Lossy Algebraic Filters With Short Tags

Benoît Libert, Chen Qian

► To cite this version:

Benoît Libert, Chen Qian. Lossy Algebraic Filters With Short Tags. PKC 2019 - 22nd International Conference on Practice and Theory of Public Key Cryptography, Apr 2019, Beijing, China. pp.34–65, 10.1007/978-3-030-17253-4_2 . hal-02124968

HAL Id: hal-02124968

<https://hal.inria.fr/hal-02124968>

Submitted on 10 May 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Lossy Algebraic Filters With Short Tags

Benoît Libert^{1,2} and Chen Qian³

¹ CNRS, Laboratoire LIP, France

² ENS de Lyon, Laboratoire LIP (U. Lyon, CNRS, ENSL, INRIA, UCBL), France

³ IRISA Rennes (France)

Abstract. Lossy algebraic filters (LAFs) are function families where each function is parametrized by a tag, which determines if the function is injective or lossy. While initially introduced by Hofheinz (Eurocrypt 2013) as a technical tool to build encryption schemes with key-dependent message chosen-ciphertext (KDM-CCA) security, they also find applications in the design of robustly reusable fuzzy extractors. So far, the only known LAF family requires tags comprised of $\Theta(n^2)$ group elements for functions with input space $\mathbb{Z}_p^n$, where p is the group order. In this paper, we describe a new LAF family where the tag size is only linear in n and prove it secure under simple assumptions in asymmetric bilinear groups. Our construction can be used as a drop-in replacement in all applications of the initial LAF system. In particular, it can shorten the ciphertexts of Hofheinz’s KDM-CCA-secure public-key encryption scheme by 19 group elements. It also allows substantial space improvements in a recent fuzzy extractor proposed by Wen and Liu (Asiacrypt 2018). As a second contribution, we show how to modify our scheme so as to prove it (almost) tightly secure, meaning that security reductions are not affected by a concrete security loss proportional to the number of adversarial queries.

Keywords. Lossy algebraic filters, efficiency, tight security, standard assumptions.

1 Introduction

As introduced by Peikert and Waters a decade ago [40], lossy trapdoor functions (LTFs) are function families where injective functions – which are efficiently invertible using a trapdoor – are computationally indistinguishable from many-to-one functions, wherein the image is drastically smaller than the domain. Since their introduction, they drew a lot of attention [19,23,25,44,46] and revealed powerful enough to imply chosen-ciphertext (IND-CCA2) security [40], deterministic public-key encryption in the standard model [9,15,42], as well as encryption schemes achieving the best possible security against selective-opening (SO) adversaries [2,5] or using imperfect randomness [1].

LOSSY ALGEBRAIC FILTERS. Lossy algebraic filters (LAFs) are a variant LTFs introduced by Hofheinz [26] as a tool enabling the design of chosen-ciphertext-secure encryption schemes with key-dependent message (KDM-CCA) security [6]. Recently, they were also used by Wen and Liu [45] in the construction of

robustly reusable fuzzy extractors. In LAF families, each function takes as arguments an input x and a tag t , which determines if the function behaves as a lossy or an injective function. More specifically, each tag $t = (t_c, t_a)$ is comprised of an auxiliary component t_a , which may consist of any public data, and a core component t_c . For any auxiliary component t_a , there should exist at least one t_c such that $t = (t_c, t_a)$ induces a lossy function $f_{\text{LAF}}(t, \cdot)$. LAFs strengthen the requirements of lossy trapdoor functions in that, for any lossy tag t , the function $f_{\text{LAF}}(t, x)$ always reveals the same information about the input x , regardless of which tag is used. In particular, for a given evaluation key ek , multiple evaluations $f_{\text{LAF}}(t_1, x), \dots, f_{\text{LAF}}(t_n, x)$ for distinct lossy tags do not reveal any more information about x than a single evaluation. On the other hand, LAFs depart from lossy trapdoor functions in that they need not be efficiently invertible using a trapdoor. For their applications to KDM-CCA security [26] and fuzzy extractors [45], lossy algebraic filters are expected to satisfy two security properties. The first one, called *indistinguishability*, requires that lossy tags be indistinguishable from random tags. The second one, named *evasiveness*, captures that lossy tags should be hard to come by without a trapdoor.

So far, the only known LAF realization is a candidate, suggested by Hofheinz [26], which relies on the Decision Linear assumption (DLIN) [12] in groups with a bilinear map. While efficient and based on a standard assumption, it incurs relatively large tags comprised of a quadratic number of group elements in the number of input symbols. More precisely, for functions admitting inputs $\mathbf{x} = (x_1, \dots, x_n)^\top \in \mathbb{Z}_p^n$, where p is the order of a pairing-friendly $\mathbb{G}$, the core components t_c contain $\Theta(n^2)$ elements of $\mathbb{G}$. For the application to KDM-CCA security [26] (where t_c should be part of ciphertexts), quadratic-size tags are not prohibitively expensive as the encryption scheme of [26, Section 4] can make do with a constant n (typically, $n = 6$). In the application to fuzzy extractors [45], however, it is desirable to reduce the tag length. In the robustly reusable fuzzy extractor of [45], the core tag component t_c is included in the public helper string P that allows reconstructing a secret key from a noisy biometric reading w . The latter lives in a metric space that should be small enough to fit in the input space $\mathbb{Z}_p^n$ of the underlying LAF family. Even if p is exponentially large in the security parameter λ , a constant n would restrict biometric readings to have linear length in λ . Handling biometric readings of polynomial length thus incurs $n = \omega(1)$, which results in large tags and longer public helper strings. This motivates the design of new LAF candidates with smaller tags.

OUR RESULTS. The contribution of this paper is two-fold. We first construct a new LAF with linear-size tags and prove it secure under simple, constant-size assumptions (as opposed to q -type assumptions, which are described using a linear number of elements in the number of adversarial queries) in bilinear groups. The indistinguishability and evasiveness properties of our scheme are implied by the Decision 3-party Diffie-Hellman assumption (more precisely, its natural analogue in asymmetric bilinear maps), which posits the pseudorandomness of tuples $(g, g^a, g^b, g^c, g^{abc})$, for random $a, b, c \in_R \mathbb{Z}_p$. For inputs in $\mathbb{Z}_p^n$, where p is the group order, our core tag components only consist of $O(n)$ group elements.

These shorter tags are obtained without inflating evaluation keys, which remain of length $O(n)$ (as in [26]).

As a second contribution, we provide a second LAF realization with $O(n)$ -size tags where the indistinguishability and evasiveness properties are both *almost tightly* related to the underlying hardness assumption. Namely, our security proofs are tight – or almost tight in the terminology of Chen and Wee [16] – in that the gap between the advantages of the adversary and the reduction only depends on the security parameter, and not on the number of adversarial queries. In the LAF suggested by Hofheinz [26], the proof of evasiveness relies on the unforgeability of Waters signatures [43]. As a result, the reduction loses a linear factor in the number of lossy tags obtained by the adversary. In our second construction, we obtain tight reductions by replacing Waters signatures with (a variant of) a message authentication code (MAC) due to Blazy, Kiltz and Pan [7]. As a result, our proof of evasiveness only loses a factor $O(\lambda)$ with respect to the Symmetric eXternal Diffie-Hellman assumption (SXDH). If our scheme is plugged into the robustly reusable fuzzy extractor of Wen and Liu [45], it immediately translates into a tight proof of robustness in the sense of the definition of [45]. While directly using our second LAF in the KDM-CCA-secure scheme of [26] does not seem sufficient to achieve tight key-dependent message security, it may still provide a building block for future constructions of tightly KDM-CCA-secure encryption schemes with short ciphertexts.

TECHNIQUES. Like the DLIN-based solution given by Hofheinz [26], our evaluation algorithms proceed by computing a matrix-vector product in the exponent, where the matrix is obtained by pairing group elements taken from the core tag t_c with elements of the evaluation key. Here, we reduce the size of t_c from $O(n^2)$ to $O(n)$ group elements using a technique suggested by Boyen and Waters [14] in order to compress the evaluation keys of DDH-based lossy trapdoor functions.

In the pairing-based LTF of [14], the evaluation key contains group elements $\{(R_i, S_i) = (g^{r_i}, (h^i \cdot u)^{r_i})\}_{i=1}^n, \{(V_j = g^{v_j}, H_j = (h^j \cdot u)^{v_j})\}_{j=1}^n$. Using a symmetric bilinear maps $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$, these make it possible to compute the off-diagonal elements of a matrix

$$M_{i,j} = e(g, h)^{r_i \cdot v_j} = \left(\frac{e(R_i, H_j)}{e(S_i, V_j)} \right)^{1/(j-i)} \quad \forall (i, j) \in [n] \times [n] \setminus \{(i, i)\}_{i=1}^n \quad (1)$$

via a “two equation” technique borrowed from the revocation system of Lewko, Sahai and Waters [34]. By including $\{D_i = e(g, g)^{r_i \cdot v_i} \cdot e(g, g)\}_{i=1}^n$ in the evaluation key, the LTF of [14] allows the evaluator to compute a matrix $(M_{i,j})_{i,j \in [n]}$ such that $M_{i,j} = e(g, g)^{r_i \cdot v_j}$ if $i \neq j$ and $M_{i,i} = e(g, g)^{r_i \cdot v_i} \cdot e(g, g)^{m_i}$ and for which $m_i = 1$ (resp. $m_i = 0$), for all $i \in [n]$, in injective (resp. lossy) functions. The indistinguishability of lossy and injective evaluation keys relies on the fact that (1) is only computable when $i \neq j$, making it infeasible to distinguish $\{D_i = e(g, h)^{r_i \cdot v_i} \cdot e(g, g)\}_{i=1}^n$ from $\{D_i = e(g, h)^{r_i \cdot v_i}\}_{i=1}^n$.

Our first LAF construction relies on the “two equation” technique of [34] in a similar way with the difference that we include $\{(V_j = g^{v_j}, H_j = (h^j \cdot u)^{v_j})\}_{j=1}^n$ in the evaluation key ek , but $\{(R_i, S_i) = (g^{r_i}, (h^i \cdot u)^{r_i})\}_{i=1}^n$ is now part of the core

tag components t_c . This makes it possible to compute off-diagonal elements of $(M_{i,j})_{i,j \in [n]}$ by pairing elements of ek with those of t_c . To enable the computation of diagonal elements $\{M_{i,i}\}_{i=1}^n$, we augment core tag components by introducing pairs $(D_i, E_i) \in \mathbb{G}^2$, which play the same role as $\{D_i = e(g, g)^{r_i \cdot v_i} \cdot e(g, g)\}_{i=1}^n$ in the LTF of [14]. In lossy tags, $\{(D_i, E_i)\}_{i=1}^n$ are of the form

$$(D_i, E_i) = (h^{r_i \cdot v_i} \cdot H_{\mathbb{G}}(\tau)^{\rho_i}, g^{\rho_i}), \quad (2)$$

for a random $\rho_i \in_R \mathbb{Z}_p$, where τ is a chameleon hashing of all tag components. Such pairs $\{(D_i, E_i)\}_{i=1}^n$ allow the evaluator to compute

$$M_{i,i} = \frac{e(D_i, g)}{e(H_{\mathbb{G}}(\tau), E_i)} = e(g, h)^{r_i \cdot v_i} \quad \forall i \in [n],$$

which results in a rank-one matrix $(M_{i,j})_{i,j \in [n]}$, where $M_{i,j} = e(g, h)^{r_i \cdot v_j}$. When computed as per (2), $\{(D_i, E_i)\}_{i=1}^n$ can be seen as “blinded” Waters signatures [43]. Namely, $(g, h, V_i = g^{v_i})$ can be seen as a verification key; h^{v_i} is the corresponding secret key; and $r_i \in \mathbb{Z}_p$ serves as a blinding factor that ensures the indistinguishability of (D_i, E_i) from random pairs. Indeed, the Decision 3-party Diffie-Hellman (D3DH) assumption allows proving that $h^{r_i \cdot v_i}$ is computationally indistinguishable from random given (g, h, g^{v_i}, g^{r_i}) . In our proof of indistinguishability, however, we need to rely on the proof technique of the Boneh-Boyen IBE [11] in order to apply a hybrid argument that allows gradually replacing pairs $\{(D_i, E_i)\}_{i=1}^n$ by random group elements.

In our proof of evasiveness, we rely on the fact that forging a pair of the form $(D_i, E_i) = (h^{r_i \cdot v_i} \cdot H_{\mathbb{G}}(\tau)^{\rho_i}, g^{\rho_i})$ on input of (g, h, g^{v_i}) is as hard as solving the 2-3-Diffie-Hellman problem [33], which consist in finding a non-trivial pair $(g^r, g^{r \cdot ab}) \in \mathbb{G}^* \times \mathbb{G}^*$ on input of (g, g^a, g^b) . In turn, this problem is known to be at least as hard as breaking the Decision 3-party Diffie-Hellman assumption.

The above techniques allow us to construct a LAF with $O(n)$ -size tags and evaluation keys made of $O(n + \lambda)$ group elements under a standard assumption. Our first LAF is actually described in terms of asymmetric pairings, but it can be instantiated in all types (i.e., symmetric or asymmetric) of bilinear groups. Our second LAF construction requires asymmetric pairing configurations and the Symmetric eXternal Diffie-Hellman (SXDH) assumption. It is very similar to our first construction with the difference that we obtain a tight proof of evasiveness by replacing Waters signatures with a variant of a MAC proposed by Blazy, Kiltz and Pan [7]. In order for the proofs to go through, we need to include n MAC instances (each with its own keys) in lossy tags, which incurs evaluation keys made of $O(n \cdot \lambda)$ group elements. We leave it is an interesting open problem to achieve tight security using shorter evaluation keys.

RELATED WORK. All-but-one lossy trapdoor functions (ABO-LTFs) [40] are similar to LAFs in that they are lossy function families where each function is parametrized by a tag that determines if the function is injective or lossy. They differ from LAFs in two aspects: (i) They should be efficiently invertible using a trapdoor; (ii) For a given evaluation key ek , there exists only one

tag for which the function is lossy. The main motivation of ABO-LTFs is the construction of chosen-ciphertext [41] encryption schemes. *All-but-many* lossy trapdoor functions (ABM-LTFs) are an extension of ABO-LTFs introduced by Hofheinz [25]. They are very similar to LAFs in that a trapdoor makes it possible to dynamically create arbitrarily many lossy tags using. In particular, each tag $t = (t_c, t_a)$ consists of an auxiliary component t_a and a core component t_c so that, by computing t_c as a suitable function of t_a , the pair $t = (t_c, t_a)$ can be made lossy, but still random-looking. The motivation of ABM-LTFs is the construction chosen-ciphertext-secure public-key encryption schemes in scenarios, such as the selective-opening setting [18,2], which involve multiple challenge ciphertexts [25]. They also found applications in the design of universally composable commitments [20]. Lossy algebraic filters differ from ABM-LTFs in that they may not have a trapdoor enabling efficient inversion but, for any lossy tag $t = (t_c, t_a)$, the information revealed by $f_{\text{LAF}}(t, x)$ is always the same (i.e., it is completely determined by x and the evaluation key ek).

LAFs were first introduced by Hofheinz [26] as a building block for KDM-CCA-secure encryption schemes, where they enable some form of “plaintext awareness”. In the security proofs of KDM-secure encryption schemes (e.g., [10]), the reduction must be able to simulate encryptions of (functions of) the secret key. When the adversary is equipped with a decryption oracle, the ability to publicly compute encryptions of the decryption key may be a problem as decryption queries could end up revealing that key. LAFs provide a way reconcile the conflicting requirements of KDM and CCA2-security by introducing in each ciphertext a LAF-evaluation of the secret key. By having the simulator encrypt a lossy function of the secret key, one can keep encryption queries from leaking too much secret information. At the same time, adversarially-generated ciphertexts always contain an injective function of the key, which prevents the adversary from learning the secret key by publicly generating encryptions of that key.

