk_{Ω} and the ciphertext is attached to the randomness. Assuming the security of the encryption scheme, this does not leak any information about the message (and notice that no decryption oracle is provided, thus IND-CPA suffices), even if the secret key sk_{ch} is known. If, however, the random coins used to generate the key material become known, an adversary can simply generate sk_{Ω} and decrypt the ciphertexts and compare the content with the message in question.

$\text{CHPG}'(1^{\lambda})$: Let $\text{pp}_{\text{ch}} \leftarrow_{\$} \text{CHPG}(1^{\lambda})$ and $\text{pp}_{\Omega} \leftarrow_{\$} \text{PG}_{\Omega}(1^{\lambda})$. Return $(\text{pp}_{\text{ch}}, \text{pp}_{\Omega})$.
 $\text{CHKG}'(\text{pp}_{\text{ch}})$: Let $(\text{sk}_{\Omega}, \text{pk}_{\Omega}) \leftarrow_{\$} \text{KG}_{\Omega}(\text{pp}_{\Omega})$, and $(\text{sk}_{\text{ch}}, \text{pk}'_{\text{ch}}) \leftarrow_{\$} \text{CHKG}(\text{pp}_{\text{ch}})$. Return $(\text{sk}_{\text{ch}}, (\text{pk}'_{\text{ch}}, \text{pk}_{\Omega}))$.
 $\text{CHash}(\text{pk}_{\text{ch}}, m)$: Parse pk_{ch} as $(\text{pk}'_{\text{ch}}, \text{pk}_{\Omega})$. Let $c \leftarrow_{\$} \text{Enc}(\text{pk}_{\Omega}, m)$. Compute $(h, r) \leftarrow_{\$} \text{CHash}(\text{pk}'_{\text{ch}}, (m, c))$. Return $(h, (r, c))$.
 $\text{CHCheck}(\text{pk}_{\text{ch}}, m, r, h)$: Parse pk_{ch} as $(\text{pk}'_{\text{ch}}, \text{pk}_{\Omega})$, and r as (r', c) . Return $\text{CHCheck}(\text{pk}'_{\text{ch}}, (m, c), r', h)$.
 $\text{CHAdapt}(\text{sk}_{\text{ch}}, m, m', r, h)$: Parse r as (r', c) , compute $r'' \leftarrow_{\$} \text{CHAdapt}(\text{sk}_{\text{ch}}, (m, c), (m', c), r', h)$ and return (r'', c) .

Construction 5: $S - \text{Ind} \not\Rightarrow E - \text{Ind}$

Clearly, in the $S - \text{Ind}$ experiment, sk_{Ω} is discarded at key generation and is thus not given to the adversary. If, however, the adversary knows the randomness used to generate the keys, it can re-create sk_{Ω} . Consequently, $E - \text{Ind}$ is trivially broken by decrypting c contained in r . \square

Theorem 11. *Strong Indistinguishability is strictly stronger than (Normal) Indistinguishability.*

Proof. We first prove that full indistinguishability implies indistinguishability and then give a counterexample showing that the other direction of the implication does not hold.

$S - \text{Ind} \implies \text{Ind}$: Assume \mathcal{A} to be an adversary who wins the indistinguishability game with non-negligible probability. Using \mathcal{A} we construct an adversary \mathcal{B} which wins the strong indistinguishability game with the same probability: \mathcal{B} receives pp_{ch} from its own challenger, receiving $(\text{sk}_{\text{ch}}, \text{pk}_{\text{ch}})$, and uses pp_{ch} and pk_{ch} to initialize \mathcal{A} . All queries to the collision-finding oracle are answered by querying \mathcal{B} 's own oracle. Whenever \mathcal{A} outputs a bit a , \mathcal{B} returns that bit to its own challenger. As the simulation is perfect, \mathcal{B} 's winning probability equals the one of \mathcal{A} .

$\text{Ind} \not\Rightarrow S - \text{Ind}$: Our scheme given in Construction 9 provides a suitable counterexample. In particular, due to the used encryption, knowledge of the secret key allows extracting the original message m . In more detail, to show that this construction is not strongly indistinguishable, consider the following strategy. The key pair $(\text{sk}_{\text{ch}}, \text{pk}_{\text{ch}})$ is generated by the challenger, but (according to the game) known to the adversary. Obtain a challenge tuple $(h, r) \leftarrow_{\$} \text{HashOrAdapt}(\text{pk}_{\text{ch}}, \text{sk}_{\text{ch}}, m, m')$, where $m \neq m'$ are random messages. Then, let $m'' \leftarrow \text{Dec}(\text{sk}_{\text{ch}}, h)$. If $m = m''$, return 0. Otherwise, return 1. Clearly, this strategy always allows learning the challenger's bit. \square

4.4. Additional Separations

We now prove some additional separations. We note that indistinguishability is strictly weaker than full indistinguishability (as formally shown in Sect. 3.3).

Theorem 12. *Even full indistinguishability and uniqueness together do not imply weak collision-resistance.*

Proof. Consider the contrived construction given in Construction 6. The basic idea is to only make one randomness valid for all messages.

$\text{CHPG}(1^\lambda)$: Return \emptyset .
 $\text{CHKG}(\text{pp}_{\text{ch}})$: Return \emptyset .
 $\text{CHash}(\text{pk}_{\text{ch}}, m)$: Return (\emptyset, \emptyset) .
 $\text{CHCheck}(\text{pk}_{\text{ch}}, m, r, h)$: Return 1, if $h = \emptyset \wedge \text{pk}_{\text{ch}} = \emptyset \wedge r = \emptyset$. Return 0.
 $\text{CHAdapt}(\text{sk}_{\text{ch}}, m, m', r, h)$: Return \emptyset , if $h = \emptyset \wedge \text{pk}_{\text{ch}} = \emptyset \wedge r = \emptyset$. Return \perp .

Construction 6: Contrived Construction 1

Clearly, this construction is fully indistinguishable and unique. Finding collisions, however, is a trivial task. \square

Theorem 13. *Even full collision-resistance and uniqueness together do not imply indistinguishability.*

Proof. Assume $\text{CH} := (\text{CHPG}, \text{CHKG}, \text{CHash}, \text{CHCheck}, \text{CHAdapt})$ to be a fully collision-resistant, unique, and fully indistinguishable chameleon-hash. In Construction 7, we construct a CH' which offers full collision-resistance and uniqueness, but is not indistinguishable. The basic idea is to manipulate the hash to contain additional information about whether an adaption took place by appending the message itself.

$\text{CHPG}'(1^\lambda)$: Return $\text{CHPG}(1^\lambda)$.
 $\text{CHKG}'(\text{pp}_{\text{ch}})$: Return $\text{CHKG}(\text{pp}_{\text{ch}})$.
 $\text{CHash}(\text{pk}_{\text{ch}}, m)$: Let $(h, r) \leftarrow_{\S} \text{CHash}(\text{pk}_{\text{ch}}, (m, m))$. Return $((h, m), r)$.
 $\text{CHCheck}(\text{pk}_{\text{ch}}, m, r, h)$: Parse h as (h', \hat{m}) . Return $\text{CHCheck}(\text{pk}_{\text{ch}}, (m, \hat{m}), r, h')$.
 $\text{CHAdapt}(\text{sk}_{\text{ch}}, m, m', r, h)$: Parse h as (h', \hat{m}) . Return $\text{CHAdapt}(\text{sk}_{\text{ch}}, (m, \hat{m}), (m', \hat{m}), r', h')$.

