Telling the Story of Chameleon Hash Functions: A 27-Year Review

Houssam DERFOUFI, Marina DEHEZ-CLEMENTI and Jean-Christophe DENEUVILLE

Fédération ENAC ISAE-SUPAERO ONERA, Université de Toulouse, Toulouse, France

Abstract. Chameleon hash functions are trapdoor hashes that allow authorized adaptations while preserving security against outsiders. They appear in chameleon signatures, sanitizable and redactable structures, and several ledger mechanisms, yet the literature remains scattered. To our knowledge, no prior work has offered a dedicated survey or SoK on CHFs. This paper provides the first unified account, covering 1998 to 2025. We build a usable overview instead of a taxonomy dump. We collect the published constructions into a compact dataset that records underlying assumptions, trapdoor arrangements, and target security notions, and we place these results on a timeline from 1998 to 2025 to show how definitions and design choices evolved. We then provide simple maps that let the reader pivot between authorization models, trapdoor structure, and functional or algebraic features, keeping comparisons focused and avoiding unnecessary parameter detail. The result is a reference for making informed choices in real deployments. We end with a short set of future roads we believe are worth exploring, drawn from the gaps we observed and aimed at aligning CHF design with practical constraints and follow-up studies.

Keywords: Chameleon Hash · Trapdoor Hash · Systematization of Knowledge

1 Introduction

Unforgeability and public verifiability are two main properties of digital signatures: respectively, only the secret key holder can produce a valid signature and anyone can check its validity given the public key. Together, they yield non-repudiation, the signer cannot deny issuing an authentic signature. As attractive as this property is, it is undesirable in settings where the signed information is classified, personal, or risk-bearing due to transferability: if a signed document leaks, anyone can prove its origin, which can expose whistleblowers or journalistic sources, reveal sealed-bid participants, deanonymize peer reviewers or political dissidents, or link individuals to sensitive health or employment records.

Consequently, undeniable signatures were introduced (Chaum and van Antwerpen [CvA90]), establishing validity requires the signer's cooperation via confirmation and disavowal protocols. Traditional undeniable signatures rely on zero-knowledge proofs; coupled with the interactive nature of verification, they induce availability and latency constraints that are impractical in many applications.

Chameleon signatures, introduced by Krawczyk and Rabin [KR98], pursue the same goal while avoiding interactive verification. The recipient can verify locally (non-interactive), yet cannot convince third parties of the signature's validity (non-transferability). The

E-mail: houssam.derfoufi@isae-supaero.fr (Houssam DERFOUFI), marina.dehez-clementi@isae-supaero.fr (Marina DEHEZ-CLEMENTI), jean-christophe.deneuville@enac.fr (Jean-Christophe DENEUVILLE)

This work is licensed under a "CC BY 4.0" license.
Date of this document: 2026-02-04.

mechanism underpinning this effect is the chameleon hash function (trapdoor hash function).

A chameleon hash function (CHF) provides collision resistance to outsiders while enabling *controlled* collisions to a trapdoor holder. More formally, for a public key pk and a trapdoor key tk, the map $CH(pk, m; r)$ is collision resistant for any adversary without access to tk (no two distinct inputs hash to the same value). With tk, for any chosen pair (m, r) one can efficiently compute (m', r') with $m \neq m'$ such that $CH(pk, m; r) = CH(pk, m'; r')$. Equivalently, collision resistance for CHFs requires that any adversary lacking tk cannot produce $m \neq m'$ and randomness r, r' with equal hash outputs.

This primitive mirrors the earlier notion of chameleon blobs [BCC88]. In that model, a blob is a commitment that is binding for outsiders (one fixed opening) yet equivocable by a party holding special knowledge; their chameleon variant (property (iv″) defined in [BCC88]) lets a simulator generate transcripts in which the same commitment opens either way without detection. Viewing $h = CH(pk, m; r)$ as the "blob" and (m, r) as its opening, outsider binding corresponds to collision resistance, while equivocation with auxiliary knowledge corresponds to trapdoor collisions. This "commitment with controlled equivocation" template is precisely how chameleon hashes are best understood.

Circling back, chameleon signatures follow the standard hash-then-sign paradigm, with the hash replaced by a CHF whose trapdoor is held by the intended recipient. Since the recipient can adapt openings of received chameleon hashes, no third party can obtain transferable evidence of validity, while the recipient still verifies locally and efficiently [KR98, BCC88].

1.1 Motivation

Despite this long history, the literature on CHFs remains fragmented. Many works have been driven by specific application needs (e.g., redactable ledgers [HZM+19a]), or motivated by broader cryptographic toolkits (e.g., lattice-based primitives [CHKP12, DM14]). As a result, ***we lack a clear systematization that connects the various threads of CHF research and clarifies their similarities and differences***, which could be used to provide guidance for future design and deployment. The community has explored questions of security, efficiency, and usability, but often in isolation and without a unifying framework.

This paper present a unified perspective on the design landscape of CHFs. Revisiting their evolution, constructions, and applications through the same lens provides not only historical perspective but also insights for the next generation of cryptographic systems.

Contributions Our contributions are twofold:

i.) We *compiled an extensive database of existing works on CHFs that captures their key design dimensions*, including underlying hardness assumptions, trapdoor structures, and security properties.

ii.) We *provided a historical narrative about the CHF research and proposed technical maps for its analysis*, proposing different focus points for extracting different insights on their construction and usage.

1.2 Scope

Systematizing almost 30 years of CHF research requires balancing breadth with depth. To frame our study, we address the following research questions:

- **RQ1**: *How have the security properties and adversarial models for CHFs evolved?* (Security frameworks)

- **RQ2**: *What are the main families of CHF constructions?* (Constructions)
- **RQ3**: *Which hardness assumptions underpin CHFs?* (Underlying hard problems)
- **RQ4**: *Which applications actually deploy CHFs and what design constraints do they impose (Applications)*?

In order to tackle these research questions, we gathered the literature following the methodology described in the next subsection.

1.3 Methodology

To systematize the literature on CHFs, we began with an extensive keyword-based search using terms including "chameleon hash", "chameleon hashing", "chameleon hash function", "trapdoor hash function", "trapdoor hash", and "trapdoor-hash". A naive Google Scholar query returned over 3,500 hits. Manual crawling would have been clearly infeasible for systematic analysis. SerpAPI is an online, real-time web-search API that returns structured JSON results from major search engines (in particular Google Scholar). We therefore relied on a combination of SerpAPI [Ser25] and a custom script to scrape the first 70 pages of results (approximately 1,200 papers), and cross-referenced these with major bibliographic databases (like the ACM Digital Library and the Semantic Scholar Graph API).

From this initial set, we performed abstract-based filtering, which reduced the pool to 784 candidate papers. Since this number remained too large, we applied a quality filter, discarding low-impact works (e.g., pre-2015 papers with very few citations or self-cited within a small research team). This yielded approximately 550 papers. We then conducted a labeling of the remaining works into two categories: on one hand, papers proposing new constructions or improvements of CHFs; and on the other hand, papers employing CHFs in applications. This resulted in approximately 70 papers for the first category, 190 papers for the second one.

Finally, we performed a backward literature search by inspecting the references of these works to capture missing but relevant contributions. We focused specifically on the DEV labeled papers as it is the main interest for our contribution. After this refinement, our corpus converged to approximately 290 papers in total (noting 78 papers labeled as DEV), which form the basis of our systematization. The whole methodology flow is summarized in Figure 1.

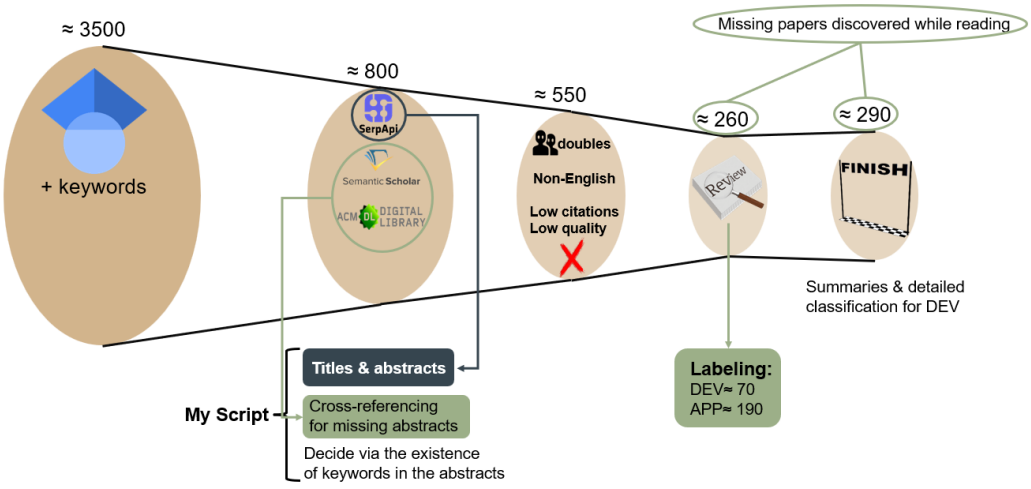

Figure 1: Illustration of the methodology applied to gather the scientific literature for the SoK.

1.4 Reader's guide

The remainder of this paper is organized as follows: Section 2 presents the necessary background and notation. The reader will not only find formal definitions in this section, but also a classification of the different mathematical assumptions underlying chameleon hash constructions. Section 3 proposes a chronological narrative of the evolutions of chameleon hash functions, detailing attempted improvements along with their respective limitations. Then a first taxonomy is proposed in Section 4. This taxonomy adopts a specific viewpoint, that is discussed in Section 5, before the conclusion of this systematization of knowledge work in Section 6.

It should also be noted that some citations refer to the same work e.g., [AdM03, AdM04a]. The oldest reference is necessary for the chronological story (see Section 3), but the most recent work might be more complete or accurate (peer-reviewed conference or journal).

2 Preliminaries

We provide the necessary background and notations in this section. We begin with general notations and reminders regarding CHFs in Section 2.1, then proceed with an overview of the underlying mathematical tools in Section 2.2. Security properties relative to CHFs are finally described and commented in Section 2.3.

2.1 Notations

Throughout this paper, $\mathcal{A}(\mathsf{args})$ will denote a Probabilistic Polynomial-Time (PPT) adversary, having knowledge of the arguments args. The security parameter is denoted λ (typically, $\lambda = 128$ bits), and negligible functions in this security parameter as $\mathsf{negl}(\lambda)$. Formally, a chameleon hash function CH is defined as follows:

Definition 1 (Chameleon Hash Function [AMVA17, CDH+17, DSS20])**.** A Chameleon Hash Function CH is a tuple of five algorithms (CHPG, CHKG, CHash, $\mathsf{CHCheck}$, $\mathsf{CHAdapt}$) defined as follows:

- CHPG: takes a security parameter λ as input, and outputs public parameters pp_{CH}.
- CHKG: takes as input the public parameters pp_{CH}, and outputs a chameleon key pair $(pk_{\mathsf{CH}}, sk_{\mathsf{CH}})$.
- CHash: takes as inputs the chameleon public key pk_{CH} and a message $m \in \mathcal{M}$, and outputs a chameleon hash value h with a randomness r.
- $\mathsf{CHCheck}$: takes as inputs the chameleon public key pk_{CH}, a message m, a chameleon hash value h and a randomness r, and outputs a bit $b \in \{0, 1\}$ (respectively for False or True depending on the validity of the hash value h with respect to the other parameters).
- $\mathsf{CHAdapt}$: takes as inputs the chameleon secret key sk_{CH}, two messages m and m', a chameleon hash value h and a randomness r; it outputs a new randomness r' that verifies $\mathsf{CHCheck}$ on inputs h, m' and r'.

Private-coin vs. public-coin hashing. Following Ateniese *et al.* [AMVA17], it is useful to distinguish whether the coins used by CHash are kept secret or made explicit. In the *private-coin* model, hashing uses hidden randomness and outputs a pair (h, ξ), where ξ is a check string; verification runs $\mathsf{CHCheck}(pk_{\mathsf{CH}}, m, h, \xi)$ without recomputing the hash from public coins. In the *public-coin* model, hashing is randomized with explicit coins r and returns (h, r); verification simply recomputes the hash,

$$\mathsf{CHCheck}(pk_{\mathsf{CH}}, m, h, r) = 1 \iff h = \mathsf{CHash}(pk_{\mathsf{CH}}, m, r).$$

Both interfaces are functionally equivalent given CHAdapt, yet they lead to different storage and verification patterns (store ξ in the private-coin case; store r in the public-coin case).

2.2 Assumptions Summary: Tradeoff Between Detail and Conciseness

In this section, we summarize the various hardness assumptions that underpin the construction of chameleon hash functions. To maintain the balance between detail and conciseness, these assumptions are consolidated into two tables: Table 1 for classical assumptions and Table 2 for post-quantum ones. Each table provides a concise reference for the assumption's full name, abbreviation, typical setting, and its corresponding citation. This format not only saves space but also avoids redundant explanations throughout the paper. By grouping them in this manner, we make the information easily accessible to the reader without losing important context. This approach allows us to maintain the flow of the paper while ensuring that all relevant assumptions are clearly presented and can be cross-referenced as needed.

Table 1: Hardness assumptions in classical CHF constructions.

Assumption (abbrev.)	**Typical setting / note**	**Citations**
Discrete Logarithm Problem (DLP)	Prime-order groups; ElGamal-style CHFs	[ElG85]
Computational Diffie–Hellman (CDH)	Prime-order groups (computational variant)	[DH76]
Decisional Diffie–Hellman (DDH)	Prime-order groups (indistinguishability)	[Bon98]
RSA Inversion (RSA)	RSA moduli; classic and ROM instantiations	[RSA78]
One-More RSA Inversion (OM-RSA)	RSA-based S-CollRes constructions	[BNPS03]
Strong RSA (SRSA)	Hidden-order groups / RSA moduli	[BP97]
Integer Factorization Problem (IFP)	RSA setting; group of quadratic residues	[LJMP90]
Quadratic Residuosity (QRP)	Quadratic residues modulo composite n	[GM82]
Decisional Composite Residuosity (DCR)	Paillier and related hidden-order settings	[Pai99]
Decisional Linear (DLIN)	Bilinear groups; k-linear family with k=2	[BBS04]
k-Linear (k-LIN)	Bilinear groups; generalized linear assumptions	[Sha07]
Symmetric External Diffie–Hellman (SXDH)	Type-III pairings; DDH holds in both sources	[Fre10]
External Diffie–Hellman (XDH)	Asymmetric pairings; DDH in G_1 only	[BGdMM05]
Power Knowledge of Exponent (PKoE)	Knowledge-type assumption for NIZKs / proofs	[Gro10]

Table 2: Hardness assumptions in post-quantum CHF constructions.

Assumption (abbrev.)	**Typical setting / note**	**Citations**
Learning Parity with Noise (LPN)	Code/linear-algebraic	[BKW03]
Learning With Errors (LWE)	Lattices; Ring-LWE variants, GPV-style trapdoors	[Reg09]
Short Integer Solution (SIS)	Lattices; trapdoor-based collisions / commitments	[Ajt96]
Syndrome Decoding (SD)	Code-based signatures / encryption family	[BMvT78]
Multivariate Quadratic (MQ)	Multivariate schemes; HFE-style instantiations	[Pat96]

2.3 CHF-specific security framework

When it comes to the security of CHFs, an experienced reader may directly refer to the framework proposed by Derler *et al.* [DSS20] and continue their reading with Section 3. For those who would like to benefit from additional context on the various security notions surrounding CHFs and their historical development, the following subsection offers a more detailed exposition.

The basic notion for CR (later labeled as weak CR) for CHFs was formalized by [KR98] (Definition 2).

Definition 2 ((Weak) Collision Resistance [KR98])**.** Consider the following experiment:

Experiment $\mathsf{Expt}_{\mathcal{A}}^{\mathsf{CollRes}}(1^\lambda)$

1. $pp_{\mathsf{CH}} \leftarrow \mathsf{CHPG}(1^\lambda)$; $(sk_{\mathsf{CH}}, pk_{\mathsf{CH}}) \leftarrow \mathsf{CHKG}(pp_{\mathsf{CH}})$.
2. $(m, r, m', r') \leftarrow \mathcal{A}(pk_{\mathsf{CH}})$.
3. Output 1 if $m \neq m'$ and $\mathsf{CH}(pk_{\mathsf{CH}}, m, r) = \mathsf{CH}(pk_{\mathsf{CH}}, m', r')$, 0 otherwise.

A CHF CH is *collision resistant* (CR) if for all PPT adversaries $\mathcal{A}$, it holds that

$$\Pr\left[\mathsf{Expt}_{\mathcal{A}}^{\mathsf{CollRes}}(1^\lambda) = 1\right] \leq \mathsf{negl}(\lambda).$$

In other words, the adversary only has access to the public key pk to carry out an attack on the scheme.

Later on, the Key Exposure Problem (KEP) was discovered in early constructions, that is, once an adversary sees a single or multiple valid collisions, they can recover the trapdoor key and thus, create collisions on their own. Note that this case is not covered by CR as defined then. Take for example the initial Discrete Logarithm (DL) based construction [KR98], once an adversary sees a valid collision

$$g^m \cdot y^r = g^{m'} \cdot y^{r'} \text{ with } y = g^x,$$

then the secret key x can easily be recovered by solving a simple linear equation $x = (m - m') \cdot (r' - r)^{-1}$. One of the first attempts to solve the KEP was [AdM03, AdM04a], integrating Identity-based Cryptography (IBC) into chameleon hashing. The authors argued that the *non-transferability* property is undermined if the scheme reveals the recipient's secret key when a signature forgery happens (by the recipient since they are the one holding the trapdoor information associated with the CHF). Namely, the signer can invalidate all subsequent legitimate issued signatures to the recipient. This is very costly, and a rational third party will think the recipient would never risk a forge that burns their key and all their other signatures. Thus, the third party is convinced by a signature the

recipient presents. Given the identity-based system properties, the CH public key can be deterministically derived per transaction from the recipient's identity and a unique tag, confining any trapdoor exposure from observed collisions to that transaction rather than revealing a reusable long-term key.

In [CZK04], Chen *et al.* observe that the idea in [AdM04a] still amounts to frequently changing the recipient's public key, which is not a true fix to the KEP. Authors therefore propose a certificate-based construction in Gap-Diffie-Hellman groups under the assumption that the **Computational Diffie-Hellman Problem (CDHP)** is hard. In their analysis, they prove *CR*, *semantic security*, and *message hiding*. The description is twofold: it presents a basic scheme that is ***Key-Exposure Free* (KEF)** but unsafe under key reuse with the same recipient; then, authors propose an improvement into a full scheme with customized identities per transaction that removes the reuse issue while preserving the aforementioned properties.

Later, Ateniese and de Medeiros [ADM04b] formally define *key-exposure freeness (KEF)* in a CH model defined with the following algorithms (KeyGen, Hash, UForge, IForge). Beyond CR, they require *semantic security*[1], and *message hiding*[2], the latter enforced by instance-forging IForge so that disputes can be settled without revealing the original opening. They introduce two design principles: one based on a single-trapdoor construction that guarantees *KEF* but not *message hiding*, and the second construction based on a double-trapdoor that supports both *KEF* and *message hiding* by tying IForge to a per-label trapdoor that becomes derivable after one collision, while the long-term secret remains safe. They also clarify that identity-based variants help localize exposure per transaction but are not *KEF* in the strict sense.

Circling back to the first notion of *CR* introduced by Krawczyk and Rabin [KR98], it was later referred to as *weak CR* (W-CollRes). In 2017, in the context of redactable blockchains, Ateniese *et al.* generalize chameleon hashing to the *secret-coin* setting, where the hashing algorithm outputs a hash/check pair and the random coins remain implicit. The authors formalize a strictly stronger security notion named *enhanced CR* (E-CollRes) [AMVA17]. The distinction is as follows: W-CollRes, together with *KEF*, fits the *public-coin* case where verification consists in recomputating from public inputs; E-CollRes captures the generalized *secret-coin* case where verification uses a check string rather than the raw randomness. In E-CollRes (Def. 3), the adversary is given access to a collision-finding oracle for adaptively chosen hashes but wins only if it outputs a valid collision $(h^*, m^*, \xi^*, m'^*, \xi'^*)$ for a target hash value h^* that was never queried to the oracle.

Definition 3 (Enhanced CR [AMVA17])**.** Consider the following experiment:

[1] "semantic security" means that the chameleon hash value does not reveal information about the underlying message; formally, the conditional entropy satisfies $H[m \mid C] = H[m]$ [AdM04a].

[2] "Message hiding" means that even if some collisions $(m', r'), (m'', r'')$ are later revealed, they do not decrease the entropy of the original opening (m, r) beyond what C already reveals, i.e., $H[(m, r) \mid C, (m', r'), (m'', r'')] = H[(m, r) \mid C]$ [AdM04a]

Experiment $\mathsf{Expt}^{\mathsf{E\text{-}CollRes}}_{\mathcal{A},\mathsf{CH}}(1^\lambda)$

1. $pp_{\mathsf{CH}} \leftarrow \mathsf{CHPG}(1^\lambda)$; $(sk_{\mathsf{CH}}, pk_{\mathsf{CH}}) \leftarrow \mathsf{CHKG}(pp_{\mathsf{CH}})$; $\mathcal{Q} \leftarrow \varnothing$.
2. Give $\mathcal{A}$ the key pk_{CH} and oracle access to $\mathcal{O}_{\text{adapt}}$ defined by: on input (m, r, m', h), if $\mathsf{CHCheck}(pk_{\mathsf{CH}}, m, r, h) = 1$ then $r' \leftarrow \mathsf{CHAdapt}(sk_{\mathsf{CH}}, m, m', r, h)$, set $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{h\}$, and return r'; otherwise return $\perp$.
3. $\mathcal{A}$ outputs $(m^*, r^*, m'^*, r'^*, h^*)$.
4. Output 1 iff all hold: $m^* \neq m'^*$ and $h^* \notin \mathcal{Q}$ and

$$\mathsf{CHCheck}(pk_{\mathsf{CH}}, m^*, r^*, h^*) = 1 = \mathsf{CHCheck}(pk_{\mathsf{CH}}, m'^*, r'^*, h^*),$$

otherwise output 0.

A CHF CH is *enhanced CR* if for all PPT adversaries $\mathcal{A}$, it holds that

$$\Pr\left[\mathsf{Expt}^{\mathsf{E\text{-}CollRes}}_{\mathcal{A},\mathsf{CH}}(1^\lambda) = 1\right] \leq \mathsf{negl}(\lambda).$$

Moving on, Camenisch *et al.* introduced the *standard CR* property ($\mathsf{S\text{-}CollRes}$) in a unified framework where the CHF is described by a tuple of five algorithms ($\mathsf{CParGen}$,CKGen,CHash,$\mathsf{CHashCheck}$,$\mathsf{CHAdapt}$) [CDH+17]. In $\mathsf{S\text{-}CollRes}$ (Definition 4), the adversary may query $\mathsf{CHAdapt}$ on messages of its choice and obtain valid openings, yet it wins only if it outputs a fresh collision $(h^*, m^*, r^*, m'^*, r'^*)$ for which the pair was not obtained from its adapt queries.

Definition 4 (Standard CR [CDH+17])**.** Consider the following experiment:

Experiment $\mathsf{Expt}^{\mathsf{S\text{-}CollRes}}_{\mathcal{A},\mathsf{CH}}(1^\lambda)$

1. $pp_{\mathsf{CH}} \leftarrow \mathsf{CHPG}(1^\lambda)$; $(sk_{\mathsf{CH}}, pk_{\mathsf{CH}}) \leftarrow \mathsf{CHKG}(pp_{\mathsf{CH}})$; $\mathcal{Q} \leftarrow \varnothing$.
2. Give $\mathcal{A}$ the key pk_{CH} and oracle access to $\mathcal{O}_{\text{adapt}}$ defined by: on input (m, r, m', h), if $\mathsf{CHCheck}(pk_{\mathsf{CH}}, m, r, h) = 1$ then $r' \leftarrow \mathsf{CHAdapt}(sk_{\mathsf{CH}}, m, m', r, h)$, set $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{m, m'\}$, and return r'; otherwise return $\perp$.
3. $\mathcal{A}$ outputs $(m^*, r^*, m'^*, r'^*, h^*)$.
4. Output 1 iff the following hold: $m^* \neq m'^*$ and $m^* \notin \mathcal{Q}$ and

$$\mathsf{CHCheck}(pk_{\mathsf{CH}}, m^*, r^*, h^*) = 1 = \mathsf{CHCheck}(pk_{\mathsf{CH}}, m'^*, r'^*, h^*),$$

otherwise output 0.