Recently, Wen and Liu [45] appealed to LAFs in the design of robustly reusable fuzzy extractors. As defined by Dodis *et al.* [17], fuzzy extractors allow one to generate a random cryptographic key R – together with some public helper string P – out of a noisy biometric reading w . The key R need not be stored as it can be reproduced from the public helper string P and a biometric reading w' which is sufficiently close to w . Reusable fuzzy extractors [13] make it possible to safely generate multiple keys $R_1, \dots, R_t$ (each with its own public helper string P_i) from correlated readings $w_1, \dots, w_t$ of the same biometric source. Wen and Liu [45] considered the problem of achieving robustness in reusable fuzzy extractors. In short, robustness prevents adversaries from covertly tampering with the public helper string P_i in order to affect the reproducibility of R_i . The Wen-Liu [45] fuzzy extractor relies on LAFs to simultaneously achieve reusability and robustness assuming a common reference string. Their solution requires the LAF to be homomorphic, meaning that function outputs should live in a group and, for any tag t and inputs x_1, x_2 , we have $f_{\text{LAF}}(t, x_1 + x_2) = f_{\text{LAF}}(t, x_1) \cdot f_{\text{LAF}}(t, x_2)$. The candidate proposed by Hofheinz [26] and ours are both usable in robustly reusable fuzzy extractors as they both satisfy this homomorphic property. Our

scheme offers the benefit of shorter public helper strings P since these have to contain a LAF tag in the fuzzy extractor of [45].

The tightness of cryptographic security proofs was first considered by Bellare and Rogaway [4] in the random oracle model [3]. In the standard model, it drew a lot of attention in digital signatures and public-key encryption the recent years (see, e.g., [29,16,7,35,36,27,21,28,22]). In the context of all-but-many lossy trapdoor functions, a construction with tight evasiveness was put forth by Hofheinz [25]. A tightly secure lattice-based ABM-LTF was described by Libert *et al.* [37] as a tool enabling tight chosen-ciphertext security from lattice assumptions. To our knowledge, the only other prior work considering tight reductions for lossy trapdoor functions is a recent result of Hofheinz and Nguyen [30]. In particular, tight security has never been obtained in the context of LAFs, nor in fuzzy extractors based on public-key techniques.

2 Background

2.1 Lossy Algebraic Filters

We recall the definition of Lossy Algebraic Filter (LAF) from [26], in which the distribution over the function domain may not be the uniform one.

Definition 1. For integers $\ell_{\text{LAF}}(\lambda), n(\lambda) > 0$, an (ℓ_{LAF}, n) -lossy algebraic filter (LAF) with security parameter λ consists of the following ppt algorithms:

Key generation. $\text{LAF.Gen}(1^\lambda)$ outputs an evaluation key ek and a trapdoor key tk . The evaluation key ek specifies an ℓ_{LAF} -bit prime p as well as the description of a tag space $\mathcal{T} = \mathcal{T}_c \times \mathcal{T}_a$, where $\mathcal{T}_c$ is efficiently samplable. The disjoint sets of injective and non-injective tags are called $\mathcal{T}_{\text{inj}}$ and $\mathcal{T}_{\text{non-inj}} = \mathcal{T} \setminus \mathcal{T}_{\text{inj}}$, respectively. We also define the subset of lossy tags $\mathcal{T}_{\text{loss}}$ to be a subset of $\mathcal{T}_{\text{non-inj}}$, which induce very lossy functions. A tag $t = (t_c, t_a)$ is described by a core part $t_c \in \mathcal{T}_c$ and an auxiliary part $t_a \in \mathcal{T}_a$. A tag may be injective, or lossy, or neither. The trapdoor tk allows sampling lossy tags.

Evaluation. $\text{LAF.Eval}(ek, t, X)$ takes as inputs an evaluation key ek , a tag $t \in \mathcal{T}$ and a function input $X \in \mathbb{Z}_p^n$. It outputs an image $Y = f_{ek,t}(X)$.

Lossy tag generation. $\text{LAF.LTag}(tk, t_a)$ takes as input the trapdoor key tk , an auxiliary part $t_a \in \mathcal{T}_a$ and outputs a core part t_c such that $t = (t_c, t_a) \in \mathcal{T}_{\text{loss}}$ forms a lossy tag.

In addition, LAF has to meet the following requirements:

Lossiness. For any $(ek, tk) \xleftarrow{R} \text{LAF.Gen}(1^\lambda)$, the following conditions should be satisfied.

- For any $t \in \mathcal{T}_{\text{inj}}$, $f_{ek,t}(\cdot)$ should behave as an injective function (note that $f_{ek,t}^{-1}(\cdot)$ is not required to be efficiently computable given tk).
- For any auxiliary tag $t_a \in \mathcal{T}_a$ and any $t_c \xleftarrow{R} \text{LAF.LTag}(tk, t_a)$, we have $t = (t_c, t_a) \in \mathcal{T}_{\text{loss}}$, meaning that $f_{ek,t}(\cdot)$ is a lossy function. Moreover, for

any input $X = (x_1, \dots, x_n) \in \mathbb{Z}_p^n$ and any $t = (t_c, t_a) \in \mathcal{T}_{\text{loss}}$, $f_{ek,t}(X)$ is completely determined by $\sum_{i=1}^n v_i \cdot x_i \bmod p$ for coefficients $\{v_i\}_{i=1}^n$ that only depend on ek .

Indistinguishability. Multiple lossy tags are computationally indistinguishable from random tags, namely:

$$\text{Adv}_Q^{\mathcal{A}, \text{ind}}(\lambda) := |\Pr[\mathcal{A}(1^\lambda, ek)^{\text{LAF.LTag}(tk, \cdot)} = 1] - \Pr[\mathcal{A}(1^\lambda, ek)^{\mathcal{O}_{\tau_c}(\cdot)} = 1]|$$

is negligible for any PPT algorithm $\mathcal{A}$, where $(ek, tk) \xleftarrow{R} \text{LAF.Gen}(1^\lambda)$ and $\mathcal{O}_{\tau_c}(\cdot)$ is an oracle that assigns a random core tag $t_c \xleftarrow{R} \mathcal{T}_c$ to each auxiliary tag $t_a \in \mathcal{T}_a$ (rather than a core tag that makes $t = (t_c, t_a)$ lossy). Here Q denotes the number of oracle queries made by $\mathcal{A}$.

Evasiveness. Non-injective tags are computationally hard to find, even with access to an oracle outputting multiple lossy tags, namely:

$$\text{Adv}_{Q_1, Q_2}^{\mathcal{A}, \text{eva}}(\lambda) := \Pr[\mathcal{A}(1^\lambda, ek)^{\text{LAF.LTag}(tk, \cdot), \text{LAF.IsInjective}(tk, \cdot)} \in \mathcal{T}_{\text{non-inj}}]$$

is negligible for legitimate adversary $\mathcal{A}$, where $(ek, ik, tk) \xleftarrow{R} \text{LAF.Gen}(1^\lambda)$ and $\mathcal{A}$ is given access to the following oracle:

- $\text{LAF.LTag}(tk, \cdot)$ which acts exactly as the lossy tag generation algorithm.
- $\text{LAF.IsInjective}(tk, \cdot)$ that takes as input a tag $t = (t_c, t_a)$. It outputs 0 if $t \in \mathcal{T}_{\text{non-inj}} = \mathcal{T} \setminus \mathcal{T}_{\text{inj}}$ and 1 if $t \in \mathcal{T}_{\text{inj}}$. If $t \notin \mathcal{T}$, the oracle outputs $\perp$.

We denote by Q_1 and Q_2 the number of queries to $\text{LAF.LTag}(tk, \cdot)$ and $\text{LAF.IsInjective}(tk, \cdot)$, respectively. By “legitimate adversary”, we mean that $\mathcal{A}$ is PPT and never outputs a tag $t = (t_c, t_a)$ such that t_c was obtained by invoking the LAF.LTag oracle on t_a .

In our construction, the tag space $\mathcal{T}$ will not be dense (i.e., not all elements of the ambient algebraic structure are potential tags). However, elements of the tag space $\mathcal{T}$ will be efficiently recognizable given ek .

We note that the above definition of evasiveness departs from the one used by Hofheinz [26] in that it uses an additional $\text{LAF.IsInjective}(tk, \cdot)$ oracle that uses the trapdoor tk to decide whether a given tag is injective or not. However, this oracle will only be used in our tightly secure LAF (and not in our first construction). Its only purpose is to enable a modular use of our tightly evasive LAF in applications to KDM security [26] or robustly reusable fuzzy extractors [45]. Specifically, by invoking the $\text{LAF.IsInjective}(tk, \cdot)$ oracle, the reduction from the security of a primitive to the underlying LAF’s evasiveness does not have to guess which adversarial query involves a non-lossy tag.

2.2 Chameleon Hash Functions

A chameleon hash function [32] is a tuple of algorithms $\text{CMH} = (\text{CMKg}, \text{CMhash}, \text{CMswitch})$ which contains an algorithm CMKg that, given a security parameter 1^λ , outputs a key pair $(hk, td) \leftarrow \mathcal{G}(1^\lambda)$. The randomized hashing algorithm

outputs $y = \text{CMhash}(hk, m, r)$ given the public key hk , a message m and random coins $r \in \mathcal{R}_{\text{hash}}$. On input of messages m, m' , random coins $r \in \mathcal{R}_{\text{hash}}$ and the trapdoor key td , the switching algorithm $r' \leftarrow \text{CMswitch}(td, m, r, m')$ computes $r' \in \mathcal{R}_{\text{hash}}$ such that $\text{CMhash}(hk, m, r) = \text{CMhash}(hk, m', r')$. The collision-resistance property mandates that it be infeasible to come up with pairs $(m', r') \neq (m, r)$ such that $\text{CMhash}(hk, m, r) = \text{CMhash}(hk, m', r')$ without knowing the trapdoor key tk . Uniformity guarantees that the distribution of hash values is independent of the message m : in particular, for all hk , and all messages m, m' , the distributions $\{r \leftarrow \mathcal{R}_{\text{hash}} : \text{CMHash}(hk, m, r)\}$ and $\{r \leftarrow \mathcal{R}_{\text{hash}} : \text{CMHash}(hk, m', r)\}$ are identical.

2.3 Hardness Assumptions

Definition 2. Let $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$ be bilinear groups of order p . The **First Decision 3-Party Diffie-Hellman (D3DH1)** assumption holds in $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$ if no PPT distinguisher can distinguish the distribution

$$\begin{aligned} D_1 &:= \{(g, \hat{g}, g^a, g^b, g^c, \hat{g}^a, \hat{g}^b, \hat{g}^c, g^{abc}) \mid g \xleftarrow{R} \mathbb{G}, \hat{g} \xleftarrow{R} \hat{\mathbb{G}}, a, b, c \xleftarrow{R} \mathbb{Z}_p\} \\ D_0 &:= \{(g, \hat{g}, g^a, g^b, g^c, \hat{g}^a, \hat{g}^b, \hat{g}^c, g^z) \mid g \xleftarrow{R} \mathbb{G}, \hat{g} \xleftarrow{R} \hat{\mathbb{G}}, a, b, c, z \xleftarrow{R} \mathbb{Z}_p\}. \end{aligned}$$

The D3DH1 assumption has a natural analogue where the pseudorandom value lives in $\hat{\mathbb{G}}$ instead of $\mathbb{G}$.

Definition 3. The **Second Decision 3-Party Diffie-Hellman (D3DH2)** assumption holds in $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$ if no PPT algorithm can distinguish between the distribution

$$\begin{aligned} D_1 &:= \{(g, \hat{g}, g^a, g^b, g^c, \hat{g}^a, \hat{g}^b, \hat{g}^c, g^{abc}) \mid g \xleftarrow{R} \mathbb{G}, \hat{g} \xleftarrow{R} \hat{\mathbb{G}}, a, b, c \xleftarrow{R} \mathbb{Z}_p\} \\ D_0 &:= \{(g, \hat{g}, g^a, g^b, g^c, \hat{g}^a, \hat{g}^b, \hat{g}^c, \hat{g}^z) \mid g \xleftarrow{R} \mathbb{G}, \hat{g} \xleftarrow{R} \hat{\mathbb{G}}, a, b, c, z \xleftarrow{R} \mathbb{Z}_p\}. \end{aligned}$$

We also need a computational assumption which is implied by D3DH2. The 2-3-CDH was initially introduced [33] in ordinary (i.e., non-pairing-friendly) discrete-logarithm hard groups. Here, we extend it to asymmetric bilinear groups.

Definition 4 ([33]). Let $(\mathbb{G}, \hat{\mathbb{G}})$ be a bilinear groups of order p with generators $g \in \mathbb{G}$ and $\hat{g} \in \hat{\mathbb{G}}$. The **2-out-of-3 Computational Diffie-Hellman (2-3-CDH)** assumption says that, given $(g, g^a, \hat{g}^a, g^b, \hat{g}^b)$ for randomly chosen $a, b \xleftarrow{R} \mathbb{Z}_p$, no PPT algorithm can find a pair $(g^r, g^{r \cdot ab})$ such that $r \neq 0$.

It is known (see, e.g., [38]) that any algorithm solving the 2-3-CDH problem can be used to break the D3DH2 assumption. On input of $(g, \hat{g}, g^a, g^b, g^c, \hat{g}^a, \hat{g}^b, \hat{g}^c, \hat{g}^z)$, where $z = abc$ or $z \in_R \mathbb{Z}_p$, the reduction can simply run a 2-3-CDH solver on input of $(g, g^a, g^b, \hat{g}^a, \hat{g}^b)$. If the solver outputs a non-trivial pair of the form $(R_1, R_2) = (g^r, g^{r \cdot ab})$, the D3DH2 distinguisher decides that $z = abc$ if and only if $e(R_1, \hat{g}^z) = e(R_2, \hat{g}^c)$.

In our constructions, we actually rely on a weaker variant of D3HD1, called wd3HD1, where $\hat{g}^a$ is not given. In our tightly secure construction (which requires asymmetric pairings), we need to rely on the following variant of wd3HD1.

Definition 5. Let $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$ be bilinear groups of order p . The **Randomized weak Decision 3-Party Diffie-Hellman (R-wD3DH1)** assumption holds in $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$ if no PPT distinguisher can distinguish the distribution

$$D_1 := \left\{ \{(g, \hat{g}, g^{a_i}, g^b, g^c, \hat{g}^b, \hat{g}^c, g^{a_i b c})\}_{i=1}^Q \mid g \xleftarrow{R} \mathbb{G}, \hat{g} \xleftarrow{R} \hat{\mathbb{G}}, \right. \\ \left. a_1, \dots, a_Q, b, c \xleftarrow{R} \mathbb{Z}_p \right\} \\ D_0 := \left\{ \{(g, \hat{g}, g^{a_i}, g^b, g^c, \hat{g}^b, \hat{g}^c, g^{z_i})\}_{i=1}^Q \mid g \xleftarrow{R} \mathbb{G}, \hat{g} \xleftarrow{R} \hat{\mathbb{G}}, \right. \\ \left. a_1, \dots, a_Q, z_1, \dots, z_Q, b, c \xleftarrow{R} \mathbb{Z}_p \right\}.$$

We do not know if D3DH1 or wD3DH1 can be tightly reduced to R-wD3DH1 (the only reduction we are aware of proceeds via a hybrid argument). In asymmetric pairings, however, we can give a tight reduction between R-wD3DH1 and a combination of wD3DH1 and SXDH.

Lemma 1. *There is a tight reduction from the wD3DH1 assumption and the DDH assumption in $\mathbb{G}$ to the R-wD3DH1 assumption. More precisely, for any R-wD3DH1 adversary $\mathcal{B}$, there exist distinguishers $\mathcal{B}_1$ and $\mathcal{B}_2$ which run in about the same time as $\mathcal{B}$ and such that*

$$\mathbf{Adv}_{\mathcal{B}}^{\text{R-wD3DH1}}(\lambda) \leq \mathbf{Adv}_{\mathcal{B}_1}^{\text{wD3DH1}}(\lambda) + \mathbf{Adv}_{\mathcal{B}_2}^{\text{DDH}_1}(\lambda),$$

where the second term denotes $\mathcal{B}_2$'s advantage as a DDH distinguisher in $\mathbb{G}$.

Proof. To prove the result, we consider the following distribution:

$$D_{int} := \left\{ \{(g, \hat{g}, g^{a \cdot \alpha_i}, g^b, g^c, \hat{g}^b, \hat{g}^c, g^{z \cdot \alpha_i})\}_{i=1}^Q \mid g \xleftarrow{R} \mathbb{G}, \hat{g} \xleftarrow{R} \hat{\mathbb{G}}, \right. \\ \left. \alpha_1, \dots, \alpha_Q, b, c, z \xleftarrow{R} \mathbb{Z}_p, a \xleftarrow{R} \mathbb{Z}_p^* \right\}$$

A straightforward reduction shows that, under the wD3DH1 assumption, D_1 is computationally indistinguishable from D_{int} . We show that, under the DDH assumption in $\mathbb{G}$, D_{int} is computationally indistinguishable from D_0 . Moreover, the reduction is tight in that the two distinguishers have the same advantage.