Construction 7: Contrived Construction 2

Clearly, CH' is still fully collision-resistant and unique, but looking at the appended messages allows deciding whether an adaption has occurred. □

Theorem 14. *Even full collision-resistance and full indistinguishability together do not imply uniqueness.*

Proof. Assume $\text{CH} := (\text{CHPG}, \text{CHKG}, \text{CHash}, \text{CHCheck}, \text{CHAdapt})$ to be a fully collision-resistant, unique, and fully indistinguishable chameleon-hash. We construct CH' as given in Construction 8. The basic idea is to append a random bit to the randomness r which is ignored during verification.

$\text{CHPG}'(1^\lambda)$: Return $\text{CHPG}(1^\lambda)$.
 $\text{CHKG}'(\text{pp}_{\text{ch}})$: Return $\text{CHKG}(\text{pp}_{\text{ch}})$.
 $\text{CHash}(\text{pk}_{\text{ch}}, m)$: Let $(h, r) \leftarrow_{\S} \text{CHash}(\text{pk}_{\text{ch}}, m)$. Return $(h, (r, 0))$.
 $\text{CHCheck}(\text{pk}_{\text{ch}}, m, r, h)$: Parse r as (r', x) . Return $\text{CHCheck}(\text{pk}_{\text{ch}}, m, r', h)$.
 $\text{CHAdapt}(\text{sk}_{\text{ch}}, m, m', r, h)$: Parse r as (r', x) . Return $\text{CHAdapt}(\text{sk}_{\text{ch}}, m, m', r', h)$.

Construction 8: Contrived Construction 3

Clearly, CH' is still fully collision-resistant and fully indistinguishable, but changing the bit in the randomness r is trivial, breaking uniqueness. □

5. Fully Collision-Resistant Chameleon-Hashes

We are now ready to present our black-box construction of fully collision-resistant chameleon-hashes.

5.1. Construction

The main idea of our construction is to encrypt a message m using an mcIND-CPA secure encryption scheme and use the ciphertext as the hash, i.e., it is very close to our “contrived” construction providing enhanced collision-resistance given in Construction 3. However, it has some important, and subtle, differences.

Namely, the randomness r is a SSE NIZK attesting membership of a tuple containing the public key used for encryption, the hash, as well as the hashed message in the following NP-language:

$$L := \{(\mathbf{pk}_\Omega, h, m) \mid \exists (\mathbf{sk}_\Omega, \xi) : h = \text{Enc}(\mathbf{pk}_\Omega, m; \xi) \vee \text{KVf}_\Omega(\mathbf{pk}_\Omega, \mathbf{sk}_\Omega) = 1\}. \quad (2)$$

Informally, this language requires the prover to demonstrate that it either knows the randomness ξ attesting that h is a well-formed encryption of m under the CH key \mathbf{pk}_Ω , or it knows a secret key \mathbf{sk}_Ω corresponding to \mathbf{pk}_Ω , instead of encrypting a signature and proving the verification relation. Our construction of a fully collision-resistant CH is presented as Construction 9. We note that compared to Ateniese et al. [5] we cannot use true-simulation extractable NIZKs (tSE-NIZKs) [32] and need SSE NIZKs.

CHPG(1^λ) : Fix a public-key encryption scheme Ω and a compatible NIZK proof system for language L in (2). Return $\text{pp}_{\text{ch}} = (\text{pp}_\Omega, \text{crs}_\Pi)$, where

$$\text{pp}_\Omega \leftarrow_{\$} \text{PG}_\Omega(1^\lambda), \text{ and } \text{crs}_\Pi \leftarrow_{\$} \text{PG}_\Pi(1^\lambda).$$

CHKG(pp_{ch}) : Return $(\mathbf{sk}_{\text{ch}}, \mathbf{pk}_{\text{ch}}) = (\mathbf{sk}_\Omega, (\text{pp}_{\text{ch}}, \mathbf{pk}_\Omega))$, where

$$(\mathbf{sk}_\Omega, \mathbf{pk}_\Omega) \leftarrow_{\$} \text{KG}_\Omega(\text{pp}_\Omega).$$

CHash($\mathbf{pk}_{\text{ch}}, m$) : Parse \mathbf{pk}_{ch} as $((\text{pp}_\Omega, \text{crs}_\Pi), \mathbf{pk}_\Omega)$, and return $(h, r) = (c, \pi)$, where

$$(c; \xi) \leftarrow_{\$} \text{Enc}(\mathbf{pk}_\Omega, m), \text{ and } \pi \leftarrow_{\$} \text{Prf}_\Pi(\text{crs}_\Pi, (\mathbf{pk}_\Omega, h, m), (\perp, \xi)).$$

CHCheck($\mathbf{pk}_{\text{ch}}, m, r, h$) : Parse \mathbf{pk}_{ch} as $((\text{pp}_\Omega, \text{crs}_\Pi), \mathbf{pk}_\Omega)$, and r as π . Return 1, if the following holds, and 0 otherwise:

$$m \in \mathcal{M} \wedge \text{Vfy}_\Pi(\text{crs}_\Pi, (\mathbf{pk}_\Omega, h, m), \pi) = 1.$$

CHAdapt($\mathbf{sk}_{\text{ch}}, m, m', r, h$) : Parse \mathbf{sk}_{ch} as \mathbf{sk}_Ω . Verify whether $m' \in \mathcal{M}$, and $\text{CHCheck}(\mathbf{pk}_{\text{ch}}, m, r, h) = 1$. Return \perp , if not. Otherwise, return $r' = \pi'$, where

$$\pi' \leftarrow_{\$} \text{Prf}_\Pi(\text{crs}_\Pi, (\mathbf{pk}_\Omega, h, m'), (\mathbf{sk}_\Omega, \perp)).$$

Construction 9: Our Construction of a Fully Collision-Resistant CH

5.2. Security

Subsequently, we prove the security of our CH in Construction 9.

Theorem 15. *If Ω is correct and Π is complete, then CH in Construction 9 is correct.*

Correctness follows from inspection and the (perfect) correctness of the used primitives.

Theorem 16. *If Ω is mcIND-CPA secure, and Π is zero-knowledge, then CH in Construction 9 is indistinguishable (N-Ind).*

In the proof, we use frameboxes and \rightsquigarrow to highlight the changes we make in the algorithms throughout a sequence of games (and we only show the changes).

Proof. To prove indistinguishability, we use a sequence of games:

Game 0: The original indistinguishability game.