A CHF CH is *standard CR* property if for all PPT adversaries $\mathcal{A}$ there exists a negligible function $\mathsf{negl}(\cdot)$ such that

$$\Pr\left[\mathsf{Expt}^{\mathsf{S\text{-}CollRes}}_{\mathcal{A}}(1^\lambda) = 1\right] \leq \mathsf{negl}(\lambda).$$

Alongside this, Camenisch *et al.* formalize *indistinguishability* (Def. 5) via a hash or adapt experiment where one must decide whether a valid opening came from CHash or from $\mathsf{CHAdapt}$ [CDH+17]. In addition, the authors isolate the *uniqueness* property as the infeasibility of finding two different openings for the same couple message/CH value (m, h). In their paper, a CHF is therefore called *secure* if it is *correct*, *S-CollRes*, and *indistinguishable*.

Definition 5 (Normal Indistinguishability [CDH+17])**.** Consider the following experiment:

Experiment $\mathsf{Expt}^{\mathsf{N\text{-}Ind}}_{\mathcal{A}}(1^\lambda)$

1. $\mathsf{pp}_{\mathsf{CH}} \leftarrow \mathsf{CParGen}(1^\lambda)$.
2. $(sk_{\mathsf{CH}}, pk_{\mathsf{CH}}) \leftarrow \mathsf{CKGen}(\mathsf{pp}_{\mathsf{CH}})$.
3. Sample $b \leftarrow \{0,1\}$ uniformly at random.
4. Run $\mathcal{A}^{\mathsf{HashOrAdapt}(\cdot,\cdot,\cdot),\,\mathsf{Adapt}(\cdot,\cdot,\cdot,\cdot,\cdot)}(pk_{\mathsf{CH}})$, where the oracles behave as follows:
   - **Oracle** $\mathsf{HashOrAdapt}(sk_{\mathsf{CH}}, m, m', b)$:
     (a) Compute $(h, r, \mathsf{td}) \leftarrow \mathsf{CHash}(pk_{\mathsf{CH}}, m')$.
     (b) Compute $(h', r', \mathsf{td}') \leftarrow \mathsf{CHash}(pk_{\mathsf{CH}}, m)$.
     (c) Let $r'' \leftarrow \mathsf{CHAdapt}(sk_{\mathsf{CH}}, m, m', r', h', \mathsf{td}')$.
     (d) If $r'' = \perp$ or $r' = \perp$, return $\perp$.
     (e) If $b = 0$, return (h, r, td).
     (f) If $b = 1$, return (h'', r, td'), where $h'' = \mathsf{CHash}(pk_{\mathsf{CH}}, m, r'')$.
   - **Oracle** Adapt is available to $\mathcal{A}$ in its standard form.
5. Let a be the bit output by $\mathcal{A}$.
6. The experiment outputs 1 if $a = b$, and 0 otherwise.

A CHF CH is *indistinguishable* if for all PPT adversaries $\mathcal{A}$, it holds that

$$\Pr\left[\mathsf{Expt}^{\mathsf{N\text{-}Ind}}_{\mathcal{A}}(1^\lambda) = 1\right] \leq \tfrac{1}{2} + \mathsf{negl}(\lambda).$$

In the same work [CDH+17], the authors extend the CH model to *ephemeral trapdoors*, namely **Chameleon Hash function with Ephemeral Trapdoor (CHET)**, where the hashing algorithm returns auxiliary data needed by the adaptation algorithm, and they define two new security properties: *outsider CR* against adversaries that only see public information and adapted openings; and *insider CR* against adversaries that even know the long term secret key but not the ephemeral data.

Two years after, in 2019, Derler *et al.* introduced the *strong indistinguishability* property (Def. 6), shortened as S-Ind: even if the adversary learns the secret key, a valid opening must not reveal whether it came from hashing or from adaptation [DSSS19]. The authors give the explicit game where the challenger flips a bit and answers a left–or–right *HashOrAdapt* query with either a fresh hash or an adapted opening, and the adversary (who knows sk) must guess the bit; they also carry the definitions over to the CHET setting, alongside *public* and *private* (outsider/insider) CR for CHET.

Definition 6 (Strong Indistinguishability [DSSS19])**.** Consider the following experiment:

Experiment $\mathsf{Expt}^{\mathsf{S\text{-}Ind}}_{\mathcal{A}}(1^\kappa)$

1. $\mathsf{pp}_{\mathsf{CH}} \leftarrow \mathsf{PPGen}_{\mathsf{CH}}(1^\kappa)$.
2. $(sk_{\mathsf{CH}}, pk_{\mathsf{CH}}) \leftarrow \mathsf{KGen}_{\mathsf{CH}}(\mathsf{pp}_{\mathsf{CH}})$.
3. Sample $b \leftarrow \{0,1\}$ uniformly at random.
4. Run $\mathcal{A}^{\mathsf{HashOrAdapt}_{\mathsf{CH}}(\cdot,\cdot,b)}(sk_{\mathsf{CH}}, pk_{\mathsf{CH}})$, where the oracle behaves as follows:
   - **Oracle** $\mathsf{HashOrAdapt}_{\mathsf{CH}}(sk_{\mathsf{CH}}, m, m', b)$:
     (a) Compute $(h, r) \leftarrow \mathsf{Hash}_{\mathsf{CH}}(pk_{\mathsf{CH}}, m')$.
     (b) Compute $(h', r') \leftarrow \mathsf{Hash}_{\mathsf{CH}}(pk_{\mathsf{CH}}, m)$.
     (c) Let $r'' \leftarrow \mathsf{Adapt}_{\mathsf{CH}}(sk_{\mathsf{CH}}, m, m', r', h')$.
     (d) If $b = 0$, return (h, r).
     (e) If $b = 1$, return (h', r'').
5. Let a be the bit output by $\mathcal{A}$.
6. The experiment outputs 1 if $a = b$, and 0 otherwise.

A CHF CH is *strongly indistinguishable* if for all PPT adversaries $\mathcal{A}$, it holds that

$$\Pr\left[\mathsf{Expt}^{\mathsf{S\text{-}Ind}}_{\mathcal{A}}(1^\kappa) = 1\right] \leq \tfrac{1}{2} + \mathsf{negl}(\kappa).$$

Shortly after, Samelin and Slamanig push this further and define the *full indistinguishability* property (Def. 7), shortened (F-Ind), where the adversary may even generate the secret keys used inside the HashOrAdapt oracle, yet must still fail to distinguish fresh from adapted openings; they formulate F-Ind both for standard CH/CHET and for policy-based CH, keeping the same left–or–right structure [SS20]. In short, S-Ind tolerates secret-key leakage, while F-Ind tolerates adversarial key generation.

Definition 7 (Full Indistinguishability [SS20])**.** Consider the following experiment:

Experiment $\mathsf{Expt}^{\mathsf{F\text{-}Ind}}_{\mathcal{A},\mathsf{CH}}(1^\kappa)$

1. $\mathsf{pp}_{\mathsf{CH}} \leftarrow \mathsf{PPGen}_{\mathsf{CH}}(1^\kappa)$.
2. Sample $b \leftarrow \{0,1\}$ uniformly at random.
3. Run $\mathcal{A}^{\mathsf{HashOrAdapt}(\cdot,\cdot,\cdot,\cdot,b)}(\mathsf{pp}_{\mathsf{CH}})$, where the oracle behaves as follows:
   - **Oracle** $\mathsf{HashOrAdapt}(sk_{\mathsf{CH}}, pk_{\mathsf{CH}}, m, m', b)$:
     (a) Compute $(h, r) \leftarrow \mathsf{Hash}_{\mathsf{CH}}(pk_{\mathsf{CH}}, m')$.
     (b) Compute $(h', r') \leftarrow \mathsf{Hash}_{\mathsf{CH}}(pk_{\mathsf{CH}}, m)$.
     (c) Let $r'' \leftarrow \mathsf{Adapt}_{\mathsf{CH}}(sk_{\mathsf{CH}}, m, m', r', h')$.
     (d) If $r = \bot$ or $r' = \bot$ or $r'' = \bot$, return $\bot$.
     (e) If $b = 0$, return (h, r).
     (f) If $b = 1$, return (h', r'').
4. Let b^* be the bit output by $\mathcal{A}$.
5. The experiment outputs 1 if $b^* = b$, and 0 otherwise.

A CHF CH is *fully indistinguishable* if for all PPT adversaries $\mathcal{A}$, it holds that

$$\Pr\left[\mathsf{Expt}^{\mathsf{F\text{-}Ind}}_{\mathcal{A}}(1^\kappa) = 1\right] \leq \tfrac{1}{2} + \mathsf{negl}(\kappa).$$

Building on previous works, Derler *et al.* systematize the landscape and introduce the *full CR* (Def. 8), shortened (F-CollRes), motivated by the observation that even S-CollRes can leave hashes malleable in practice [DSS20].

Definition 8 (Full CR [DSS20])**.** Consider the following experiment:

Experiment $\mathsf{Expt}^{\mathsf{F\text{-}CollRes}}_{\mathcal{A},\mathsf{CH}}(1^\lambda)$

1. $pp_{\mathsf{CH}} \leftarrow \mathsf{CHPG}(1^\lambda)$; $(sk_{\mathsf{CH}}, pk_{\mathsf{CH}}) \leftarrow \mathsf{CHKG}(pp_{\mathsf{CH}})$.
2. Initialize $\mathcal{Q} \leftarrow \varnothing$.
3. Give $\mathcal{A}$ the key pk_{CH} and oracle access to $\mathcal{O}_{\text{adapt}}$ defined by: on input (m, r, m', h), if $\mathsf{CHCheck}(pk_{\mathsf{CH}}, m, r, h) = 1$ then $r' \leftarrow \mathsf{CHAdapt}(sk_{\mathsf{CH}}, m, m', r, h)$, set $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(h, m); (h, m')\}$, and return r'; otherwise return $\perp$.
4. $\mathcal{A}$ outputs $(m^*, r^*, m'^*, r'^*, h^*)$.
5. Output 1 iff all hold: $m^* \neq m'^*$ and $(h^*, m^*) \notin \mathcal{Q}$ and
$$\mathsf{CHCheck}(pk_{\mathsf{CH}}, m^*, r^*, h^*) = 1 = \mathsf{CHCheck}(pk_{\mathsf{CH}}, m'^*, r'^*, h^*),$$
otherwise output 0.

A CHF CH is *fully CR* if for all PPT adversaries $\mathcal{A}$, it holds that

$$\Pr\left[\mathsf{Expt}^{\mathsf{F\text{-}CollRes}}_{\mathcal{A}}(1^\kappa) = 1\right] \leq \mathsf{negl}(\kappa).$$

In [DSS20], the authors formalize the family of collision experiments side by side and prove strict relations: S-CollRes implies W-CollRes, F-CollRes strictly implies S-CollRes, while E-CollRes and S-CollRes are incomparable. The separations are witnessed by concrete counterexamples, for instance the RSA based scheme of [CDH+17] satisfies S-CollRes but fails F-CollRes due to rerandomizability [DSS20]. They also revisit *indistinguishability* in a unified syntax, introducing *enhanced indistinguishability* (E-Ind) to express that the adversary also learns the key generation randomness, and establishing the strict hierarchy F-Ind $\Rightarrow$ E-Ind $\Rightarrow$ S-Ind $\Rightarrow$ Ind. They give compact left or right games for all variants. On the constructive side, they present a simple black box design that achieves F-CollRes using mcIND-CPA encryption and an SSE-NIZK, which simultaneously provides *normal indistinguishability*; this contrasts with earlier random oracle based E-CollRes templates that stack heavier primitives [DSS20, AMVA17]. Finally, they map the notions to applications, highlight where weaker properties can break redactable blockchains or limit signature upgrades, and argue that F-CollRes captures the intuitive requirement of no new collisions beyond what is already revealed.

Since 2020, new security properties have emerged, for instance the *revocation CR* [PVM21], *adaptive-identity based CR* [LXD+22], or even *forgery indistinguishability* [WCZ+24]. For more insights on the current landscape with respect to the security properties of CHFs, we refer the reader to Sub-section 5.1.

3 Chronological story

In this section, we present the complete and updated story of CHFs. We review the subsequent evolution of their design and the resulting diversity of construction techniques. Our goal is to provide the first unified and chronological view of the historical trajectory and the current landscape of CHFs.

Note: we tell the story of chameleon hash functions in a chronological manner, each paragraph represents a separate paper. For conciseness, we keep our description of each contribution modular and direct, which leads to some new notions being introduced without properly defining them (e.g. Σ protocol in [BR08]); thus, we encourage the reader to consult the individual papers if details are more of interest.

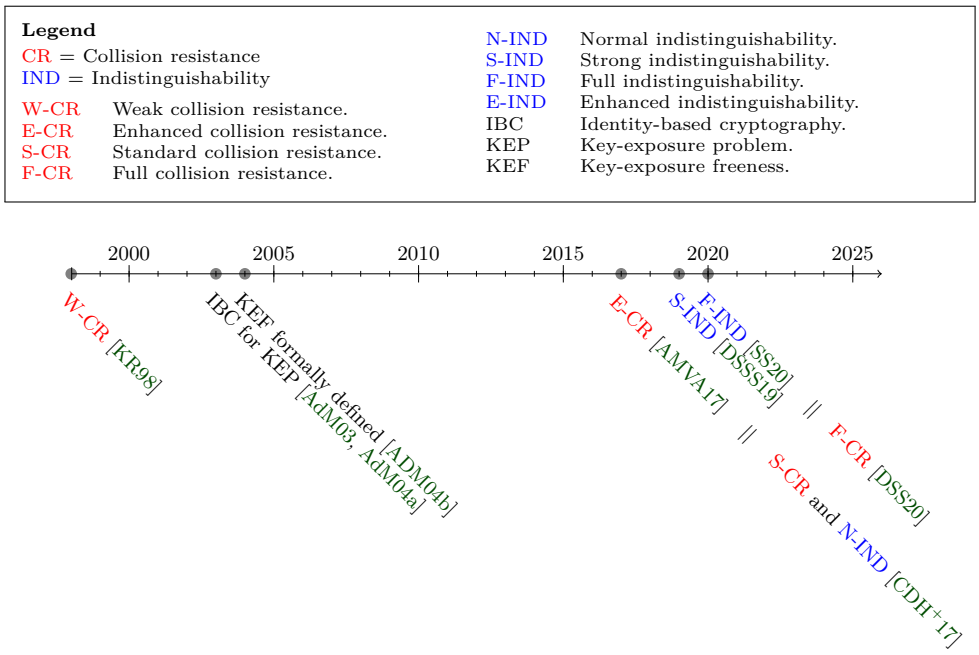

Figure 2: Security properties timeline

1998. Krawczyk and Rabin [KR98] introduce CHFs and formalize three core properties: *weak collision resistance*, *trapdoor collisions*, and *uniformity*. They give a black-box construction from claw-free trapdoor permutations, and two direct instantiations, a *factoring-based* scheme using $f_0(x) = x^2 \bmod n$ and $f_1(x) = 4x^2 \bmod n$ over the Jacobi $+1$ domain with hashing by repeated squaring and conditional multiplication by 4, and a *discrete-logarithm-based* scheme in a prime order group with public key $y = g^x$ and hash $h = g^m y^r$, with adaptation solving $m + xr = m' + xr' \bmod q$. *Weak collision resistance* reduces to *factoring* or *discrete logarithm* according to the instantiation, *trapdoor collisions* follow from the secret key, *uniformity* holds by permutation structure, and standard domain extension applies by first compressing the message with a Collision-Resistant Hash Function (CRHF). The native application is chameleon signatures via hash-then-sign with the recipient holding the trapdoor, which preserves unforgeability of the base scheme while yielding *non-transferability* for third parties.

2000. Cramer and Shoup exhibit an explicit trapdoor-hash pattern inside their strong-RSA signature scheme: from a signature (e, y, y') on m, set $x' \leftarrow (y')^{e'} h^{-H(m)}$ and obtain $y^e = x\, h^{H(x')}$, so $H(x')$ functions as a trapdoor hash tied to x' and simulatable by a party holding the signing trapdoor [CS00, Sec. 5]. They do not define new CHFs nor state security properties; we retain this work for the first explicit *RSA*-setting appearance of the trapdoor-hash mechanism that later papers isolate as a CH construction.

2001. Shamir and Tauman [ST01] formalize trapdoor hash families and give three concrete constructions: a *factoring*-based scheme with $f_0(x) = x^2 \bmod n$ and $f_1(x) = 4x^2 \bmod n$ over QR_n, a *discrete-logarithm* scheme $h(m; r) = g^m y^r$ in a prime-order subgroup, and a new *factoring*-based scheme with $HK = (n, g)$ and $h(m; r) = g^{m \circ r} \bmod n$ with $r \in \mathbb{Z}_{\lambda(n)}$. They prove *weak collision resistance* under *factoring* or discrete logarithms according to the instantiation, *trapdoor collisions* with randomness indistinguishability, and identify an inversion property for the first construction that supports lifting security from random-message to adaptive attacks. Native use is minimal: trapdoor hashing drives the hash–sign–switch paradigm to obtain efficient on-line/off-line signatures.

2003. Ateniese and de Medeiros [AdM03, AdM04a] carry the idea to identities, the trapdoor is issued per identity by a trusted extractor so key exposure is localized and provisioning is simple. The formulation targets *weak collision resistance* under RSA style unforgeability in the random oracle model, *trapdoor collisions* are per identity, message hiding is retained, and the intended use is chameleon signatures and a sealed bid sketch where one signed message suffices.

In the same year, Zhang *et al.* [ZSNS03] shift the identity based line to bilinear pairings, a direct realization in Gap Diffie Hellman (GDH) groups with proofs in the random oracle model. Security is framed as resistance to collision forgery under active attacks with semantic security, assumptions are the Computational Diffie Hellman (CDH) for the first scheme and a weak CDH variant for the second. The trapdoor is single holder per identity, and the use case is identity based chameleon signatures with compact parameters.

2004. Later, Chen *et al.* [CZK04] revisit the KEP and point out that the identity based variants that appeared the same year, Ateniese and de Medeiros, only localize exposure by issuing per identity trapdoors, they do not eliminate it. They then give a direct certificate based construction in Gap Diffie Hellman groups that is exposure free, a revealed collision does not leak the secret and does not enable further collisions, with *weak collision resistance* and semantic security proven under CDH. The trapdoor is single holder, and the intended use remains chameleon signatures with a simple denial step.

Ateniese and de Medeiros [ADM04b], finally put the KEP at the center, give a clean model with *weak CR*, message hiding, and key exposure freeness, and relate chameleon hashes to trapdoor commitments by isolating when stateful designs fail and when a double–trapdoor template is necessary. They show that the identity–based line only localizes exposure [AdM03] and that the later pairing-based proposal gives the first fully exposure–free instance [CZK04]; then they build exposure–free schemes under non–pairing assumptions and note that the RSA commitment used in the 2003 identity–based construction is already exposure–free in a PKI setting. Concretely, they give three exposure–free constructions, two under RSA and *factoring* (label–bound RSA trapdoor, plus a Paillier–based trapdoor commitment), and one under SDH in Gap groups, and also a single–trapdoor variant that drops perfect message hiding (generic–group style proof) [ADM04b, Secs. 4–5]. they conclude their work with an open question that steers later work: double–trapdoor yields identity–based or exposure–free, not both. Are there efficient schemes that are simultaneously identity–based and exposure–free, perhaps with more than two consecutive trapdoors? [ADM04b, Sec. 7].

Susilo and Mu [SM04] fold chameleon hashing into deniable ring authentication and, along the way, give three identity based chameleon hashes under *factoring* and RSA with a PKG extracting per identity trapdoors. Security follows the identity-based line, *weak collision resistance* against passive attacks with semantic security in the random oracle model, assumptions are RSA or *factoring* depending on the variant, the trapdoor is single holder per identity with easy re-randomization, and the proofs show reductions from CHF collisions to RSA signatures [SM04, §2.1–3]. They compare parameters and report shorter keys than Ateniese–de Medeiros for the same security level, which fixes the practical side of identity based provisioning in this branch [SM04, Table 1]. The application thread stays light, any ring signature can be turned *non-transferable* by hashing to the verifier's identity and letting the verifier simulate a collision if challenged.

2005. Schmidt-Samoa and Takagi [SST05] take the Paillier–Okamoto–Uchiyama path and give two factorization–based CH families over moduli $n = p^2q$ [SST05]. The first offers extremely fast weak altering and also strong altering, with *weak collision resistance* under the p^2q *factoring* assumption and *uniformity* inherited from the trapdoor map; the second adds an on–line/off–line profile. They note a practical caveat, some variants require the hasher to sample from $\mathbb{Z}_{pq}$ which would reveal the trapdoor, and resolve it by an alleviated form that samples from a public set and proves relaxed *uniformity* from Decisional Composite Residuosity Assumption (DCRA), preserving security for chameleon

signatures [SST05, Def. 8, Lem. 1]. Trapdoor structure is single–holder; the application thread is brief, these families support efficient on–line/off–line chameleon signatures.

Contini *et al.* [CLS05, CLS06] introduce Very Smooth Hash (VSH) and, for our line, define and instantiate a *randomized trapdoor hash* under the very–smooth–residue assumption with a factoring trapdoor. Security is collision resistance in the message input and *random trapdoor collisions* with *uniformity*, plus an *inversion* trapdoor property; proofs are in the standard model. The trapdoor is single holder; evaluation is fast while trapdoor collision finding is slower. The application note is brief, these hashes serve to tighten signature reductions and support designated–verifier primitives.

2007. Ren *et al.* [RMS07] give an *identity-based* CH *without key exposure*, avoiding the prior "identity customization" technique. The construction is pairing based in the standard bilinear setting, the trapdoor is single and long term per identity issued by a PKG (private key generator), and hashing includes a per-transaction event tag to isolate openings across different messages. They prove *weak collision resistance* from discrete logarithm in the target group and *semantic security*, and they show that revealing a collision does not leak the recipient's secret key, so keys remain usable after repudiation. They extend to a multi-party variant that binds a set of identities with the same guarantees, and they apply the scheme to deniable email authentication to mitigate phishing, including a mailing-list setting.

Gao *et al.* [GWX07] give a *factoring/RSA-based* CH that removes the key-exposure flaw of earlier *discrete-logarithm*-based designs (e.g., [KR98]): even after revealing many chosen collisions, the secret key remains hidden. The trapdoor is *single and long-term* (the RSA secret), there is no ephemeral component, and adaptation uses the secret exponent to re-open the same digest to a new message. They prove *weak collision resistance* from the RSA/*factoring* assumption and argue *key-exposure freeness* and *semantic security* directly in their model. Compared with the contemporaneous pairing/identity-based KE-free approach, this construction stays in the RSA setting while achieving the same "no key exposure" guarantee with a simpler assumption and standard algorithms.

Bresson *et al.* [BCG07] improve a previous online/offline signatures scheme [CMTW06] by *modifying* the chameleon hash to carry *two independent long-term trapdoors* (two discrete-log secrets), concretely $\mathsf{CH}(m, r, s) = g^m h_1^r h_2^s$. This double-trapdoor CHF lets the simulator program commitments without the one-more DL oracle, so the full threshold scheme reduces to the standard *discrete logarithm* assumption. They do not define CHF security notions; trapdoor structure is *dual, long-term, threshold-shared* among signers; no per-hash ephemeral trapdoor; adaptation is computed from shares in the online phase. The contribution is thus an efficiency and assumption upgrade for threshold online/offline signatures via a tweaked CHF component rather than a new CHF model.

Later that year, Chen *et al.* [CZSM07] crystallize the double–trapdoor pattern on elliptic curves, separating a long–term secret from a one–time key so revealed collisions do not leak the master secret. The security headline is *weak collision resistance* with *trapdoor collisions* under the *discrete-logarithm* assumption, proofs for the surrounding reduction use the random–oracle model, and the result is a compact, exposure–free CHF that fits cleanly with fast online phases.

2008. Chow [Cho08] introduces coins multi–CH and shows how to lift a key–exposure–free discrete–log CHF to a cooperative variant where collisions require all designated parties. Security is stated as collision resistance and key–exposure freeness in the Multi–Hash / Multi–Forge syntax, with proofs in the discrete–log setting. The neat move is a Schnorr–style aggregation: each verifier blinds with a one–shot exponent, an arbitrary combiner forms the new randomness and challenge, each verifier returns a share, and the combiner sums shares to obtain a valid collision, no single key is leaked [Cho08, §5.1–5.3]. This pins the point in the story where trapdoor power is distributed without secure multi–party computation and where group–held trapdoors enter the CHF line.