First, we show that, under the wD3DH1 assumption, D_{int} is computationally indistinguishable from D_1 .

We can build a wD3DH1 distinguisher $\mathcal{B}_1$ from any distinguisher for D_1 and D_{int} . With $(g, \hat{g}, g^a, g^b, g^c, \hat{g}^b, \hat{g}^c, T)$ as input where $g \xleftarrow{R} \mathbb{G}$, $\hat{g} \xleftarrow{R} \hat{\mathbb{G}}$ and $a, b, c \xleftarrow{R} \mathbb{Z}_p$, $\mathcal{B}_1$ uniformly draws $\alpha_i, \dots, \alpha_Q \xleftarrow{R} \mathbb{Z}_p$ and computes

$$D_{B_1} := \left\{ \{(g, \hat{g}, g^{a \cdot \alpha_i}, g^b, g^c, \hat{g}^b, \hat{g}^c, T^{\alpha_i})\}_{i=1}^Q \mid \alpha_1, \dots, \alpha_Q \xleftarrow{R} \mathbb{Z}_p \right\}.$$

It is easy to see that if $T = g^{abc}$, then D_{B_1} is identical to D_1 . If $T \in_R \mathbb{G}$, then D_{B_1} is distributed as D_{int} . Hence, any distinguisher between D_1 and D_{int} with D_{B_1} implies a distinguisher $\mathcal{B}_1$ for the wD3DH1 problem.

Next, we show that, under the DDH assumption in $\mathbb{G}$, D_{int} is computationally indistinguishable from D_0 . In order to build a DDH distinguisher $\mathcal{B}_2$ out of a distinguisher between D_{int} and D_0 , we use the random self-reducibility of the DDH assumption.

Lemma 2 (Random Self-Reducibility [39]). *Letting $\mathbb{G}$ be a group of prime order p , there exists a PPT algorithm R that takes as input $(g, g^a, g^b, g^c) \in \mathbb{G}^4$, for any $a, b, c \in \mathbb{Z}_p$, and returns a triple $(g^a, g^{b'}, g^{c'}) \in \mathbb{G}^3$ such that:*

- If $c = ab \pmod{p}$, then b' is uniformly random in $\mathbb{Z}_p$ and $c' = ab'$.
- If $c \neq ab \pmod{p}$, then $b', c' \in_R \mathbb{Z}_p$ are independent and uniformly random.

On input of $(g, g^z, g^\alpha, T) \in \mathbb{G}^4$, where $g \xleftarrow{R} \mathbb{G}$ and $z, \alpha \xleftarrow{R} \mathbb{Z}_p$, $\mathcal{B}_2$ uses algorithm R to generate Q instances $\{(g^z, g^{\alpha_i}, T_i)\}_{i=1}^Q$. Next, $\mathcal{B}_2$ draws $\hat{g} \xleftarrow{R} \hat{\mathbb{G}}$, $a, b, c \xleftarrow{R} \mathbb{Z}_p$ and defines the following distribution:

$$D_{\mathcal{B}_2} := \left\{ \{(g, \hat{g}, (g^{\alpha_i})^a, g^b, g^c, \hat{g}^b, \hat{g}^c, T_i)\}_{i=1}^Q \mid \hat{g} \xleftarrow{R} \hat{\mathbb{G}}, a, b, c \xleftarrow{R} \mathbb{Z}_p \right\}.$$

We observe that, if $T = g^{z \cdot \alpha}$, we have $T_i = g^{z \cdot \alpha_i}$ for all $i \in [Q]$. In this case, $D_{\mathcal{B}_2}$ is identical to D_{int} . In contrast, if $T \in_R \mathbb{G}$, the random self-reducibility ensures that $T_1, \dots, T_Q \in_R \mathbb{G}$ are i.i.d, meaning that $D_{\mathcal{B}_2}$ is identical to D_0 . Using a distinguisher between D_{int} and D_0 and feeding it with $D_{\mathcal{B}_2}$, we obtain a distinguisher $\mathcal{B}_2$ for the DDH problem in $\mathbb{G}$. $\square$

3 A Lossy Algebraic Filter With Linear-Size Tags

We present a LAF based on DDH-like assumptions with tags of size $O(n)$, where n is the number of input symbols when the input is viewed as a vector over $\mathbb{Z}_p$. Our tags are comprised of $4n$ elements of $\mathbb{G}$, which outperforms the construction of [26] for $n > 4$. In his application to KDM-CCA security [26], Hofheinz uses a LAF scheme with $n = 6$, in which case we decrease the tag size from 43 to 24 group elements⁴ and thus shorten ciphertexts by 19 group elements.

The construction is inspired by the lossy TDF of [14] and relies on the revocation technique of Lewko, Sahai and Waters [34] (LSW) in the same way. In asymmetric pairings $e : \mathbb{G} \times \hat{\mathbb{G}} \rightarrow \mathbb{G}_T$, the evaluation key contains a set of LSW ciphertexts $\{(\hat{V}_j = \hat{g}^{v_j}, \hat{H}_j = (\hat{h}^j \cdot \hat{u})^{v_j})\}_{j=1}^n$, while each core tag component t_c can be seen as containing a set of LSW secret keys $\{(R_i, S_i) = (g^{r_i}, (h^i \cdot u)^{r_i})\}_{i=1}^n$, allowing the evaluator compute $M_{i,j} = e(g, \hat{h})^{r_i \cdot v_j}$ for any pairwise distinct indices $i \neq j$. In lossy tags (t_c, t_a) , diagonal elements $\{M_{i,i}\}_{i=1}^n$ are handled by having t_c contain Waters signatures $(D_i, E_i) = (h^{r_i \cdot v_i} \cdot H_{\mathbb{G}}(\tau)^{\rho_i}, g^{\rho_i})$, where $\rho_i \in_R \mathbb{Z}_p$ and $H_{\mathbb{G}} : \{0, 1\}^L \rightarrow \mathbb{G}$ is an algebraic hash function mapping the output τ of a chameleon hash function to the group $\mathbb{G}$. For indistinguishability purposes, pairs

⁴ The LAF of [26] was described in terms of symmetric pairings but it extends to asymmetric pairings $e : \mathbb{G} \times \hat{\mathbb{G}} \rightarrow \mathbb{G}_T$ where tags are comprised of elements in $\mathbb{G}$.

$\{(D_i, E_i)\}_{i=1}^n$ are not immediately recognizable as Waters signatures because the underlying secret key h^{v_i} is blinded by a random exponent $r_i = \log_g(R_i)$. Still, running the verification algorithm of Waters signatures on (D_i, E_i) allows the evaluation algorithm to derive $M_{i,i} = e(g, \hat{h})^{r_i \cdot v_i}$, so that $(M_{i,j})_{i,j \in [n]}$ forms a rank-1 matrix. In injective tags, $\{(D_i, E_i)\}_{i=1}^n$ are uniformly distributed in $\mathbb{G}$, so that $(M_{i,j})_{i,j \in [n]}$ is the sum of a rank-1 matrix and a diagonal matrix.

3.1 Description

Key generation. $\text{LAF.Gen}(1^\lambda)$ conducts the following steps.

1. Choose bilinear groups $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$ of prime order $p > 2^\lambda$ with random generators $g, h, u \xleftarrow{R} \mathbb{G}$ and $\hat{g}, \hat{h}, \hat{u} \xleftarrow{R} \hat{\mathbb{G}}$ subject to the constraints $\log_g(h) = \log_{\hat{g}}(\hat{h})$ and $\log_g(u) = \log_{\hat{g}}(\hat{u})$.
2. Choose a chameleon hash function $\text{CMH} = (\text{CMKg}, \text{CMhash}, \text{CMswitch})$, where the hashing algorithm $\text{CMhash} : \{0, 1\}^* \times \mathcal{R}_{\text{hash}} \rightarrow \{0, 1\}^L$ has output length $L \in \text{poly}(\lambda)$. Generate a pair $(hk_{\text{CMH}}, td_{\text{CMH}}) \leftarrow \text{CMKg}(1^\lambda)$ made of a hashing key hk_{CMH} and a trapdoor td_{CMH} .
3. Choose random exponents $w_0, \dots, w_L \xleftarrow{R} \mathbb{Z}_p$ and define

$$W_k = g^{w_k}, \quad \hat{W}_k = \hat{g}^{w_k} \quad \forall k \in [0, L]$$

that will be used to instantiate two hash functions $H_{\mathbb{G}} : \{0, 1\}^L \rightarrow \mathbb{G}$, $H_{\hat{\mathbb{G}}} : \{0, 1\}^L \rightarrow \hat{\mathbb{G}}$ which map any string $\mathbf{m} \in \{0, 1\}^L$ to

$$H_{\mathbb{G}}(\mathbf{m}) = W_0 \cdot \prod_{k=1}^L W_k^{m[k]}, \quad H_{\hat{\mathbb{G}}}(\mathbf{m}) = \hat{W}_0 \cdot \prod_{k=1}^L \hat{W}_k^{m[k]},$$

respectively. Note that $e(g, H_{\hat{\mathbb{G}}}(\mathbf{m})) = e(H_{\mathbb{G}}(\mathbf{m}), \hat{g})$ for any $\mathbf{m} \in \{0, 1\}^L$.

4. Let $n \in \text{poly}(\lambda)$ be the desired input length. For each $j \in [n]$, choose $v_j \xleftarrow{R} \mathbb{Z}_p$ and define

$$\hat{V}_j = \hat{g}^{v_j}, \quad \hat{H}_j = (\hat{h}^j \cdot \hat{u})^{v_j} \quad \forall j \in [n].$$

5. Output the evaluation key ek and the lossy tag generation key tk , which consist of

$$ek := \left(hk_{\text{CMH}}, g, h, u, \hat{g}, \hat{h}, \hat{u}, \{W_k, \hat{W}_k\}_{k=0}^L, \{\hat{V}_j, \hat{H}_j\}_{j=1}^n \right),$$

$$tk := (td_{\text{CMH}}, \{v_j\}_{j=1}^n).$$

The tag space $\mathcal{T} = \mathcal{T}_c \times \mathcal{T}_{aux}$ is defined as a product of $\mathcal{T}_a = \{0, 1\}^*$ and

$$\mathcal{T}_c := \left\{ (\{R_i, S_i, D_i, E_i\}_{i=1}^n, r_{\text{hash}}) \mid r_{\text{hash}} \in \mathcal{R}_{\text{CMH}} \wedge \right.$$

$$\left. \forall i \in [n] : (R_i, S_i, D_i, E_i) \in \mathbb{G}^{*4} \wedge e(R_i, \hat{h}^i \cdot \hat{u}) = e(S_i, \hat{g}) \right\},$$

where $\mathbb{G}^* := \mathbb{G} \setminus \{1_{\mathbb{G}}\}$. The range of the function family is $\text{Rng}_\lambda = \mathbb{G}_T^{n+1}$ and its domain is $\mathbb{Z}_p^n$.

Lossy tag generation. $\text{LAF.LTag}(tk, t_a)$ takes in an auxiliary tag component $t_a \in \{0, 1\}^*$ and uses $tk = (td_{\text{CMH}}, \{v_j\}_{j=1}^n, \{w_k\}_{k=0}^L)$ to generate a lossy tag as follows.

1. For each $i \in [n]$, choose $r_i \xleftarrow{R} \mathbb{Z}_p^*$ and compute

$$R_i = g^{r_i}, \quad S_i = (h^i \cdot u)^{r_i} \quad \forall i \in [n]. \quad (3)$$

2. For each $i \in [n]$, choose $\rho_i \xleftarrow{R} \mathbb{Z}_p$ and compute

$$D_i = h^{r_i \cdot v_i} \cdot H_{\mathbb{G}}(\tau)^{\rho_i}, \quad E_i = g^{\rho_i} \quad \forall i \in [n],$$

where $\tau \in \{0, 1\}^L$ is chosen uniformly in the range of CMhash .

3. Use the trapdoor td_{CMH} to find $r_{\text{hash}} \in \mathcal{R}_{\text{hash}}$ such that

$$\tau = \text{CMhash}(hk_{\text{hash}}, (t_a, \{R_i, S_i, D_i, E_i\}_{i=1}^n), r_{\text{hash}}) \in \{0, 1\}^L$$

and output the tag $t = (t_c, t_a)$, where $t_c = (\{R_i, S_i, D_i, E_i\}_{i=1}^n, r_{\text{hash}})$.

Each lossy tag is associated with a matrix $(M_{i,j})_{i,j \in [n]} = (e(g, \hat{h})^{r_i \cdot v_j})_{i,j}$, which is a rank-1 matrix in the exponent. Its diagonal entries consist of

$$M_{i,i} = \frac{e(D_i, \hat{g})}{e(E_i, H_{\hat{\mathbb{G}}}(\tau))} = e(g, \hat{h})^{r_i \cdot v_i} \quad \forall i \in [n], \quad (4)$$

while its non-diagonal entries

$$M_{i,j} = \left(\frac{e(R_i, \hat{H}_j)}{e(S_i, \hat{V}_j)} \right)^{1/(j-i)} = e(g, \hat{h})^{r_i \cdot v_j} \quad \forall (i, j) \in [n] \times [n] \setminus \{(i, i)\}_{i=1}^n, \quad (5)$$

are obtained by pairing tag component (R_i, S_i) with evaluation key components $(\hat{V}_j, \hat{H}_j)$.

Random Tags. A random tag can be publicly sampled as follows.

1. For each $i \in [n]$, choose $r_i \xleftarrow{R} \mathbb{Z}_p^*$ and compute $\{R_i, S_i\}_{i=1}^n$ as in (3).
2. For each $i \in [n]$, choose $(D_i, E_i) \xleftarrow{R} \mathbb{G}^* \times \mathbb{G}^*$ uniformly at random.
3. Choose $r_{\text{hash}} \xleftarrow{R} \mathcal{R}_{\text{hash}}$.

Finally, output the tag $t = (t_c, t_a)$, where $t_c = (\{R_i, S_i, D_i, E_i\}_{i=1}^n, r_{\text{hash}})$.

We note that, in both random and lossy tags, we have $e(R_i, \hat{u}^i \cdot \hat{h}) = e(S_i, \hat{g})$ for all $i \in [n]$, so that elements of $\mathcal{T}$ are publicly recognizable.

Evaluation. $\text{LAF.Eval}(ek, t, \mathbf{x})$ takes in the input $\mathbf{x} \in \mathbb{Z}_p^n$ and the tag $t = (t_c, t_a)$. It parses t_c as $(\{R_i, S_i, D_i, E_i\}_{i=1}^n, r_{\text{hash}})$ and proceeds as follows.

1. Return $\perp$ if there exists $i \in [n]$ such that $e(R_i, \hat{h}^i \cdot \hat{u}) \neq e(S_i, \hat{g})$.

2. Compute the matrix $(M_{i,j})_{i,j \in [n]} \in \mathbb{G}_T^{n \times n}$ as

$$M_{i,i} = \frac{e(D_i, \hat{g})}{e(E_i, H_{\hat{\mathbb{G}}}(\tau))} \quad \forall i \in [n], \quad (6)$$

where $\tau = \text{CMhash}(hk_{hash}, (t_a, \{R_i, S_i, D_i, E_i\}_{i=1}^n), r_{hash})$, and

$$M_{i,j} = \left(\frac{e(R_i, \hat{H}_j)}{e(S_i, \hat{V}_j)} \right)^{1/(j-i)} \quad \forall (i,j) \in [n] \times [n] \setminus \{(i,i)\}_{i=1}^n, \quad (7)$$

Note that, since $R_i = g^{r_i}$ and $S_i = (h^i \cdot u)^{r_i}$ for some $r_i \in \mathbb{Z}_q$, we have

$$\begin{aligned} M_{i,i} &= e(g, \hat{h})^{r_i \cdot v_i + \omega_i}, & \forall i \in [n] \\ M_{i,j} &= e(g, \hat{h})^{r_i \cdot v_j}, & \forall i \neq j, \end{aligned} \quad (8)$$

for some vector $(\omega_1, \dots, \omega_n)^\top \in \mathbb{Z}_p^n$ that only contains non-zero entries if $t = (t_c, t_a)$ is injective.