Game 1: As Game 0, but we modify the algorithms CHPG, CHash, and CHAdapt used inside the game:

CHPG'(1 $^\lambda$) :

$$\text{crs}_\Pi \leftarrow_{\S} \text{PG}_\Pi(1^\lambda) \rightsquigarrow (\text{crs}_\Pi, \tau) \leftarrow_{\S} \text{SIM}_1(1^\lambda).$$

CHash'(pk_{ch}, m) :

$$\pi \leftarrow_{\S} \text{Prf}_\Pi(\text{crs}_\Pi, (\text{pk}_\Omega, h, m), (\perp, \xi)) \rightsquigarrow \pi \leftarrow_{\S} \text{SIM}_2(\text{crs}_\Pi, \tau, (\text{pk}_\Omega, h, m))$$

CHAdapt'(sk_{ch}, m, m', r, h) :

$$\pi' \leftarrow_{\S} \text{Prf}_\Pi(\text{crs}_\Pi, (\text{pk}_\Omega, h, m'), (\text{sk}_\Omega, \perp)) \rightsquigarrow \pi' \leftarrow_{\S} \text{SIM}_2(\text{crs}_\Pi, \tau, (\text{pk}_\Omega, h, m')).$$

Transition – Game 0 \rightarrow Game 1 : We bound the probability for an adversary to detect this game change by presenting a hybrid game, which, depending on a zero-knowledge challenger \mathcal{C}^{Zk} , either produces the distribution in Game 0 or Game 1, respectively. In particular, assume that we use the following changes:

CHPG''(1 $^\lambda$) :

$$(\text{crs}_\Pi, \tau) \leftarrow_{\S} \text{SIM}_1(1^\lambda) \rightsquigarrow \text{crs}_\Pi \leftarrow_{\S} \mathcal{C}^{\text{Zk}}.$$

CHash''(pk_{ch}, m) :

$$\pi \leftarrow_{\S} \text{SIM}_2(\text{crs}_\Pi, \tau, (\text{pk}_\Omega, h, m)) \rightsquigarrow \pi \leftarrow_{\S} \mathcal{C}^{\text{Zk}}.P_b((\text{pk}_\Omega, h, m), (\perp, \xi)).$$

CHAdapt''(sk_{ch}, m, m', r, h) :

$$\pi' \leftarrow_{\S} \text{SIM}_2(\text{crs}_\Pi, \tau, (\text{pk}_\Omega, h, m')) \rightsquigarrow \pi' \leftarrow_{\S} \mathcal{C}^{\text{Zk}}.P_b((\text{pk}_\Omega, h, m'), (\text{sk}_\Omega, \perp)).$$

Clearly, if the challenger's internal bit is 0, we simulate the distribution in Game 0, whereas we simulate the distribution in Game 1 otherwise. We have that $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{zk}(\lambda)$.

Game 2: As Game 1, but we further modify the CHash algorithm as follows:

CHash''(pk_{ch}, m) :

$$(c; \xi) \leftarrow_{\S} \text{Enc}(\text{pk}_{\Omega}, m) \rightsquigarrow (c; \xi) \leftarrow_{\S} \text{Enc}(\text{pk}_{\Omega}, 0)$$

Transition – Game 1 → Game 2 : We bound the probability for an adversary to distinguish between two consecutive games by introducing a hybrid game which uses a multi-challenge IND-CPA challenger to interpolate between two consecutive games.

CHKG(pp_{ch})' : Return $(\perp, \text{pk}_{\text{ch}}) = (\perp, (\text{pp}_{\text{ch}}, \text{pk}_{\Omega}))$, where

$$(\text{sk}_{\Omega}, \text{pk}_{\Omega}) \leftarrow_{\S} \text{KG}_{\Omega}(\text{pp}_{\Omega}) \rightsquigarrow \text{pk}_{\Omega} \leftarrow_{\S} \mathcal{C}^{\text{mc-cpa}}$$

CHash'''(pk_{ch}, m) :

$$(c; \xi) \leftarrow_{\S} \text{Enc}(\text{pk}_{\Omega}, 0) \rightsquigarrow (c; \perp) \leftarrow_{\S} \mathcal{C}^{\text{mc-cpa}}.\text{Enc}'(m, 0)$$

Now, depending on the challenger's bit, we either simulate Game 1 or Game 2. Thus we have that $|\Pr[S_1] - \Pr[S_{2_i}]| \leq \nu_{\text{mc-cpa}}(\lambda)$

Now, the indistinguishability game is independent of the bit b , proving indistinguishability. \square

Theorem 17. *If Ω is perfectly correct and mcIND-CPA secure and Π is zero-knowledge as well as simulation-sound extractable, then CH in Construction 9 is fully collision-resistant.*

Proof. To prove full collision-resistance, we use a sequence of games.

Game 0: The original full collision-resistance game.

Game 1: As Game 0, but we modify the CHPG and the CHAdapt algorithm as follows:

CHPG'(1^λ) :

$$\text{crs}_{\Pi} \leftarrow_{\S} \text{PG}_{\Pi}(1^{\lambda}) \rightsquigarrow (\text{crs}_{\Pi}, \tau) \leftarrow_{\S} \text{SIM}_1(1^{\lambda})$$

CHAdapt'(sk_{ch}, m, m', r, h) :

$$\pi' \leftarrow_{\S} \text{Prf}_{\Pi}(\text{crs}_{\Pi}, (\text{pk}_{\Omega}, h, m'), (\text{sk}_{\Omega}, \perp)) \rightsquigarrow \pi' \leftarrow_{\S} \text{SIM}_2(\text{crs}_{\Pi}, \tau, (\text{pk}_{\Omega}, h, m'))$$

Transition – Game 0 → Game 1 : We bound the probability for an adversary to detect this game change by presenting a hybrid game, which, depending on a

zero-knowledge challenger \mathcal{C}^{zk} , either produces the distribution in Game 0 or Game 1, respectively.

CHPG''(1^λ):

$$(\text{crs}_\Pi, \tau) \leftarrow_{\S} \text{SIM}_1(1^\lambda) \rightsquigarrow \boxed{\text{crs}_\Pi \leftarrow_{\S} \mathcal{C}^{\text{zk}}}.$$

CHAdapt''($\text{sk}_{\text{ch}}, m, m', r, h$):

$$\pi' \leftarrow_{\S} \text{SIM}_2(\text{crs}_\Pi, \tau, (\text{pk}_\Omega, h, m')) \rightsquigarrow \boxed{\pi' \leftarrow_{\S} \mathcal{C}^{\text{zk}}.P_b((\text{pk}_\Omega, h, m'), \text{sk}_\Omega)}.$$

Clearly, if the challenger's internal bit is 0 we simulate the distribution in Game 0, whereas we simulate the distribution in Game 1 otherwise. We have that $|\Pr[S_0] - \Pr[S_1]| \leq \nu_{\text{zk}}(\lambda)$.