Bellare and Ristov [BR08, BR14] shift the narrative from case studies to a general theory, they show that the $\Sigma 2H$ transform turns any suitable Σ protocol into a keyed hash that is chameleon, and conversely that every CH arises from a Σ protocol. *Weak collision resistance* is proved via the special soundness of the underlying Σ protocol, *trapdoor collisions* follow from possessing a witness for the relation proved by that protocol, and *uniformity* comes from the honest–verifier zero–knowledge and challenge distribution. Proofs are stated in the standard model relative to the Σ protocol and the reduction is tight. The trapdoor structure is single holder and equals the witness, for Schnorr and Okamoto it is a discrete logarithm (or pair), for Guillou–Quisquater it is an RSA root, for Strong Fiat–Shamir it relies on *factoring*, for the modified Micali–Shamir it relies on the square–roots–of–prime–products assumption, and for the Hess pairing identification it is a discrete log in the source group.

Chen *et al.* [CZT+08] introduce a double-trapdoor CH to solve the KEP in Shamir–Tauman online/offline signatures. The trapdoor structure has a single long-term secret x and a per-message one-time secret k; both are required to adapt, and reusing the one-time key would enable universal forgeries, hence the enforced "one-time" design. The construction is concrete over elliptic-curve discrete-logarithm groups, proved *key-exposure-free* with *weak collision resistance* under the *discrete-logarithm* assumption in the random-oracle model. As an application, they give an efficient generic online/offline signature scheme whose online cost is essentially one modular multiplication, and, further, the first *threshold* online/offline scheme without a trusted dealer by distributing the long-term trapdoor via DKG and combining $t+1$ collision shares in the exponent.

2009. Hohenberger and Waters [HW09] needed a plain-RSA CH for their hash-and-sign and, to this end, they adopt and simplify Ateniese–de Medeiros' identity-based CHF [AdM03] by removing the identity layer. The result is an RSA-based CHF in the standard model with a *single long-term* trapdoor (factorization and the RSA inverse), standard *weak collision resistance* under *RSA*; they present it explicitly as a more efficient drop-in for their RSA line.

Chen *et al.* [CZW+09, CZW+11] identify a concrete gap and then fill it: prior discrete–log key–exposure–free CHFs existed only in GDH groups with pairings, so they ask whether efficient constructions without GDH are possible and give two answers. The first is in GDH groups under CDH with explicit exposure freeness; the second is in plain *discrete-logarithm* groups, dispensing with the GDH oracle by certifying DH tuples via proofs of equality of discrete logs (proofs of knowledge). Security provided by *weak collision resistance*, *semantic security*, and explicit *key exposure freeness* in the ROM; the trapdoor is single holder, with labels binding collisions through a full–domain hash.

2010. Chen *et al.* [CTZD10] reexamine the 2007 *factoring* line and pinpoint three flaws in Gao–Wang–Xie's key–exposure–free CHF [GWX07], collisions were defined only up to a hidden sign, the exposure–free proof relied on an unrealistically strong signing oracle, and the collision method yielded essentially one collision where many exist. They fix all three and give an improved *factoring*–based CHF: make the sign bit explicit so collisions are standard equality, tighten exposure–freeness to EUF–CMA of a Rabin–style signature in the ROM, and show how to derive a family of collisions. Security remains *weak collision resistance* with *semantic security* and explicit *key exposure–freeness*; the trapdoor is single holder, and labels bind via a hash.

Harn *et al.* [HHL10] define multiple collision trapdoor hash families so a fixed CH value can be safely reused across many messages, this is essentially a double trapdoor CHF. They give direct discrete logarithm and *factoring* instantiations that add a one time secret to enable many controlled collisions without leaking the long term trapdoor; security is *weak collision resistance* with uniform *trapdoor collisions* in the random oracle model under the *discrete-logarithm* or the *factoring* hardness assumption, the trapdoor is single holder and adaptations re randomize naturally.

Chen *et al.* [CZS+10, CZW+14] resolve the open problem raised by Ateniese and de Medeiros [AdM03, AdM04a]: they give the first *identity–based* CH that is *key–exposure free*. The construction sits in GDH groups with pairings, uses a master secret x, a per–identity secret $SID = xH(ID)$, and an ephemeral trapdoor bound to a transaction label L (per-identity trapdoor); revealing a collision leaks only the ephemeral value and does not enable ephemerals for other labels. Security is proved in the random oracle model under the Bilinear Diffie-Hellman (BDH) assumption, *weak collision resistance* against active attackers with *semantic security* and explicit *key exposure–freeness*. Native use stays light, identity–based chameleon signatures with the verifier's identity as label and standard denial when collisions are shown.

Chandrasekhar *et al.* [CLLS10] revisit trapdoor hashing for settings with multiple trapdoors, formalize *collision-forgery* and *ephemeral-collision-forgery resistance*, and give a precise *key-exposure* notion in that model. They propose a discrete-logarithm based, exposure-free CHF (DL-MTH) that binds the hash to the public key and computes collisions via an ephemeral key pair, yielding a double-trapdoor structure (long-term secret and per-use ephemeral) with *weak collision resistance*, *semantic security*, and explicit *key exposure-freeness* in the random oracle model under the *discrete-logarithm* assumption. Application use is a clean proxy-signature framework where a hashed warrant is fixed by the delegator and the proxy releases only an ephemeral hash key with each signature.

2012. Cash *et al.* [CHKP12] give the first lattice–based (post–quantum) CH: over $\mathbb{Z}_q$, fix $A_0 \in \mathbb{Z}_q^{n\times k}$ and $A_1 \in \mathbb{Z}_q^{n\times m}$ and define $h_A(m;r) = A_0 m + A_1 r$ with r drawn from a discrete Gaussian; the trapdoor is a short basis S for $\Lambda^{\perp}(A_1)$ enabling, for any target y and message m, efficient sampling of r conditioned on $h_A(m;r) = y$. Security is in the weak sense under SIS with standard–model proofs (collision resistance via SIS; *trapdoor collisions* via Gaussian preimage sampling; *uniformity* from the sampling distribution); trapdoor structure is single holder (the short basis). Native use is light but pivotal: this CHF plugs into their stateless hash–and–sign in the standard model and underlies the bonsai–tree framework.

2013. Pan *et al.* [PLQZ13] build chameleon hashes from *inner automorphism groups*, the first CHF on noncommutative groups The key is $E = \mathrm{Inn}(g^x)$; hashing maps (m, r) to a universal hash of $\mathrm{Inn}(g^{m+xr})$; the trapdoor x yields $r' = x^{-1}(m - m') + r \bmod p$. Security is in the weak sense under the discrete logarithm problem over $\langle \mathrm{Inn}(g) \rangle$; chameleon property and (computational) *uniformity* are proved; the trapdoor is single holder. Native use is one-time signatures instantiated in the same setting, with parameter-size and runtime comparisons against *factoring*-based variants.

Lai *et al.* [LDW13] introduce **Accountable CH (ACH)**, a two–party variant between plain CH and **Identity-based CH (IBCH)**, and use it to realize accountable trapdoor sanitizable signatures. Accountability means that any revealed collision carries verifiable evidence of *who* enabled it: per–transaction trapdoors $td_{i,TID}$ are issued by one party to the other, so if a collision on TID appears, the issuer can cryptographically link it to the holder or clear themselves. They give a compact pairing–based construction in the ROM under CDH with *weak collision resistance against active attacks* and a *forgery–indistinguishability* property; trapdoors are single–holder per user and per TID, and no long–term key leaks. The construction then plugs into trapdoor sanitization to restore traceability without losing the delegation flexibility.

2014. Ducas and Micciancio [DM14] touch the chameleon-hash line only as a vehicle: in Appendix B they recall the standard lift from static to adaptive security via CHFs and give a lattice instantiation within their ring-SIS framework. They define CHFs with *uniformity*, preimage-sampling correctness (with key-independent output distribution), and *collision resistance*, then instantiate these via the algorithms GenTrap/SampleD and prove *collision resistance* under *RingSIS*, all in the standard model. The role is connective tissue: it reaffirms that lattice CHFs plug cleanly into hash-and-sign without random oracles and

serve as a stock upgrade from non-adaptive to adaptive security, as already leveraged in Cash *et al.* work [CHKP12].

Chandrasekhar and Singhal [CS14] formalize multi trapdoor hashing, a CH interface where several parties, each with its own trapdoor, can generate shares that combine into one collision. The construction adapts the earlier *discrete-logarithm*-based multiple collision idea rather than introducing a new *discrete-logarithm*-based core, and scales to a multi party variant (DL-μTH) with hash size and cost independent of the number of participants. Security is given in the weak sense, with an explicit ephemeral collision notion and key exposure resistance in the ROM under discrete logarithms; trapdoors are long term per participant plus per use ephemeral. Applications are sketched for aggregate style and sanitizable settings.

Wang and Tanaka [WT14] introduce the **Leakage–Resilient Chameleon Hash Function (LR–CHF)**, meaning a CH that remains *collision-resistant* even if an attacker learns bounded partial information about the trapdoor; the trapdoor still enables collisions with uniform output, but any leakage up to a stated bound does not help to forge a collision. They build LR–CHF from a leakage–resilient hard relation, concretely an Okamoto–style SPR relation under the n–representation assumption in a *discrete-logarithm* group, and prove in the bounded–leakage model that collision resistance and reversibility hold with a single–holder trapdoor. As a brief use, composing LR–CHF with the GBSW transform yields a generic lift from weakly to strongly unforgeable signatures while preserving leakage tolerance.

Boneh and Corrigan–Gibbs [BCG14] derive a chameleon–hash candidate from low–degree bivariate polynomials modulo an RSA composite (e.g., $H(m,r) = m^7 + 3r^7 \bmod N$), pointing to the hardness of finding collisions but giving no formal CHF definition or proof of collision resistance or *trapdoor collisions*. The trapdoor is implicitly the factorization of N, yet the collision–finding procedure and security theorems are not developed.

2015. Wang and Tanaka [WT15] extend leakage resilience to the continual model and define a CH with an *updatable* single–holder trapdoor, *weak collision resistance* holds despite unbounded total leakage as long as each round leaks at most a bounded amount. Besides *weak collision resistance*, they prove *trapdoor collisions* with *uniformity*, *reversibility*, and *update correctness/indistinguishability*; the instantiation is under SXDH with Groth–Sahai proofs in the standard bounded–leakage model. As a brief use, this CHF yields a generic lift from weakly to strongly unforgeable signatures under continual leakage via a refined GBSW transform.

Liu *et al.* [LZZ+16] introduces a variant of *Weak collision resistance*, they call it *Universally collision resistant CH* (UCH), where collision resistance must hold even when the adversary chooses the hash key (public key). They give a generic UCH from any CHF and a standard hash, setting $\mathsf{UCHash}(chk, m, r) = \mathsf{Hash}(\mathsf{CHash}(chk, m, r) \,\|\, chk)$; *trapdoor collisions* and (random) output distribution carry over, while "universal" CR follows from the base CHF or Hash colliding. This UCH is then the pivot for their CCA2 online/offline PIPE transform, preventing key-swapping malleation during ciphertext validation.

2017. Ateniese *et al.* [AMVA17] introduce the public–coin versus secret–coin distinction for chameleon hashes and define *enhanced collision resistance*, hardness that persists even after many oracle–generated collisions. They give a generic lift from any public–coin CHF to a secret–coin CHF with enhanced CR using a CPA–secure PKE and a true–simulation–extractable NIZK, and they instantiate it both in the ROM from DDH via ElGamal with Fiat–Shamir, and in the standard model from pairings using Groth–Sahai–style tSE NIZKs under SXDH/DLIN–type assumptions. The trapdoor is a single long-term key composed of 2 parts (tk, sk), where tk is the underlying public-coin CHF trapdoor and sk is the PKE decryption key used to recover the encrypted randomness during adaptations. On use, they show centralized and Multi-Party Computation (MPC)

governance and a Bitcoin sketch where the inter–block link is a CHF value (instead of an usual collision resistant hash function) that supports auditable redactions.

Xie *et al.* [XZWZ17] define a **Homomorphic CH (HCHF)**: from hashes of inputs (m_i, r_i) one can publicly compute a new hash that equals the CH of $f(m_1, \ldots, m_k)$ under some derived randomness, for any polynomial–depth circuit f. The lattice instantiation is under SIS in the standard model with a single short–basis trapdoor; security includes *weak collision resistance*, *trapdoor collisions* via Gaussian preimage sampling, and statistical *uniformity* of outputs across homomorphic evaluation. This HCHF then drives leveled fully homomorphic signatures whose evaluated signature size grows polynomially with circuit depth.

Chandrasekhar *et al.* [CIS17] refine the double–trapdoor CHF work in [CIS17] by proposing *DL–mDTH*, a *discrete–logarithm* variant of DL–DTH that *removes the zero–knowledge proof of knowledge for the ephemeral trapdoor* while preserving *weak collision resistance*, *key–exposure freeness*, and *semantic security*. The trapdoor structure remains long–term plus per–use ephemeral, proofs are in the random oracle model under the *discrete-logarithm* assumption, and the construction is used inside their proxy–signature–driven access–control protocol for cloud HIE.

Camenisch *et al.* [CDH+17] introduce **Chameleon Hashes with Ephemeral Trapdoors (CHET)**: each hash comes with per–hash trapdoor data chosen by the hasher, and even the long–term trapdoor holder cannot adapt a hash without that ephemeral value. They give a full model with *normal indistinguishability* of randomness (hash vs. adapt), *public collision resistance* against outsiders and *private collision resistance* even against the long–term key holder lacking the ephemeral, and discuss *uniqueness*; a CHET is secure if it is *correct*, *indistinguishable*, *publicly CR*, and *privately CR* [CDH+17, Defs. 10–14]. Constructions include a black–box "bootstrapped" CHET from any secure *unique* CHF (Theorem 1), two direct *discrete-logarithm*–based schemes (one in known–order groups using NIZK proofs and IND–CCA2 encryption for state–free recovery of randomness, one in pairing/gap groups), and an RSA–like scheme in the random–oracle model that also achieves uniqueness; assumptions and proof hooks are spelled out (DL, IND–CCA2, CR of the base hash, NIZK extractability; plus ROM and one–more RSA where needed) [CDH+17, §4.1–4.2; App. E]. As a light touch on applications, CHET resolves *invisible* sanitizable signatures with a full framework and proof that hide which blocks are admissible from outsiders [CDH+17, §5].

2018. Zheng *et al.* [ZLL+18] give a lattice CH tailored for hash–sign–switch: over $\mathbb{Z}_q$ they set $\mathsf{CH}_{(A',B')}(m, r) = B'm + A'r \bmod q$ with a single short–basis trapdoor (a G–trapdoor for A') enabling adaptations via Gaussian preimage sampling. Security is proved in the standard model under the *ISIS* assumption: *weak collision resistance* from ISIS, *trapdoor collisions*, and (computational) *uniformity* of the adapted randomness; the trapdoor is single–holder. The construction then drives their lattice online/offline signature, using programmable hash functions to keep the offline work heavy and the online step light.

Krenn *et al.* [KMS18] generalize CHET to **Chameleon Hashes with Dual Long–Term Trapdoors (CHDLTT)**: besides a regular trapdoor sk_{ret}, a reusable pair $(\mathsf{ptd}, \mathsf{std})$ enables adaptations for all hashes tied to ptd. Security is recast for the two–key setting with *normal indistinguishability*, *public collision resistance* (hard even after Adapt queries) and *private collision resistance* (hard for the holder of sk_{ret} without std), and optional *uniqueness*. Constructions: a black–box lift from any CHF by hashing both $(pk_{\mathsf{ret}}, \mathsf{ptd}, m)$ in two independent instances, and a direct scheme in hidden–order groups proven in the ROM under the *one–more RSA* assumption. Conceptually, CHDLTT removes CHET's per–hash ephemeral constraint and restores "global" collision resistance over the full message space; briefly, this supports efficient multi–party sanitization patterns (e.g., three–party sanitizable signatures).

2019. Blazy *et al.* [BGM+19] give the first *code–based* CH, built from the KKS signature assumption together with coding–theory decision problems (DSD and its extended form).

Security is proved in the standard model: *weak collision resistance*, *trapdoor collisions* via the KKS signing trapdoor, and *uniformity* in a computational–indistinguishability sense; the trapdoor is single holder and inherited from KKS. As a light touch on use, they derive a chameleon–hash–based two–tier signature and a binary–tree lift, yielding a code–based signature in the standard model; they also note current limitations, large sizes and the weaker *uniformity* notion.

Derler *et al.* [DSSS19] introduce *policy based chameleon hashes* (PCH) so a hash carries an access policy and only holders of attributes that satisfy the policy can find collisions; this moves redactability from block level to transaction level and makes rewriting fine grained. They strengthen the foundation with *strong indistinguishability* (even if the adversary learns the secret key, a valid opening must not reveal whether it came from hashing or from adaptation) and give two security axes for collisions, *outsider* collision resistance that remains hard even after Adapt oracle queries, and *insider* collision resistance that enforces the policy against key holders. The trapdoor structure is a single long–term CHET key together with a per–hash *ephemeral* trapdoor etd; hashing outputs (h, r, etd) and publishes an ABE ciphertext of etd under the policy, while Adapt requires both the long–term key and the recovered etd—policy gating without threshold sharing [CDH+17]. The construction is generic, combine a CHET with a CCA–secure CP–ABE, and they also give a more efficient CHET than the earlier template; an efficient instantiation uses the RSA–based CHF from [CDH+17] for CHET and FAME CP–ABE made CCA2 via a Fujisaki–Okamoto style transform, with a concrete implementation and performance figures. The application snapshot is direct, transaction–level rewriting under attribute policies and more expressive delegations for sanitizable signatures.

Shu *et al.* [SHL+19] take the ECC branch and package a double–trapdoor CH (long–term plus ephemeral) under *elliptic curve discrete-logarithm* assumption, then lift it to a *multi–trapdoor* variant where several trapdoors can jointly yield collisions and, crucially, enable batch collision computation. Security is given in the random–oracle model with *weak collision resistance*, *trapdoor collisions*, and explicit *key–exposure freeness* for the long–term secret; the structure remains the familiar double–trapdoor pattern on elliptic curves. As a brief note on use, these CHFs back an aggregate signature whose length is constant in the number of signers.

Huang *et al.* [HZM+19a] push redactability in blockchains toward automation by introducing a *revocable CH* (RCH) and using it to build a *self–redactable blockchain*. The trapdoor picture is long–term x plus a *published, time–bound* ephemeral trapdoor etd tied to a label CID and epoch t; anyone holding etd can adapt within that epoch, and both the committed hash and etd *expire* at $t+\Delta t$, with an $\mathsf{Un\text{-}Revoke}$ algorithm to re–activate a hash using fresh ephemerals. Security is formalized with (i) *normal indistinguishability*, and two collision notions aligned to revocation: (ii) collision resistance between *nonrevoked* hashes, and (iii) collision resistance between a *revoked* and a nonrevoked hash; proofs are in the ROM under DDH/CDH in pairing groups, and the ephemeral trapdoor's hardness reduces to the *discrete-logarithm* assumption. The construction keeps a single–holder long–term trapdoor, pushes adaptation power into per–epoch ephemerals to limit abuse, and then plugs into a BLS–based chameleon–signature wrapper to realize transaction–level, policy–driven self–redaction in IoT–oriented ledgers.

Lu *et al.* [LAZ19] introduce *Chameleon Hash Plus (CH+)* as a CHF–like primitive tailored to an RST–style ring–signature framework: given a hash key and trapdoor, one can, for any target message m' and only the hash value C, compute r' with $\mathrm{Hash}(hk, m', r') = C$ (no need for the original (m, r)). They formalize properties one wayness, collision resistance (weak), near–uniform randomness after adaptation, and (near) uniform key distribution—and give two lattice instantiations: (i) a standard–model CH+ from SIS/ISIS via GPV sampling (trapdoor T as a short basis), and (ii) an efficient NTRU/Falcon–based CH+ where inversion is Falcon preimage sampling. The trapdoor is single holder; outputs

are statistically close to the target sampling distribution. As a brief pointer to use, CH+ is the core building block enabling their practical (linkable) lattice ring signatures.

Huang *et al.* [HZM+19b] introduce what they call "first *threshold CH* (TCH)", pushing collision control from a single holder to k cooperating parties. In GDH groups, each party holds x_i and the public hash key is $hk = g^{\prod_i x_i}$; hashing uses H_0, H_1 and randomness $r = (g^\alpha, hk^\alpha)$, while Forge composes per–party exponents in a ring so only the full set can adapt (m, r) to m'. Security notions are stated explicitly for TCH—*semantic security*, *weak collision resistance*, and *key–exposure freeness* with reductions to CDH in a GDH group (ROM). The trapdoor structure is long–term, *threshold–held*, and adaptations remain label–bound via a customized identity. On use, TCH underpins an accountable–and–sanitizable chameleon signature and a redactable consortium blockchain where authorized sensors jointly compute collisions and update signatures without hard forks.

Samelin and Slamanig [SS19, SS20] do not give a new CH; they tighten CH/CHET and policy-based CH by allowing adversary-generated keys and by formalizing *full indistinguishability* via a left–right Hash/Adapt oracle, then show that the RSA-based CH of Camenisch *et al.* [CDH+17] achieves full indistinguishability in the ROM, lift this to CHET, and re-prove the NDSS'19 [DSSS19] policy-based CH under the stronger model (full indistinguishability, insider/public collision-resistance, and uniqueness). As an application, they instantiate these guaranties inside policy-based sanitizable signatures with public verifiability and accountability; the trapdoor remains single-holder long-term with a per-hash ephemeral component (CHET).

2020. Lin *et al.* [LLLS20] give a *linearly (additively) homomorphic CH* over lattices: for public matrices $C = (C_0, C_1)$ over $\mathbb{Z}_q$, define $h_C(v; r) = C_0 v + C_1 r \bmod q$ with $v \in \mathbb{Z}_p^\ell$ and $r \in \mathbb{Z}_q^m$, so $h_C(v; r) + h_C(v'; r') = h_C(v+v';\ r+r')$. Security includes *standard collision resistance*, *trapdoor collisions* via Gaussian preimage sampling (GPV/SamplePre) using a single short basis T_{C_1} for $\Lambda_q^\perp(C_1)$, and computational *uniformity* of adapted randomness, under SIS/ISIS in the standard model. The trapdoor is *single-holder*; parameters pick $p < q$ to control norm growth under additions. This CHF then upgrades their lattice LHS to EUF–CMA while keeping public matrices constant-size.

Khalili *et al.* [KDS20] take *enhanced collision resistance* (ECR) from the 2017 redactable-blockchain [AMVA17] line and make it practical: they pinpoint three issues in the instantiation (an inefficient $\phi : \mathbb{Z}_p \to G$ map, an incomplete witness-link, and heavy sizes) and then give *two* efficient secret-coin CHFs with ECR . First, they design a public–coin CHF in pairings under DCDH (Divisible Computational Diffie–Hellman), then a lift to *secret–coin + ECR* using Cramer–Shoup (CCA2) and Groth–Sahai NIZKs; a shorter variant swaps GS+CS for a se–SNARK (Groth–Maller). Security: standard collision resistance for the public–coin core, ECR after the lift. *Trapdoor structure:* single–holder; for the core CHF the trapdoor is a single exponent $x \in \mathbb{Z}_p$ (discrete log witness), and in the ECR lift the trapdoor becomes $(x, \mathsf{sk}_{\mathrm{CS}})$. Efficiency improves over [AMVA17] in key size and inversion cost, matching redactable–ledger needs.

Zhang *et al.* [ZLL+20] introduce what they call "*threshold trapdoor CH* (TTCH)" and use it to power Re-chain, a redactable blockchain with fixed storage. TTCH means collisions can be produced only by a threshold t of n authorized parties: a master secret x yields a shared trapdoor as Shamir shares of $d = x^{-1} \pmod p$; each party returns a partial credential and Lagrange interpolation derives the new randomness R' that validates the adapted hash. Security is given for *weak collision resistance* and a *strong "unforgeability"* notion (stronger than *key-exposure freeness*) in bilinear groups under DCDH; the trapdoor structure is *threshold-held* (no single holder), with no ephemeral component.