3. Compute $(V_{T,j})_{j \in [n]}$ as $V_{T,j} = e(h, \hat{V}_j) = e(g, \hat{h})^{v_j}$ for each $j \in [n]$.

4. Use the input $\mathbf{x} = (x_1, \dots, x_n)^\top \in \mathbb{Z}_p^n$ to compute

$$\begin{aligned} Y_0 &= \prod_{j=1}^n V_{T,j}^{x_j} \\ Y_i &= \prod_{j=1}^n M_{i,j}^{x_j} \quad \forall i \in [n] \end{aligned} \quad (9)$$

and output $\mathbf{Y} = (Y_0, Y_1, \dots, Y_n)^\top \in \mathbb{G}_T^{n+1}$.

While the above construction inherits the $\Theta(\lambda)$ -size public keys of Waters signatures [43], we believe that it can be adapted to other signature schemes in the standard model (e.g., [8,31]) so as to obtain shorter evaluation keys.

INJECTIVITY AND LOSSINESS. For any injective tag, all entries of the vector $(\omega_1, \dots, \omega_n)^\top$ are non-zero in (8). We can use Y_0 to ensure that the function is injective. As long as $\omega_i \neq 0$ for all $i \in [n]$, the evaluation algorithm (9) yields a vector $\mathbf{Y} = (Y_0, Y_1, \dots, Y_n) \in \mathbb{G}_T^{n+1}$ of the form

$$\begin{aligned} Y_0 &= e(g, \hat{h})^{\sum_{j=1}^n v_j \cdot x_j} \\ Y_i &= e(g, \hat{h})^{\omega_i \cdot x_i + r_i \cdot \sum_{j=1}^n v_j \cdot x_j} \quad \forall i \in [n], \end{aligned}$$

meaning that $x_i \in \mathbb{Z}_p$ is uniquely determined by (Y_0, Y_i) and (R_i, D_i, E_i) (note that the triple (R_i, D_i, E_i) uniquely defines ω_i).

For any lossy tag, the evaluation outputs $\mathbf{Y} = (Y_0, Y_1, \dots, Y_n) \in \mathbb{G}_T^{n+1}$ such that

$$\begin{aligned} Y_0 &= e(g, \hat{h})^{\sum_{j=1}^n v_j \cdot x_j} \\ Y_i &= e(g, \hat{h})^{r_i \cdot \sum_{j=1}^n v_j \cdot x_j} \quad \forall i \in [n], \end{aligned}$$

which always reveals the same information $\sum_{j=1}^n v_j \cdot x_j \bmod p$ about the input vector $\mathbf{x} = (x_1, \dots, x_n)^\top$, no matter which tag is used.

3.2 Security

The proof of indistinguishability relies on the wD3DH1 assumption via a hybrid argument over the queries to the $\text{LAF.LTag}(tk, \cdot)$ oracle and over the pairs $\{(D_i, E_i)\}_{i=1}^n$ produced by $\text{LAF.LTag}(tk, \cdot)$ at each query. Using the R-wD3DH1 assumption, it is possible to modify the proof so as to use a hybrid argument over the pairs $\{(D_i, E_i)\}_{i=1}^n$ only (meaning that all queries to $\text{LAF.LTag}(tk, \cdot)$ are processed in parallel at each game transition). However, this proof would require the SXDH assumption – which only holds in asymmetric pairings – to apply the result of Lemma 1. In contrast, the proof of Theorem 1 allows instantiations in all bilinear group configurations, even in symmetric pairings.

The proof of Theorem 1 uses a hybrid argument to gradually replace pairs $\{(D_i, E_i)\}_{i=1}^n$ by truly random group elements in outputs of the lossy tag generation oracle. To this end, it relies on the proof technique of the Boneh-Boyen IBE [11] in the proof of Lemma 3. Namely, in order to embed a D3DH1 instance $(g, h, g^{v_k}, g^{r_k}, T \stackrel{?}{=} h^{r_k \cdot v_k})$ in the k -th pair (D_k, E_k) , for indexes $i > k$, the reduction has to simulate $h^{r_i \cdot v_k}$ for a known $r_i \in \mathbb{Z}_p$ and an unknown h^{v_k} .

Theorem 1. *The above LAF provides indistinguishability under the wD3DH1 assumption in $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$.*

Proof. We first recall that, for any injective or non-injective tag $t = (t_c, t_a)$, the core component $t_c = (\{R_i, S_i, D_i, E_i\}_{i=1}^n, r_{hash})$ imply a matrix $(M_{i,j})_{i,j \in [n]}$ where the off-diagonal entries are $M_{i,j} = e(g, \hat{h})^{r_i \cdot v_j}$ and the diagonal entries are of the form (8). In injective tags, the vector $(\omega_1, \dots, \omega_n)^\top \in \mathbb{Z}_p^n$ only contains non-zero entries. In lossy tags, we have $(\omega_1, \dots, \omega_n)^\top = \mathbf{0}^n$. We define a sequence of hybrid games. In $\text{Game}_{(0,0)}$, the adversary has access to the real oracle $\text{LAF.LTag}(tk, \cdot)$ oracle that always outputs lossy tags. In $\text{Game}_{(Q,n)}$, the adversary is given access to an oracle $\mathcal{O}_{\mathcal{T}_c}(\cdot)$ that always outputs random tags.

Game $_{(\ell,k)}$ ($1 \leq \ell \leq Q, 1 \leq k \leq n$): In this game, the adversary interacts with a hybrid oracle $\text{LAF.LTag}^{(\ell,k)}(tk, \cdot)$. At the μ -th query, this oracle outputs tags $t^{(\mu)} = (t_c^{(\mu)}, t_a^{(\mu)})$ such that

- If $\mu < \ell$, the tag $t_c^{(\mu)} = (\{R_i, S_i, D_i, E_i\}_{i=1}^n, r_{hash})$ implies a matrix $(M_{i,j}^{(\mu)})_{i,j \in [n]}$ of the form (8) where $(\omega_1^{(\mu)}, \dots, \omega_n^{(\mu)})^\top$ is uniform over $\mathbb{Z}_p^n$
- If $\mu = \ell$, $t_c^{(\mu)} = (\{R_i, S_i, D_i, E_i\}_{i=1}^n, r_{hash})$ implies a matrix $(M_{i,j}^{(\mu)})_{i,j \in [n]}$ of the form (8) where the first k entries of $(\omega_1^{(\mu)}, \dots, \omega_n^{(\mu)})^\top$ are uniform over $\mathbb{Z}_p$ and its last $n - k$ entries are zeroes.
- If $\mu > \ell$, the matrix $(M_{i,j}^{(\mu)})_{i,j \in [n]}$ implied by the core tag component $t_c^{(\mu)} = (\{R_i, S_i, D_i, E_i\}_{i=1}^n, r_{hash})$ is a rank-1 matrix in the exponent since $(\omega_1^{(\mu)}, \dots, \omega_n^{(\mu)})^\top = \mathbf{0}^n$.

Lemma 3 shows that, for all pairs $(\ell, k) \in [Q] \times [n]$, these games are computationally indistinguishable from one another, which yields the stated result. $\square$

Lemma 3. For each $k \in [n]$ and $\ell \in [Q]$, $\text{Game}_{(\ell,k)}$ is computationally indistinguishable from $\text{Game}_{(\ell,k-1)}$ if the wD3DH1 assumption holds. Under the same assumption, $\text{Game}_{(\ell,1)}$ is computationally indistinguishable from $\text{Game}_{(\ell-1,n)}$.

Proof. For the sake of contradiction, assume that there exists $\ell \in [Q]$, $k \in [n]$ such that the adversary $\mathcal{A}$ can distinguish $\text{Game}_{(\ell,k)}$ from $\text{Game}_{(\ell,k-1)}$ with noticeable advantage (the indistinguishability of $\text{Game}_{(\ell-1,n)}$ and $\text{Game}_{(\ell,1)}$ can be proved in a completely similar way). We build a wD3DH1 distinguisher $\mathcal{B}$ that inputs $(g, \hat{g}, g^a, g^b, g^c, \hat{g}^b, \hat{g}^c, T)$ with the goal of deciding if $T = g^{abc}$ or $T \in_R \mathbb{G}$.

To this end, $\mathcal{B}$ defines $h = g^b$, $\hat{h} = \hat{g}^b$ and $\hat{V}_k = \hat{g}^c$. It picks $\alpha \xleftarrow{R} \mathbb{Z}_p$ and defines $\hat{u} = \hat{h}^{-k} \cdot \hat{g}^\alpha$ as well as $u = h^{-k} \cdot g^\alpha$, which implicitly sets $v_k = c$. This allows defining

$$\hat{H}_k = (\hat{h}^k \cdot \hat{u})^c = (\hat{g}^c)^\alpha,$$

In addition, $\mathcal{B}$ defines $(W_0, W_1, \dots, W_L) \in \mathbb{G}^{L+1}$ and $(\hat{W}_0, \hat{W}_1, \dots, \hat{W}_L) \in \hat{\mathbb{G}}^{L+1}$ by setting

$$W_i = (g^b)^{\alpha_i} \cdot g^{\beta_i}, \quad \hat{W}_i = (\hat{g}^b)^{\alpha_i} \cdot \hat{g}^{\beta_i} \quad \forall i \in \{0, \dots, L\}$$

for randomly chosen $\alpha_0, \dots, \alpha_L \xleftarrow{R} \mathbb{Z}_p$, $\beta_0, \dots, \beta_L \xleftarrow{R} \mathbb{Z}_p$. Then, $\mathcal{B}$ chooses $v_i \xleftarrow{R} \mathbb{Z}_p$ for each $i \in [n] \setminus \{k\}$ and defines the rest of the evaluation key ek by setting

$$\hat{V}_i = \hat{g}^{v_i}, \quad \hat{H}_i = (\hat{h}^i \cdot \hat{u})^{v_i}, \quad \forall i \in [n] \setminus \{k\}$$

Then, at each invocation of the $\text{LAF.LTag}(tk, \cdot)$ oracle, $\mathcal{B}$ responds as follows. At the μ -th query $t_a^{(\mu)}$, it generates a core tag $t_c^{(\mu)}$ such that

- If $\mu < \ell$, $t_c^{(\mu)} = (\{R_i, S_i, D_i, E_i\}_{i=1}^n, r_{hash})$ contains $\{\hat{D}_i, \hat{E}_i\}_{i=1}^n$ uniformly random pairs whereas $\{R_i, \hat{S}_i\}_{i=1}^n$ are chosen as in the real algorithm sampling random tags.
- If $\mu = \ell$, $t_c^{(\mu)} = (\{R_i, S_i, D_i, E_i\}_{i=1}^n, r_{hash})$ is generated as follows. It sets

$$R_k = g^a, \quad S_k = (g^a)^\alpha.$$

As for indexes $i \neq k$, it chooses $r_1, \dots, r_{k-1}, r_{k+1}, \dots, r_n \xleftarrow{R} \mathbb{Z}_p$ and sets

$$R_i = g^{r_i}, \quad S_i = (h^i \cdot u)^{r_i} \quad \forall i \in [n] \setminus \{k\}.$$

It generates the pairs $\{D_i, E_i\}_{i=1}^n$ by choosing $(D_i, E_i) \xleftarrow{R} \mathbb{G}^2$ at random for each $i \in [k-1]$. The k -th pair (D_k, E_k) is defined as

$$D_k = T \cdot H_{\mathbb{G}}(\tau)^{\rho_k}, \quad E_k = g^{\rho_k}. \quad (10)$$

for a randomly chosen $\rho_k \xleftarrow{R} \mathbb{Z}_p$. As for $\{D_i, E_i\}_{i=k+1}^n$, they are obtained by choosing a random $\tau = \tau[1] \dots \tau[L] \in \{0, 1\}^L$ in the range of CMhash and choosing $\rho_i \xleftarrow{R} \mathbb{Z}_p$ before setting

$$\begin{aligned} D_i &= H_{\mathbb{G}}(\tau)^{\rho_i} \cdot (g^c)^{-r_i \cdot \frac{\beta_0 + \sum_{i=1}^L \beta_i \cdot \tau[i]}{\alpha_0 + \sum_{i=1}^L \alpha_i \cdot \tau[i]}} \\ E_i &= g^{\rho_i} \cdot (g^c)^{-\frac{r_i}{\alpha_0 + \sum_{i=1}^L \alpha_i \cdot \tau[i]}} \end{aligned} \quad (11)$$

which can be written

$$\begin{aligned} D_i &= g^{bc \cdot r_i} \cdot H_{\mathbb{G}}(\tau)^{\tilde{\rho}_i} = h^{v_k \cdot r_i} \cdot H_{\mathbb{G}}(\tau)^{\tilde{\rho}_i} \\ E_i &= g^{\tilde{\rho}_i} \end{aligned}$$

if we define $\tilde{\rho}_i = \rho_i - \frac{c \cdot r_i}{\alpha_0 + \sum_{i=1}^L \alpha_i \cdot \tau[i]}$. Note that the reduction $\mathcal{B}$ fails if $\alpha_0 + \sum_{i=1}^L \alpha_i \cdot \tau[i] = 0$ but this only occurs with negligible chance since the coordinates $(\alpha_0, \dots, \alpha_L) \in \mathbb{Z}_p^L$ are independent of $\mathcal{A}$'s view. Finally, $\mathcal{B}$ uses the trapdoor td_{CMH} of the chameleon hash function to find coins $r_{hash} \in \mathcal{R}_{\text{CMH}}$ such that $\tau = \text{CMhash}(hk_{hash}, (t_a, \{R_i, S_i, D_i, E_i\}_{i=1}^n), r_{hash})$.

- If $\mu > \ell$, the tags are generated as lossy tags. To this end, $\mathcal{B}$ proceeds as in the previous case, except that all elements $\{D_i, E_i\}_{i=1}^n$ (and not only the last $n - k$ ones) are generated as per (11).

It is easy to see that, if $T = g^{abc}$, the pair (D_k, E_k) of (10) can be written

$$D_k = h^{v_k \cdot r_k} \cdot H_{\mathbb{G}}(\tau)^{\rho_k}, \quad E_k = g^{\rho_k},$$

meaning that $\mathcal{A}$'s view is the same as in $\text{Game}_{(\ell, k-1)}$. In contrast, if $T \in_R \mathbb{G}$, then (D_k, E_k) can be written

$$D_k = h^{\omega_k + v_k \cdot r_k} \cdot H_{\mathbb{G}}(\tau)^{\rho_k}, \quad E_k = g^{\rho_k},$$

for some uniformly random $\omega_k \in_R \mathbb{Z}_p$. In this case, $\mathcal{A}$'s view corresponds to $\text{Game}_{(\ell, k)}$. $\square$

The evasiveness property is established by Theorem 2 for which a proof is given in Appendix A.

Theorem 2. *The above LAF provides evasiveness assuming that: (i) CMH is a collision-resistant chameleon hash function; (ii) The wD3DH1 and 2-3-CDH assumptions both hold in $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$.*

Recall that the wD3DH1 and 2-3-CDH assumptions are implied by the D3DH1 and D3DH2 assumptions, respectively. Theorem 1 and Theorem 2 thus guarantee the D3DH1 and D3DH2 assumptions suffice to ensure the indistinguishability and evasiveness properties of our LAF construction (indeed, chameleon hash functions also exist under these assumptions).

3.3 Towards All-But-Many Lossy Trapdoor Functions

Our LAF construction can be modified to construct an all-but-many lossy trapdoor function [25]. Recall that ABM-LTFs do not require evaluations on lossy tags to always output the same information about the input: on any lossy tag, the image size is only required to be much smaller. On the other hand, ABM-LTFs require that, for any injective tag, the function be efficiently invertible using a trapdoor.

Our construction can be turned into an ABM-LTF in the following way. In the evaluation algorithm, a binary input vector $\mathbf{x} = (x_1, \dots, x_n)^\top \in \{0, 1\}^n$ is mapped to the output $(Y_0, \dots, Y_n) \in \mathbb{G}_T^{n+1}$, where

$$Y_0 = \prod_{i=1}^n e(R_i, \hat{h})^{x_i}$$

$$Y_j = \prod_{i=1}^n M_{i,j}^{x_i} \quad \forall j \in [n],$$

which can be written

$$Y_0 = e(g, \hat{h})^{\sum_{i=1}^n r_i \cdot x_i}$$

$$Y_j = e(g, \hat{h})^{\omega_j \cdot x_j + v_j \cdot \sum_{i=1}^n r_i \cdot x_i} \quad \forall j \in [n].$$

Using $ik = (v_1, \dots, v_n) \in \mathbb{Z}_p^n$ as an inversion key, one can decode the j -th input bit as $x_j = 0$ (resp. $x_j = 1$) if $Y_j/Y_0^{v_j} = 1_{\mathbb{G}_T}$ (resp. $Y_j/Y_0^{v_j} \neq 1_{\mathbb{G}_T}$).