Game 2: As Game 1, but we further modify the CHPG algorithm as follows:

CHPG'''(1^λ):

$$(\text{crs}_\Pi, \tau) \leftarrow_{\S} \text{SIM}_1(1^\lambda) \rightsquigarrow \boxed{(\text{crs}_\Pi, \tau, \zeta) \leftarrow_{\S} \mathcal{E}_1(1^\lambda)}.$$

Transition – Game 1 \rightarrow Game 2 : Under simulation-sound extractability, Game 1 and Game 2 are indistinguishable. That is, $|\Pr[S_1] - \Pr[S_2]| = 0$.

Game 3: As Game 2, but we keep a list \mathcal{Q} of all tuples (h, r, m) previously submitted to the collision-finding oracle which are accepted by the CHCheck algorithm, where h was never submitted to the collision-finding oracle before.

Transition – Game 2 \rightarrow Game 3 : This change is conceptual, i.e., $|\Pr[S_2] - \Pr[S_3]| = 0$.

Game 4: As Game 3, but for every valid collision $(m^*, r^*, m'^*, r'^*, h^*)$ output by the adversary we observe that either (m^*, r^*) or (m'^*, r'^*) must be a “fresh” collision, i.e., one that was never output by the collision-finding oracle. We assume, without loss of generality, that (m'^*, r'^*) is the “fresh” collision. We run $(\text{sk}', \xi') \leftarrow_{\S} \mathcal{E}_2(\text{crs}_\Pi, \zeta, (\text{pk}_\Omega, h^*, m'^*), r'^*)$ and abort if the extraction fails. We call this event E_1 .

Transition – Game 3 \rightarrow Game 4 : Game 3 and Game 4 proceed identically, unless E_1 occurs. Assume, toward contradiction, that event E_1 occurs with non-negligible probability. We now construct an adversary \mathcal{B} which breaks the simulation-sound extractability property of the NIZK proof system with non-negligible probability. We engage with a simulation-sound extractability challenger \mathcal{C}^{sse} and modify the algorithms as follows:

CHPG'''(1 λ) :

$$(\text{crs}_\Pi, \tau, \zeta) \leftarrow_{\S} \mathcal{E}_1(1^\lambda) \rightsquigarrow \boxed{\text{crs}_\Pi \leftarrow_{\S} \mathcal{C}^{\text{sse}}}.$$

CHAdapt'''(sk_{ch}, m, m', r, h) :

$$\pi' \leftarrow_{\S} \text{SIM}_2(\text{crs}_\Pi, \tau, (\text{pk}_\Omega, h, m')) \rightsquigarrow \boxed{\pi' \leftarrow_{\S} \mathcal{C}^{\text{sse}}.\text{SIM}(\text{pk}_\Omega, h, m')}.$$

In the end we output $((\text{pk}_\Omega, h^*, m'^*), r'^*)$ to the challenger. This shows that we have $|\Pr[S_3] - \Pr[S_4]| \leq \nu_{\text{sse}}(\lambda)$.

Game 5: As Game 4, but we observe that if (m^*, r^*) does not correspond to a fresh collision for h^* in the above sense, then we will have an entry $(h^*, r, m) \in \mathcal{Q}$ where (m, r) is a “fresh” collision, i.e., one computed by the adversary. We run the extractor for the fresh collision, i.e., either obtain $(\text{sk}'', \xi'') \leftarrow_{\S} \mathcal{E}_2(\text{crs}_\Pi, \zeta, (\text{pk}_\Omega, h^*, m^*), r^*)$ or $(\text{sk}'', \xi'') \leftarrow_{\S} \mathcal{E}_2(\text{crs}_\Pi, \zeta, (\text{pk}_\Omega, h^*, m), r)$, respectively. In case the extraction fails, we abort. We call the abort event E_2 .

Transition – Game 4 \rightarrow Game 5 : Analogously to the transition between Game 3 and Game 4, we argue that Game 4 and Game 5 proceed identically unless E_2 occurs which is why we do not restate the reduction to simulation-sound extractability here. We have that $|\Pr[S_4] - \Pr[S_5]| \leq \nu_{\text{sse}}(\lambda)$.

Reduction to mcIND-CPA: We are now ready to construct an adversary \mathcal{B} which breaks the mcIND-CPA security of the underlying Ω . Our adversary \mathcal{B} proceeds as follows. It receives pp_Ω and pk_Ω from its own challenger. It embeds them straightforwardly as pp_{ch} and pk_{ch} to initialize \mathcal{A} . Now we know that we have extracted two witnesses (sk, ξ) as well as (sk'', ξ'') where one attests membership of $(\text{pk}_\Omega, h^*, m'^*)$ in L and one attests membership of $(\text{pk}_\Omega, h^*, m'')$ for some $m'' \neq m'^*$ in L . By the perfect correctness of the encryption scheme, we know that at most one of them can be consistent with the ciphertext contained in h^* , which implies that either sk or sk'' will be the key for the underlying encryption scheme (which of them we figure out by using KVf_Ω). With knowledge of the key, \mathcal{B} trivially breaks the mcIND-CPA security of the underlying Ω by randomly sending two distinct messages to its own challenger (for encryption), simply decrypting the returned ciphertext, and answering with the correct bit. We have that $\Pr[S_5] \leq \nu_{\text{mc-cpa}}(\lambda)$. This concludes the proof. \square

5.3. Concrete Instantiation

A suitable instantiation for Ω is ElGamal [37]. The algorithm KVf_Ω is simply checking whether $g^{\text{sk}_\Omega} = g^x = \text{pk}_\Omega$. Note that for Π we only need to extract a bounded number of times (i.e., twice). To this end one may use Fiat-Shamir transformed Σ -protocols for DLOG relations in the random-oracle model [36] when additionally applying the compiler by Faust et al. [35]. In particular, Faust et al. show that such proofs are simulation-sound extractable when additionally including the statement x upon hashing in the

challenge computation and if the Σ -protocol provides a property called quasi-unique responses. The latter is straightforward for the statements which need to be proven in our context. See, e.g., [30], for a detailed discussion of this transformation.

For the sake of completeness and to demonstrate how efficiently our approach can be instantiated, we provide this concrete instantiation as Construction 10. Therefore, let $(\mathbb{G}, g, q) \leftarrow_{\$} \mathbf{GGen}(1^\lambda)$ be an instance generator which returns a prime-order, and multiplicatively written, group \mathbb{G} where the DDH problem is hard, along with a generator g such that $\langle g \rangle = \mathbb{G}$. Note that an SSE NIZK for the required L in (3) can easily be obtained as an *equality* proof of two discrete logarithms together with an *or* composition of a proof of a discrete logarithm [24] of Fiat-Shamir transformed Σ -protocols discussed above.

$$L := \{(y, h, m) \mid \exists (x, \xi) : h = (g^\xi, m \cdot y^\xi) \vee y = g^x\}. \quad (3)$$

CHPG(1^λ): Outputs the public parameters $\mathbf{pp}_{\text{ch}} = (\mathbb{G}, g, q, H)$, where $(\mathbb{G}, g, q) \leftarrow_{\$} \mathbf{GGen}(1^\lambda)$ is a group \mathbb{G} of prime order q generated by g , and $H : \{0, 1\}^* \rightarrow \mathbb{Z}_q$ is a hash function (which we assume to behave like a random oracle and to be implicitly available to all algorithms below).