Tian *et al.* [TZZ+20] extend policy–based chameleon hashing to *policy–based CH with black–box accountability* (PCHBA) for permissioned ledgers: the CHET backbone remains, but each hash embeds an *ephemeral* trapdoor R encrypted under an attribute–based system so only policy–authorized modifiers can recover it, and a homomorphic–signature binding

ensures all adaptations of a transaction tie to the same R for public linking. The trapdoor structure is *single–holder long–term* CH key per modifier (distributed by an attribute authority) plus per–transaction *ephemeral* R. Security lifts CHET's *normal indistinguishability* and *P-collision resistance* (the set Q_3 is being fed tuples (h, m, Λ) with Λ being a policy, this format look very similar to *full collision resistance* introduced in a later work) and introduces *accountability*: given a dispute, an authority can cryptographically identify the responsible modifier, while any user can extract an accused set via ABET's black–box traceability. The instantiation combines CHET with an anonymous CP–ABE/HIBE hybrid (FAME–style under DLIN) and new asymmetric–pairing assumptions DLE and eDDH; Schnorr–style homomorphic signatures provide the binding layer, yielding fine–grained, accountable rewriting at the transaction level .

Derler *et al.* [DSS20] tighten the target for chameleon hashes by *introducing full collision-resistance (F-CollRes)* (the adversary is not allowed to use tuples (h, m) that they have seen collisions for) and *Enhanced Indistinguishability* (The adversary gets the randomness involved in the generation of the secret key). On the constructive side, they give a clean *black-box* CHF with a single long-term trapdoor built from a perfectly correct, key-verifiable mcIND-CPA PKE together with an SSE–NIZK; informally, the digest commits to m via the PKE and the opening is a proof of the disjunction "either this digest encodes m under the CHF key or I know the matching secret key," which is what enables adaptation by switching witnesses. This compiler *achieves F-CollRes and normal indistinguishability (N-Ind).* Concretely, they instantiate the PKE with *ElGamal* in prime-order groups under the *DDH* assumption, and they obtain the needed *simulation-sound extractable* NIZKs in the ROM from standard Σ-protocols via known compilers; key-verifiability is a lightweight check tied to the same *discrete-logarithm* setting. A perfectly binding commitment would already suffice to reach F-CollRes, but choosing an mcIND-CPA encryption as the "commitment" additionally yields indistinguishability so fresh and adapted hashes are distributionally aligned. In addition, they also give a minimal CHF that attains only *E-CollRes* while still meeting *N-Ind.* Finally, they sketch the application angle: stronger CHF binding directly stabilizes redactable/rewritable blockchains and strengthens online/offline signatures by allowing safe reuse of offline work without prohibitive overheads.

Derler *et al.* [DSSS20] remove the PKE layer from [DSS20] by using a non-interactive commitment together with an SSE–NIZK over a disjunctive relation tied to a one-way function, yielding a black-box CHF with a single long-term trapdoor that *achieves full collision resistance (F-CollRes)* and *strong indistinguishability (S-Ind).* They also *introduce randomness unforgeability*, a weaker alternative to uniqueness forbidding outsiders from producing new valid randomness for hashes they did not create. Concretely, they instantiate over *discrete-logarithm* settings and, separately, from LPN using Stern-type Fiat–Shamir proofs with SSE and suitable commitments. The LPN path gives, to my knowledge within the paper, the first *post-quantum* CHF achieving *F-CollRes*, improving over prior post-quantum directions that effectively matched weaker collision levels (e.g. Cash *et al.* [CHKP12]). The trapdoor is single and long-term in both lines, and the stated F-CollRes and S-Ind guarantees are proved for each instantiation.

2021. Xie *et al.* [XSL+21] give the first *identity-based* CH provable in the *standard model*, redefining collision resistance to an *adaptive-identity, basic* notion (ID-B-CollRes) (no collision finding oracle is given, still in the older model since it is an identity-based CHF) and pairing it with a *semantic security* requirement for hash outputs. The construction is pairing-based without random oracles, uses Waters-style key material, and proves ID-B-CollRes under *CDH* via a reduction to Waters EU-CMA security; the setting is public-coin and the trapdoor is single and long-term per identity. They also give a direct application to *online/offline* identity-based signatures by plugging the IB-CHF into the standard transform, with a fixed-cost online phase and experimental evidence that the overhead remains small on mobile-class hardware.

Thanalakshmi *et al.* [TPAS21] propose a hash-based (graph/DAG) CH together with a chameleon signature composition aimed at post-quantum security. The CHF is defined over DAGs using a homomorphic hash H and a homomorphic pseudorandom generator PG; the public key has a permanent part and an *ephemeral* part tied to a cut, and the trapdoor is a *double* trapdoor (sk, r) that enables adaptations without leaking long-term secrets. They formalize three CHF properties in this setting, *weak collision resistance* via a reduction to preimage resistance of H, *semantic security*, and *key-exposure freeness* (no key disclosure even with oracle access to the collision computation); proofs are given by direct reductions in their model. Instantiation uses only standard hash-function assumptions and PRG security, hence their claim of quantum resistance, with no number-theoretic or pairing assumptions. On top of the CHF, they build a chameleon signature by combining the verifier's graph for hashing and a graph-based one-time signature for signing, and they argue *unforgeability*, *non-transferability*, *non-repudiation*, and a denial protocol, plus a basic cost analysis in terms of DAG width and evaluations.

Zhang *et al.* [ZZLD21] design a *multi-authority policy-based CH* (MAPCH) that combines chameleon hashes with *ephemeral* trapdoors and multi-authority CP-ABE so that only users whose attribute keys satisfy a policy can adapt a given digest; the trapdoor structure is long-term *tk* plus per-hash ephemeral *etd* encrypted under the access policy, which also removes the single-authority bottleneck present in prior PCH designs. They formalize three guarantees for MAPCH—indistinguishability of hash versus adapted openings, outsider collision resistance, and insider collision resistance—is proven by reduction to CHET's *private-collision resistance* together with MA-ABE IND-CCA2, while *normal indistinguishability* and *outsider-CR* are defined and argued in the same framework (the paper calls the latter "easy to prove"). Concretely, their proof-of-concept instantiates MAPCH with the RSA-based CHET of [CDH+17] and the RW15 multi-authority ABE, implements the full transaction flow in Hyperledger Fabric with minimal structural changes, and reports about 120 tps throughput and ≈ 8 ms latency at 20,000 concurrent redactable transactions.

Xu [XZL+21] introduce *revocable policy-based CH* (RPCH), formalize security for it, *insider collision resistance*, and *uniqueness*), and realize a concrete RPCH by replacing the FAME layer in PCH with a new revocable ABE that is *adaptively secure* in prime-order pairings under *DLIN* in the random oracle model and resistant to decryption key exposure. Revocation uses the tree-based KUNodes method with public key updates per period; a modifier holds an attribute-based long-term secret and combines it with the period update to derive a decryption key. Hashing encrypts the *CHET* ephemeral trapdoor under the access policy and timestamp, so adaptation needs both the attribute-derived decryption key and the CHET long-term trapdoor, which enables revocation of rewriting privileges. They integrate RPCH into mutable blockchains and sanitizable signatures and implement on MNT224 pairings.

Xu *et al.* introduce *revocable policy-based CH* (RPCH), i.e., PCH augmented with revocation so that only policy-satisfying, *non-revoked* holders in a given period can adapt a hash [XZL+21]. They formalize security for *RPCH—full indistinguishability*, *insider collision resistance*, and *uniqueness* and realize it by replacing FAME in PCH with an *adaptively secure* revocable ABE (RABE) under DLIN on prime-order pairings (tree-based KUNodes for public key updates). The trapdoor structure is *double*: a single long-term trapdoor together with a per-hash *ephemeral* trapdoor encrypted under the access policy and timestamp; adaptation thus needs both the period decryption key (from RABE) and the CHET long-term trapdoor, and revocation simply withholds the update. They integrate RPCH into mutable blockchains and sanitizable signatures and implement on MNT224 pairings with modest overhead.

Panwar *et al.* [PVM21] define and build a **Revocable Chameleon Hash with Ephemeral Trapdoor (RCHET)**: a single long-term trapdoor together with a per-hash

ephemeral trapdoor, both required to adapt; revocation rotates the ephemeral part so past holders lose rewriting power. They formalize and strengthen the CHF guarantees with *outsider* and *insider (private) collision resistance* and add *revocation collision resistance* (it states that an adversary is unable to find a collision even if they have access to both the long-term trapdoor and the ephemeral one after the later one has been updated), giving the adversary message- and trapdoor-adaptation oracles, and they include *normal indistinguishability* in the same model. RCHET is instantiated in prime-order groups under *DL* and *DDH* with NIZK proofs of knowledge and an *IND-CCA2* PKE used to protect randomness; security relies on *public/private collision resistance* and they introduce *revocation collision resistance*, and standard NIZK properties. For policy control they introduce *RFAME*, a revocable CP-ABE built under *DLIN* on Type-III pairings, and combine it with dynamic group signatures to obtain anonymous yet traceable rewrites.

Wu [WKD21] give two single trapdoor, key exposure free lattice chameleon hashes for redactable blockchains. In the first scheme (without a lattice basis trapdoor), the long term trapdoor is a matrix X with $AX = qI \bmod 2q$; computing adaptations follows a twin signature style and enhanced collision resistance reduces to *ISIS*, with a distributed key generation and MPC sum to share X and compute collisions jointly. In the second scheme (with a lattice basis), the trapdoor is a GPV short basis S in a Pedersen like form $h = Ar + Tc$; *enhanced collision resistance* reduces to *SIS* in the random oracle model and collisions do not leak S. Both schemes also prove *semantic security* of hash outputs, avoid PKE and NIZK, and target post quantum security. The trapdoor structure is single; control of redaction can be by threshold sharing or by voting.

Jia *et al.* [JLL+21] propose a redactable blockchain with a semi-trusted regulator and introduce *stateful CH with revocable subkey* (sCHRS). The trapdoor structure has a *master* key (regulator) and per-user *subordinate* keys; both can adapt, and a per-hash *state* binds which subordinate key is authorized so the master can revoke that key's power for that specific hash. They define security for sCHRS (correctness, revocability, *secret-key collision resistance* for original and new sub-keys, and master-key revocation-resistance) and adopt an what they call it *enhanced collision resistance* (but in fact it is *standard collision resistance*) baseline for the underlying CHF. Construction is black-box from two standard CHFs by hashing the inner digest together with the state under the outer CHF; instantiation uses *RSA-based* chameleon hashes. On-chain, "witness transactions" record modifications and state, yielding supervision and user self-management at transaction-field granularity with small overheads.

2022. Li *et al.* [LXD+22] give an efficient identity-based CH (IB-CH) in the *standard model*, targeting mobile settings and online/offline IBS. The trapdoor is single-holder per identity (no ephemeral component), and security follows their framework: *adaptive-identity basic collision resistance* (ID-B-CollRes) via a reduction to the unforgeability of the underlying Gentry signature under q-SDH, and *semantic security*; no treatment of key-exposure freeness. They report substantial efficiency gains over the only prior standard-model IB-CH, including markedly smaller public parameters and a 50% reduction in the online phase when plugged into the generic OO-IBS transform.

Xu *et al.* [XHY+22] introduce a new design of *policy-based CH with blackbox accountability* (PCHA) to add linkability and traitor tracing to PCH while keeping CHET's fine-grained control. The generic design composes CHET, ABTT, and IBS: the transaction owner hashes with CHET and ABE-seals the *ephemeral* trapdoor via ABTT under an access policy, and the modifier's act is publicly linked through an IBS. The trapdoor structure remains single and long term for CHET with a per-hash ephemeral trapdoor, and adaptation requires both, which enables revocation and tracing without changing the digest. They formalize four targets for PCHA, *full indistinguishability* for the pair (h, σ), collision resistance against insider capabilities, uniqueness, and accountability, and prove that FIND reduces to CHET *full indistinguishability*, CR to *strong private CR of CHET* plus IND-CCA security of the ABE layer, UNI to CHET uniqueness, and ACT to EUF-CMA of

IBS together with collusion-resistant traceability of ABTT. Concretely, they build an *adaptively secure* ABTT on prime-order pairings with IND-CPA under *DLIN* and CRT under a modified *q-BDHE* plus fingerprint codes, then lift to IND-CCA via Fujisaki–Okamoto; the CHET instantiation uses the RSA-based CHF, and the implementation shows modest overhead relative to PCH.

Wei *et al.* [WZZ+22] propose a CH with a changeable trapdoor (CHCT) for redactable blockchains in federated learning. The trapdoor structure is single and long term at the owner, each hash includes an ephemeral trapdoor that can be rotated under the supervision of a threshold set of supervisors, and adaptation requires the current ephemeral trapdoor together with an update parameter; the owner can abolish and reissue the trapdoor after a collision, and the framework supports a flexible expiration period with public supervision of updates without revealing the trapdoor. They adopt the security properties from the earlier work of Ateniese and Medeiros [ADM04b], *collision resistance* and *key-exposure freeness* (but they define them differently from the original work) with explicit oracles and prove the former from CDH and the latter from ℓ-SDH in bilinear groups, using pairing-based algorithms with public verification and no reliance on public-key encryption or NIZKs. They instantiate CHCT and build a redactable medical blockchain, reporting lower cost than prior threshold and rotating-trapdoor designs and stable runtime across different underlying hash functions.

Ma *et al.* [MZZ+22] introduce *decentralized policy-based CH* (DPCH) for transaction-level rewrites without a fully trusted authority. The trapdoor structure uses a single long-term chameleon secret key held by the modifier together with a per-transaction *ephemeral* trapdoor sealed under a multi-authority ABE policy; adaptation requires both, which enables fine-grained control while preventing the owner from unilateral edits. They formalize three targets for DPCH, *indistinguishability*, *outsider collision resistance*, and *insider collision resistance* under authority corruption, and reduce them to *strong indistinguishability* and *standard collision resistance* of the underlying CHF, static IND-CCA security of MA-ABE via KEM and Fujisaki–Okamoto, and EUF-CMA of the signature layer. The generic composition uses CH, MA-ABE, and digital signatures; a practical instantiation combines the RSA-based CH and BLS signatures with prime-order pairing MA-ABE, and an evaluation shows favorable overhead compared to centralized PCH constructions.

2023. Tian *et al.* [TZZ+23] formalize *revocable policy-based CH* (RPCH) and give a practical construction. The trapdoor structure combines a single long-term CHF secret key with a per-hash *ephemeral* trapdoor encrypted under a revocable ABE; adaptation requires both a valid, time-updated decryption key and the CHF trapdoor, which enables public revocation of rewriting privileges. Security includes *indistinguishability* and *collision resistance* for RPCH, with reductions to CHF *normal indistinguishability* and the RABE's *semantic security*. Concretely, they instantiate the CHF over *discrete-logarithm* groups and design a revocable ABE from FAME in prime-order pairings under *DLIN*, using KUNodes for logarithmic key-update cost; they avoid re-encryption during adaptation to prevent policy overwriting and report favorable performance versus RSA-based variants. They acknowledge closely related 2021 work [XZL+21] on revocable PCH and regulator-style CH, and position RPCH with an explicit IND model and a more practical revocation mechanism.

Duan *et al.* [DWL+23] define *policy-based CH with black-box traceability* (PCHT) and give a generic, black-box composition from CHET and ABET equipped with robust fingerprinting codes. The trapdoor structure follows CHET: a single long-term CHF secret together with a per-hash ephemeral trapdoor; adaptation requires both, and the ephemeral part is ABET-encrypted under an access policy. Security reduces *indistinguishability* to CHET *indistinguishability*, *collision resistance* to CHET *collision resistance* plus adaptive security of ABET, and traitor tracing to the robustness of the fingerprinting code. Concretely, they instantiate CHET in *discrete-logarithm* groups and ABET from FAME on

Type-III pairings under DLIN with constant-size ciphertexts, and implement a redactable IoT blockchain on Hyperledger (MNT224), arguing moderate overhead while enabling authority-side black-box tracing, which improves over earlier accountability-only PCH lines.

Li *et al.* [LXZ+23] define non-interactive threshold CH (NITCH) with (t, n) *collision resistance* and robustness, they give a concrete instantiation in type-III bilinear groups under the co-CDH assumption in the random-oracle model. The trapdoor is threshold-distributed by a Pedersen-style DKG, each party produces a non-interactive adaptation share equipped with a NIZK of *discrete-logarithm* equality, and any $t+1$ shares combine by Lagrange interpolation in the exponent to adapt, so there is no single long-term holder and no per-hash ephemeral trapdoor. *Indistinguishability* is omitted from the CHF model. They use NITCH to build a decentralized random beacon and a committee-evolution protocol, then compose these into Wolverine, a permissionless redactable blockchain with public accountability and transaction consistency.

Li and Liu [LL23, LL24] introduce *tagged CH* (tCH), where a public tag τ is bound to both Hash and Adapt and outputs from fresh hashing and adaptation are required to be *statistically indistinguishable*. They formalize two tagged collision notions—*tCR* and *tagged full collision resistance (tF-CR)*—and prove that in the *one-time tag* regime (each settled block uses a unique tag) *tCR* and *tF-CR* are equivalent. Intuitively, binding all edits of a block to its fixed tag lets them obtain "full" non-malleability without the usual SSE-NIZK machinery needed in untagged CHFs. Concretely, they give two lattice instantiations *without NIZKs*: a standard-model scheme with an almost-tight reduction to *SIS* (using a PRF, e.g., LWE-based), and a tighter, more efficient ROM variant reduced to *SIS*. In one-time tag mode these achieve the effective strength of *tF-CR*, yielding (by their claim) the first post-quantum path to secure redactable blockchains in the standard model without random oracles or heavy zero-knowledge. The trapdoor is single and long-term from lattice trapdoor generation ; collisions are sampled directly via lattice trapdoor algorithms.

2024. Duan *et al.* [DWWG24] introduce *t-times CH* (t-CH) and *t-times chameleon signature* (t-CS) to enforce *controlled redactions*: once more than t collisions are created for the same digest, a public *extraction* algorithm recovers the trapdoor. They formalize the t-CH/CS syntax and games, adding *extractability* (recover sk from $t+1$ openings) to the usual *correctness*, *indistinguishability*, and *collision resistance*; the trapdoor is single and long term, structured as a vector of per-slot secrets that enables linear extraction, and they also give *multi-modifier* variants that split the trapdoor among parties. Concretely, they instantiate in prime-order discrete-log groups with hash-in-the-exponent and an Adapt rule that updates the randomness linearly; security reduces *weak collision resistance* to *DL*, *strong indistinguishability* to *uniformity* in the *ROM*, and t-CS unforgeability to *DL* via forking. As an application, they build *controlled redactable blockchain* protocols in permissioned and permissionless settings where exceeding the limit triggers public key recovery and rollback.

Li *et al.* [LLM+24] introduce *forward/backward-secure policy-based CH* (FB-PCH). They define two new guarantees. *Indistinguishability against full key exposure* means the adversary learns all randomness used in key generation and in forward updates yet still cannot tell hash from adapt. *FB collision resistance* means that, even with present key compromise or after revocation, the adversary cannot produce valid collisions for past hashes (forward security) nor for future hashes (backward security). The construction combines a DL-based CHF with SSE-NIZKs and a new *FB-ABE* supporting non-interactive forward updates and epoch-based revocation. Trapdoor structure is single and long term in the underlying CHF, together with a per-hash *ephemeral* trapdoor that is FB-ABE encrypted under a policy and time; authorized, non-revoked modifiers recover it to adapt. Security is proved from *discrete-logarithm* and pairing-based assumptions for the CHF and FB-ABE components, and an implementation on type-III pairings reports overhead comparable to prior PCH while adding forward and backward security.

Wang and Ismail [WI24] build a *code-based* CH in the rank metric and use it to obtain *tightly secure two-tier signatures*. On the security side, they adopt the classic model with *weak collision resistance*, *uniformity* of outputs, and *trapdoor-collision correctness*. The construction combines KKS-style rank-metric primitives with an IPHFE mapping (the first CHF based on both KKS and HFE schemes); the trapdoor is single and long term, essentially the KKS secret together with the linear maps used in the encoding. Security reduces to the hardness of rank-syndrome decoding and to the underlying KKS one-time signature assumption, targeting post-quantum parameters. There is no threshold split and no policy or ABE layer. The CHF then composes with a code-based signature to yield two-tier signatures with a tight reduction to the CHF's collision resistance.

Yusoff *et al.* [YAH24] refine policy-based CHF by *replacing* CHET with Krenn *et al.*'s [KMS18] dual long-term trapdoor CHF (CHDLTTF). The CHF trapdoor structure becomes master key plus a *reusable* second long-term trapdoor, so transactions that share a policy or modifier reuse the same second trapdoor rather than generating and ABE-encrypting a fresh ephemeral one each time; this preserves fine-grained control while reducing per-transaction overhead. They also strengthen the ABE layer to CCA security using a Boneh–Katz style transform with a lightweight MAC/KEM wrapper. Security properties remain the usual same as in [KMS18] for the CHF, and the implementation reports lower hashing cost and linear scalability versus CHET-based PCH.

Zhao *et al.* [ZGL+24] define *targeted policy-based CH* (TPBCh) and *tokenized redact privileges*. TPBCh binds rewriting power to a designated new message m' *and* an access policy: the trapdoor structure uses a long-term CHF secret together with a per-hash *ephemeral* trapdoor that is ABE-encrypted and bound to m'; Adapt succeeds only if the modifier first obtains a key matching the policy *and* m', then recovers the ephemeral trapdoor and uses it with the CHF secret. Tokenized privileges issue a one-time, time-bounded authorization recorded on-chain, so permission is checked *before* any rewrite and consumed after one use. Security targets *enhanced collision resistance* (outsider with access to the Adapt oracle, not the usual E-CollRes) and introduces a system-level *re-enhanced collision resistance* experiment; the ABE layer is assumed *CCA2*-secure. Their prototype shows *moderate overhead* (approximately 21.3% vs. PBCH) while preventing uncontrolled rewrites via ex-ante token checks.

Yang *et al.* [YCL+24] present CHSTS, a decentralized CH built from Schnorr threshold signatures so that a collision can be computed only when the transaction proposer cooperates with a t-out-of-n regulator set. The trapdoor is threshold distributed among regulators, the proposer's signing key is required during adaptation, and there is no per-hash ephemeral trapdoor. They adopt *key-exposure freeness* together with *public* and *private collision resistance* (no collision finding oracle is given, parallel to *weak collision resistance*) and prove them under discrete logarithm and computational Diffie–Hellman in prime-order groups. The design uses only DLP primitives and hash-to-challenge techniques, avoiding pairings and public-key encryption. For redactable IoV chains they add a vector-commitment field to each block to enable consistency checking and stateful redaction across dependent transactions, and they integrate CHSTS into Hyperledger Fabric with experiments showing lower overhead than pairing-based thresholds and throughput close to an immutable Fabric baseline.

Wu *et al.* [WDY+24] develop a *verifiable distributed CH* (VDCH) and a matching consensus protocol CVTSS for consortium chains. The trapdoor is *threshold distributed* among nodes via verifiable key sharing with secure MPC, there is no per-hash ephemeral trapdoor, and adaptation is produced from t verified collision shares. Each node proves its share correct with a Schnorr NIZK so any party can check shares before combining them; this adds fault tolerance and removes the single-holder trust point. They formalize verifiability and adopt *standard collision resistance* and reduce it to *CDH* in a Gap-DH group; *correctness* follows from their hashing and share-combination rules. CVTSS

coordinates redactions with verifiable threshold signatures, aggregates the network's decision, and prevents abuse of one-time chameleon keys while avoiding Nakamoto-style voting. A full consortium-chain design built from VDCH and CVTSS is evaluated with malicious nodes and shows acceptable time cost.

Wang *et al.* [WCZ+24] propose a *hierarchical identity-based CH with revocable subkey* (HIBCH–RS) over lattices, aimed at quantum resistance and trapdoor *updates*. Trapdoors are delegated along an identity tree via a new *SlaveKey* mechanism so a parent cannot forge a child's key; per-identity trapdoors are *revocable* and *updatable* through an explicit Update algorithm, and there is no per-hash ephemeral trapdoor (collisions are produced directly via lattice trapdoor sampling). The syntax comprises Setup, SlaveKey, DelKey, Update, CH, and Forge. On security, they formalize correctness, a revocability property, and *resistance to collision forgery under active attacks* (what they call it *enhanced-CR* experiment with adaptation queries), and they further define a *forgery indistinguishability* property (outputs of CH vs. forged collisions). The construction relies on standard lattice tools (TrapGen, SampleD, DelTrap with FRD encoding) and proves collision security under LWE/ISIS-style assumptions; a Hyperledger Fabric prototype demonstrates practicality.