Unfortunately, the above ABM-LTF does not seem immediately usable in the application to selective-opening chosen-ciphertext security, which was suggested in [25]. The reason is that our tags have a special and publicly recognizable structure, where (R_i, S_i) both depend on the same exponent $r_i \in \mathbb{Z}_p$. In the selective-opening setting, the problem arises when the adversary chooses to corrupt some senders, at which point the reduction should reveal the random coins used to create lossy/injective tags. In our construction, this would entail to reveal $r_i \in \mathbb{Z}_p$, which is incompatible with our proofs of indistinguishability and evasiveness. In the ABM-LTF constructions of [25,37], lossy tags are explainable because they are pseudorandom, which allows the reduction to pretend that they have been randomly sampled in their ambient space. Here, the special structure of lossy/injective tags prevents us from explaining the generation of lossy tags in the same way for corrupted senders. The only apparent way to sample a pair (R_i, S_i) satisfying $e(R_i, \hat{h}^i \cdot \hat{u}) = e(S_i, \hat{g})$ is to choose $r_i \in \mathbb{Z}_p$ and compute $(R_i, S_i) = (g^{r_i}, (h^i \cdot u)^{r_i})$.

We thus leave it as an open problem to build an ABM-LTF with explainable linear-size tags under DDH-like assumptions.

4 A Lossy Algebraic Filter With Tight Security

In this section, we modify our first LAF construction in such a way that we can prove it tightly secure under constant-size assumptions.⁵ To this end, we replace Waters signatures by a variant of the MAC described by Blazy, Kiltz and Pan [7], which is itself inspired by the Naor-Reingold PRF [39].

⁵ While the assumption of Definition 5 is described using $O(Q)$ group elements, it tightly reduces to wD3DH1 and DDH which both take a constant number of group elements to describe.

4.1 A Variant of the BKP MAC

The MAC construction below is identical to the signature scheme implied by the IBE scheme of [7, Appendix D] with two differences which prevent public verification in order to obtain a pseudo-random MAC instead of a digital signature. The signature scheme of [7] was actually designed by transposing a pseudo-random MAC from standard DDH-hard groups to bilinear groups in order to enable public verification. Here, we cannot immediately use the MAC of [7] because we need bilinear maps in the evaluation algorithm of our LAF.

In order to obtain a pseudo-random MAC, we thus modify the signature scheme of [7] by introducing an additional randomizer $r \in \mathbb{Z}_p$ and an extra group element h , of which the discrete logarithm $\log_g(h)$ serves as a private verification key.

Keygen($1^\lambda, 1^L$): Given a security parameter λ and a message length $L \in \text{poly}(\lambda)$, choose asymmetric bilinear groups $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$ of prime order $p > 2^\lambda$ with generators $g, h \xleftarrow{R} \mathbb{G}$, $\hat{g} \xleftarrow{R} \hat{\mathbb{G}}$.

1. Choose $\theta, \alpha, \beta \xleftarrow{R} \mathbb{Z}_p$ and compute $\hat{g}^\theta \in \hat{\mathbb{G}}$. For each $\mu \in \{0, 1\}$, choose vectors $\mathbf{x}_\mu = (x_{1,\mu}, \dots, x_{L,\mu}) \xleftarrow{R} \mathbb{Z}_p^L$, $\mathbf{y}_\mu = (y_{1,\mu}, \dots, y_{L,\mu}) \xleftarrow{R} \mathbb{Z}_p^L$.
2. Set $v = \alpha + \theta \cdot \beta$ and $\mathbf{z}_\mu = \mathbf{x}_\mu + \theta \cdot \mathbf{y}_\mu \in \mathbb{Z}_p^L$. Compute $\hat{V} = \hat{g}^v$ and, for each $\mu \in \{0, 1\}$, define $\hat{\mathbf{Z}}_\mu = (\hat{Z}_{1,\mu}, \dots, \hat{Z}_{L,\mu}) = \hat{g}^{\mathbf{z}_\mu}$.

Output a secret key $\text{sk}_{\text{mac}} = (\alpha, \beta, \mathbf{x}_0, \mathbf{x}_1, \mathbf{y}_0, \mathbf{y}_1, \eta)$, where $\eta = \log_g(h)$, and public parameters consisting of $\text{pp} = ((\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T), g, \hat{g}, h, \hat{g}^\theta, (\hat{V}, \hat{\mathbf{Z}}_0, \hat{\mathbf{Z}}_1))$.

Mac.Sig($\text{pp}, \text{sk}_{\text{mac}}, M$): To generate a MAC for $M = m[1] \dots m[L] \in \{0, 1\}^L$ using $\text{sk}_{\text{mac}} = (x, y, \mathbf{x}_0, \mathbf{x}_1, \mathbf{y}_0, \mathbf{y}_1, \eta)$, choose $r, \rho \xleftarrow{R} \mathbb{Z}_p$ and compute

$$\begin{aligned}\sigma_1 &= h^{\alpha \cdot r} \cdot g^{\rho \cdot (\sum_{k=1}^L x_{k,m[k]})} \\ \sigma_2 &= h^{\beta \cdot r} \cdot g^{\rho \cdot (\sum_{k=1}^L y_{k,m[k]})} \\ \sigma_3 &= g^\rho \\ \sigma_4 &= g^r\end{aligned}$$

Mac.Ver($\text{pp}, \text{sk}_{\text{mac}}, M, \sigma$): Given $\text{sk}_{\text{mac}} = (\alpha, \beta, \mathbf{x}_0, \mathbf{x}_1, \mathbf{y}_0, \mathbf{y}_1, \eta)$ and an L -bit message $M = m[1] \dots m[L]$, a purported MAC $\sigma = (\sigma_1, \sigma_2, \sigma_3, \sigma_4)$ is accepted if and only if

$$e(\sigma_1, \hat{g}) \cdot e(\sigma_2, \hat{g}^\theta) = e(\sigma_4, \hat{V})^\eta \cdot e(\sigma_3, \prod_{k=1}^L \hat{Z}_{k,m[k]}). \quad (12)$$

We note that the verification algorithm can be modified in such a way that it does not require any pairing evaluation. The above description is just meant to simplify the presentation of the security proof of our LAF construction.

The proof is essentially identical to that of [7] but we give it for completeness. We note that, in the security definitions of MACs, the adversary is generally allowed to make verification queries. Here, for simplicity, we prove unforgeability in a game where the adversary knows $\eta = \log_g(h)$, which allows it to run the verification oracle itself. This dispenses us with the need for a verification oracle.

Lemma 4. *The above construction is an unforgeable MAC assuming that the SXDH assumption holds in $(\mathbb{G}, \hat{\mathbb{G}})$. Namely, any forger $\mathcal{A}$ making Q MAC queries within running time $t_{\mathcal{A}}$ has advantage at most*

$$\mathbf{Adv}_{\mathcal{A}}^{\text{uf-mac}}(\lambda) \leq \mathbf{Adv}_{\mathcal{B}_1}^{\text{DDH}_2}(\lambda) + 2L \cdot \mathbf{Adv}_{\mathcal{B}_2}^{\text{DDH}_1}(\lambda),$$

where $\mathcal{B}_1$ and $\mathcal{B}_2$ are PPT distinguishers against the DDH assumption in $\mathbb{G}_1$ and $\mathbb{G}_2$, respectively, which run in time $t_{\mathcal{A}} + Q \cdot \text{poly}(\lambda)$.

Proof. To prove the result, we consider a sequence of games. For each index i , we call W_i the event that the challenger outputs 1 in **Game_i**.

Game₀: This is the real game MAC security game, where the adversary $\mathcal{A}$ is additionally given $\eta = \log_g(h)$ in such a way that it can run the verification algorithm (and test whether equation (12) holds) by itself. The challenger outputs 1 if and only if $\mathcal{A}$ eventually outputs a pair $(M^*, \sigma^* = (\sigma_1^*, \sigma_2^*, \sigma_3^*, \sigma_4^*))$ satisfying

$$e(\sigma_1^*, \hat{g}) \cdot e(\sigma_2^*, \hat{g}^\theta) = e(\sigma_4^*, \hat{V})^\eta \cdot e(\sigma_3^*, \prod_{k=1}^L \hat{Z}_{k, m^*[k]}), \quad (13)$$

where $M^* = m^*[1] \dots m^*[L] \in \{0, 1\}^L$, although M^* was not previously queried to the MAC oracle. By definition, $\Pr[W_0] = \mathbf{Adv}_{\mathcal{A}}^{\text{uf-mac}}(\lambda)$.

Game₁: In this game, we modify again the verification oracle as follows. When $\mathcal{A}$ outputs a pair $(M^*, \sigma^* = (\sigma_1^*, \sigma_2^*, \sigma_3^*, \sigma_4^*))$ such that M^* was not queried to the MAC oracle but (M^*, σ^*) still satisfies (13), the challenger checks if

$$\begin{aligned} \sigma_1^* &= \sigma_4^{\star \eta \cdot \alpha} \cdot \sigma_3^{\star \sum_{k=1}^L x_{k, m^*[k]}} \\ \sigma_2^* &= \sigma_4^{\star \eta \cdot \beta} \cdot \sigma_3^{\star \sum_{k=1}^L y_{k, m^*[k]}}. \end{aligned} \quad (14)$$

We call E_1 the event that equalities (14) are satisfied. If they are not satisfied, the challenger outputs 0. Otherwise, it outputs 1 as it did in **Game₀**. If we denote by E_0 the analogue of event E_1 in **Game₀**, we have

$$\begin{aligned} \Pr[W_0] &= \Pr[W_0 \wedge E_0] + \Pr[W_0 \wedge \neg E_0] \\ &= \Pr[W_1 \wedge E_1] + \Pr[W_0 \wedge \neg E_0] = \Pr[W_1] + \Pr[W_0 \wedge \neg E_0] \end{aligned}$$

since $\Pr[W_1 \wedge \neg E_1] = 0$. Lemma 5 shows that event $W_0 \wedge \neg E_0$ would contradict the DDH assumption in $\hat{\mathbb{G}}$: namely, we have $\Pr[W_0 \wedge \neg E_0] \leq \mathbf{Adv}^{\text{DDH}_2}(\lambda)$, which implies $|\Pr[W_1] - \Pr[W_0]| \leq \mathbf{Adv}^{\text{DDH}_2}(\lambda)$.

We now use a sub-sequence of L hybrid games over the input bits of queried messages. For convenience, we define **Game_{2,0}** to be identical to **Game₁**.

Game_{2,i} ($1 \leq i \leq L$): In this sub-sequence of games, we modify the key generation phase and the MAC oracle in the following way.

- At the beginning of the game, the challenge defines $\hat{V} = \hat{g}^v$ for a random $v \xleftarrow{R} \mathbb{Z}_p$.
- MAC queries are handled as follows. Let $R : \{0, 1\}^i \rightarrow \mathbb{Z}_p$ be a truly random function mapping i -bit input to $\mathbb{Z}_p$. At each message M queried by $\mathcal{A}$, the challenger computes $(\sigma_3, \sigma_4) = (g^\rho, g^r)$ for random $\rho, r \xleftarrow{R} \mathbb{Z}_p$. Then, it outputs $(\sigma_1, \sigma_2, \sigma_3, \sigma_4)$, where

$$\begin{aligned}\sigma_1 &= h^{(v - \theta \cdot R(m[1] \dots m[i])) \cdot r} \cdot g^{\rho \cdot (\sum_{k=1}^L x_{k, m^*[k]})} \\ \sigma_2 &= h^{R(m[1] \dots m[i]) \cdot r} \cdot g^{\rho \cdot (\sum_{k=1}^L y_{k, m^*[k]})}\end{aligned}$$

When the adversary outputs $(M^*, \sigma^* = (\sigma_1^*, \sigma_2^*, \sigma_3^*, \sigma_4^*))$ satisfying (13) for a new message M^* , the challenger checks if the following equalities are satisfied:

$$\begin{aligned}\sigma_1^* &= \sigma_4^{*\eta \cdot (v - \theta \cdot R(m^*[1] \dots m^*[i]))} \cdot \sigma_3^{*\sum_{k=1}^L x_{k, m^*[k]}} \\ \sigma_2^* &= \sigma_4^{*\eta \cdot R(m^*[1] \dots m^*[i])} \cdot \sigma_3^{*\sum_{k=1}^L y_{k, m^*[k]}}.\end{aligned}\tag{15}$$

If so, the challenger outputs 1. Otherwise, it outputs 0. Lemma 6 shows that $\text{Game}_{2,i}$ is indistinguishable from $\text{Game}_{2,(i-1)}$ under the DDH assumption in $\mathbb{G}$. Namely, $|\Pr[W_{2,i}] - \Pr[W_{2,(i-1)}]| \leq \mathbf{Adv}^{\text{DDH}_1}(\lambda)$.

In $\text{Game}_{2,L}$, we claim that $\Pr[W_{2,L}] = 1/p$. Indeed, the equalities (15) can only hold by pure chance when $i = L$ because $m^*[1] \dots m^*[L]$ was never involved in an output of the MAC oracle. Hence, the random function output $R(m^*[1] \dots m^*[L])$ is perfectly independent of $\mathcal{A}$'s view. Since $\Pr[W_{2,0}] = \Pr[W_1]$, we obtain the claimed upper bound for $\Pr[W_0]$. $\square$

Lemma 5. *In Game_0 , we have $\Pr[W_0 \wedge \neg E_0] \leq \mathbf{Adv}^{\text{DDH}_2}(\lambda)$.*

Proof. Towards a contradiction, let us assume that, in Game_1 , the adversary $\mathcal{A}$ can output a pair $(M^*, \sigma^* = (\sigma_1^*, \sigma_2^*, \sigma_3^*, \sigma_4^*))$ satisfying (13) but not (14). We construct a distinguisher $\mathcal{B}$ for the DDH assumption in $\hat{\mathbb{G}}$. Our distinguisher $\mathcal{B}$ takes as input $(\hat{g}, \hat{g}^\theta, \hat{g}^\omega, \hat{T}) \in \hat{\mathbb{G}}^4$ and decides if $\hat{T} = \hat{g}^{\alpha \cdot \omega}$ or $\hat{T} \in_R \hat{\mathbb{G}}$. To this end, $\mathcal{B}$ will compute a pair of the form $(w, w^\theta) \in \mathbb{G}^2$ with $w \neq 1_{\mathbb{G}}$, which allows solving the given DDH instance in $\hat{\mathbb{G}}$ by testing if $e(w, \hat{T}) = e(w^\theta, \hat{g}^\omega)$. Indeed, the latter equality holds if and only if $\hat{T} = \hat{g}^{\alpha \cdot \omega}$.

The reduction $\mathcal{B}$ runs the real key generation algorithm and answers all MAC and verification queries exactly as in Game_1 . By hypothesis, $\mathcal{B}$ has non-negligible probability of outputting a pair $(M^*, \sigma^* = (\sigma_1^*, \sigma_2^*, \sigma_3^*, \sigma_4^*))$ satisfying (13) although

$$\sigma_1^* \neq \sigma_4^{*\eta \cdot \alpha} \cdot \sigma_3^{*\sum_{k=1}^L x_{k, m^*[k]}}, \quad \sigma_2^* \neq \sigma_4^{*\eta \cdot \beta} \cdot \sigma_3^{*\sum_{k=1}^L y_{k, m^*[k]}}.$$

At this point, $\mathcal{B}$ uses sk_{mac} to construct a different valid MAC $(\sigma'_1, \sigma'_2, \sigma_3^*, \sigma_4^*)$ satisfying (13) and such that $(\sigma'_1, \sigma'_2) \neq (\sigma_1^*, \sigma_2^*)$. Namely, $\mathcal{B}$ computes

$$\sigma'_1 = \sigma_4^{*\eta \cdot \alpha} \cdot \sigma_3^{*\sum_{k=1}^L x_{k, m^*[k]}}, \quad \sigma'_2 = \sigma_4^{*\eta \cdot \beta} \cdot \sigma_3^{*\sum_{k=1}^L y_{k, m^*[k]}}.$$

By dividing the two verification equations for $(\sigma'_1, \sigma'_2, \sigma_3^*, \sigma_4^*)$ and $(\sigma_1^*, \sigma_2^*, \sigma_3^*, \sigma_4^*)$, we get

$$e(\sigma_1^*/\sigma'_1, \hat{g}) \cdot e(\sigma_2^*/\sigma'_2, \hat{g}^\theta) = 1_{\mathbb{G}_T}$$

meaning that $\sigma_1^*/\sigma'_1 = (\sigma_2^*/\sigma'_2)^\theta$. Since $\sigma_1^* \neq \sigma'_1$, this provides $\mathcal{B}$ with a non-trivial pair $(w, w^\theta) = (\sigma_2^*/\sigma'_2, \sigma_1^*/\sigma'_1)$, which is sufficient to solve DDH in $\hat{\mathbb{G}}$. $\square$

Lemma 6. *Under the DDH assumption in $\mathbb{G}$, the challenger outputs 1 with about the same probabilities in $\text{Game}_{3,(i-1)}$ and $\text{Game}_{3,i}$. We have*

$$|\Pr[W_{2,i}] - \Pr[W_{2,(i-1)}]| \leq 2 \cdot \text{Adv}^{\text{DDH}_1}(\lambda).$$

(The proof is given in Appendix B.1.)