CHKG(\mathbf{pp}_{ch}): Return $(\mathbf{sk}_{\text{ch}}, \mathbf{pk}_{\text{ch}}) = (x, y)$, where $x \leftarrow_{\$} \mathbb{Z}_q$ and $y \leftarrow g^x$.

CHash($\mathbf{pk}_{\text{ch}}, m$): Parse \mathbf{pk}_{ch} as y , choose $(\xi, k_1, e_2, s_2) \leftarrow_{\$} \mathbb{Z}_q^4$, set $u_{1,1} \leftarrow g^{k_1}$, $u_{1,2} \leftarrow y^{k_1}$, $u_2 \leftarrow g^{s_2} \cdot y^{-e_2}$, $e \leftarrow H((y, h, m), (u_{1,1}, u_{1,2}, u_2))$ and $e_1 \leftarrow e - e_2 \bmod q$. Then compute $s_1 \leftarrow k_1 + e_1 \xi \bmod q$ and finally, return $(h, r) = (c, \pi)$, where

$$c \leftarrow (c_1, c_2) = (g^\xi, m \cdot y^\xi), \text{ and } \pi \leftarrow (e_1, e_2, s_1, s_2).$$

CHCheck($\mathbf{pk}_{\text{ch}}, m, r, h$): Parse \mathbf{pk}_{ch} as y and r as (e_1, e_2, s_1, s_2) , and h as (c_1, c_2) . Return 1 if the following holds, and 0 otherwise:

$$m \in \mathbb{G} \wedge e_1 + e_2 = H((y, h, m), (g^{s_1} \cdot c_1^{-e_1}, y^{s_1} \cdot (c_2/m)^{-e_1}, g^{s_2} \cdot y^{-e_2})).$$

CHAdapt($\mathbf{sk}_{\text{ch}}, m, m', r, h$): Parse \mathbf{sk}_{ch} as x , and h as (c_1, c_2) . Verify whether $m' \in \mathbb{G}$, and

CHCheck($\mathbf{pk}_{\text{ch}}, m, r, h$) = 1. Return \perp if not. Otherwise, choose $(k_2, e_1, s_1) \leftarrow_{\$} \mathbb{Z}_q^3$, set $u_{1,1} \leftarrow g^{s_1} \cdot c_1^{-e_1}$, $u_{1,2} \leftarrow y^{s_1} \cdot (c_2/m')^{-e_1}$, $u_2 \leftarrow g^{k_2}$, $e \leftarrow H((y, h, m'), (u_{1,1}, u_{1,2}, u_2))$, and $e_2 \leftarrow e - e_1 \bmod q$. Finally compute $s_2 \leftarrow k_2 + e_2 x \bmod q$, and return $r' = \pi'$, where

$$\pi' \leftarrow (e_1, e_2, s_1, s_2).$$

Construction 10: Concrete instantiation of a Fully Collision-Resistant CH

5.4. Comparison

Subsequently, in Table 1 we compare existing constructions of chameleon-hashes providing the *W-CollRes*, *E-CollRes* and *S-CollRes* notions with instantiations of our approach (in the random oracle and standard model) providing the stronger *F-CollRes* notion. Here E denotes an exponentiation in the respective algebraic structure, “?” denotes that it is unclear how efficient this can be realized due to requirement of an invertible

Table 1. Comparison of different chameleon-hash functions.

| Scheme | CR | $ h $ | $ h _{\text{bit}}$ | $ r $ | $ r _{\text{bit}}$ | CHash | CHAdapt | Ass. | Model |
|----------|----|-----------------|--------------------|---------------------------------|--------------------|----------------------|---------------------|--------|-------|
| [46] | W | $1\mathbb{G}$ | 256 | $1\mathbb{Z}_q$ | 256 | $2E_{\mathbb{G}}$ | $0E_{\mathbb{G}}$ | DLOG | SM |
| [5] (1) | E | $1\mathbb{G}$ | 256 | $12\mathbb{G}+7\mathbb{Z}_q$ | 4876 | $17E_{\mathbb{G}}$ | ? | DDH | ROM |
| [5] (2) | E | $1\mathbb{G}_1$ | 382 | $6\mathbb{G}_1+13\mathbb{G}_2$ | 12211 | $51E_{\mathbb{G}_1}$ | ? | SXDH | SM |
| [45] (1) | E | $1\mathbb{G}_1$ | 382 | $9\mathbb{G}_1+4\mathbb{G}_2$ | 6490 | $25E_{\mathbb{G}_1}$ | $1E_{\mathbb{Z}_q}$ | SXDH | SM |
| [45] (2) | E | $1\mathbb{G}_1$ | 382 | $3\mathbb{G}_1$ | 1164 | $6E_{\mathbb{G}_1}$ | $1E_{\mathbb{Z}_q}$ | PKoE | SM |
| [18] | S | $1\mathbb{Z}_N$ | 3072 | $1\mathbb{Z}_N$ | 3072 | $1E_{\mathbb{Z}_N}$ | $1E_{\mathbb{Z}_N}$ | OM-RSA | ROM |
| Ours | F | $2\mathbb{G}$ | 514 | $4\mathbb{Z}_q$ | 1024 | $6E_{\mathbb{G}}$ | $5E_{\mathbb{G}}$ | DDH | ROM |
| Ours | F | $2\mathbb{G}_1$ | 764 | $\approx 1-2k \mathbb{G}_{1/2}$ | – | – | – | SXDH | SM |

$|\cdot|_{\text{bit}}$ refers to the bit size of the respective value which is currently believed to provide 128 bit security. We use 256bit elliptic curves for standard known order groups ($|\mathbb{G}| = 257$, $|\mathbb{Z}_q| = 256$), 3072bit RSA modulus for the RSA setting ($|\mathbb{Z}_N| = 3072$), and 381bit BLS12 curves for the SXDH setting ($|\mathbb{G}_1| = 382$, $|\mathbb{G}_2| = 763$, $|\mathbb{Z}_q| = 256$)

onto mapping into the used group (cf. the discussion in [45]). **SM** and **RO** denote the standard and the random oracle model, respectively.

Furthermore, **DDH**, **SXDH**, **PKoE**, and **OM-RSA** denote the decisional Diffie–Hellman, the symmetric DDH, the power knowledge of exponent [42], and the one-more RSA inversion [10] assumptions. We also stress that for constructions relying on **SXDH**, for typical instantiations of type-III bilinear groups, we have that $|\mathbb{G}_2| = 2(|\mathbb{G}_1| - 1) + 1$ (where $|\cdot|$ denotes the size of the representation of a group element). Regarding our construction in the standard model, e.g., using **SSE NIZKs** based on Groth–Sahai **NIZKs**, one can use the compiler in [27] to efficiently achieve simulation-sound extractability. We, however, note that a naive instantiation of our template in the standard model would still require to include bit-wise proofs of the parts of the witness which are in \mathbb{Z}_q , which would, all in all, require a number of group elements in the order of $1k - 2k$ (a very rough estimate; thus we also omit the remaining costs which is indicated by “–” in Table 1). It seems that switching to a variant of ElGamal in the target group (and maybe some other tweaks) would help to work around the requirement of having bit-wise proofs. While we are not able to provide a more efficient instantiation, we hope that future work will be able to do so. Finally, we note that we omit comparing our scheme given in Construction 3 as it is contrived and its sole purpose is to prove a separation result.