Chan *et al.* [CCTY24] "reconstruct" chameleon hashing to be the first that attains *both* full indistinguishability and full collision-resistance in one scheme. The core idea is to make the digest independent of the public key and move the binding into the opening: the randomness is an OR–NIZK that either binds the digest to the message or proves knowledge of the trapdoor, which preserves indistinguishability even against a secret-key holder while still allowing simulation for full collision-resistance reductions. They also point out a *key-replacement* attack in multi-key settings and fix it by additionally hashing the (authenticated) public-key set into the recorded value to prevent re-binding under a different key. Beyond single-owner security, they formalize a multi-party setting (creator vs. owner) and introduce *(un)claimability* and *(un)deniability*—showing how parties can (or provably cannot) assert or deny authorship of a digest/opening—and sketch a multi-owner variant by extending the OR-proof across many keys. The trapdoor structure is a single long-term secret per owner (multiple in the multi-owner variant), with no per-hash ephemeral trapdoor. They give two concrete instantiations: an ECC/DL scheme (full IND and full CollRes in the ROM) and a lattice scheme under Module-LWE and Module-SIS, both implemented to demonstrate practicality; security relies on OWF, collision-resistant hashing, and simulation-sound extractable NIZKs. Finally, they discuss applications to redactable blockchains and why the new multi-party notions matter for accountability and privacy.

2025. Li *et al.* [LSW25] revisit decentralized CHF–based redactable blockchains and show that Jia *et al.* DCH [JCH+22] does *not* reach its claimed security: once a block at height h has been redacted, any party can derive further collisions for the same label without the threshold trapdoor, enabling unilateral redactions. They trace the root cause to using a KEF-level guarantee and then fix it by designing a direct pairing-based CHF with *standard collision resistance* and *indistinguishability*. Concretely, hashing is $h = H(m) \cdot g^{\xi}$ with randomness $r = y^{\xi}$, verification checks $e(h/H(m), y) = e(g, r)$, and adaptation uses the secret exponent x to map r to r'; collision resistance reduces to the EUF–CMA security of an underlying BLS-style signature, and the scheme avoids SSE–NIZKs. Their improved DCH keeps a (t, n) threshold long-term trapdoor via Shamir shares and eliminates any per-hash ephemeral trapdoor; adaptation is a share-based re-randomization that preserves decentralization. Because h is no longer recomputable from (m, r), the block header stores a "current-hash&rand" field; theory and experiments show comparable costs to [JCH+22] with smaller randomness.

Liu and Zhao [LZ25] propose *time-modifiable CH* (TMCH): hashing and verification are bound to a time label t, and besides the usual Adapt (message update) they add Update to change t while preserving the digest. They formalize two *indistinguishability* notions, *typical* (T-IND, hash vs. adapt) and *new* (N-IND, hash vs. update), and matching

collision resistance notions (Typical-CR and New-CR), and argue how relaxing uniqueness is compatible with these goals. The concrete instantiation is pairing-based from Boneh–Boyen: a single long-term trapdoor, standard per-hash randomness, and proofs that T-/N-IND reduce to *DDH* and ℓ-*SDH*, while T-/N-CR reduce to ℓ-*SDH*; they report an implementation and comparisons indicating low, policy-independent overhead.

Building on the line of tagged chameleon hash functions [LL23], Clermont *et al.* [CDJ+25] remove tags and construct a lattice chameleon hash with *full* collision resistance in the random–oracle model, the first chameleon hash function achieving full collision resistance without the use of NIZK protocols: they commit to the hashing randomness using a trapdoor commitment that satisfies *DEE*, and then adapt by pre–sampling a matching randomness and *equivocating* the commitment; *DEE* ("distributional equivalence of equivocation") means an honest commit–open pair is computationally indistinguishable from one produced by committing to a fixed value and later equivocating with the trapdoor, which lets the proof answer Adapt queries without the secret key. The trapdoor remains single–holder and long–term (GPV–style), and they note that uniqueness is not achieved (they prove a generic attack on uniqueness for their design), which matters if the CHFs is plugged into some SSS frameworks (they address this via a VRS-based transform, but that part uses NIZK—outside the CHF itself).

4 Classification of the constructions, a first taxonomy

As Section 3 demonstrates, there exists a wide variety of constructions for CHFs, each built on different hardness assumptions and cryptographic paradigms. To make sense of this diversity, we present a taxonomy that organizes existing constructions into coherent families, as illustrated in Figure 3. In the following subsections, we detail each branch of the taxonomy, namely the authorization model, the trapdoor structure, and the functional or algebraic enhancements.

Scope and organizing principle. The chronological narrative motivates the design space; the present section complements it with a compact, type-driven taxonomy. We classify along three orthogonal axes that every construction instantiates, independently of concrete assumptions or proof flavors: authorization model, trapdoor structure, and functional enhancements. Consequently, a given scheme is located by its coordinates on these axes without conflating mechanism, policy, and capability.

Reading guide. Each subsection proceeds definition first, then minimal context, then precise pointers to the canonical formulations. The items are self-contained to make the map of chameleon-hash types explicit: a reader identifies the relevant authorization model, trapdoor structure, and enhancement set, respectively, and then consults the cited items for full formal detail.

4.1 Authorization model (who is able to adapt)

This axis specifies the set of legitimate adaptors and the semantics by which adaptation rights are granted, respectively per identity, per policy, or per quorum. Identity-based variants authorize holders of per-identity secrets. Policy- or attribute-based variants authorize precisely those holding keys that satisfy an access predicate, including accountable, traceable, revocable, multi-authority, decentralized, and targeted subtypes. Threshold or committee settings distribute authorization to a t-of-n set. The model answers *who*, not *how*; it composes with any trapdoor arrangement below.

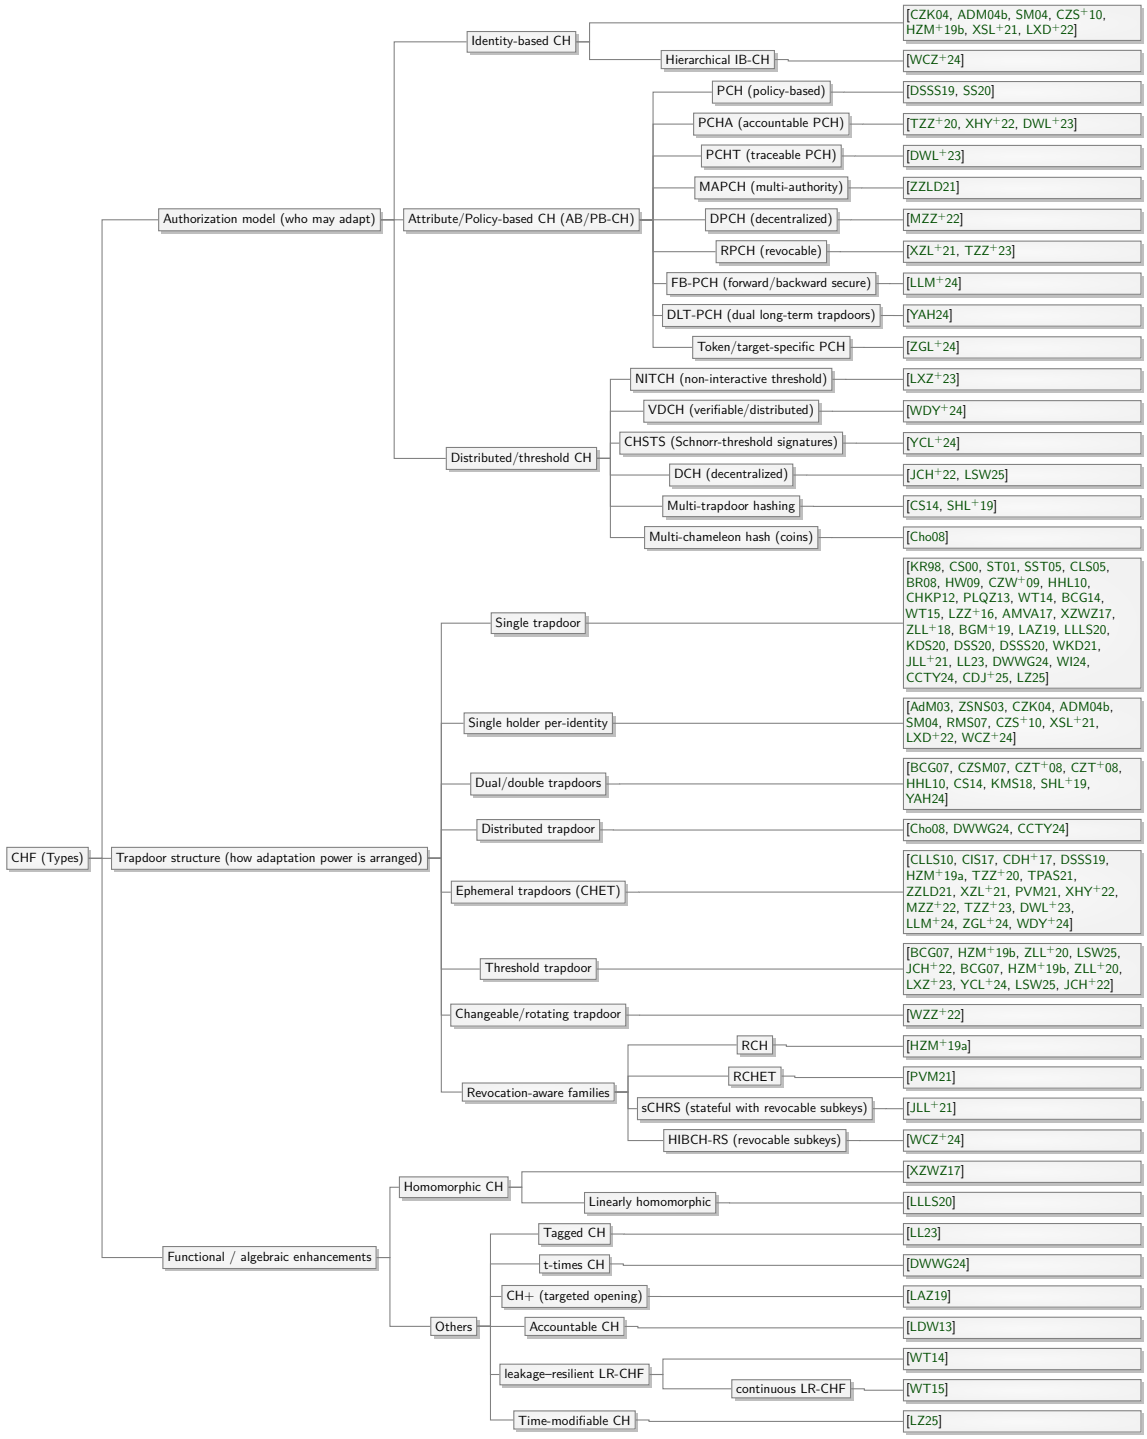

Figure 3: Classification of CHF constructions (note: the main branches are not mutually exclusive)

4.1.1 Identity-Based Chameleon Hash (IB-CH).

An authority issues a per-identity trapdoor. The holder of that identity can adapt openings, outsiders cannot. The earliest identity-based formulations appear in 2003, carrying trapdoor

hashing and chameleon signatures to the identity setting and showing compact realizations under standard number-theoretic assumptions [AdM03, AdM04a, ZSNS03]. Follow-ups clarify exposure and denial, and study issuance and revocation details in identity systems [CZK04, ADM04b, SM04, RMS07, CZS+10].

Hierarchical Identity-Based Chameleon Hash (HIB-CH). Identities are organized hierarchically, respectively as vectors along a tree, and a higher level authority derives per-identity trapdoors for its descendants. A holder of the secret for identity $ID = (ID_1, \ldots, ID_\ell)$ can adapt openings, while delegation proceeds by deriving child keys from parent keys without reinitializing public parameters. This improves scalability for large organizations and enables structured revocation policies; concrete hierarchical and revocation-aware realizations are given in [WCZ+24].

4.1.2 Policy-based Chameleon Hash (PCH)

A policy-based chameleon hash binds an access policy to the digest so that an opening can be adapted if, and only if, the presented attributes satisfy that policy. Formal syntax and collision/indistinguishability experiments are defined together with concrete realizations that compose a CHF with a policy layer, typically via attribute-based mechanisms; these works serve as the canonical definitions for PCH [DSSS19, SS20].

Accountable Policy-based Chameleon Hash (APCH). Accountable policy-based CH augments PCH with black-box attribution: whenever an adaptation is produced, verifiable evidence links the event to the enabling party, without revealing trapdoors. This is achieved by layering accountability over the policy substrate, and instantiated with efficient constructions and system considerations [TZZ+20, XHY+22, DWL+23].

Traceable Policy-based Chameleon Hash (PCHT). Traceable policy-based CH strengthens accountability to full tracing, enabling identification of who authorized or performed an adaptation rather than merely proving that some authorized party existed. The model and constructions formalize traceability guarantees and show how to integrate tracing into policy-bound hashing [DWL+23].

Multi-Authority Policy-based Chameleon Hash (MAPCH). Multi-authority PCH removes reliance on a single issuer by distributing attribute issuance and verification across several authorities while preserving PCH semantics. The model prevents a single point of failure or control and provides concrete realizations with policy enforcement under multiple issuers [ZZLD21].

Decentralized Policy-based Chameleon Hash (DPCH). Decentralized PCH eliminates central authorities entirely: policy satisfaction and authorization are enforced in a decentralized setting, so that adaptation rights arise from protocol-driven policy evaluation rather than a trusted issuer. Definitions and constructions demonstrate feasibility and articulate the security interface with the underlying CHF [MZZ+22].

Revocable Policy-based Chameleon Hash (RPCH). Revocable policy-based CH equips the policy layer with revocation so that previously valid keys or epochs can be disabled without reissuing all credentials. Revocation is integrated into the hashing and adaptation interface, yielding concrete mechanisms and precise security targets for revocable authorization [XZL+21, TZZ+23].

Forward/Backward-secure Policy-based Chameleon Hash (FB-PCH). Forward / backward-secure PCH controls the temporal impact of compromise: forward security

prevents retroactive adaptations after present compromise, and backward security prevents future adaptations from past exposures. Constructions realize these epoch-evolving guarantees on top of the policy substrate [LLM+24].

Dual Long-Term trapdoors for Policy-based Chameleon Hash (DLT-PCH). Dual-long-term PCH realizes policy-bound adaptation using two long-lived trapdoors that must be combined, separating roles such as long-term authority and operational capability while retaining policy checks at adaptation time. This variant clarifies how policy enforcement composes with a dual-trapdoor architecture [YAH24].

Targeted/token-specific Policy-based Chameleon Hash. Targeted policy-based CH restricts even authorized parties to predesignated openings, binding adaptation power to a specific target (for example via a token). The definition and constructions show how tokenization narrows the image of adaptation within a policy-bound framework [ZGL+24].

4.1.3 Distributed and threshold families

Distribution dictates that the adaptation power is shared by a set of parties. Threshold chameleon hashing requires any t of n shares to adapt; non-interactive variants aggregate shares without on-line interaction; verifiable variants add public checks. System-oriented designs for ledgers use regulator or committee structures [HZM+19b]; formal threshold definitions and concrete schemes appear in [ZLL+20]; recent work gives non-interactive or publicly verifiable threshold constructions with implementations [WDY+24, YCL+24]. When this capability is expressed through policies it aligns with PCH (Sub-section 4.1); when it is expressed through secret sharing it aligns with threshold trapdoors in Subsection 4.2.

4.2 Trapdoor structure (how adaptation power is managed)

This axis captures the mechanics required at adaptation time. Single long-term trapdoors, and their per-identity refinements, represent the baseline. Dual or double long-term trapdoors separate roles so two persistent secrets must be combined. Ephemeral designs bind a per-hash or per-epoch token to the digest; more formally, adaptation requires the long-term key together with an auxiliary value scoped to the target. Threshold sharing treats the trapdoor itself as a (t, n) secret. Changeable or rotating forms allow supervised refresh of the auxiliary component. Revocation-aware structures incorporate state that enables selective disabling. The distinction between dual and ephemeral is structural: dual requires two long-lived secrets for any adaptation; ephemeral requires a target-scoped auxiliary, hence per-hash or per-epoch behavior.

Single trapdoor (baseline). A single secret suffices to adapt under a public key. This structure is present from the beginning and remains the pivot across many discrete-log, RSA, lattice, and code-based instantiations [KR98, ST01, SST05, CLS05, BR08, HW09, BR14]. Later generic frameworks still evaluate security with the adaptation algorithm consuming exactly one secret at adaptation time [DSS20, AMVA17, WKD21, WI24, CCTY24].

Single holder per identity. Mostly related to identity-based chameleon hashing: each identity receives its own long-term secret from a PKG, so identities adapt independently. The model and its exposure or denial nuances are established by the 2003–2004 lines [AdM03, AdM04a, ZSNS03, CZK04, ADM04b, SM04] with additional identity-centric developments and revocation in later works [RMS07, CZS+10, XSL+21, LXD+22, WCZ+24].

Dual or double trapdoors. Two separately generated secrets must be combined to adapt, which separates authority from operational capability. Concrete patterns and instantiations include double-trapdoor structures and hidden-order settings [BCG07, CZSM07, CZT+08, HHL10, CS14, KMS18, SHL+19, YAH24]. The defining feature is that both components are long-term and any adaptation requires their joint use.

Distributed trapdoor. No single entity holds the whole secret; shares are combined, often with public checks, to adapt. This captures key distribution and combination protocols beyond mere threshold counting and is used in committee-style deployments [WDY+24, CCTY24]. The emphasis here is how the secret is stored and evaluated.

Ephemeral trapdoors (CHET). Besides a long-term key, each hash or epoch binds an auxiliary ephemeral that is required for adaptation. This prevents the long-term holder from unilaterally modifying a fixed digest. The idea is developed across multiple platforms, with explicit per-use ephemerals and clean applications to access control and redactable ledgers [CLLS10, CHKP12, WT14, CIS17, CDH+17, DSSS19, HZM+19a, TZZ+20, ZZLD21, XZL+21, PVM21, XHY+22, MZZ+22, TZZ+23, DWL+23, LLM+24, ZGL+24]. Some designs are per-epoch by default; rotating the ephemeral per hash yields per-transaction granularity.

Threshold trapdoor. The trapdoor itself is (t, n)-shared and any t valid partial openings combine into an adaptation. Several entries are non-interactive or publicly verifiable. Formal and implementation-oriented works include [HZM+19b, ZLL+20, YCL+24, LSW25, JCH+22].

Changeable or rotating trapdoor. The ephemeral component can be rotated under public supervision so that subsequent adaptations require fresh tokens, enabling time-scoped rewriting without global rekeying. Formal treatments and concrete realizations analyze rotation cost and public verifiability [WZZ+22, MZZ+22].

Revocation-aware structures. Revocation is embedded into the structure. Time-bound ephemerals suit consortium chains [HZM+19a]; CHET with the ephemeral encrypted under a revocable policy enables fine-grained disabling [PVM21, XZL+21]; a stateful master–subordinate arrangement (sCHRS) binds, per hash, which subordinate may adapt and lets the master revoke that power case-by-case [JLL+21]. Although two keys exist in sCHRS, adaptation consumes a single authorized secret for that hash.

Note on dual/double versus CHET and on per-epoch ephemerals. Dual or double structures and CHET both prevent the long-term holder from adapting a fixed digest alone, but they differ architecturally: dual requires two long-lived secrets for any adaptation [CZSM07, CZT+08, HHL10]; CHET requires a per-hash or per-epoch token bound to the target [CLLS10, CHKP12, WT14]. Several systems are per-epoch by default; stating this explicitly matters for applications that require per-transaction isolation.

Note on the differences between dual/double, distributed,threshold trapdoors. double trapdoor might look like a distributed/threshold trapdoor with $n = 2$. However, a double/dual trapdoor uses two independently generated secrets, whereas a distributed/threshold trapdoor splits one master secret into shares with t-out-of-n semantics. They coincide only in the special AND case (both secrets required), which behaves like $t = n = 2$.

4.3 Functional and algebraic enhancements

Finally, this axis records capabilities layered on top of the basic interface. Homomorphic variants support algebraic operations on digests that mirror message operations. Tagged

variants bind public context to Hash and Adapt. t-times variants bound the number of permissible adaptations. Chameleon Hash Plus enables targeted openings to designated messages. Accountable variants provide evidence of who enabled a change in the plain setting. Leakage-resilient lines preserve guarantees under bounded or continual leakage. These features are orthogonal to both authorization and trapdoor structure.

4.3.1 Homomorphic chameleon hashing.

Homomorphic CH supports algebraic operations on hashes that mirror operations on messages, enabling structured updates and efficient aggregation without rehashing from scratch. Instances include group-homomorphic and linearly homomorphic realizations with lattice sampling and public-coin templates [XZWZ17, CIS17, LLLS20].

Tagged chameleon hash. A public tag (τ) is bound to Hash and Adapt; tagged collision notions relate hashing and adaptation across tags. In the one-time-tag regime, the tagged notions collapse in the model reported in the literature and tags help inhibit malleability without heavy zero-knowledge [LL23, LL24]. Tags are natural in ledger contexts where blocks or fields carry unique labels.

t-times chameleon hash. At most t successful adaptations are permitted under a public key; after $t+1$ openings a public extractor recovers the secret or blocks further rewrites. This gives controlled rewriting that composes with governance policies [DWWG24].

Chameleon Hash Plus (CH+). Given a digest, an opening to a designated target message can be computed under stated constraints. From a digest h, an adaptor holding the trapdoor can compute an opening to a designated target message m' using only h, i.e., without the original pair (m, r); admissible targets and any bound public context are fixed by the scheme's interface. CH+ is formalized via targeted one–wayness and (weak) collision resistance, with lattice instantiations (standard-model via GPV/SIS and an efficient NTRU/Falcon preimage-sampling design), making it a natural fit for chameleon-signature wrappers and redactable logs that require precise control of the image of Adapt [LAZ19].

Accountable chameleon hash (ACH). Accountability is added to basic CH so that revealed collisions carry verifiable evidence of who authorized the change. The model and protocols show how to graft accountability without altering the base trapdoor structure [LDW13]; policy-based systems reuse the same idea at the authorization layer.

Leakage-resilient and continual-leakage CH. Security is maintained when bounded information about secrets leaks, and extended across many rounds in continual-leakage models. Constructions derive from leakage-resilient building blocks in both discrete-log and factoring settings [WT14, BCG14, WT15]. These refinements compose with the structural and authorization choices above.

Time-modifiable CHF. Binds temporal metadata to the digest and defines an adaptation interface that can update both the message and its time label under explicit time rules (for example, within a bounded window or only forward in time). The authorization predicate is unchanged (per identity, per policy, or per quorum), but the admissible openings are restricted by a temporal policy that is enforced at verification. Constructions realize this by coupling the hash with epoch tokens or time-indexed auxiliary data so that a valid collision must respect the prescribed time relation, consequently enabling redactable blockchains with auditable "grace periods" and post-hoc corrections while preventing unconstrained retroactive edits. See the formalization and instantiation in the time-modifiable line [LZ25].

5 Discussions

We now step back from the details and read the results as a whole. The aim is to make clear what the figures and timeline actually support, what they leave out, and how to use them safely when choosing a scheme. We then draw a few simple lessons for practice and point out gaps that are worth exploring next.

5.1 Taxonomies

Our primary taxonomy (Section 4) was designed to address the research question **"What are the main families of CHF constructions?" (RQ.2)**. We detailed this classification and explained the rationale behind its design (Section 4).

However, the database we built covers topics ranging from fundamental works on CHFs, their constructions, and up to their applications. We would like to stress out that it therefore enables the construction of alternative taxonomies that answer different research questions.

5.1.1 Underlying hardness assumptions

For instance, by analyzing the underlying assumptions of each contribution, we can derive another taxonomy (Figure 4) that answers the question **"Which hardness assumptions underpin CHFs?" (RQ.3)**. This reveals how the literature maps onto distinct branches and the recent attempts to propose quantum-resistant alternatives.

5.1.2 The evolution of security properties

Similarly, one can shift the focus from constructions to security properties and ask: **"How have the security properties and adversarial models for CHFs evolved?" (RQ.1)**. We have shown in Section 2 that the security framework has evolved throughout the years, reaching a point of consensus with Derler *et al.*'s work [DSS20]. However, the taxonomy presented in Figure 5 shows that there are still efforts to make in order to unify the research. Indeed, we can observe on the graph that recent papers do not fall in the categorization of the authors, and that depending on applications, new security properties are defined (e.g., traceability in [DWL+23]).