4.2 The LAF Construction

In order to apply a hybrid argument in our proof of indistinguishability, we need to use n instances of the MAC of Section 4.1, each of which has its own secret key $\text{sk}_{\text{mac},j}$ and its own set of public parameters $\text{pp}_j = (g, \hat{g}, h, \hat{g}^{\theta_j}, (\hat{V}_j, \hat{\mathbf{Z}}_{j,0}, \hat{\mathbf{Z}}_{j,1}))$. As a result, we need an evaluation key containing $\Theta(n \cdot L)$ group elements. We leave it as an open problem to shorten the evaluation while retaining tight security and short tags.

Key generation. $\text{LAF.Gen}(1^\lambda)$ conducts the following steps.

1. Choose asymmetric bilinear groups $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$ of prime order $p > 2^\lambda$ with generators $g, h \xleftarrow{R} \mathbb{G}$, $\hat{g} \xleftarrow{R} \hat{\mathbb{G}}$ and let $\eta = \log_g(h)$.
2. Choose a chameleon hash function $\text{CMH} = (\text{CMKg}, \text{CMhash}, \text{CMswitch})$, where the hashing algorithm $\text{CMhash} : \{0, 1\}^* \times \mathcal{R}_{\text{hash}} \rightarrow \{0, 1\}^L$ has output length $L \in \text{poly}(\lambda)$. Generate a pair $(hk_{\text{CMH}}, td_{\text{CMH}}) \leftarrow \text{CMKg}(1^\lambda)$ made of a hashing key hk_{CMH} and a trapdoor td_{CMH} .
3. Generate n keys for the MAC of Section 4.1 which all share the same parameters $g, h \in \mathbb{G}$, $\hat{g} \in \hat{\mathbb{G}}$. Namely, for each $j \in [n]$, conduct the following steps.
 - a. Choose $\theta_j \xleftarrow{R} \mathbb{Z}_p$ and compute $\hat{g}^{\theta_j} \in \hat{\mathbb{G}}$.
 - b. For each $\mu \in \{0, 1\}$, choose vectors $\mathbf{x}_{j,\mu} = (x_{j,1,\mu}, \dots, x_{j,L,\mu}) \xleftarrow{R} \mathbb{Z}_p^L$ and $\mathbf{y}_{j,\mu} = (y_{j,1,\mu}, \dots, y_{j,L,\mu}) \xleftarrow{R} \mathbb{Z}_p^L$.
 - c. Compute $\mathbf{z}_{j,\mu} = \mathbf{x}_{j,\mu} + \theta_j \cdot \mathbf{y}_{j,\mu}$ and $\hat{\mathbf{Z}}_{j,\mu} = \hat{g}^{\mathbf{z}_{j,\mu}} = (\hat{g}^{z_{j,1,\mu}}, \dots, \hat{g}^{z_{j,L,\mu}})$ for each $\mu \in \{0, 1\}$.
 - d. Choose $\alpha_j, \beta_j \xleftarrow{R} \mathbb{Z}_p$ and compute $\hat{V}_j = \hat{g}^{\alpha_j + \theta_j \cdot \beta_j}$.
 - e. Define $\text{sk}_{\text{mac},j} = (\alpha_j, \beta_j, \mathbf{x}_{j,0}, \mathbf{x}_{j,1}, \mathbf{y}_{j,0}, \mathbf{y}_{j,1})$.
4. Choose $u \xleftarrow{R} \mathbb{G}$ and $\hat{u} \xleftarrow{R} \hat{\mathbb{G}}$ subject to the constraints $\log_g(h) = \log_{\hat{g}}(\hat{h})$ and $\log_g(u) = \log_{\hat{g}}(\hat{u})$.

5. Define

$$\hat{H}_j = (\hat{h}^j \cdot \hat{u})^{\alpha_j + \theta_j \cdot \beta_j} \quad \forall j \in [n].$$

6. Output the evaluation key ek and the lossy tag generation key tk , which consist of

$$\begin{aligned} ek &:= (g, h, u, \hat{g}, \hat{h}, \hat{u}, \{\hat{g}^{\theta_j}\}_{j=1}^n, \{\hat{Z}_{j,\mu}\}_{j \in [n], \mu \in \{0,1\}}, \{\hat{V}_j, \hat{H}_j\}_{j=1}^n, hk_{\text{CMH}}), \\ tk &:= (\{\text{sk}_{\text{mac},j}\}_{j=1}^n, \eta, td_{\text{CMH}}). \end{aligned}$$

The tag space $\mathcal{T} = \mathcal{T}_c \times \mathcal{T}_{aux}$ is defined as a product of $\mathcal{T}_a = \{0,1\}^*$ and

$$\begin{aligned} \mathcal{T}_c &:= \{(\{R_i, S_i, D_i, E_i, F_i\}_{i=1}^n, r_{hash}) \mid r_{hash} \in \mathcal{R}_{hash} \wedge \\ &\quad \forall i \in [n] : (R_i, S_i, D_i, E_i, F_i) \in \mathbb{G}^5 \wedge e(R_i, \hat{h}^i \cdot \hat{u}) = e(S_i, \hat{g})\}. \end{aligned}$$

The range of the function family is $\text{Rng}_\lambda = \mathbb{G}_T^{n+1}$ and its domain is $\mathbb{Z}_p^n$.

Lossy tag generation. $\text{LAF.LTag}(tk, t_a)$ takes in an auxiliary tag component $t_a \in \{0,1\}^*$ and uses $tk = (\{\text{sk}_{\text{mac},j}\}_{j=1}^n, \eta)$ to generate a lossy tag as follows.

1. For each $i \in [n]$, choose $r_i \xleftarrow{R} \mathbb{Z}_p$ and compute

$$R_i = g^{r_i}, \quad S_i = (h^i \cdot u)^{r_i} \quad \forall i \in [n]. \quad (16)$$

2. Choose a random string $\tau \in \{0,1\}^L$ in the range of CMhash . Then, for each $i \in [n]$, choose $\rho_i \xleftarrow{R} \mathbb{Z}_p$ and compute

$$\begin{aligned} D_i &= h^{\alpha_i \cdot r_i} \cdot g^{\rho_i \cdot (\sum_{k=1}^L x_{i,k,\tau[k]})}, \\ E_i &= h^{\beta_i \cdot r_i} \cdot g^{\rho_i \cdot (\sum_{k=1}^L y_{i,k,\tau[k]})}, \\ F_i &= g^{\rho_i}. \end{aligned} \quad \forall i \in [n] \quad (17)$$

3. Use the trapdoor td_{CMH} of the chameleon hash function to find random coins $r_{hash} \in \mathcal{R}_{hash}$ such that

$$\tau = \text{CMhash}(hk_{\text{CMH}}, (t_a, \{R_i, S_i, D_i, E_i, F_i\}_{i=1}^n), r_{hash}) \in \{0,1\}^L.$$

4. Output the tag $t = (t_c, t_a)$, where $t_c = (\{R_i, S_i, D_i, E_i, F_i\}_{i=1}^n, r_{hash})$.

Each lossy tag corresponds to a matrix $(M_{i,j})_{i,j \in [n]} = (e(g, \hat{h})^{r_i \cdot (\alpha_j + \theta_j \cdot \beta_j)})_{i,j}$, which forms a rank-1 matrix in the exponent. Its diagonal entries consist of

$$M_{i,i} = \frac{e(D_i, \hat{g}) \cdot e(E_i, \hat{g}^{\theta_i})}{e(F_i, \prod_{k=1}^L \hat{Z}_{i,k,\tau[k]})} = e(g, \hat{h})^{r_i \cdot (\alpha_i + \theta_i \cdot \beta_i)} \quad \forall i \in [n], \quad (18)$$

while its non-diagonal entries

$$\begin{aligned} M_{i,j} &= \left(\frac{e(R_i, \hat{H}_j)}{e(S_i, \hat{V}_j)} \right)^{1/(j-i)} \\ &= e(g, \hat{h})^{r_i \cdot (\alpha_j + \theta_j \cdot \beta_j)} \quad \forall (i,j) \in [n] \times [n] \setminus \{(i,i)\}_{i=1}^n, \end{aligned} \quad (19)$$

are obtained by pairing tag component (R_i, S_i) with evaluation key components $(\hat{V}_j, \hat{H}_j)$.

Random Tags. A random tag can be publicly sampled as follows.

1. For each $i \in [n]$, choose $r_i \xleftarrow{R} \mathbb{Z}_p$ and compute $\{R_i, S_i\}_{i=1}^n$ as in (16).
2. For each $i \in [n]$, choose $(D_i, E_i, F_i) \xleftarrow{R} \mathbb{G}^3$ uniformly at random.
3. Choose $r_{hash} \xleftarrow{R} \mathcal{R}_{hash}$.

Output the tag $t = (t_c, t_a)$, where $t_c = (\{R_i, S_i, D_i, E_i, F_i\}_{i=1}^n, r_{hash})$.

We note that, in both random and lossy tags, we have $e(R_i, \hat{u}^i \cdot \hat{h}) = e(S_i, \hat{g})$ for all $i \in [n]$, so that elements of $\mathcal{T}$ are publicly recognizable.

Evaluation. $\text{LAF.Eval}(ek, t, \mathbf{x})$ takes in the input $\mathbf{x} \in \mathbb{Z}_p^n$ and the tag $t = (t_c, t_a)$.

It parses t_c as $(\{R_i, S_i, D_i, E_i, F_i\}_{i=1}^n, r_{hash})$ and does the following.

1. Return $\perp$ if there exists $i \in [n]$ such that $e(R_i, \hat{h}^i \cdot \hat{u}) \neq e(S_i, \hat{g})$.
2. Compute the matrix $(M_{i,j})_{i,j \in [n]} \in \mathbb{G}_T^{n \times n}$ as

$$M_{i,i} = \frac{e(D_i, \hat{g}) \cdot e(E_i, \hat{g}^{\theta_i})}{e(F_i, \prod_{k=1}^L \hat{Z}_{i,k,\tau[k]})} \quad \forall i \in [n], \quad (20)$$

where $\tau = \text{CMhash}(hk_{\text{CMH}}, (t_a, \{R_i, S_i, D_i, E_i, F_i\}_{i=1}^n), r_{hash}) \in \{0, 1\}^L$, and

$$M_{i,j} = \left(\frac{e(R_i, \hat{H}_j)}{e(S_i, \hat{V}_j)} \right)^{1/(j-i)} \quad \forall (i, j) \in [n] \times [n] \setminus \{(i, i)\}_{i=1}^n, \quad (21)$$

Since $R_i = g^{r_i}$ and $S_i = (h^i \cdot u)^{r_i}$ for some $r_i \in \mathbb{Z}_q$, we have

$$\begin{aligned} M_{i,i} &= e(g, \hat{h})^{r_i \cdot (\alpha_i + \theta_i \cdot \beta_i) + \omega_i}, & \forall i \in [n] \\ M_{i,j} &= e(g, \hat{h})^{r_i \cdot (\alpha_j + \theta_j \cdot \beta_j)}, & \forall i \neq j, \end{aligned} \quad (22)$$

for some vector $(\omega_1, \dots, \omega_n)^\top \in \mathbb{Z}_p^n$ that only contains non-zero entries if $t = (t_c, t_a)$ is injective.

3. Compute the vector $(V_{T,j})_{j \in [n]}$ as $V_{T,j} = e(h, \hat{V}_j) = e(g, \hat{h})^{\alpha_j + \theta_j \cdot \beta_j}$ for each $j \in [n]$.
4. Use the input $\mathbf{x} = (x_1, \dots, x_n)^\top \in \mathbb{Z}_p^n$ to compute

$$\begin{aligned} Y_0 &= \prod_{j=1}^n V_{T,j}^{x_j} \\ Y_i &= \prod_{j=1}^n M_{i,j}^{x_j} \quad \forall i \in [n] \end{aligned} \quad (23)$$

and output $\mathbf{Y} = (Y_0, Y_1, \dots, Y_n)^\top \in \mathbb{G}_T^{n+1}$.

The lossiness/injectivity properties can be analyzed exactly in the same way as in the construction of Section 3. Indeed, by defining $v_j = \alpha_j + \theta_j \cdot \beta_j$ for each $j \in [n]$, we find that $\{\hat{V}\}_{j=1}^n$ and $(M_{ij})_{i,j \in [n]}$ are distributed as in Section 3.

4.3 Security

Theorem 3. *The above LAF provides indistinguishability assuming that the $wD3DH1$ assumption holds in $(\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T)$ and that the DDH assumptions holds in $\mathbb{G}$. The advantage of any PPT distinguisher $\mathcal{A}$ making Q queries within time $t_{\mathcal{A}}$ is bounded by*

$$\mathbf{Adv}^{\text{indist}}(\lambda) \leq n \cdot (\mathbf{Adv}_{\mathcal{B}_1}^{wD3DH1}(\lambda) + \mathbf{Adv}_{\mathcal{B}_2}^{DDH1}(\lambda))$$

for PPT algorithm $\mathcal{B}_1, \mathcal{B}_2$ running in time $t_{\mathcal{A}} + Q \cdot \text{poly}(\lambda)$.

Proof. We define a sequence of hybrid games. In Game_0 , the adversary has access to the real oracle $\text{LAF.LTag}(tk, \cdot)$ oracle that always outputs lossy tags. In Game_n , the adversary is given access to an oracle $\mathcal{O}_{\mathcal{T}}(\cdot)$ that always outputs random tags in the tag space $\mathcal{T}$.

Game $_{\xi}'$ ($1 \leq \xi \leq n$): The adversary interacts with an oracle $\text{LAF.LTag}^{(\ell, k)}(tk, \cdot)$ that outputs tags $t = (t_c, t_a)$ with the following hybrid distribution. In the core component $t_c = (\{R_i, S_i, D_i, E_i, F_i\}_{i=1}^n, r_{\text{hash}})$, the first $\xi - 1$ tuples $\{(R_i, S_i, D_i, E_i, F_i)\}_{i=1}^{\xi}$ of t_c are random group elements satisfying the equality $e(R_i, \hat{h}^i \cdot \hat{u}) = e(S_i, \hat{g})$. The last $n - \xi$ tuples $\{(R_i, S_i, D_i, E_i, F_i)\}_{i=\xi+1}^n$ are generated exactly as in lossy tags. The ξ -th tuple $(R_{\xi}, S_{\xi}, D_{\xi}, E_{\xi}, F_{\xi})$ has a special distribution where $e(R_{\xi}, \hat{h}^{\xi} \cdot \hat{u}) = e(S_{\xi}, \hat{g})$, D_{ξ} is completely random in $\mathbb{G}$ and

$$\begin{aligned} E_{\xi} &= h^{\beta_{\xi} \cdot \log_g(R_{\xi})} \cdot g^{\rho_{\xi} \cdot \sum_{k=1}^L y_{\xi, k, \tau[k]}}, \\ F_{\xi} &= g^{\rho_{\xi}} \end{aligned}$$

Game $_{\xi}$ ($1 \leq \xi \leq n$): The adversary interacts with an oracle $\text{LAF.LTag}^{(\ell, k)}(tk, \cdot)$ that outputs $t = (t_c, t_a)$ such that the first ξ tuples $\{(R_i, S_i, D_i, E_i, F_i)\}_{i=1}^{\xi}$ of t_c are random subject to the constraint $e(R_i, \hat{h}^i \cdot \hat{u}) = e(S_i, \hat{g})$ while $\{(R_i, S_i, D_i, E_i, F_i)\}_{i=\xi+1}^n$ are generated as in lossy tags.