6. Applications

In this section we discuss (stronger) collision-resistance notions of chameleon-hashes in context of two applications, namely redactable blockchains as well as online/offline signatures.

6.1. Redactable Blockchains

While one of the major goals of blockchains is their immutability and in particular their use as an immutable append-only log, recently, starting with the work of Ateniese et al.

[5], there has been an increasing interest in blockchains that allow some controlled after-the-fact modification of their content. This is motivated by illegal content that was shown to be included into the Bitcoin blockchain [48], which represents a significant challenge for law enforcement agencies [55], as well as legislations like the European General Data Protection Regulation (GDPR) and the associated “right to be forgotten”. Solutions to this problem may either be for the permissioned- or permissionless-blockchain setting and cryptographic in nature [5,26,51] or non-cryptographic, where in the latter case it is based on the consensus layer of the blockchain [31].

We are considering the former and focus on block-level rewriting (change entire blocks) of blockchains instead of transaction-level rewriting (change single transactions within a block) in a permissionless setting (such as Bitcoin), as this illustrates the problem with much wider implications. In the following we are using the notation used in [5], and describe a block as triple of the form $B = \langle s, x, \text{ctr} \rangle$, where $s \in \{0, 1\}^\lambda$, $x \in \{0, 1\}^*$ and $\text{ctr} \in \mathbb{N}$ and a block is valid if

$$\mathbf{validblock}_q^D(B) := (H(\text{ctr}, G(s, x)) < D) \wedge (\text{ctr} < q) = 1.$$

Here, $H : \{0, 1\}^* \rightarrow \{0, 1\}^{2\lambda}$ and $G : \{0, 1\}^* \rightarrow \{0, 1\}^{2\lambda}$ are collision-resistant hash functions, and the parameters $D \in \mathbb{N}$ and $q \in \mathbb{N}$ are the difficulty level of the block and the maximum number of hash queries that a user is allowed to make in any given round of the protocol, respectively. The chaining of blocks is now done by requiring that when attaching a (valid) block $B' = \langle s', x', \text{ctr}' \rangle$ we have that $s' = H(\text{ctr}, G(s, x))$. Now to make blocks redactable, one changes the description of blocks to $B = \langle s, x, \text{ctr}, (h, r) \rangle$ where the new component is a chameleon-hash (h, r) and the validation predicate changes to

$$\mathbf{validblock}_q^D(B) := (H(\text{ctr}, h) < D) \wedge \mathbf{CHCheck}(\text{pk}_{\text{ch}}, (s, x), r, h) = 1 \wedge (\text{ctr} < q) = 1.$$

Chaining is now done by requiring that when attaching a (valid) block $B' = \langle s', x', \text{ctr}' \rangle$ we have that $s' = H(\text{ctr}, h)$. Observe that now computing a collision in the chameleon-hash gives very much power as it basically allows to rewrite the entire history of the blockchain.

Ateniese et al. in [5] discuss different ways to control this power to actually compute collisions (i.e., run $\mathbf{CHAdapt}$) where (1) either sk_{ch} may be available to some fully trusted single party only, or (2) sk_{ch} is generated using a multi-party computation (MPC) protocol and $\mathbf{CHAdapt}$ is also performed in a distributed way by some set of parties. We will discuss the implications of different collision-resistance notions to this setting, which is independent of which of these two approaches is going to be used.

We recall that Ateniese et al. [5], who introduced this application, rely on $\mathbf{E-CollRes}$ and Derler et al. in more recent work in [26] rely on $\mathbf{S-CollRes}$. Now, note that in such a permissionless setting as discussed above, where everybody is allowed to participate, it is reasonable to assume that an adversary sees the collisions computed for any blocks over some time in the system (as they will be broadcasted). Now let us discuss the single notions:

Weak Collision-Resistance (W-CollRes) A chameleon-hash providing this notion of collision-resistance provides absolutely no guarantees, as after seeing a single collision all guarantees are lost. A prime example is the Pedersen CH due to Krawczyk and Rabin [46] (cf. Sect. 4.1),

where a single seen collision exposes the secret key sk_{CH} to everybody. Clearly, this has significant consequences in the above scenario as then everybody can arbitrarily alter the blockchain.

Enhanced Collision-Resistance (E-CollRes) Recall that an adversary, when attacking some hash h^* , must have never input h^* to $\text{CHAdapt}'$. Now, this means that if an adversary targets a specific hash and then happens to see a collision for this hash (for some reason), suddenly all guarantees are lost and arbitrary collisions could be computed. Note that our construction in Sect. 4 clearly demonstrates potential problems with CHs only satisfying this notion. This still represents a significant problem with this application.

Standard Collision-Resistance (S-CollRes) Recall, that an adversary is only restricted to not query message m^* (which is associated with the computed collision h^*) was never queried to the collision-finding oracle. While this still might be problematic in the redactable blockchain setting, messages can very likely be made unique by perpending a large enough random tag/nonce (note that in this could easily be done in the block format of, e.g., the Bitcoin block structure). So, this notion seems suitable if the aforementioned constrained may, under certain circumstances, be guaranteed to be met, but is far away from being ideal.

Full Collision-Resistance (F-CollRes) We recall that, here, only the collision (h^*, m^*) was not generated by the collision-finding oracle, but there is no other restriction whatsoever. Consequently, this collision-resistance notion seems the “right” notion as no issues on higher levels need to be considered and very strong guarantees are already provided by the notion itself.

6.2. Online/Offline Signatures

Online/offline signatures (OOS) [33,34] are signatures which run in two phases, a potentially computationally expensive offline phase and a more efficient online phase. Latter clearly should be more efficient than the full signing algorithm. Thus, if the online phase is then run by a resource constrained signer, this allows such signers to compute signatures even if it might be too expensive to run the full signing algorithm of the respective signature scheme.

6.2.1. Hash-sign-switch OOS

In [52], Shamir and Tauman introduced the so called hash-sign-switch paradigm for OOS. Here, the key pair of *any* signature scheme is extended by the key pair of a chameleon-hash. The offline phase represents computing a signature on a chameleon-hash value h of a random message m' (the hash part). The online phase then represents computing a collision for h with the message m to be signed (the switch part). Shamir and Tauman in [52] propose (among an instantiation based on factoring) the use of the W-CollRes by Krawczyk and Rabin [46]. Note that this requires that for every offline signature, a new signature for a fresh chameleon-hash needs to be computed. Otherwise,

due to the key-exposure of the chameleon-hash the so obtained OOS gets insecure, i.e., one can forge signatures for arbitrary messages after seeing two signatures.