5.1.3 List of applications

Finally, when the focus changes to applications, we believe that one can find the answer to: **"Which applications actually deploy CHFs, and what design constraints do they impose? (RQ.4)**. Here, our classification groups contributions by their application domains, such as redactable blockchains, signature schemes, identity-based cryptography, or access-control systems (Figure 6).

Reading guide For readability we group applications into four families: (A) Authorized editing after issuance, (B) Non-transferability / designated verification, (C) Privacy with accountability, and (D) Security/tightness & performance wrappers. These families are not mutually exclusive; each work is placed where its primary intent is most evident.

A. Authorized editing after issuance (edit rights are the point).

Sanitizable signatures (authorized edits after issuance). Ateniese *et al.* introduced sanitizable signatures: the signer designates which parts of a document are editable, digests those regions with a chameleon hash, and signs the global structure [ACDMT05]. A designated sanitizer, holding the corresponding trapdoor, can later update an authorized

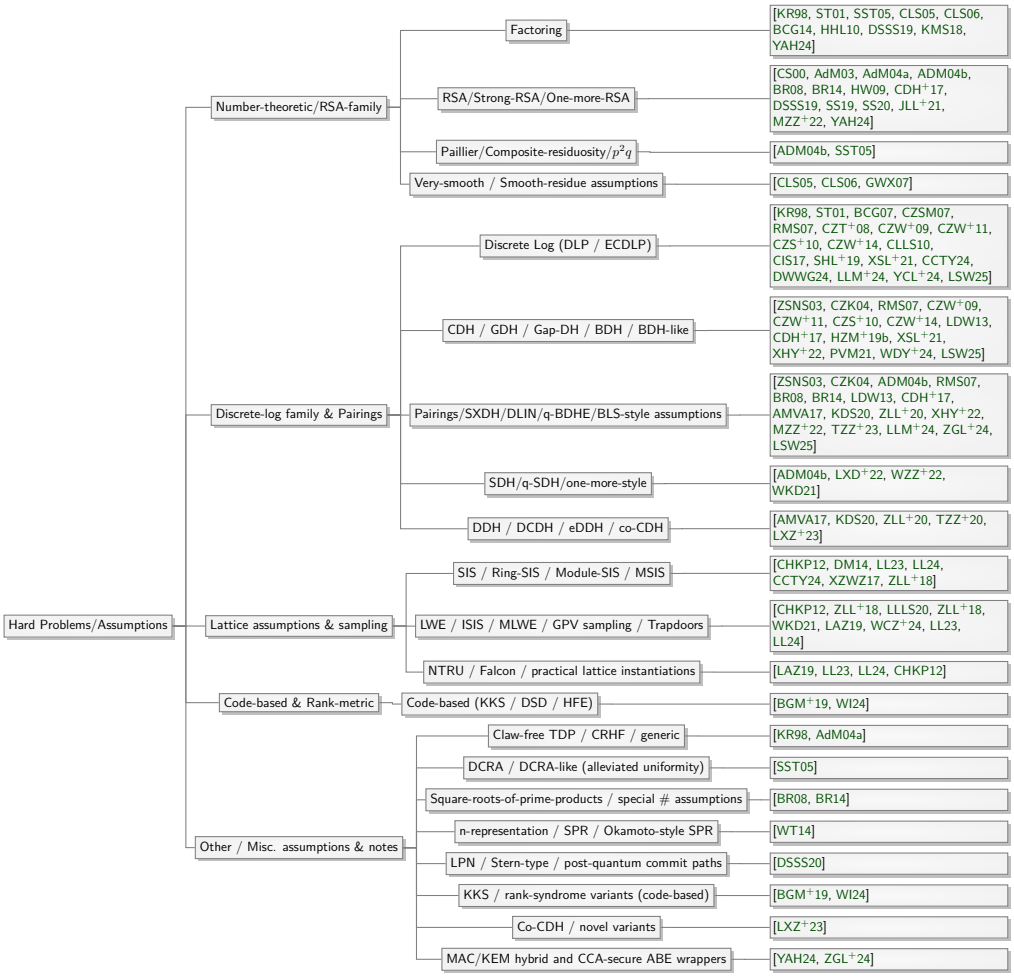

Figure 4: Classification of CHF Underlying Hard Problems

region and adjust the CHF opening so the original signature remains valid, without contacting the signer. This fits editorial, medical, and compliance workflows where certain fields must change after issuance, for example redacting PII, normalizing formats, or removing sensitive fragments, while preserving end-to-end authenticity and a clear record of *who* may change *what*. Compared to redactable signatures, which permit public removal of marked content, sanitizable signatures keep edit rights explicit and accountable.

Trapdoor-sanitizable signatures (delegating fine grained edit rights). Canard *et al.* introduced trapdoor-sanitizable signatures: the signer can later issue scoped trapdoors that authorize specific parties to modify declared regions of a signed document, while downstream verification remains unchanged [CLM08]. In practice, trapdoors can be tied to identities, attributes, or roles (via identity-based chameleon hashing), enabling per-message or per-field edit capabilities, with natural support for expiration and revocation. This preserves end-to-end authenticity, keeps governance explicit about *who* may change *what*, and scales administrative control without altering the public verification interface.

Edit-friendly image/document authentication. Xu *et al.* introduce verifiable image revision from chameleon hashes: the signer designates editable regions, digests them with a chameleon hash, and signs the overall structure; an authorized editor can later update content and adjust the opening so the original signature still verifies, without interacting

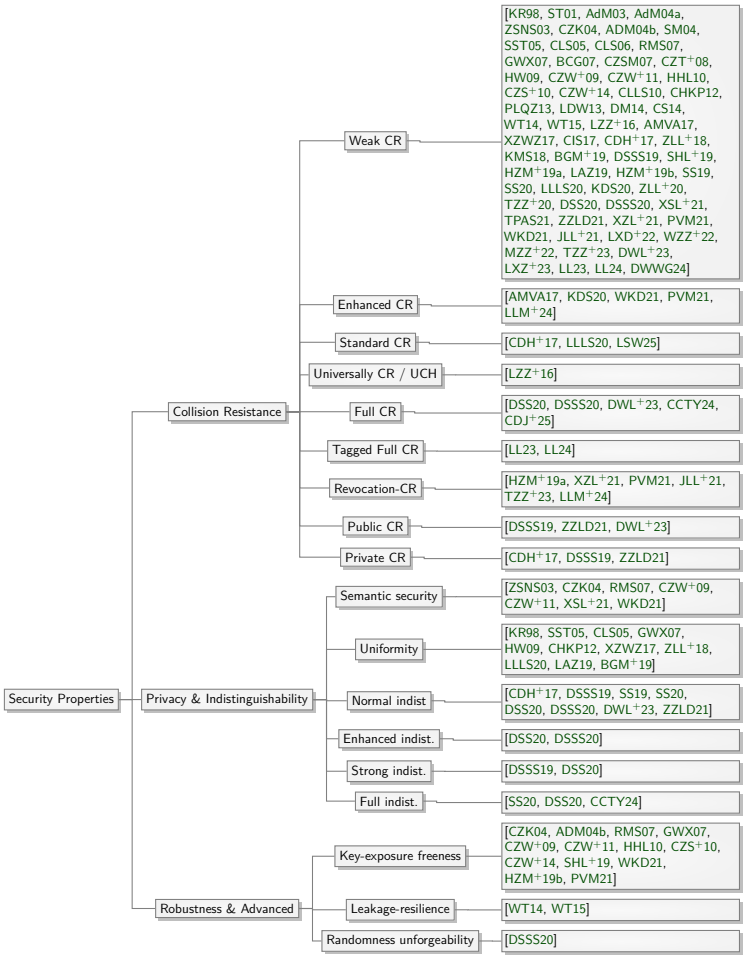

Figure 5: A Taxonomy of Chameleon Hash Security Properties (note: the main branches are not mutually exclusive)

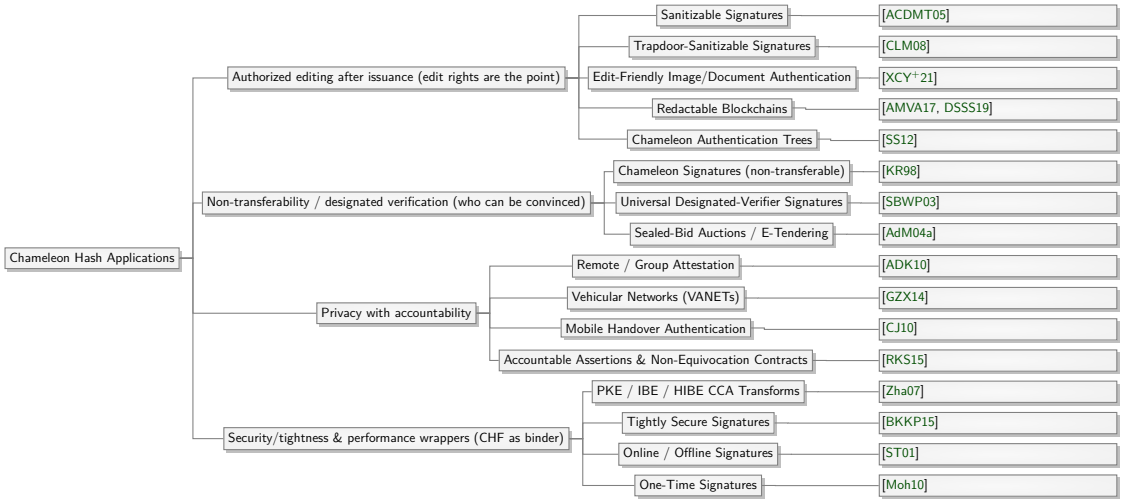

Figure 6: Chameleon Hash Applications

with the signer [XCY+21]. Practical designs embed the editing rule inside the signed data and hash non-editable regions with ordinary hashes, consequently unauthorized changes invalidate verification while permitted, privacy-aware edits remain publicly checkable.

Redactable blockchains. Ateniese *et al.* introduced redactable blockchains: the predecessor pointer $\mathsf{H}(prev||\cdot)$ is replaced with a chameleon hash $\mathsf{CH}(pk, \cdot)$, so that when a block's contents must be altered or pruned, an authorized editor holding tk adapts the opening while preserving the stored link value [AMVA17]. Consequently, edits such as removing illicit payloads, excising personal data, compressing historical blocks, or correcting metadata can be performed without forking or re-mining the suffix, while light clients and new nodes continue to validate against the public headers they see; the "hash-chain" shape and validation logic remain largely intact, which explains the clean integration with Bitcoin-like designs and the small measured overheads shown by a Bitcoin-core prototype. Governance is orthogonal and flexible: centralized editors in permissioned ledgers, threshold-held trapdoors among designated authorities, or consensus-gated edits in permissionless settings, in each case the edit right becomes an explicit policy object that can be audited. Follow-ups refine policy scoping via identity- or attribute-tied trapdoors (policy-based chameleon hashes) to enable fine-grained, transaction-level rewriting [DSSS19].

Chameleon authentication trees. The Schröder introduced chameleon authentication trees: a Merkle-style structure where selected internal nodes use a chameleon hash so appends or authorized edits are absorbed by adapting openings, keeping the published root unchanged [SS12]. Verifiers still check ordinary authentication paths against the stable root, yielding $\mathcal{O}(\log n)$ updates, short proofs, and a fixed public anchor suitable for streaming or archival settings where the root is embedded in a certificate, ledger, or transparency log. Applications include verifiable data streaming and accountable publishing.

B. Non-transferability / designated verification (who can be convinced).

Chameleon signatures (non-transferable proof to outsiders). Krawczyk and Rabin [KR98] keep the standard hash-then-sign pipeline and replace the digest with a chameleon hash whose trapdoor is held by the recipient. Because the recipient can generate an alternative message that yields the same chameleon value, they can verify the signature for themselves but cannot produce transferable evidence for a third party. Any transcript could be reconstructed after receipt using the trapdoor. The construction preserves the familiar shape of digital signatures while changing the evidentiary outcome that matters in practice: authenticity for the intended party and plausible deniability, i.e., non-transferability, to everyone else.

Universal designated-verifier signatures. Steinfeld *et al.* introduced universal designated-verifier signatures: a public signature can be converted by any holder into a form that convinces only a selected verifier, using the verifier's public key [SBWP03]. The designated verifier, who holds the matching secret key, can always simulate what they see, consequently outsiders cannot obtain transferable evidence, while the verifier's local check remains as lightweight as in the underlying scheme. This fits receipt, acknowledgement, and accountability settings where convincing power must be restricted to a single party, for example confirmations or closed-door tenders.

Sealed-bid auctions and e-tendering (commit privately, open under policy). Ateniese and de Medeiros used identity-based chameleon hashing to design sealed-bid auctions with minimal communication: a bidder submits one chameleon-signed bid, the auctioneer issues a receipt, and bids are revealed at close [AdM04a]. The chameleon layer binds the bid to an identity or auction tag, so commitments are binding to everyone

yet openable by the authorized party at the defined time, and the non-transferability property prevents insiders from providing transferable proofs of others' bids. In practice, this supports e-tendering where bids remain private during the bidding phase, can be opened collectively at closing, and can be re-bound in narrow, auditable exceptions such as dispute resolution, while fitting naturally into PKI or IBE deployments.

C. Privacy with accountability (hide today, resolve if needed).

Remote (group-based) attestation with privacy. Alsouri *et al.* introduced a chameleon-hash based, group-oriented attestation that hides exact platform configurations while keeping verifiers confident that a device operates within a whitelisted class [ADK10].A device proves membership in this whitelisted class, while a trusted authority can later resolve the exact state or extract accountability if it equivocates. Updates appear as authorized "retargetings" to the same anchor, consequently public artifacts need not change when patches land. The result balances configuration privacy for honest devices with a credible recourse path for operators.

VANETs (anonymous beacons with authority traceability). Guo *et al.* use chameleon hashing to authenticate safety beacons without linkability over time [GZX14]. In their design, messages are signed under pseudonyms whose chameleon bindings are resolvable only by a designated authority, consequently honest drivers remain unlinkable to peers and roadside units, while misbehaving ones can be deanonymized after the fact. The protocol supports both V2R and V2V mutual authentication with conditional privacy, and keeps verification light enough for resource-constrained on-board units (no bilinear pairings, fast public checks). This yields short handover latencies and a clear, auditable policy for *who* may unmask *whom*.

Mobile handover authentication. Choi and Jung use chameleon-hash–based credentials for fast handovers: a device presents a non-linkable commitment that the new base station validates locally, while only core components with the appropriate trapdoor can later recover or check the underlying identity if needed [CJ10]. The device simultaneously runs an authenticated ephemeral Diffie–Hellman, consequently no AAA round-trip [3] is required during handover and long-term identifiers remain hidden. In practice, this keeps latency and energy low and integrates with standard challenge–response flows; their prototype reports multi-millisecond handovers with simple RSA verification rather than bilinear pairings.

Accountable assertions and non-equivocation contracts. Ruffing, Kate, and Schröder introduced accountable assertions and non-equivocation contracts: parties bind statements to contexts via a chameleon hash, and if two inconsistent openings appear, a public extractor recovers the equivocator's key, enabling a prearranged deposit to be forfeited [RKS15]. In effect, "do not say two different things" becomes an enforceable, auditable rule with low on-ledger overhead, since verification is a standard public check and penalties are automatic through the deposit mechanism. This captures slashing, exclusion, or public blame in a single mechanism that fits append-only or blockchain-backed settings.

D. Security/tightness & performance wrappers (CHF as binder).

[3] AAA stands for *Authentication, Authorization, and Accounting.* An "AAA round-trip" is a full request/response exchange between the access point and a remote AAA server for user authentication, which adds extra latency during handover.

PKE/(H)IBE transforms (hardening encryption against active tampering). Zhang proposed a simple methodology for lifting IND-CPA encryption to IND-CCA security, for public-key, identity-based, and hierarchical settings [Zha07]. The idea is to place a chameleon hash as a lightweight binder on the tweakable parts of the ciphertext, for example randomness, a tag, or an identity, and to have the scheme record and check that value during decryption. Any tampering must preserve the bound value, which is out of reach without the trapdoor, so decryption oracles no longer help an attacker craft useful modifications. The attraction is practical: the transform fits into existing IBE and HIBE frameworks with minor format adjustments, keeps ciphertexts and keys compact, and provides resistance to active manipulation that plain IND-CPA designs do not. In short, Zhang cast chameleon hashes as a clean upgrade path from tag-based or identity-based encryption in the CPA setting to CCA security, with the same reasoning extending naturally to hierarchical variants.

Tightly secure signatures (standard model, tight reductions). Blazy *et al.* and Pan showed that suitable chameleon hashes yield signature schemes with tight reductions in the standard model [BKKP15]. Practically, tightness means the provable loss does not blow up with the number of signing queries, consequently parameters need not be inflated to compensate for weak reductions. The payoff is predictable security margins, smaller keys and signatures at a given target level, and fast verification under familiar assumptions (RSA, CDH, SIS). For operators, this translates to lower HSM load and cheaper high-throughput deployments; for auditors and regulated environments, it avoids reliance on the random-oracle heuristic and simplifies assurance narratives. In short, CHF is used as the lever for tightness rather than new functionality, so integration is routine and the verification interface remains standard.

Online/Offline signatures (precompute now, adapt later). Shamir and Tauman showed that chameleon hash functions enable an on-line/off-line signing method that front-loads the computational cost [ST01]. The signer prepares off-line by producing signatures on chameleon hash images before the message is known; when the message arrives, they adjust the opening so the pre-signed image matches the hash of the actual message and release the small on-line component. The result is near-instant on-line signing for constrained devices or high-throughput settings, together with a clean separation between the signing interface and message content (the oracle only observes chameleon images). Keys remain short, verification is simple, and the overall effect is a practical "front-load the cost" signing workflow.

Weak→Strong signature upgrade (sUF from wUF). Steinfeld *et al.* [SPW07] give a generic transform that strengthens any weakly unforgeable signature into a strongly unforgeable one by wrapping it with two randomized trapdoor (chameleon) hash functions that are strongly collision-resistant (it is hard to find two randomness numbers $r = r'$ and a message m such that (m, r) and (m, r') are a collision-pair for the CHF) and satisfy the random trapdoor collision property [SPW07]. Concretely, the message is hashed together with signature material through one CHF, then "switched" via the trapdoor of a second CHF so that the underlying scheme can be treated as a black box; the result is sUF-CMA security without abandoning the existing construction or assumptions (illustrated by a DSA instantiation) at the cost of small key/signature overhead and a few extra exponentiations. This casts CHFs as a security wrapper: they bind the tweakable parts of the signing workflow so post-hoc modifications reduce to breaking collision resistance, while integration keeps the familiar verify interface of the underlying scheme.

One-time signatures. Mohassel introduced a general transformation that turns any *weakly CR* chameleon hash function into a strongly unforgeable one-time signature [Moh10]. In practice, the signer commits once to a chameleon image, then binds that commitment

to the chosen message by revealing the appropriate opening; since alternative messages would require different openings that only the holder can compute, reusing the signature on a new message is prevented. The payoff is small signatures, lightweight verification, and very fast signing in concrete instantiations, including options with no exponentiations that are attractive for constrained devices and high-rate telemetry.

5.1.4 From properties to applications

Reading the map with care. Linking security notions to applications is not static, but historically contingent. As definitions evolve and attacks surface, the "right" level often shifts upward. A classic example is redactable blockchains: the initial proposal [AMVA17] argued that *enhanced collision resistance* sufficed for safe header preservation under authorized edits. Subsequent analyses and constructions indicate that weaker collision notions are malleable in multi–edit or cross–context settings, and that *full collision resistance* better matches the ledger's safety goals [LL23, LL24]. Consequently, the mapping we present should be read as a current best assignment grounded in today's understanding, not as a timeless prescription.

A natural question that may arise is "If I have a CHF with these properties, in which applications can I use it?", and the answer is more subtle than a static lookup table might suggest: as the literature on redactable blockchains illustrates, our understanding of the "right" notion can evolve over time. Early designs for redactable ledgers were instantiated under enhanced collision resistance, which was later shown to be malleable in this setting; more recent work argues that full collision resistance (against an adversary with access to an adaptation oracle) is the appropriate baseline for predecessor links in permissionless ledgers [AMVA17, DSSS19, DSS20, LL23]. Consequently, any mapping from applications to CR/IND levels should be read as a *snapshot of current best practice*, not as a fixed prescription.

At a high level, our study suggests the following patterns. Primitives whose whole point is *authorized editing after issuance*, such as sanitizable and trapdoor-sanitizable signatures, edit-friendly image and document authentication, redactable blockchains, and chameleon authentication trees [ACDMT05, CLM08, XCY+21, AMVA17, DSSS19, SS12], primarily rely on the collision-resistance ladder: weak collision resistance plus key-exposure freeness suffices for basic "edit without forking" workflows, while stronger CR levels (up to full) become necessary once edits are performed in adversarially rich environments such as public ledgers. In contrast, applications driven by *non-transferability or designated verification* (chameleon signatures, universal designated-verifier signatures, sealed-bid auctions) and by *privacy with accountability* (group-based attestation, VANET beacons, mobile handover authentication) [KR98, SBWP03, AdM04a, ADK10, GZX14, CJ10] mostly hinge on having *some* robust CR notion and carefully designed protocols around it; indistinguishability of hash versus adapted openings is rarely modelled explicitly and is often implicit in the surrounding encryption or anonymization layer. Finally, when CHFs are used as a technical lever to harden existing primitives, for example PKE/(H)IBE transforms, tightly secure signatures, on-line/off-line signatures, or one-time signatures [Zha07, BKKP15, ST01, Moh10], the CHF serves as a binder that inherits its security goals from the wrapped primitive: weak or standard CR with key-exposure freeness typically suffices, and the transform does not fundamentally depend on a particular IND level of the CHF itself.

We therefore see the CR level (weak(key-exposure freeness), standard, enhanced, full), as the main axis on which most current applications position themselves, while indistinguishability variants become explicit mainly in a narrower line of work where the CHF itself is the privacy layer, such as invisible or policy-based sanitizable signatures instantiated from CH/CHET and policy-based chameleon hashes with (strong or full) indistinguishability guarantees [CDH+17, SS20]; in these settings, one requires that adapted openings (and sometimes sanitized views) be computationally indistinguishable from fresh

hashes, so that edits or sanitization do not leak which version is the "real" one.

5.2 Addressed research questions and open questions

RQ1 (Security frameworks). What should a CHF be secure against, and how do the definitions relate? We organize the landscape along two independent axes: collision resistance and indistinguishability. For collision resistance, four baseline notions appear in the literature, ordered by adversarial power: weak (W-CollRes), standard (S-CollRes), enhanced (E-CollRes), and full (F-CollRes). The ladder reflects which oracles the adversary receives and what auxiliary information is available. The order is partial: S-CollRes and E-CollRes are incomparable in general, while F-CollRes strictly implies both. Indistinguishability forms a separate ladder (Ind, S-Ind, E-Ind, F-Ind) that is orthogonal to collision resistance; uniqueness is also orthogonal and can be required when applications benefit from canonical encodings. Consequently, the target collision resistance notion and the desired indistinguishability level should be fixed by the application threat model and auditability requirements, not by the algebraic setting or the hardness assumption alone. Many refinements exist in the literature, often tailored to use cases: key-exposure-free variants, outsider versus insider formulations, tagged and tagged-full flavors, selective versus adaptive patterns, and interface-driven tweaks. A complete catalog would add complexity with limited comparative value, so we retain only the variants that recur across applications and that materially affect design choices in our taxonomy and maps.

RQ2 (Constructions). Constructions separate cleanly along three axes. Authorization: identity–based single holder, policy/attribute–based, and threshold/committee. Trapdoor structure: single long–term, per–hash ephemeral (e.g., CHET–style), dual long–term or delegated variants. Functional/algebraic features: homomorphic or structure–preserving hashing, tagged and time–modifiable variants, accountable and leakage–resilient refinements. This taxonomy explains how policy or threshold layers compose with a base CHF without collapsing the collision notion, and how ephemeral or dual trapdoors scope adaptation while protecting long–term keys.

RQ3 (Underlying hard problems). Classical schemes predominantly rely on discrete–log and RSA families, with pairing assumptions (e.g., DLIN, SXDH) when additional algebraic structure is required. Post–quantum lines exist over lattices (LWE/SIS, and LPN in particular) and, to a lesser extent, codes; they remain comparatively sparse and heavier in parameters. Black–box routes to F-CollRes often pass through commitments or PKE plus succinct proofs; lattice and LPN instantiations reach F-CollRes, but typically with NIZKs. A tagged approach achieves a tag–scoped full collision resistance (tF-CR); under a single–tag mode this coincides with full CR for that tag domain, and avoids NIZKs. A 2025 design attains full CR without NIZKs [CDJ+25] by removing tags and employing commitments satisfying DEE, at the cost of uniqueness. The trade–off suggests that removing proof systems shifts burden to commitment programmability and may constrain uniqueness and IND levels.