For each index $\xi \in [n]$, Lemma 7 shows that Game_{ξ}' is computationally indistinguishable from $\text{Game}_{\xi-1}$ if the R- $wD3DH1$ assumption holds. In a second step, Lemma 8 shows that Game_{ξ}' is indistinguishable from Game_{ξ} under the DDH assumption in $\mathbb{G}$. By applying Lemma 1, we obtain that the scheme provides indistinguishability under tight reductions from the hardness of $wD3DH1$ and that of the DDH problem in $\mathbb{G}$. $\square$

Lemma 7. *Game $_{\xi}'$ is computationally indistinguishable from Game $_{\xi-1}$ under the R- $wD3DH1$ assumption. The advantage of any PPT distinguisher between the two games can be bounded by $\mathbf{Adv}^{\xi' - (\xi-1)}(\lambda) \leq \mathbf{Adv}^{\text{R-}wD3DH1}(\lambda)$.*

Proof. Let us assume that there exists $\xi \in [n]$ such that the adversary $\mathcal{A}$ can distinguish Game_{ξ}' from $\text{Game}_{\xi-1}$ with non-negligible advantage. We build a R- $wD3DH1$ distinguisher $\mathcal{B}$ that takes as input $\{(g, \hat{g}, g^{a_i}, g^b, g^c, \hat{g}^b, \hat{g}^c, T_i)\}_{i=1}^Q$ with

the goal of deciding if $T_i = g^{a_i bc}$ for each $i \in [Q]$ or if $\{T_i\}_{i=1}^Q$ are all independent and uniformly distributed over $\mathbb{G}$.

To this end, $\mathcal{B}$ defines $h = g^b$, $\hat{h} = \hat{g}^b$. It also picks $\theta'_\xi, \beta'_\xi \xleftarrow{R} \mathbb{Z}_p$ uniformly and sets

$$\hat{g}^{\theta_\xi} = (\hat{g}^b)^{\theta'_\xi}, \quad \hat{V}_\xi = (\hat{g})^c \cdot \hat{g}^{\theta'_\xi \cdot \beta'_\xi},$$

which implicitly defines

$$\alpha_\xi = c, \quad \beta_\xi = \beta'_\xi / b, \quad \theta_\xi = b \cdot \theta'_\xi.$$

It chooses $\nu \xleftarrow{R} \mathbb{Z}_p$ and defines $\hat{u} = \hat{h}^{-\xi} \cdot \hat{g}^\nu$ as well as $u = h^{-\xi} \cdot g^\nu$. This allows defining

$$\hat{H}_\xi = (\hat{h}^\xi \cdot \hat{u})^{c + \theta'_\xi \cdot \beta'_\xi} = (\hat{V}_\xi)^\nu,$$

For all indexes $j \in [n] \setminus \{\xi\}$, it chooses $\alpha_j, \beta_j, \theta_j \xleftarrow{R} \mathbb{Z}_p$ and faithfully computes $\hat{V}_j = \hat{g}^{\alpha_j + \theta_j \cdot \beta_j}$ and

$$\hat{H}_j = (\hat{h}^j \cdot \hat{u})^{\alpha_j + \theta_j \cdot \beta_j}.$$

Then, it constructs the MAC secret keys $\{\mathbf{x}_{j,\mu}, \mathbf{y}_{j,\mu}\}_{j=1}^n$ for randomly chosen vectors $\mathbf{x}_{j,\mu} = (x_{j,1,\mu}, \dots, x_{j,L,\mu}) \xleftarrow{R} \mathbb{Z}_p^L$, $\mathbf{y}_{j,\mu} = (y_{j,1,\mu}, \dots, y_{j,L,\mu}) \xleftarrow{R} \mathbb{Z}_p^L$. For each $j \in [n]$, it defines

$$\begin{aligned} \hat{\mathbf{Y}}_{j,\mu} &= (\hat{Y}_{j,1,\mu}, \dots, \hat{Y}_{j,L,\mu}) = \hat{g}^{\mathbf{y}_{j,\mu}}, & \mathbf{Y}_{j,\mu} &= (Y_{j,1,\mu}, \dots, Y_{j,L,\mu}) = g^{\mathbf{y}_{j,\mu}} \\ \hat{\mathbf{X}}_{j,\mu} &= (\hat{X}_{j,1,\mu}, \dots, \hat{X}_{j,L,\mu}) = \hat{g}^{\mathbf{x}_{j,\mu}}, & \mathbf{X}_{j,\mu} &= (X_{j,1,\mu}, \dots, X_{j,L,\mu}) = g^{\mathbf{x}_{j,\mu}}. \end{aligned}$$

Then, it computes

$$\begin{aligned} \hat{\mathbf{Z}}_{j,\mu} &= \hat{\mathbf{X}}_{j,\mu} \cdot \hat{\mathbf{Y}}_{j,\mu}^{\theta_j} & \forall j \in [n] \setminus \{\xi\} \\ \hat{\mathbf{Z}}_{\xi,\mu} &= \hat{\mathbf{X}}_{\xi,\mu} \cdot (\hat{g}^b)^{\mathbf{y}_{\xi,\mu} \cdot \theta'_\xi} \end{aligned}$$

At the t -th invocation of the $\text{LAF.LTag}(tk, \cdot)$ oracle, $\mathcal{B}$ sets

$$R_\xi = g^{a_t}, \quad S_\xi = (g^{a_t})^\nu = (h^\xi \cdot u)^{a_t},$$

where g^{a_t} is fetched from the t -th input tuple $(g, \hat{g}, g^{a_t}, g^b, g^c, \hat{g}^b, \hat{g}^c, T_t)$. For all indexes $i \neq \xi$, it chooses $r_1, \dots, r_{\xi-1}, r_{\xi+1}, \dots, r_n \xleftarrow{R} \mathbb{Z}_p$ and sets

$$R_i = g^{r_i}, \quad S_i = (h^i \cdot u)^{r_i} \quad \forall i \in [n] \setminus \{\xi\}.$$

It generates the triples $\{D_i, E_i, F_i\}_{i=1}^n$ by choosing $(D_i, E_i, F_i) \xleftarrow{R} \mathbb{G}^3$ at random for each $i \in [\xi - 1]$. The ξ -th triple (D_ξ, E_ξ, F_ξ) is defined as

$$\begin{aligned} D_\xi &= T_t \cdot \left(\prod_{k=1}^L \hat{Y}_{\xi,k,\tau[k]} \right)^{\rho_\xi}, \\ E_\xi &= (g^{a_t})^{\beta'_\xi} \cdot \left(\prod_{k=1}^L \hat{Y}_{\xi,k,\tau[k]} \right)^{\rho_\xi}, \\ F_\xi &= g^{\rho_\xi}. \end{aligned}$$

for a randomly chosen $\rho_\xi \xleftarrow{R} \mathbb{Z}_p$ and $\tau \xleftarrow{R} \{0, 1\}^L$. As for $\{D_i, E_i, F_i\}_{i=\xi+1}^n$, they are obtained by choosing $\rho_i, r_i \xleftarrow{R} \mathbb{Z}_p$ before setting

$$D_i = (g^b)^{\alpha_i \cdot r_i} \left(\prod_{k=1}^L X_{\xi, k, \tau[k]} \right)^{\rho_i}, \quad E_i = (g^b)^{\beta_i \cdot r_i} \left(\prod_{k=1}^L Y_{\xi, k, \tau[k]} \right)^{\rho_i}, \quad F_i = g^{\rho_i}.$$

Then, it uses the trapdoor td_{CMH} of the chameleon hash function to find coins $r_{\text{hash}} \in \mathcal{R}_{\text{hash}}$ such that $\tau = \text{CMhash}(hk_{\text{CMH}}, (t_a, \{R_i, S_i, D_i, E_i, F_i\}_{i=1}^n), r_{\text{hash}})$.

It is easy to see that, if $T_t = g^{a_t bc}$, the triple (D_ξ, E_ξ, F_ξ) can be written

$$\begin{aligned} D_\xi &= h^{\alpha_\xi \cdot r_\xi} \cdot \left(\prod_{k=1}^L \hat{X}_{\xi, k, \tau[k]} \right)^{\rho_\xi}, \\ E_\xi &= h^{\beta_\xi \cdot r_\xi} \cdot \left(\prod_{k=1}^L \hat{Y}_{\xi, k, \tau[k]} \right)^{\rho_\xi} \\ F_\xi &= g^{\rho_\xi}, \end{aligned}$$

meaning that $\mathcal{A}$'s view is the same as in $\text{Game}_{\xi-1}$. In contrast, if $T_t \in_R \mathbb{G}$, it can be written $T_t = g^{a_t bc + z_t}$ for some uniformly random $z_t \in_R \mathbb{Z}_p$. In this case, (D_ξ, E_ξ, F_ξ) can be written

$$\begin{aligned} D_\xi &= h^{z_t + \alpha_\xi \cdot r_\xi} \cdot \left(\prod_{k=1}^L \hat{X}_{\xi, k, \tau[k]} \right)^{\rho_\xi}, \\ E_\xi &= h^{\beta_\xi \cdot r_\xi} \cdot \left(\prod_{k=1}^L \hat{Y}_{\xi, k, \tau[k]} \right)^{\rho_\xi}, \\ F_\xi &= g^{\rho_\xi}, \end{aligned}$$

for some random $z_t \in_R \mathbb{Z}_p$ that does not appear anywhere else. In this case, $\mathcal{A}$'s view corresponds to Game'_ξ . $\square$

Lemma 8. *Game_ξ is computationally indistinguishable from Game'_ξ under the DDH assumption in $\mathbb{G}$. The advantage of any PPT distinguisher between the two games can be bounded by $\text{Adv}^{\xi-\xi'}(\lambda) \leq \text{Adv}^{\text{DDH}_1}(\lambda)$.*

Proof. We assume that there exists $\xi \in [n]$ such that $\mathcal{A}$ can tell apart Game'_ξ from Game_ξ with noticeable advantage. We build a distinguisher $\mathcal{B}$ that takes as input Q tuples $\{(g, g^{a_i}, g^{a_i \cdot b}, g^b, T_i)\}_{i=1}^Q$ in $\mathbb{G}^5$ with the goal of deciding if $T_i = g^{a_i b}$ for each $i \in [Q]$ or if $\{T_i\}_{i=1}^Q$ are independent and uniformly distributed over $\mathbb{G}$. This assumption is known (see, e.g., [39, Lemma 4.4]) to have a tight reduction from the DDH assumption.

To this end, $\mathcal{B}$ defines $h = g^\eta$, $\hat{h} = \hat{g}^\eta$ for a random $\eta \xleftarrow{R} \mathbb{Z}_p$. It also computes $\hat{g}^{\theta_\xi}$ for a randomly chosen $\theta_\xi \xleftarrow{R} \mathbb{Z}_p$. Then, it picks $v_\xi \xleftarrow{R} \mathbb{Z}_p$ uniformly and sets

$$\hat{V}_\xi = \hat{g}^{v_\xi}.$$

Implicitly, $\mathcal{B}$ will define

$$\beta_\xi = b, \quad \alpha_\xi = v_\xi - b \cdot \theta_\xi$$

although it does not know (α_ξ, β_ξ) . It chooses $\hat{u} \in \hat{\mathbb{G}}$ and $u \in \mathbb{G}$ by setting $u = g^\nu$ and $\hat{u} = \hat{g}^\nu$ for a random $\nu \xleftarrow{R} \mathbb{Z}_p$. Then, $\mathcal{B}$ defines

$$\hat{H}_\xi = (\hat{h}^\xi \cdot \hat{u})^{v_\xi}.$$

For all indexes $j \in [n] \setminus \{\xi\}$, it chooses $\alpha_j, \beta_j, \theta_j \xleftarrow{R} \mathbb{Z}_p$ and faithfully computes $\hat{V}_j = \hat{g}^{\alpha_j + \theta_j \cdot \beta_j}$ and

$$\hat{H}_j = (\hat{h}^j \cdot \hat{u})^{\alpha_j + \theta_j \cdot \beta_j}.$$

Then, it constructs the MAC secret keys $\{\mathbf{x}_{j,\mu}, \mathbf{y}_{j,\mu}\}_{j=1}^n$ by for randomly chosen vectors $\mathbf{x}_{j,\mu} = (x_{j,1,\mu}, \dots, x_{j,L,\mu}) \xleftarrow{R} \mathbb{Z}_p^L$, $\mathbf{y}_{j,\mu} = (y_{j,1,\mu}, \dots, y_{j,L,\mu}) \xleftarrow{R} \mathbb{Z}_p^L$. For each $j \in [n]$, it defines

$$\begin{aligned} \hat{\mathbf{Y}}_{j,\mu} &= (\hat{Y}_{j,1,\mu}, \dots, \hat{Y}_{j,L,\mu}) = \hat{g}^{\mathbf{y}_{j,\mu}}, & \mathbf{Y}_{j,\mu} &= (Y_{j,1,\mu}, \dots, Y_{j,L,\mu}) = g^{\mathbf{y}_{j,\mu}} \\ \hat{\mathbf{X}}_{j,\mu} &= (\hat{X}_{j,1,\mu}, \dots, \hat{X}_{j,L,\mu}) = \hat{g}^{\mathbf{x}_{j,\mu}}, & \mathbf{X}_{j,\mu} &= (X_{j,1,\mu}, \dots, X_{j,L,\mu}) = g^{\mathbf{x}_{j,\mu}}. \end{aligned}$$

Then, it computes

$$\hat{\mathbf{Z}}_{j,\mu} = \hat{\mathbf{X}}_{j,\mu} \cdot \hat{\mathbf{Y}}_{j,\mu}^{\theta_j} \quad \forall j \in [n].$$

For each $t \in [Q]$, the t -th invocation of the $\text{LAF.LTag}(tk, \cdot)$ oracle is handled by setting

$$R_\xi = g^{a_t}, \quad S_\xi = (g^{a_t})^{\eta \cdot \xi + \nu} = (h^\xi \cdot u)^{a_t},$$

where g^{a_t} is fetched from the t -th input tuple $(g, g^{a_t}, g^{a_t \cdot b}, g^b, T_t)$. For all indexes $i \neq \xi$, it chooses $r_1, \dots, r_{\xi-1}, r_{\xi+1}, \dots, r_n \xleftarrow{R} \mathbb{Z}_p$ and sets

$$R_i = g^{r_i}, \quad S_i = (h^i \cdot u)^{r_i} \quad \forall i \in [n] \setminus \{\xi\}.$$

It generates the triples $\{D_i, E_i, F_i\}_{i=1}^n$ by choosing $(D_i, E_i, F_i) \xleftarrow{R} \mathbb{G}^3$ at random for each $i \in [\xi - 1]$. The ξ -th triple (D_ξ, E_ξ, F_ξ) is defined by sampling $D_\xi \xleftarrow{R} \mathbb{G}$ uniformly and setting

$$\begin{aligned} E_\xi &= T_t^\eta \cdot \left(\prod_{k=1}^L \hat{Y}_{\xi,k,\tau[k]} \right)^{\rho_\xi}, \\ F_\xi &= g^{\rho_\xi}. \end{aligned}$$

for randomly chosen $\rho_\xi \xleftarrow{R} \mathbb{Z}_p$ and $\tau \xleftarrow{R} \{0, 1\}^L$. As for $\{D_i, E_i, F_i\}_{i=\xi+1}^n$, they are obtained by choosing choosing $\rho_i, r_i \xleftarrow{R} \mathbb{Z}_p$ before setting

$$D_i = h^{\alpha_i \cdot r_i} \left(\prod_{k=1}^L X_{\xi,k,\tau[k]} \right)^{\rho_i}, \quad E_i = h^{\beta_i \cdot r_i} \left(\prod_{k=1}^L Y_{\xi,k,\tau[k]} \right)^{\rho_i}, \quad F_i = g^{\rho_i}.$$

Then, it uses the trapdoor td_{CMH} of the chameleon hash function to obtain coins $r_{\text{hash}} \in \mathcal{R}_{\text{hash}}$ such that $\tau = \text{CMhash}(hk_{\text{CMH}}, (t_a, \{R_i, S_i, D_i, E_i, F_i\}_{i=1}^n), r_{\text{hash}})$.