6.2.2. Key-exposure in OOS

Chen et al. in [21] observe that this key-exposure problem in OOS following this “hash-sign-switch” paradigm might impose a huge storage overhead due to the number of precomputed signatures in the offline phase. They then suggest to fix this problem by introducing a special double-trapdoor hash family based on the discrete logarithm assumption combined with a one-time trapdoor/hash key pair for each message signing. Although this removes a part of the problem, this is still not entirely generic and imposes an additional overhead.

We want to stress, that besides the storage overhead pointed out by Chen et al. [21], constructing such OOS using a chameleon-hash providing only **W-CollRes** might be even more problematic when it comes to what we informally call robustness. Imagine that due to a fault or some behavior triggered by an adversary, one of the signatures precomputed in the offline phase gets reused in the online phase. Then, the OOS is immediately completely broken. Note that this is somewhat reminiscent of the problem of secret key leakage when reusing the randomness in Schnorr-type signatures as repeatedly seen in case of ECDSA in practice (cf. [44]).

6.2.3. F-CollRes *CH* in OOS

Now, when instantiating OOS on the “hash-sign-switch” paradigm based on a **F-CollRes** chameleon-hash instead, this immediately resolves the above robustness issue and yields a completely generic solution. More so, in the offline phase only a *single* signature needs to be precomputed, which can be reused for all online signing operations while allowing the adversary to query signatures for arbitrary messages. Clearly, when it comes to concrete efficiency, it needs to be guaranteed that the online part remains more efficient than the signing operation of the underlying signature scheme. Taking for instance the concrete instantiation in Sect. 5.3, precomputing all the message-independent values of the **Adapt** algorithm except for u_2 (which is critical to robustness) in the offline phase, then the online phase requires two exponentiations. So while this does not yield a benefit when building OOS on Schnorr-type signatures, it will so for instance when using the BLS signature scheme [15] (where an estimate of signing including the hashing to the curve [56] requires a cost of strictly more than two exponentiations).⁹

Relying on a **F-CollRes** chameleon-hash thus provides a fully generic construction of OOS with this robustness feature (in contrast to [21] which is based on the discrete logarithm assumption), and using the recent results in [28] even immediately yields a construction from post-quantum assumptions.

⁹We note that even without any precomputation for the chameleon-hash and only using pre-computation tables for the static generator g_1 and public key y of 16 \mathbb{G}_1 elements each (as used in libraries such as RELIC [3]) and multi-base scalar multiplications this amounts to the cost of three exponentiations. Finally, one could even reduce the cost to a single exponentiation, but then one needs to ensure that the precomputed u_2 value never gets reused.

Acknowledgements

We want to thank the anonymous reviewers for their helpful feedback. The work of D.S. was done while with AIT Austrian Institute of Technology. This work was supported by the EU's Horizon 2020 ECSEL Joint Undertaking under grant agreement n°783119 (SECREDas), from the European Union's Horizon 2020 research and innovation programme under grant agreement n°871473 (KRAKEN) and by the Austrian Science Fund (FWF) and netidee SCIENCE under grant agreement P31621-N38 (PROFET).

Funding Open Access funding enabled and organized by Projekt DEAL.

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

References

- [1] M. Abe, B. David, M. Kohlweiss, R. Nishimaki, M. Ohkubo, Tagged one-time signatures: Tight security and optimal tag size, in *PKC*. (2013), pp. 312–331
- [2] S. Alsouri, Ö. Dagdelen, S. Katzenbeisser, Group-based attestation: Enhancing privacy and management in remote attestation, in *Trust*. (2010), pp. 63–77
- [3] D.F. Aranha, C.P.L. Gouvêa, T. Markmann, R.S. Wahby, K. Liao, RELIC is an Efficient Library for Cryptography. <https://github.com/relic-toolkit/relic>
- [4] G. Ateniese, D.H. Chou, B. de Medeiros, G. Tsudik, Sanitizable signatures, in *ESORICS*. (2005), pp. 159–177
- [5] G. Ateniese, B. Magri, D. Venturi, E.R. Andrade, Redactable blockchain - or - rewriting history in bitcoin and friends, in *EuroS&P*. (2017), pp. 111–126
- [6] G. Ateniese, B. de Medeiros, Identity-based chameleon hash and applications, In *FC*. (2004), pp. 164–180
- [7] G. Ateniese, B. de Medeiros, On the key exposure problem in chameleon hashes, in *SCN*. (2004), pp. 165–179
- [8] F. Bao, R.H. Deng, X. Ding, J. Lai, Y. Zhao, Hierarchical identity-based chameleon hash and its applications, in *ACNS*. (2011), pp. 201–219
- [9] M. Bellare, A. Boldyreva, S. Micali, Public-key encryption in a multi-user setting: Security proofs and improvements, in *Eurocrypt*. (2000), pp. 259–274
- [10] M. Bellare, C. Namprempe, D. Pointcheval, M. Semanko, The one-more-rsa-inversion problems and the security of chaum's blind signature scheme. *J. Cryptol.* **16**(3), 185–215 (2003)
- [11] M. Bellare, T. Ristov, Hash functions from sigma protocols and improvements to VSH, in *Asiacrypt*. (2008), pp. 125–142
- [12] M. Bellare, T. Ristov, A characterization of chameleon hash functions and new, efficient designs. *J. Cryptol.* **27**(4), 799–823 (2014)
- [13] M. Bellare, D. Riepel, L. Shea, Highly-effective backdoors for hash functions and beyond. *Cryptology ePrint Archive, Paper 2024/536* (2024). <https://eprint.iacr.org/2024/536>
- [14] O. Blazy, S.A. Kakvi, E. Kiltz, J. Pan, Tightly-secure signatures from chameleon hash functions, in *PKC*. (2015), pp. 256–279