RQ4 (Applications and constraints). Our application overview shows that real systems using CHFs cluster into a few recurring families: *authorised editing after issuance* (sanitizable and trapdoor-sanitizable signatures, edit-friendly image/document authentication, redactable blockchains, chameleon authentication trees), *non-transferability and privacy-with-accountability* (chameleon signatures, universal designated-verifier signatures, sealed-bid auctions, group-based attestation, VANET beacons, mobile handover), and *wrapper uses* that harden existing primitives (PKE/(H)IBE transforms, tightly secure signatures, online/offline and one-time signatures). Across these families, the main knob is

the collision-resistance level together with key-exposure freeness: weak or standard CR is enough for local editing and wrapper constructions, whereas enhanced or full CR becomes necessary in highly adversarial environments such as public ledgers. Indistinguishability levels show up explicitly only in a narrower class of strongly privacy-aware designs; for most deployed protocols, once the CR level and coin/authorization model are fixed, the compatible application types are essentially those captured by our applications taxonomy.

5.3 Open directions

1) Post quantum CHFs beyond lattices. There is a clear gap in code–based CHFs at strong collision levels (even construction wise). A concrete road is: identify CHF–friendly code–based trapdoors that admit efficient controlled adaptation; pair them with commitments or PKE whose security reduces to SD (or rank–metric analogues) without random oracles; and prototype F-CollRes instances with competitive key and hash sizes.

2) Full collision resistance without NIZKs, and the uniqueness gap. Recent lattice-based CHFs achieve full collision resistance without NIZKs by relying on trapdoor commitments with distributional equivalence of equivocation (DEE); normal indistinguishability follows from DEE, and full collision resistance from SIS. However, these designs do not achieve uniqueness, and the evidence points to an inherent obstacle for PreSample/Gaussian-based approaches. [CDJ+25]A concrete open problem is to recover uniqueness while keeping full collision resistance and normal indistinguishability under lattice assumptions. Two possible directions (both speculative) are: exploring commitment or trapdoor tools that allow equivocation without enabling multiple openings for the same digest, and rethinking the CHF interface so that adaptation randomness is tied to a committed state in a one-way fashion. We treat both as open leads rather than claims.

3) Raising indistinguishability in NIZK–free designs. Tagged schemes avoid NIZKs but confine guarantees to a tag domain; DEE–based schemes drop tags but currently settle for weaker IND (normal) or lose uniqueness. A systematic target is S-/E-/F-Ind in NIZK–free settings. Candidate tools include simulation–extractable commitments with restricted programmability, domain–separated coins derived from PRFs under public keys, and verifiable re–randomization that does not leak linkability. The goal is a hierarchy of indistinguishability implications that survives the removal of proofs.

4) Tag–to–untag lifting with preserved properties. It is natural to ask whether a generic transformation can lift tF-CR (single–tag or few–tag regimes) to untagged F-CR while preserving uniqueness and at least S-Ind. Necessary conditions likely include pseudorandom domain separation for tags, non–malleable commitment to coins, and independence between tag space and message space.

5) Governance–aware CHFs. Applications increasingly require key rotation, delegation, revocation, and accountable adaptation. Formal models should capture dynamic authorization (policy updates, quorum changes) and audit requirements without weakening the chosen collision notion. Threshold and multi–authority CHFs with accountable transcripts, and leakage–resilient trapdoor sharing, are promising design points.

6) Benchmarks and standard interfaces. A common evaluation suite would improve comparability across models: reference APIs for public–coin and secret–coin interfaces, test vectors, and performance harnesses measuring hashing, checking, and adaptation separately. Reporting should separate fixed–cost (key generation) and per–edit costs, and include memory footprints for r or ξ storage. Such artifacts would support reproducible claims and clarify practical trade–offs between classical and PQ lines.

Summary. The landscape supports precise design choices per application, yet key frontiers remain: PQ strength beyond lattices (especially code–based), full CR with uniqueness, and higher indistinguishability without NIZKs. Progress on any of these would

materially improve deployability in redactable ledgers, sanitizable signatures, accountable data structures and any application case that may arise.

6 Conclusion

Twenty-seven years after their introduction, there was no dedicated survey or SoK on chameleon hash functions. This work fills that gap by reading the literature through a single lens: who may adapt, which trapdoor structure enables adaptation, and which functional or algebraic features are actually needed by applications. The result is a curated corpus, a compact taxonomy, and a chronological map that place definitions and relations in one place and make the main trade-offs explicit. Our contributions are twofold. First, we compile and annotate a corpus of CHF works along three axes (authorization model, trapdoor structure, functional and algebraic features), together with the underlying assumptions and the target collision and indistinguishability notions. Second, we provide a unified historical narrative and technical maps that connect these axes to concrete use cases, so choices follow from application constraints rather than from isolated constructions. The scope is broad but not exhaustive. We favor comparability and clear design choices over full parameter accounting; adjacent primitives are included only when required to explain recent lines. The dataset and taxonomy are intended to be reusable and to support replication. In practice, the design path is simple. Choose the collision notion first. Select the authorization layer and the trapdoor arrangement that match governance and audit requirements. Add features only when they provide measurable benefit for the target application. The discussion records the open directions that remain.

References

[ACDMT05] Giuseppe Ateniese, Daniel H Chou, Breno De Medeiros, and Gene Tsudik. Sanitizable signatures. In *European Symposium on Research in Computer Security*, pages 159–177. Springer, 2005. URL: URL:https://ics.uci.edu/~gts/paps/esorics-2005.pdf,doi:10.1007/11555827_10.

[ADK10] Sami Alsouri, Özgür Dagdelen, and Stefan Katzenbeisser. Group-based attestation: Enhancing privacy and management in remote attestation. In *Trust and Trustworthy Computing – TRUST 2010*, volume 6101 of *Lecture Notes in Computer Science*, pages 63–77, Berlin, Heidelberg, 2010. Springer. URL: https://download.hrz.tu-darmstadt.de/pub/FB20/Dekanat/Publikationen/CASED/Group-Based_Attestation_Alsouri-Dagdelen-Katzenbeisser.pdf, doi:10.1007/978-3-642-13869-0_5.

[AdM03] Giuseppe Ateniese and Breno de Medeiros. Identity-based chameleon hash and applications. Technical Report 2003/184, Cryptology ePrint Archive, 2003. URL: https://eprint.iacr.org/2003/167.pdf.

[AdM04a] Giuseppe Ateniese and Breno de Medeiros. Identity-based chameleon hash and applications. In *Financial Cryptography – FC 2004*, volume 3110 of *Lecture Notes in Computer Science*, pages 164–180, Berlin, Heidelberg, 2004. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-540-27809-2_19, doi:10.1007/978-3-540-27809-2_19.

[ADM04b] Giuseppe Ateniese and Breno De Medeiros. On the key exposure problem in chameleon hashes. In *International Conference on Security in Communication Networks*, pages 165–179. Springer, 2004. URL: https://link.springer.

com/chapter/10.1007/978-3-540-30598-9_12, doi:10.1007/978-3-540-30598-9_12.

[Ajt96] Miklós Ajtai. Generating hard instances of lattice problems. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing (STOC '96)*, pages 99–108, Philadelphia, Pennsylvania, USA, 1996. ACM. URL: https://dl.acm.org/doi/10.1145/237814.237838, doi:10.1145/237814.237838.

[AMVA17] Giuseppe Ateniese, Bernardo Magri, Daniele Venturi, and Ewerton R. Andrade. Redactable blockchain – or – rewriting history in bitcoin and friends. In *2017 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 111–126, Paris, France, April 2017. IEEE. doi:10.1109/EuroSP.2017.37.

[BBS04] Dan Boneh, Xavier Boyen, and Hovav Shacham. Short group signatures. In *Advances in Cryptology – CRYPTO 2004*, volume 3152 of *Lecture Notes in Computer Science*, pages 41–55, Berlin, Heidelberg, 2004. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-540-28628-8_3, doi:10.1007/978-3-540-28628-8_3.

[BCC88] Gilles Brassard, David Chaum, and Claude Crépeau. Minimum disclosure proofs of knowledge. *Journal of Computer and System Sciences*, 37(2):156–189, 1988. URL: https://www.sciencedirect.com/science/article/pii/0022000088900050.

[BCG07] Emmanuel Bresson, Dario Catalano, and Rosario Gennaro. Improved on-line/off-line threshold signatures. In *Public Key Cryptography – PKC 2007*, volume 4450 of *Lecture Notes in Computer Science*, pages 217–232, Berlin, Heidelberg, 2007. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-540-71677-8_15, doi:10.1007/978-3-540-71677-8_15.

[BCG14] Dan Boneh and Henry Corrigan-Gibbs. Bivariate polynomials modulo composites and their applications. In *Advances in Cryptology – ASIACRYPT 2014*, volume 8873 of *Lecture Notes in Computer Science*, pages 42–62, Berlin, Heidelberg, 2014. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-319-19962-7_13, doi:10.1007/978-3-319-19962-7_13.

[BGdMM05] Lucas Ballard, Matthew Green, Breno de Medeiros, and Fabian Monrose. Correlation-resistant storage via keyword-searchable encryption. Technical Report 2005/417, Cryptology ePrint Archive, 2005. URL: https://eprint.iacr.org/2005/417.

[BGM+19] Olivier Blazy, Philippe Gaborit, Dang Truong Mac, Ayoub Otmani, and Jean-Pierre Tillich. A code-based signature scheme in the standard model. In *Proceedings of the 11th International Workshop on Coding and Cryptography (WCC 2019)*, Saint-Jacut-de-la-Mer, France, 2019. Extended abstract. URL: https://www.lebesgue.fr/sites/default/files/proceedings_WCC/WCC_2019_paper_62.pdf.

[BKKP15] Olivier Blazy, Saqib A. Kakvi, Eike Kiltz, and Jiaxin Pan. Tightly-secure signatures from chameleon hash functions. In *Public-Key Cryptography – PKC 2015*, volume 9020 of *Lecture Notes in Computer Science*, pages 256–279, Berlin, Heidelberg, 2015. Springer. URL: https://www.iacr.org/archive/pkc2015/90200149/90200149.pdf, doi:10.1007/978-3-662-46447-2_12.

[BKW03] Avrim Blum, Adam Kalai, and Hal Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *Journal of the ACM (JACM)*, 50(4):506–519, July 2003. doi:10.1145/792538.792543.

[BMvT78] Elwyn R. Berlekamp, Robert J. McEliece, and Henk C. A. van Tilborg. On the inherent intractability of certain coding problems. *IEEE Transactions on Information Theory*, 24(3):384–386, May 1978. doi:10.1109/TIT.1978.1055873.

[BNPS03] Mihir Bellare, Chanathip Namprempre, David Pointcheval, and Michael Semanko. The one-more-RSA-inversion problems and the security of chaum's blind signature scheme. *Journal of Cryptology*, 16(3):185–215, 2003. URL: https://link.springer.com/article/10.1007/s00145-002-0120-1, doi:10.1007/s00145-002-0120-1.

[Bon98] Dan Boneh. The decision diffie-hellman problem. In *Algorithmic Number Theory – ANTS-III*, volume 1423 of *Lecture Notes in Computer Science*, pages 48–63, Berlin, Heidelberg, 1998. Springer. URL: https://link.springer.com/chapter/10.1007/BFb0054851, doi:10.1007/BFb0054851.

[BP97] Niko Barić and Birgit Pfitzmann. Collision-free accumulators and fail-stop signature schemes without trees. In *Advances in Cryptology – EUROCRYPT '97*, volume 1233 of *Lecture Notes in Computer Science*, pages 480–494, Berlin, Heidelberg, 1997. Springer. URL: https://link.springer.com/chapter/10.1007/3-540-69053-0_33, doi:10.1007/3-540-69053-0_33.

[BR08] Mihir Bellare and Todor Ristov. Hash functions from sigma protocols and improvements to VSH. In *Advances in Cryptology – ASIACRYPT 2008*, volume 5350 of *Lecture Notes in Computer Science*, pages 324–341, Berlin, Heidelberg, 2008. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-540-89255-7_9, doi:10.1007/978-3-540-89255-7_9.

[BR14] Mihir Bellare and Todor Ristov. A characterization of chameleon hash functions and new, efficient designs. *Journal of Cryptology*, 27(4):799–823, October 2014. URL: https://link.springer.com/article/10.1007/s00145-013-9155-8, doi:10.1007/s00145-013-9155-8.

[CCTY24] Kwan Yin Chan, Liqun Chen, Yangguang Tian, and Tsz Hon Yuen. Reconstructing chameleon hash: Full security and the multi-party setting. In *Proceedings of the 19th ACM Asia Conference on Computer and Communications Security (ASIA CCS '24)*, pages 1066–1081, Singapore, July 2024. ACM. URL: https://dl.acm.org/doi/10.1145/3634737.3656291, doi:10.1145/3634737.3656291.

[CDH+17] Jan Camenisch, Maria Dubovitskaya, Kai-Hsin Huang, Gregory Neven, and Ivan Osipkov. Chameleon-hashes with ephemeral trapdoors: And applications to invisible sanitizable signatures. In *Public-Key Cryptography – PKC 2017*, volume 10174 of *Lecture Notes in Computer Science*, pages 152–182, Berlin, Heidelberg, 2017. Springer. URL: https://eprint.iacr.org/2017/011.pdf, doi:10.1007/978-3-662-54388-7_6.

[CDJ+25] Sebastian Clermont, Samed Düzlü, Christian Janson, Laurens Porzenheim, and Patrick Struck. Lattice-based sanitizable signature schemes: Chameleon hash functions and more. In *Post-Quantum Cryptography – PQCrypto 2025*, volume 15577 of *Lecture Notes in Computer Science*, pages 278–311, Cham, 2025. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-031-86599-2_10, doi:10.1007/978-3-031-86599-2_10.

[CHKP12] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. *Journal of Cryptology*, 25(4):601–639, October 2012. URL: https://link.springer.com/chapter/10.1007/978-3-642-13190-5_27, doi:10.1007/978-3-642-13190-5_27.

[Cho08] Sherman S. M. Chow. Multi-designated verifiers signatures revisited. *International Journal of Network Security (IJNS)*, 7(3):348–357, May 2008. URL: http://ijns.jalaxy.com.tw/contents/ijns-v7-n3/ijns-2008-v7-n3-p348-357.pdf.

[CIS17] Santosh Chandrasekhar, Ahmed Ibrahim, and Mukesh Singhal. A novel access control protocol using proxy signatures for cloud-based health information exchange. *Computers & Security*, 67:73–88, July 2017. URL: https://www.sciencedirect.com/science/article/abs/pii/S0167404817300342, doi:10.1016/j.cose.2017.02.008.

[CJ10] Jaeduck Choi and Souhwan Jung. A handover authentication using credentials based on chameleon hashing. *IEEE Communications Letters*, 14(1):54–56, January 2010. URL: https://ieeexplore.ieee.org/document/5370272, doi:10.1109/LCOMM.2010.01.091607.

[CLLS10] Santosh Chandrasekhar, Shengbao Li, Rui Li, and Willy Susilo. Efficient proxy signatures based on trapdoor hash functions. *IET Information Security*, 4(4):322–332, December 2010. URL: https://digital-library.theiet.org/doi/abs/10.1049/iet-ifs.2009.0204, doi:10.1049/iet-ifs.2009.0204.

[CLM08] Sébastien Canard, Fabien Laguillaumie, and Michel Milhau. Trapdoor sanitizable signatures and their application to content protection. In *Applied Cryptography and Network Security – ACNS 2008*, volume 5037 of *Lecture Notes in Computer Science*, pages 258–276, Berlin, Heidelberg, 2008. Springer. URL: https://www.lirmm.fr/~laguillaum/acns2008web.pdf, doi:10.1007/978-3-540-68914-0_16.

[CLS05] Scott Contini, Arjen K. Lenstra, and Ron Steinfeld. VSH, an efficient and provable collision resistant hash function. Technical Report 2005/193, Cryptology ePrint Archive, 2005. URL: https://eprint.iacr.org/2005/193.

[CLS06] Scott Contini, Arjen K. Lenstra, and Ron Steinfeld. Vsh, an efficient and provable collision-resistant hash function. In *Advances in Cryptology – EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, pages 165–182, Berlin, Heidelberg, 2006. Springer. URL: https://link.springer.com/chapter/10.1007/11761679_11, doi:10.1007/11761679_11.

[CMTW06] Chris Crutchfield, David Molnar, David Turner, and David Wagner. Generic on-line/off-line threshold signatures. In *Public Key Cryptography – PKC 2006*, volume 3958 of *Lecture Notes in Computer Science*, pages 58–74, Berlin, Heidelberg, 2006. Springer. URL: https://link.springer.com/chapter/10.1007/11745853_5, doi:10.1007/11745853_5.

[CS00] Ronald Cramer and Victor Shoup. Signature schemes based on the strong RSA assumption. *ACM Transactions on Information and System Security (TISSEC)*, 3(3):161–185, August 2000. URL: https://dl.acm.org/doi/pdf/10.1145/357830.357847, doi:10.1145/357830.357847.

[CS14] Santosh Chandrasekhar and Mukesh Singhal. Multi-trapdoor hash functions and their applications in network security. In *2014 IEEE Conference on Communications and Network Security (CNS)*, pages 209–217, San Francisco, CA, USA, October 2014. IEEE. URL: https://ieeexplore.ieee.org/document/6997516, doi:10.1109/CNS.2014.6997516.

[CTZD10] Xiaofeng Chen, Haibo Tian, Fangguo Zhang, and Yong Ding. Comments and improvements on key-exposure free chameleon hashing based on factoring. In *Information Security and Cryptology – Inscrypt 2010*, volume 6584 of *Lecture Notes in Computer Science*, pages 415–426, Berlin, Heidelberg, 2010. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-642-21518-6_29, doi:10.1007/978-3-642-21518-6_29.

[CvA90] David Chaum and Hans van Antwerpen. Undeniable signatures. In *Advances in Cryptology – CRYPTO '89*, volume 435 of *Lecture Notes in Computer Science*, pages 212–216, New York, NY, USA, 1990. Springer. URL: https://link.springer.com/chapter/10.1007/0-387-34805-0_20, doi:10.1007/0-387-34805-0_20.

[CZK04] Xiaofeng Chen, Fangguo Zhang, and Kwangjo Kim. Chameleon hashing without key exposure. In *Information Security – ISC 2004*, volume 3225 of *Lecture Notes in Computer Science*, pages 87–98, Berlin, Heidelberg, 2004. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-540-30185-7_8, doi:10.1007/978-3-540-30185-7_8.

[CZS+10] Xiaofeng Chen, Fangguo Zhang, Willy Susilo, Haibo Tian, Jin Li, and Kwangjo Kim. Identity-based chameleon hash scheme without key exposure. In Ron Steinfeld and Philip Hawkes, editors, *ACISP 10: 15th Australasian Conference on Information Security and Privacy*, volume 6168 of *Lecture Notes in Computer Science*, pages 200–215, Sydney, NSW, Australia, July 5–7, 2010. Springer Berlin Heidelberg, Germany. doi:10.1007/978-3-642-14081-5_13.

[CZSM07] Xiaofeng Chen, Fangguo Zhang, Willy Susilo, and Yi Mu. Efficient generic On-Line/Off-Line signatures without key exposure. In *Applied Cryptography and Network Security – ACNS 2007*, volume 4521 of *Lecture Notes in Computer Science*, pages 18–30, Berlin, Heidelberg, 2007. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-540-72738-5_2, doi:10.1007/978-3-540-72738-5_2.

[CZT+08] Xiaofeng Chen, Fangguo Zhang, Haibo Tian, Baodian Wei, Willy Susilo, Yi Mu, Hyunrok Lee, and Kwangjo Kim. Efficient generic On-Line/Off-Line (threshold) signatures without key exposure. *Information Sciences*, 178(21):4192–4203, November 2008. URL: https://www.sciencedirect.com/science/article/pii/S0020025508002569, doi:10.1016/j.ins.2008.06.022.

[CZW+09] Xiaofeng Chen, Fangguo Zhang, Baodian Wei, Willy Susilo, and Yi Mu. Key-exposure free chameleon hashing and signatures based on discrete logarithm systems. Technical Report 2009/035, Cryptology ePrint Archive, 2009. URL: https://eprint.iacr.org/2009/035.pdf.

[CZW+11] Xiaofeng Chen, Fangguo Zhang, Baodian Wei, Willy Susilo, and Yi Mu. Discrete logarithm based chameleon hashing and signatures without key exposure. *Computers & Electrical Engineering*, 37(4):614–623, July 2011.

URL: https://www.sciencedirect.com/science/article/abs/pii/S0045790611000383, doi:10.1016/j.compeleceng.2011.03.011.

[CZW+14] Xiaofeng Chen, Fangguo Zhang, Baodian Wei, Willy Susilo, and Yi Mu. Identity-based chameleon hashing and signatures without key exposure. *Information Sciences*, 265:198–210, May 2014. URL: https://www.sciencedirect.com/science/article/abs/pii/S0020025513008700?via%3Dihub, doi:10.1016/j.ins.2013.12.020.

[DH76] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. In *IEEE Transactions on Information Theory*, volume 22, pages 644–654, 1976. URL: https://ieeexplore.ieee.org/document/1055638, doi:10.1109/TIT.1976.1055638.

[DM14] Léo Ducas and Daniele Micciancio. Improved short lattice signatures in the standard model. In *Advances in Cryptology – CRYPTO 2014*, volume 8616 of *Lecture Notes in Computer Science*, pages 335–352, Berlin, Heidelberg, 2014. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-662-44371-2_19, doi:10.1007/978-3-662-44371-2_19.

[DSS20] David Derler, Kai Samelin, and Daniel Slamanig. Bringing order to chaos: The case of collision-resistant chameleon-hashes. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020: 23rd International Conference on Theory and Practice of Public Key Cryptography, Part I*, volume 12110 of *Lecture Notes in Computer Science*, pages 462–492, Edinburgh, UK, May 4–7, 2020. Springer, Cham, Switzerland. doi:10.1007/978-3-030-45374-9_16.

[DSSS19] David Derler, Kai Samelin, Daniel Slamanig, and Christoph Striecks. Fine-grained and controlled rewriting in blockchains: Chameleon-hashing gone attribute-based. In *ISOC Network and Distributed System Security Symposium – NDSS 2019*, San Diego, CA, USA, February 24–27, 2019. The Internet Society. doi:10.14722/ndss.2019.23066.

[DSSS20] David Derler, Kai Samelin, Daniel Slamanig, and Christoph Striecks. Fully collision-resistant chameleon-hashes from simpler and post-quantum assumptions. In *Security and Cryptography for Networks – SCN 2020*, volume 12238 of *Lecture Notes in Computer Science*, pages 213–233, Cham, 2020. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-030-57990-6_21, doi:10.1007/978-3-030-57990-6_21.

[DWL+23] Pengfei Duan, Hao Wang, Zhi Li, Yi Mu, and Robert H. Deng. Policy-based chameleon hash with black-box traceability for redactable blockchain in iot. *Electronics*, 12(7):1646, April 2023. URL: https://www.mdpi.com/2079-9292/12/7/1646, doi:10.3390/electronics12071646.

[DWWG24] Junke Duan, Wei Wang, Licheng Wang, and Lize Gu. Controlled redactable blockchain based on t-times chameleon hash and signature. *IEEE Transactions on Information Forensics and Security*, 19:7560–7572, 2024. URL: https://ieeexplore.ieee.org/document/10620354, doi:10.1109/TIFS.2024.3436925.

[ElG85] Taher ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31(4):469–472, July 1985. URL: https://ieeexplore.ieee.org/document/1057074, doi:10.1109/TIT.1985.1057074.

[Fre10] David Mandell Freeman. Converting pairing-based cryptosystems from composite-order groups to prime-order groups. In *Advances in Cryptology – EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 44–61, Berlin, Heidelberg, 2010. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-642-13190-5_3, doi:10.1007/978-3-642-13190-5_3.

[GM82] Shafi Goldwasser and Silvio Micali. Probabilistic encryption & how to play mental poker keeping secret all partial information. In *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing (STOC)*, page 365–377. Association for Computing Machinery, 1982. URL: https://dl.acm.org/doi/10.1145/3335741.3335749, doi:10.1145/3335741.3335749.