- [15] D. Boneh, B. Lynn, H. Shacham, Short signatures from the weil pairing, in C. Boyd, editors, *Advances in Cryptology - ASIACRYPT 2001, 7th International Conference on the Theory and Application of Cryptology and Information Security, Gold Coast, Australia, December 9-13, 2001, Proceedings. Lecture Notes in Computer Science*, vol. 2248 (Springer, 2001), pp. 514–532
- [16] G. Brassard, D. Chaum, C. Crépeau, Minimum disclosure proofs of knowledge. *J. Comput. Syst. Sci.* **37**(2), 156–189 (1988)
- [17] C. Brzuska, M. Fischlin, T. Freudenreich, A. Lehmann, M. Page, J. Schelbert, D. Schröder, F. Volk, Security of sanitizable signatures revisited, in *PKC*. (2009), pp. 317–336
- [18] J. Camenisch, D. Derler, S. Krenn, H.C. Pöhls, K. Samelin, D. Slamanig, Chameleon-hashes with ephemeral trapdoors - and applications to invisible sanitizable signatures, in *PKC*. (2017), pp. 152–182
- [19] D. Cash, D. Hofheinz, E. Kiltz, C. Peikert, Bonsai trees, or how to delegate a lattice basis, in *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Monaco / French Riviera, May 30 - June 3, 2010. Proceedings.* (2010), pp. 523–552
- [20] X. Chen, F. Zhang, K. Kim, Chameleon hashing without key exposure, in *ISC*. (2004), pp. 87–98
- [21] X. Chen, F. Zhang, W. Susilo, Y. Mu, Efficient generic on-line/off-line signatures without key exposure, in *ACNS*. (2007), pp. 18–30
- [22] J. Choi, S. Jung, A handover authentication using credentials based on chameleon hashing. *IEEE Commun. Lett.* **14**(1), 54–56 (2010)
- [23] A. Cingolani, *Bitcoin as an Ideal Redactable Transaction Ledger*. Master’s thesis, Sapienza University of Rome (2020)
- [24] R. Cramer, I. Damgård, B. Schoenmakers, Proofs of partial knowledge and simplified design of witness hiding protocols, in *Crypto*. (1994), pp. 174–187
- [25] R. Cramer, V. Shoup, A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack, in *Crypto*. (1998), pp. 13–25
- [26] D. Derler, K. Samelin, D. Slamanig, C. Striecks, Fine-grained and controlled rewriting in blockchains: Chameleon-hashing gone attribute-based, in *NDSS* (2019)
- [27] D. Derler, D. Slamanig, Key-homomorphic signatures: definitions and applications to multiparty signatures and non-interactive zero-knowledge. *Des. Codes Cryptogr.* **87**(6), 1373–1413 (2019)
- [28] D. Derler, S. Krenn, K. Samelin, D. Slamanig, Fully collision-resistant chameleon-hashes from simpler and post-quantum assumptions, in C. Galdi, V. Kolesnikov, editors, *Security and Cryptography for Networks - 12th International Conference, SCN2020, Amalfi, Italy, September 14-16, 2020, Proceedings. Lecture Notes in Computer Science*, vol. 12238 (Springer, 2020), pp. 427–447
- [29] D. Derler, K. Samelin, D. Slamanig, Bringing order to chaos: The case of collision-resistant chameleon-hashes, in A. Kiayias, M. Kohlweiss, P. Wallden, V. Zikas, editors, *Public-Key Cryptography - PKC 2020*. (2020), pp. 462–492
- [30] D. Derler, D. Slamanig, Highly-efficient fully-anonymous dynamic group signatures, in *AsiaCCS*. (2018), pp. 551–565
- [31] D. Deuber, B. Magri, S.A.K. Thyagarajan Redactable blockchain in the permissionless setting, in *IEEE S&P*. (2019), pp. 124–138
- [32] Y. Dodis, K. Haralambiev, A. López-Alt, D. Wichs Efficient public-key cryptography in the presence of key leakage, in *Asiacrypt*. (2010), pp. 613–631
- [33] S. Even, O. Goldreich, S. Micali, On-line/off-line digital signatures. *J. Cryptol.* **9**(1), 35–67 (1996)
- [34] S. Even, O. Goldreich, S. Micali, On-line/off-line digital schemes, in G. Brassard, editors, *Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 1989, Proceedings. Lecture Notes in Computer Science*, vol. 435 (Springer, 1989), pp. 263–275
- [35] S. Faust, M. Kohlweiss, G.A. Marson, D. Venturi, On the non-malleability of the fiat-shamir transform, in *Indocrypt*. (2012), pp. 60–79
- [36] A. Fiat, A. Shamir, How to prove yourself: Practical solutions to identification and signature problems, in *Crypto*. (1986), pp. 186–194
- [37] T.E. Gamal, A public key cryptosystem and a signature scheme based on discrete logarithms, in *Crypto*. (1984), pp. 10–18
- [38] J. Groth, Simulation-sound NIZK proofs for a practical language and constant size group signatures, in *Asiacrypt*. (2006), pp. 444–459

- [39] J. Groth, Efficient fully structure-preserving signatures for large messages, in *Asiacrypt*. (2015), pp. 239–259
- [40] J. Groth, A. Sahai, Efficient non-interactive proof systems for bilinear groups, in *Eurocrypt*. (2008), pp. 415–432
- [41] S. Guo, D. Zeng, Y. Xiang, Chameleon hashing for secure and privacy-preserving vehicular communications. *IEEE Trans. Parallel Distrib. Syst.* **25**(11) (2014)
- [42] S. Hada, T. Tanaka, On the existence of 3-round zero-knowledge protocols, in *Crypto*. (1998), pp. 408–423
- [43] S. Hohenberger, B. Waters, Short and stateless signatures from the RSA assumption, in *Crypto*. (2009), pp. 654–670
- [44] J. Jancar, V. Sedlacek, P. Svenda, M. Šýs, Minerva: The curse of ECDSA nonces systematic analysis of lattice attacks on noisy leakage of bit-length of ECDSA nonces. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* **2020**(4), 281–308 (2020). <https://doi.org/10.13154/tches.v2020.i4.281-308>
- [45] M. Khalili, M. Dakhilalian, W. Susilo, Efficient chameleon hash functions in the enhanced collision resistant model. *Inf. Sci.* **510**, 155–164 (2020)
- [46] H. Krawczyk, T. Rabin, Chameleon signatures, in *NDSS*. (2000), pp. 143–154
- [47] Y. Li, S. Liu, Tagged chameleon hash from lattices and application to redactable blockchain. *Cryptology ePrint Archive, Paper 2023/774 (to appear at PKC 2024)* (2023). <https://eprint.iacr.org/2023/774>
- [48] R. Matzutt, J. Hiller, M. Henze, J.H. Ziegeldorf, D. Müllmann, O. Hohlfeld, K. Wehrle, A quantitative analysis of the impact of arbitrary blockchain content on bitcoin, in *FC*. (2018), pp. 420–438
- [49] P. Mohassel, One-time signatures and chameleon hash functions, in *SAC*. (2010), pp. 302–319
- [50] T.P. Pedersen, Non-interactive and information-theoretic secure verifiable secret sharing, in *Crypto*. (1991), pp. 129–140
- [51] K. Samelin, D. Slamanig, Policy-based sanitizable signatures, in *CT-RSA*. (2020), pp. 538–563
- [52] A. Shamir, Y. Tauman, Improved online/offline signature schemes, in *Crypto*. (2001), pp. 355–367
- [53] R. Steinfeld, L. Bull, H. Wang, J. Pieprzyk, Universal designated-verifier signatures, in *Asiacrypt*. (2003), pp. 523–542
- [54] R. Steinfeld, H. Wang, J. Pieprzyk, Efficient extension of standard schnorr/rsa signatures into universal designated-verifier signatures, in *PKC*. (2004), pp. 86–100
- [55] G. Tziakouris, Cryptocurrencies - A forensic challenge or opportunity for law enforcement? an INTER-POL perspective. *IEEE S&P* **16**(4) (2018)
- [56] R.S. Wahby, D. Boneh, Fast and simple constant-time hashing to the BLS12-381 elliptic curve. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* **20**(4), 154–179 (2019). <https://doi.org/10.13154/tches.v2019.i4.154-179>
- [57] R. Zhang, Tweaking TBE/IBE to PKE transforms with chameleon hash functions, in *ACNS*. (2007), pp. 323–339