[Gro10] Jens Groth. Short pairing-based non-interactive zero-knowledge arguments. In *Advances in Cryptology – ASIACRYPT 2010*, volume 6477 of *Lecture Notes in Computer Science*, pages 321–340, Berlin, Heidelberg, 2010. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-642-17373-8_19, doi:10.1007/978-3-642-17373-8_19.

[GWX07] Wei Gao, Xue-Li Wang, and Dong-Qing Xie. Chameleon hashes without key exposure based on factoring. *Journal of Computer Science and Technology*, 22(1):109–113, January 2007. URL: https://link.springer.com/article/10.1007/s11390-007-9015-9, doi:10.1007/s11390-007-9015-9.

[GZX14] Song Guo, Deze Zeng, and Yang Xiang. Chameleon hashing for secure and privacy-preserving vehicular communications. *IEEE Transactions on Parallel and Distributed Systems*, 25(11):2794–2803, November 2014. URL: https://ieeexplore.ieee.org/document/6654169, doi:10.1109/TPDS.2013.277.

[HHL10] Lein Harn, Wen-Jung Hsin, and Changlu Lin. Efficient On-Line/Off-Line signature schemes based on multiple-collision trapdoor hash families. *The Computer Journal*, 53(9):1478–1484, November 2010. URL: https://academic.oup.com/comjnl/article/53/9/1478/363044, doi:10.1093/comjnl/bxp044.

[HW09] Susan Hohenberger and Brent Waters. Realizing hash-and-sign signatures under standard assumptions. In *Advances in Cryptology – EUROCRYPT 2009*, volume 5479 of *Lecture Notes in Computer Science*, pages 333–350, Berlin, Heidelberg, 2009. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-642-01001-9_19, doi:10.1007/978-3-642-01001-9_19.

[HZM+19a] Ke Huang, Xiaosong Zhang, Yi Mu, Fatemeh Rezaeibagha, Xiaojiang Du, and Nadra Guizani. Achieving intelligent trust-layer for internet-of-things via self-redactable blockchain. *IEEE Transactions on Industrial Informatics*, 2019. Early access. URL: https://ieeexplore.ieee.org/document/8848417, doi:10.1109/TII.2019.2943331.

[HZM+19b] Ke Huang, Xiaosong Zhang, Yi Mu, Xiaofen Wang, Guomin Yang, Xiaojiang Du, Fatemeh Rezaeibagha, Qi Xia, and Mohsen Guizani. Building redactable consortium blockchain for industrial internet-of-things. *IEEE Transactions on Industrial Informatics*, 15(6):3670–3679, June 2019. doi:10.1109/TII.2019.2901011.

[JCH+22] Meng Jia, Jing Chen, Kun He, Ruiying Du, Li Zheng, Mingxi Lai, Donghui Wang, and Fei Liu. Redactable blockchain from decentralized chameleon hash

functions. *IEEE Transactions on Information Forensics and Security*, 17:2771–2783, July 2022. URL: https://ieeexplore.ieee.org/document/9833529, doi:10.1109/TIFS.2022.3192716.

[JLL+21] Yanxue Jia, Liwei Liu, Ximeng Liu, Robert H. Deng, and Jianfeng Ma. Redactable blockchain supporting supervision and self-management. In *Proceedings of the 2021 ACM Asia Conference on Computer and Communications Security (ASIACCS '21)*, pages 345–356, Hong Kong, China, June 2021. ACM. URL: https://dl.acm.org/doi/10.1145/3433210.3453091, doi:10.1145/3433210.3453091.

[KDS20] Mojtaba Khalili, Mohammad Dakhilalian, and Willy Susilo. Efficient chameleon hash functions in the enhanced collision resistant model. *Information Sciences*, 510:155–164, 2020. URL: https://www.sciencedirect.com/science/article/abs/pii/S002002551930831X, doi:10.1016/j.ins.2019.09.001.

[KMS18] Stephan Krenn, Daniel Müller, and Kai Samelin. Chameleon-hashes with dual long-term trapdoors and their applications. In *Progress in Cryptology – AFRICACRYPT 2018*, volume 10831 of *Lecture Notes in Computer Science*, pages 111–129, Cham, 2018. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-319-89339-6_2, doi:10.1007/978-3-319-89339-6_2.

[KR98] Hugo Krawczyk and Tal Rabin. Chameleon hashing and signatures. In *Proceedings of the Symposium on Network and Distributed System Security*, 1998. URL: https://www.ndss-symposium.org/wp-content/uploads/2017/09/Chameleon-Signatures-paper-Hugo-Krawczyk.pdf.

[LAZ19] Xingye Lu, Man Ho Au, and Zhenfei Zhang. Raptor: A practical lattice-based (linkable) ring signature. In *Applied Cryptography and Network Security – ACNS 2019*, volume 11464 of *Lecture Notes in Computer Science*, pages 110–130, Cham, 2019. Springer. Bogotá, Colombia, June 5–7, 2019. URL: https://link.springer.com/chapter/10.1007/978-3-030-21568-2_6, doi:10.1007/978-3-030-21568-2_6.

[LDW13] Junzuo Lai, Xuhua Ding, and Yongdong Wu. Accountable trapdoor sanitizable signatures. In *Information Security Practice and Experience – ISPEC 2013*, volume 7863 of *Lecture Notes in Computer Science*, pages 17–33, Berlin, Heidelberg, 2013. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-642-38033-4_9, doi:10.1007/978-3-642-38033-4_9.

[LJMP90] Arjen K. Lenstra, Hendrik W. Lenstra Jr., Mark S. Manasse, and John M. Pollard. The number field sieve. In *Proceedings of the Twenty-Second Annual ACM Symposium on Theory of Computing (STOC '90)*, pages 564–572, Baltimore, Maryland, USA, 1990. ACM. URL: https://dl.acm.org/doi/10.1145/100216.100295, doi:10.1145/100216.100295.

[LL23] Yiming Li and Shengli Liu. Tagged chameleon hash from lattices and application to redactable blockchain. Technical Report 2023/774, Cryptology ePrint Archive, 2023. URL: https://eprint.iacr.org/2023/774.

[LL24] Yiming Li and Shengli Liu. Tagged chameleon hash from lattices and application to redactable blockchain. In *Public-Key Cryptography – PKC 2024*, 2024. URL: https://link.springer.com/chapter/10.1007/978-3-031-57725-3_10, doi:10.1007/978-3-031-57725-3_10.

[LLLS20] Cheng-Jun Lin, Baocang Liu, Xiaoyuan Liu, and Willy Susilo. Linearly homomorphic signatures from lattices. *The Computer Journal*, 63(12):1871–1885, December 2020. URL: https://academic.oup.com/comjnl/article/63/12/1871/5874290, doi:10.1093/comjnl/bxaa034.

[LLM+24] Nan Li, Yingjiu Li, Mark Manulis, Yangguang Tian, and Guomin Yang. Practical and secure policy-based chameleon hash for redactable blockchains. *The Computer Journal*, 67(11):3128–3139, November 2024. URL: https://academic.oup.com/comjnl/article-abstract/67/11/3128/7723759, doi:10.1093/comjnl/bxae075.

[LSW25] Cong Li, Qingni Shen, and Zhonghai Wu. Redactable blockchain from decentralized chameleon hash functions, revisited. *IEEE Transactions on Computers*, 74(6):1911–1920, 2025. URL: https://ieeexplore.ieee.org/document/10899848, doi:10.1109/TC.2025.3544878.

[LXD+22] Cong Li, Zhikang Xie, Jisheng Dong, Qingni Shen, and Yuejian Fang. Efficient identity-based chameleon hash for mobile devices. In *ICASSP 2022 – IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3089–3093, Singapore, May 2022. IEEE. URL: https://ieeexplore.ieee.org/document/9746617, doi:10.1109/ICASSP43922.2022.9746617.

[LXZ+23] Jianhao Li, Qiuliang Xu, Zhenfeng Zhang, Zhiming Zheng, and Robert H. Deng. Wolverine: A scalable and transaction-consistent redactable permissionless blockchain. *IEEE Transactions on Information Forensics and Security*, 18:1653–1666, March 2023. URL: https://ieeexplore.ieee.org/document/10044703, doi:10.1109/TIFS.2023.3245406.

[LZ25] Zikai Liu and Zhen Zhao. Time-modifiable chameleon hash for building redactable blockchains. volume 14, 2025. URL: https://www.mdpi.com/2079-9292/14/15/2995, doi:10.3390/electronics14152995.

[LZZ+16] Weiran Liu, Rui Zhang, Jiawei Zhang, Dong Zheng, Chun Chen, and Yixian Yang. Online/offline public-index predicate encryption for fine-grained mobile access control. In *Computer Security – ESORICS 2016*, Lecture Notes in Computer Science, Cham, 2016. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-319-45741-3_30, doi:10.1007/978-3-319-45741-3_30.

[Moh10] Payman Mohassel. One-time signatures and chameleon hash functions. In *Selected Areas in Cryptography – SAC 2010*, volume 6544 of *Lecture Notes in Computer Science*, pages 302–319, Berlin, Heidelberg, 2010. Springer. URL: https://link.springer.com/content/pdf/10.1007/978-3-642-19574-7_21.pdf, doi:10.1007/978-3-642-19574-7_21.

[MZZ+22] Jinhua Ma, Rongzuo Zhang, Zhenfeng Zhang, Qi Li, Shundong Li, and Yixian Yang. Redactable blockchain in decentralized setting. *IEEE Transactions on Information Forensics and Security*, 17, 2022. URL: https://ieeexplore.ieee.org/document/9729109, doi:10.1109/TIFS.2022.3156808.

[Pai99] Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *Advances in Cryptology – EUROCRYPT '99*, volume 1592 of *Lecture Notes in Computer Science*, pages 223–238, Berlin, Heidelberg, 1999. Springer. URL: https://link.springer.com/chapter/10.1007/3-540-48910-X_16, doi:10.1007/3-540-48910-X_16.

[Pat96] Jacques Patarin. Hidden fields equations (HFE) and isomorphisms of polynomials (IP): Two new families of asymmetric algorithms. In *Advances in Cryptology – EUROCRYPT '96*, volume 1070 of *Lecture Notes in Computer Science*, pages 33–48, Berlin, Heidelberg, 1996. Springer. URL: https://link.springer.com/chapter/10.1007/3-540-68339-9_4, doi:10.1007/3-540-68339-9_4.

[PLQZ13] Ping Pan, Shaohui Lv, Guohua Qian, and Fangguo Zhang. Chameleon hash functions and one-time signature schemes from inner automorphism groups. *Fundamenta Informaticae*, 126(1):103–119, 2013. URL: https://dl.acm.org/doi/abs/10.5555/2636469.2636473, doi:10.5555/2636469.2636473.

[PVM21] Gaurav Panwar, Roopa Vishwanathan, and Satyajayant Misra. Retrace: Revocable and traceable blockchain rewrites using attribute-based cryptosystems. In *Proceedings of the 26th ACM Symposium on Access Control Models and Technologies (SACMAT '21)*, pages 103–114. ACM, 2021. URL: https://dl.acm.org/doi/10.1145/3450569.3463565, doi:10.1145/3450569.3463565.

[Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6):1–40, September 2009. URL: https://dl.acm.org/doi/10.1145/1568318.1568324, doi:10.1145/1568318.1568324.

[RKS15] Tim Ruffing, Aniket Kate, and Dominique Schröder. Liar, liar, coins on fire!: Penalizing equivocation by loss of bitcoins. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS'15)*, pages 219–230, Denver, CO, USA, October 2015. ACM. URL: https://dl.acm.org/doi/10.1145/2810103.2813686, doi:10.1145/2810103.2813686.

[RMS07] Qiong Ren, Yi Mu, and Willy Susilo. Mitigating phishing by a new id-based chameleon hash without key exposure. In *AusCERT Asia Pacific Information Technology Security Conference 6th*, 2007. URL: https://ro.uow.edu.au/articles/conference_contribution/Mitigating_phishing_by_a_new_ID-based_chameleon_hash_without_key_exposure/27689556?file=50426505, doi:10779/uow.27689556.v1.

[RSA78] Ronald L. Rivest, Adi Shamir, and Leonard M. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2):120–126, February 1978. URL: https://dl.acm.org/doi/10.1145/359340.359342, doi:10.1145/359340.359342.

[SBWP03] Ron Steinfeld, Laurence Bull, Huaxiong Wang, and Josef Pieprzyk. Universal designated-verifier signatures. In *Advances in Cryptology – ASIACRYPT 2003*, volume 2894 of *Lecture Notes in Computer Science*, pages 523–542, Berlin, Heidelberg, 2003. Springer. URL: https://iacr.org/archive/asiacrypt2003/13_Session12/32_077/UDVS_asia03_final_1.pdf, doi:10.1007/978-3-540-40061-5_33.

[Ser25] SerpAPI Team. Serpapi: Search engine results api, 2025. Accessed: 2025-11-14. URL: https://serpapi.com.

[Sha07] Hovav Shacham. A cramer–shoup encryption scheme from the linear assumption and from progressively weaker linear variants. Technical Report 2007/074, Cryptology ePrint Archive, 2007. URL: https://eprint.iacr.org/2007/074.

[SHL+19] Hong Shu, Yining Hu, Jingyu Lai, Chunxiang Xu, and Rongxing Lu. An aggregate signature scheme based on a trapdoor hash function for the internet of things. *Sensors*, 19(19):4239, October 2019. URL: https://www.mdpi.com/1424-8220/19/19/4239, doi:10.3390/s19194239.

[SM04] Willy Susilo and Yi Mu. Deniable ring authentication revisited. In *Applied Cryptography and Network Security – ACNS 2004*, volume 3089 of *Lecture Notes in Computer Science*, pages 149–163, Berlin, Heidelberg, 2004. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-540-24852-1_11, doi:10.1007/978-3-540-24852-1_11.

[SPW07] Ron Steinfeld, Josef Pieprzyk, and Huaxiong Wang. How to strengthen any weakly unforgeable signature into a strongly unforgeable signature. In *Cryptographers' Track at the RSA Conference*, pages 357–371. Springer, 2007. URL: https://link.springer.com/chapter/10.1007/11967668_23, doi:10.1007/11967668_23.

[SS12] Dominique Schröder and Heike Schröder. Verifiable data streaming. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM CCS 2012: 19th Conference on Computer and Communications Security*, pages 953–964, Raleigh, NC, USA, October 16–18, 2012. ACM Press. doi:10.1145/2382196.2382297.

[SS19] Kai Samelin and Daniel Slamanig. Policy-based sanitizable signatures. Technical Report 2019/410, Cryptology ePrint Archive, 2019. URL: https://eprint.iacr.org/2019/410.

[SS20] Kai Samelin and Daniel Slamanig. Policy-based sanitizable signatures. In *Topics in Cryptology – CT-RSA 2020*, volume 12006 of *Lecture Notes in Computer Science*, pages 538–563, Cham, 2020. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-030-40186-3_23, doi:10.1007/978-3-030-40186-3_23.

[SST05] Katja Schmidt-Samoa and Tsuyoshi Takagi. Paillier's cryptosystem modulo p^2q and its applications to trapdoor commitment schemes. In *Progress in Cryptology – Mycrypt 2005*, volume 3810 of *Lecture Notes in Computer Science*, pages 296–313, Berlin, Heidelberg, 2005. Springer. URL: https://link.springer.com/chapter/10.1007/11554868_21, doi:10.1007/11554868_21.

[ST01] Adi Shamir and Yael Tauman. Improved Online/Offline signature schemes. In *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 355–367, Berlin, Heidelberg, 2001. Springer. URL: https://iacr.org/archive/crypto2001/21390353.pdf, doi:10.1007/3-540-44647-8_21.

[TPAS21] Perumal Thanalakshmi, S. Priyanka, K. Arputharaj, and D. Suresh. A hash-based quantum-resistant chameleon signature scheme. *Sensors*, 21(24):8417, December 2021. URL: https://www.mdpi.com/1424-8220/21/24/8417, doi:10.3390/s21248417.

[TZZ+20] Yangguang Tian, Jianting Zhang, Zhenfeng Zhang, Shengli Liu, and Jiang Zhang. Policy-based chameleon hash for blockchain rewriting with black-box accountability. In *Proceedings of the 36th Annual Computer Security Applications Conference (ACSAC 2020)*, pages 388–399, Austin, TX, USA, December 2020. ACM. URL: https://dl.acm.org/doi/10.1145/3427228.3427247, doi:10.1145/3427228.3427247.

[TZZ+23] Yangguang Tian, Jianting Zhang, Zhenfeng Zhang, Shengli Liu, and Jiang Zhang. Revocable policy-based chameleon hash for blockchain rewriting. *The Computer Journal*, 66(10):2365–2378, October 2023. URL: https://academic.oup.com/comjnl/article/66/10/2365/6761540, doi:10.1093/comjnl/bxac083.

[WCZ+24] Xiayu Wang, Yineng Chen, Xinghui Zhu, Cheng Li, and Kui Fang. A redactable blockchain scheme supporting quantum-resistance and trapdoor updates. *Applied Sciences*, 14(2):832, January 2024. URL: https://www.mdpi.com/2076-3417/14/2/832, doi:10.3390/app14020832.

[WDY+24] Xiangyu Wu, Xuehui Du, Qiantao Yang, Na Wang, and Wenjuan Wang. Redactable consortium blockchain based on verifiable distributed chameleon hash functions. *Journal of Parallel and Distributed Computing*, 183:104777, 2024. URL: https://www.sciencedirect.com/science/article/abs/pii/S0743731523001478, doi:10.1016/j.jpdc.2023.104777.

[WI24] Yong Wang and Eddie Shahril Ismail. Tightly-secure two-tier signatures on code-based digital signatures with chameleon hash functions. *Mathematics*, 12(15):2375, 2024. URL: https://www.mdpi.com/2227-7390/12/15/2375, doi:10.3390/math12152375.

[WKD21] Chunhui Wu, Lishan Ke, and Yusong Du. Quantum resistant key-exposure free chameleon hash and applications in redactable blockchain. *Information Sciences*, 548:438–449, February 2021. URL: https://www.sciencedirect.com/science/article/abs/pii/S0020025520309956, doi:10.1016/j.ins.2020.10.008.

[WT14] Yuyu Wang and Keisuke Tanaka. Generic transformation to strongly existentially unforgeable signature schemes with leakage resiliency. In *Provable Security – ProvSec 2014*, volume 8782 of *Lecture Notes in Computer Science*, pages 117–129, Cham, 2014. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-319-12475-9_9, doi:10.1007/978-3-319-12475-9_9.

[WT15] Yuyu Wang and Keisuke Tanaka. Generic transformation to strongly existentially unforgeable signature schemes with continuous leakage resiliency. In *Information Security and Privacy – ACISP 2015*, volume 9144 of *Lecture Notes in Computer Science*, pages 216–233, Cham, 2015. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-319-19962-7_13, doi:10.1007/978-3-319-19962-7_13.

[WZZ+22] Jiannan Wei, Shengli Zhang, Huawei Zhang, Jianwei Liu, Hongwei Li, and Rongxing Lu. A redactable blockchain framework for secure federated learning in industrial internet of things. *IEEE Internet of Things Journal*, 9(18):17901–17911, 2022. URL: https://ieeexplore.ieee.org/document/9743331, doi:10.1109/JIOT.2022.3162499.

[XCY+21] Junpeng Xu, Haixia Chen, Xu Yang, Wei Wu, and Yongcheng Song. Verifiable image revision from Chameleon hashes. *Cybersecurity*, 4(1):34, October 2021. URL: https://cybersecurity.springeropen.com/articles/10.1186/s42400-021-00097-3, doi:10.1186/s42400-021-00097-3.

[XHY+22] Shengmin Xu, Xinyi Huang, Jiaming Yuan, Yingjiu Li, and Robert H Deng. Accountable and fine-grained controllable rewriting in blockchains. *IEEE Transactions on Information Forensics and Security*, 18:101–116, 2022. URL:

https://ieeexplore.ieee.org/document/9931163, doi:10.1109/TIFS.2022.3217742.

[XSL+21] Zhikang Xie, Qingni Shen, Cong Li, Jisheng Dong, and Yuejian Fang. Identity-based chameleon hash without random oracles and application in the mobile internet. In *ICC 2021 – IEEE International Conference on Communications*, pages 1–6, Montreal, QC, Canada, June 2021. IEEE. URL: https://ieeexplore.ieee.org/document/9500446, doi:10.1109/ICC42927.2021.9500446.

[XZL+21] Shengmin Xu, Zhixin Zhang, Weijie Liu, Chao Li, and Huaqun Wang. Revocable policy-based chameleon hash. In *Computer Security – ESORICS 2021*, volume 12973 of *Lecture Notes in Computer Science*, pages 3–22, Cham, 2021. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-030-88418-5_16, doi:10.1007/978-3-030-88418-5_16.

[XZWZ17] Dong Xie, Wei Zhang, Xuejia Wang, and Zhengan Zhang. Homomorphic signatures from chameleon hash functions. 46(2):274–286, 2017. URL: https://itc.ktu.lt/index.php/ITC/article/view/14320, doi:10.5755/j01.itc.46.2.14320.

[YAH24] Mohd Najwadi Yusoff, Shams Mhmood Abd Ali, and Hasan Falah Hasan. Enhancing efficiency and fine-grained control in redactable blockchains with epbchf. *Iraqi Journal for Computer Science and Mathematics*, 5(3):194–212, 2024. URL: https://ijcsm.researchcommons.org/ijcsm/vol5/iss3/15/, doi:10.52866/ijcsm.2024.05.03.010.

[YCL+24] Yuxiang Yang, Yuling Chen, Zhiquan Liu, Chaoyue Tan, and Yun Luo. Verifiable and redactable blockchain for internet of vehicles data sharing. *IEEE Internet of Things Journal*, 2024. URL: https://ieeexplore.ieee.org/document/10726583, doi:10.1109/JIOT.2024.3483809.

[ZGL+24] Liushun Zhao, Deke Guo, Lailong Luo, Junjie Xie, Yulong Shen, and Bangbang Ren. Tiger tally: A secure iot data management approach based on redactable blockchain. *Computer Networks*, 248:110500, 2024. URL: https://www.sciencedirect.com/science/article/abs/pii/S1389128624003323, doi:10.1016/j.comnet.2024.110500.

[Zha07] Rui Zhang. Tweaking TBE/IBE to PKE transforms with chameleon hash functions. In *Applied Cryptography and Network Security, 5th International Conference (ACNS 2007), Zhuhai, China, June 5–8, 2007, Proceedings*, volume 4521 of *Lecture Notes in Computer Science*, pages 323–339. Springer, 2007. URL: https://link.springer.com/chapter/10.1007/978-3-540-72738-5_21, doi:10.1007/978-3-540-72738-5_21.

[ZLL+18] Mingmei Zheng, Zhihua Liu, Jiejie Liu, Zhenfeng Zhang, and Dawu Gu. A new design of online/offline signatures based on lattice. In *Information Security Practice and Experience – ISPEC 2018*, volume 11125 of *Lecture Notes in Computer Science*, pages 3–22, Cham, 2018. Springer. URL: https://link-springer-com.demo.remotlog.com/chapter/10.1007/978-3-319-99807-7_12, doi:10.1007/978-3-319-99807-7_12.

[ZLL+20] Jingning Zhang, Haibin Liu, Yang Liu, Jun Shao, and Rongxing Lu. Serving at the edge: A redactable blockchain with fixed storage. In *Web Information Systems and Applications – WISA 2020*, volume 12323 of *Lecture Notes in Computer Science*, pages 3–15, Cham, 2020. Springer. URL: https:

//link.springer.com/chapter/10.1007/978-3-030-60029-7_58, doi:10.1007/978-3-030-60029-7_58.

[ZSNS03] Fangguo Zhang, Reihaneh Safavi-Naini, and Willy Susilo. Id-based chameleon hashes from bilinear pairings. Technical Report 2003/208, Cryptology ePrint Archive, 2003. URL: https://eprint.iacr.org/2003/208.

[ZZLD21] Zongyang Zhang, Jin Zhang, Jiguo Li, and Robert H. Deng. Redactable transactions in consortium blockchain: Controlled by multi-authority CP-ABE. In *Information Security and Privacy – ACISP 2021*, volume 12702 of *Lecture Notes in Computer Science*, pages 473–493, Cham, 2021. Springer. URL: https://link.springer.com/chapter/10.1007/978-3-030-90567-5_21, doi:10.1007/978-3-030-90567-5_21.