We observe that, if $T_t = g^{a_t \cdot b}$ for each $t \in [Q]$, the triples (D_ξ, E_ξ, F_ξ) are distributed as

$$\begin{aligned} D_\xi &\in_R \mathbb{G}, \\ E_\xi &= h^{\beta_\xi \cdot \log_g(R_\xi)} \cdot \left(\prod_{k=1}^L \hat{Y}_{\xi, k, \tau[k]} \right)^{\rho_\xi} \\ F_\xi &= g^{\rho_\xi}, \end{aligned}$$

so that $\mathcal{A}$'s view is the same as in Game'_ξ . In contrast, if $T_t \in_R \mathbb{G}$, it can be written $T_t = g^{a_t b + z_t}$ for some uniformly random $z_t \in_R \mathbb{Z}_p$ that does not appear anywhere else. In this case, (D_ξ, E_ξ, F_ξ) is just a triple of uniformly random group elements, meaning that $\mathcal{A}$'s view is the same as in Game_ξ . $\square$

Theorem 4. *The above LAF provides evasiveness under the SXDH and wD3DH1 assumptions, assuming that CMH is a collision-resistant chameleon hash function. Namely, for any PPT evasiveness adversary, there exist efficient algorithms $\mathcal{B}_0, \mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ with comparable running time and such that*

$$\begin{aligned} \text{Adv}_Q^{\mathcal{A}, \text{eva}} &\leq \text{Adv}_{\mathcal{B}_0}^{\text{CMH-CR}}(\lambda) + n \cdot \text{Adv}_{\mathcal{B}_1}^{\text{wD3DH1}}(\lambda) \\ &\quad + n \cdot \text{Adv}_{\mathcal{B}_2}^{\text{DDH}_2}(\lambda) + 2n \cdot (L+1) \cdot \text{Adv}_{\mathcal{B}_3}^{\text{DDH}_1}(\lambda), \end{aligned}$$

(The proof is given in Appendix C.)

Acknowledgements

We thank the reviewers for their careful reading. This work was funded in part by the French ANR ALAMBIC project (ANR-16-CE39-0006).

References

1. M. Bellare, Z. Brakerski, M. Naor, T. Ristenpart, G. Segev, H. Shacham, and S. Yilek. Hedged public-key encryption: How to protect against bad randomness. In *Asiacrypt*, 2009. 1
2. M. Bellare, D. Hofheinz, and S. Yilek. Possibility and impossibility results for encryption and commitment secure under selective opening. In *Eurocrypt*, 2009. 1, 1
3. M. Bellare and P. Rogaway. Random oracles are practical: a paradigm for designing efficient protocols. In *ACM-CCS*, 1993. 1
4. M. Bellare and P. Rogaway. The exact security of digital signatures: How to sign with RSA and Rabin. In *Eurocrypt*, 1996. 1
5. M. Bellare and S. Yilek. Encryption schemes secure under selective opening attack. Cryptology ePrint Archive: Report 2009/101, 2009. 1
6. J. Black, P. Rogaway, and T. Shrimpton. Encryption-scheme security in the presence of key-dependent messages. In *SAC*, 2002. 1

7. O. Blazy, E. Kiltz, and J. Pan. (Hierarchical) Identity-Based Encryption from Affine Message Authentication. In *Crypto*, 2014. 1, 1, 4, 4.1, 4.1
8. F. Böhl, D. Hofheinz, T. Jager, J. Koch, J.-H. Seo, C. Striecks. Practical Signatures from Standard Assumptions. In *Eurocrypt*, 2013. 3.1
9. S. Boldyreva, S. Fehr, and A. O’Neill. On notions of security for deterministic encryption, and efficient constructions without random oracles. In *Crypto*, 2008. 1
10. D. Boneh, S. Halevi, M. Hamburg, R. Ostrovsky. Circular-Secure Encryption from Decision Diffie-Hellman. In *Crypto*, 2008. 1
11. D. Boneh, X. Boyen. Efficient Selective Identity-Based Encryption Without Random Oracles. In *Eurocrypt*, 2004. 1, 3.2
12. D. Boneh, X. Boyen, and H. Shacham. Short Group Signatures. In *Crypto*, 2004. 1
13. X. Boyen. Reusable cryptographic fuzzy extractors. In *ACM-CCS*, 2004. 1
14. X. Boyen and B. Waters. Shrinking the keys of discrete-log-type lossy trapdoor functions. In *ACNS*, 2010. 1, 1, 3
15. Z. Brakerski and G. Segev. Better security for deterministic public-key encryption: The auxiliary-input setting. In *Crypto*, 2011. 1
16. J. Chen and H. Wee. Fully, (almost) tightly secure IBE and dual system groups. In *Crypto*, 2013. 1, 1
17. Y. Dodis, R. Ostrovsky, L. Reyzin, A. Smith. Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data. In *Eurocrypt*, 2004. 1
18. C. Dwork, M. Naor, O. Reingold, and L. Stockmeyer. Magic functions. *J. of the ACM*, 50(6), 2003. 1
19. D. Freeman, O. Goldreich, E. Kiltz, and G. Rosen, A. and Segev. More constructions of lossy and correlation-secure trapdoor functions. *J. of Cryptology*, 26(1), 2013. 1
20. E. Fujisaki. All-but-many encryption - a new framework for fully-equipped UC commitments. In *Asiacrypt*, 2014. 1
21. R. Gay, D. Hofheinz, E. Kiltz, and H. Wee. Tightly CCA-secure encryption without pairings. In *Eurocrypt*, 2016. 1
22. R. Gay, D. Hofheinz, L. Kohl. Kurosawa-Desmedt Meets Tight Security. In *Crypto*, 2017. 1
23. B. Hemenway and R. Ostrovsky. Extended-DDH and lossy trapdoor functions. In *PKC*, 2012. 1
24. D. Hofheinz and E. Kiltz. Programmable hash functions and their applications. In *Crypto*, 2008. A
25. D. Hofheinz. All-but-many lossy trapdoor functions. In *Eurocrypt*, 2012. 1, 1, 3.3
26. D. Hofheinz. Circular chosen-ciphertext security with compact ciphertexts. In *Eurocrypt*, 2013. Cryptology ePrint Archive: Report 2012/150. 1, 1, 2.1, 2.1, 3, 4
27. D. Hofheinz. Algebraic partitioning: Fully compact and (almost) tightly secure cryptography. In *TCC-A*, 2016. 1
28. D. Hofheinz. Adaptive partitioning. In *Eurocrypt*, 2017. 1
29. D. Hofheinz and T. Jager. Tightly secure signatures and public-key encryption. In *Crypto*, 2012. 1
30. D. Hofheinz and N.-K. Nguyen. On Tightly Secure Primitives in the Multi-Instance Setting. Cryptology ePrint Archive: Report 2018/958. 1
31. C. Jutla, A. Roy. Shorter Quasi-Adaptive NIZK Proofs for Linear Subspaces. In *Asiacrypt*, 2013. 3.1
32. H. Krawczyk and T. Rabin. Chameleon Signatures. In *NDSS*, 2000. 2.2
33. S. Kunz-Jacques, D. Pointcheval. About the Security of MTI/C0 and MQV. In *SCN*, 2006. 1, 2.3, 4

34. A. Lewko, A. Sahai, B. Waters. Revocation Systems with Very Small Private Keys. *IEEE Symposium on Security and Privacy*, 2010. 1, 3
35. B. Libert, M. Joye, T. Peters, and M. Yung. Concise Multi-Challenge CCA-secure Encryption and Signatures with Almost Tight Security. In *Asiacrypt*, 2014. 1
36. B. Libert, T. Peters, M. Joye, and M. Yung. Compactly hiding linear spans - tightly secure constant-size simulation-sound QA-NIZK proofs and applications. In *Asiacrypt*, 2015. 1
37. B. Libert, A. Sakzad, D. Stehlé, and R. Steinfeld. All-But-Many Lossy Trapdoor Functions and Selective-Opening Chosen-Ciphertext Security. In *Crypto*, 2017. 1, 3.3
38. B. Libert, D. Vergnaud. Multi-Use Unidirectional Proxy Re-Signatures. In *ACM-CCS*, 2008. 2.3
39. M. Naor and O. Reingold. Number-theoretic constructions of efficient pseudo-random functions. In *FOCS*, 1997. 2, 4, 4.3
40. C. Peikert and B. Waters. Lossy trapdoor functions and their applications. In *STOC*, 2008. 1, 1
41. C. Rackoff and D. Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In *Crypto*, 1991. 1
42. A. Raghunathan, G. Segev, and S. Vadhan. Deterministic public-key encryption for adaptively chosen plaintext distributions. In *Eurocrypt'13*, 2013. 1
43. B. Waters. Efficient Identity-Based Encryption Without Random Oracles. In *Eurocrypt*, 2005. 1, 1, 3.1
44. H. Wee. Dual projective hashing and its applications - lossy trapdoor functions and more. In *Eurocrypt*, 2012. 1
45. Y. Wen, S. Liu. Robustly Reusable Fuzzy Extractor from Standard Assumptions. In *Asiacrypt*, 2018. 1, 1, 2.1
46. M. Zhandry. The magic of ELFs. In *Crypto*, 2016. 1

A Proof of Theorem 2

Proof. Let us assume that a PPT adversary $\mathcal{A}$ can break the evasiveness property with noticeable advantage. We show that it contradicts either: (i) The indistinguishability of the scheme; (ii) The collision-resistance of the chameleon hash function; (iii) The 2-3-CDH assumption. We will prove this claim via a sequence of hybrid games.

In Game_0 , the adversary $\mathcal{A}$ proceeds as in the real evasiveness security experiment. In the final game, we show that, if the adversary can output a lossy tag, we can easily construct an algorithm breaking the 2-3-CDH assumption with non-negligible advantage.

For each i , we denote by bad_i the event that $\mathcal{A}$ manages to output a non-trivial lossy tag in Game_i .

Game₀: In this game, the adversary $\mathcal{A}$ has access to the lossy tag generation oracle $\text{LAF.LTag}(tk, \cdot)$ that always outputs lossy tags. By definition,

$$\Pr[\text{bad}_0] = \Pr[\mathcal{A}(1^\lambda, ek)^{\text{LAF.LTag}(tk, \cdot)}]. \quad (24)$$

Game₁: In this game, we define bad_{hash} to be the event that the adversary $\mathcal{A}$ outputs a tag $t = ((\{R_i, S_i, D_i, E_i\}_{i=1}^n, r_{hash}))$ for which the corresponding chameleon hash collides with that of some tag produced by the oracle $\text{LAF.LTag}(tk, \cdot)$. The only difference between Game_1 and Game_0 is that Game_1 aborts when bad_{hash} occurs. It is straightforward that

$$|\Pr[\text{bad}_1] - \Pr[\text{bad}_0]| = \Pr[\text{bad}_{hash} \text{ in Game}_1]. \quad (25)$$

We want to use the collision resistance property of the underlying chameleon hash function to bound the probability $\Pr[\text{bad}_{hash} \text{ in Game}_1]$. However, the lossy key generation oracle uses the trapdoor td_{CMH} to create lossy tags. To avoid a circularity, we consider Game'_1 , where the lossy key generation oracle always outputs injective tags instead of lossy ones. Using the indistinguishability between lossy and injective tags (established by Theorem 1), we have

$$|\Pr[\text{bad}_{hash} \text{ in Game}_1] - \Pr[\text{bad}_{hash} \text{ in Game}'_1]| = nQ \cdot \text{Adv}^{\text{wD3DH1}}(\lambda). \quad (26)$$

Since Game_0 and Game_1 only differ when bad_{hash} occurs in Game_1 , we can bound the probability (25) as

$$\begin{aligned} |\Pr[\text{bad}_1] - \Pr[\text{bad}_0]| &= |\Pr[\text{bad}_{hash} \text{ in Game}_1]| \\ &\leq |\Pr[\text{bad}_{hash} \text{ in Game}'_1]| + nQ \cdot \text{Adv}^{\text{wD3DH1}}(\lambda) \end{aligned}$$

In Game'_1 , we clearly have $\Pr[\text{bad}_{hash} \text{ in Game}'_1] \leq \text{Adv}^{\text{CR}}_{\text{CMH}}(\lambda)$: since the trapdoor of CMH is not used, we can readily build a reduction that breaks the collision-resistance of CMH out of an adversary for which bad_{hash} occurs with noticeable probability. This immediately implies

$$|\Pr[\text{bad}_1] - \Pr[\text{bad}_0]| \leq \text{Adv}^{\text{CR}}_{\text{CMH}}(\lambda) + nQ \cdot \text{Adv}^{\text{wD3DH1}}(\lambda)$$

We now proceed to bound $\Pr[\text{bad}_1]$ by showing that, using the adversary $\mathcal{A}$ in Game_1 , we can build an algorithm $\mathcal{B}$ breaking the 2-3-CDH assumption.

Algorithm $\mathcal{B}$ takes as input $(g^a, g^b, \hat{g}^a, \hat{g}^b)$ with the goal of computing $g^r, g^{r \cdot ab}$. To this end, $\mathcal{B}$ defines $h = g^a$. It randomly chooses a $J \xleftarrow{R} [n]$ and sets $V_J = g^b$, which implicitly sets $v_J = b$. In addition, $\mathcal{B}$ defines $(W_0, W_1, \dots, W_L) \in \mathbb{G}^{L+1}$ and $(\hat{W}_0, \hat{W}_1, \dots, \hat{W}_L) \in \hat{\mathbb{G}}^{L+1}$ as

$$W_i = (g^a)^{\alpha_i} \cdot g^{\beta_i} \quad \hat{W}_i = (\hat{g}^a)^{\alpha_i} \cdot \hat{g}^{\beta_i}$$

where $\alpha_0 = -1$ and $\alpha_1, \dots, \alpha_L \xleftarrow{R} \{-1, 0, 1\}$ and $\beta_0, \dots, \beta_L \xleftarrow{R} \mathbb{Z}_p$.

In order to simulate the LAF.LTag oracle on input of t_a , $\mathcal{B}$ proceeds as follows:

1. For each $i \in [n]$, $\mathcal{B}$ uniformly samples $r_i \xleftarrow{R} \mathbb{Z}_p^*$ and computes $\{R_i, S_i\}_{i=1}^n$ as in (3).

2. $\mathcal{B}$ samples a random τ in the range of CMhash . For each $i \in [n] \setminus \{J\}$, it chooses $\rho_i \xleftarrow{R} \mathbb{Z}_p$ and computes

$$D_i = h^{r_i \cdot v_i} \cdot H_{\mathbb{G}}(\tau)^{\rho_i}, \quad E_i = g^{\rho_i}.$$

3. For $i = J$, $\mathcal{B}$ aborts if $\alpha_0 + \sum_{k=1}^L \alpha_k \cdot \tau[k] = 0$. Otherwise, $\mathcal{B}$ chooses $\rho_J \xleftarrow{R} \mathbb{Z}_p$ and computes (D_J, E_J) as in (11):

$$\begin{aligned} D_J &= H_{\mathbb{G}}(\tau)^{\rho_J} \cdot (V_J)^{-r_J \cdot \frac{\beta_0 + \sum_{k=1}^L \beta_k \cdot \tau[k]}{\alpha_0 + \sum_{k=1}^L \alpha_k \cdot \tau[k]}} \\ E_J &= g^{\rho_J} \cdot (V_J)^{-\frac{r_J}{\alpha_0 + \sum_{k=1}^L \alpha_k \cdot \tau[k]}} \end{aligned} \quad (27)$$

4. Next, $\mathcal{B}$ uses the trapdoor td_{CMH} of the chameleon hash function in order to find random coins $r_{\text{hash}} \in \mathcal{R}_{\text{CMH}}$ such that

$$\tau = \text{CMhash}(hk_{\text{hash}}, (t_a, \{R_i, S_i, D_i, E_i\}_{i=1}^n), r_{\text{hash}}).$$

5. Finally, $\mathcal{B}$ outputs (t_c, t_a) with $t_c = (\{R_i, S_i, D_i, E_i\}_{i=1}^n, r_{\text{hash}})$.

As in (11), if we define $\tilde{\rho}_J = \rho_J - \frac{b \cdot r_J}{\alpha_0 + \sum_{k=1}^L \alpha_k \cdot \tau[k]}$, we observe that (27) can be written as $D_J = h^{b \cdot r_J} \cdot H_{\mathbb{G}}(\tau)^{\tilde{\rho}_J}$ and $E_J = g^{\tilde{\rho}_J}$. Hence, if $\mathcal{B}$ does not abort in any query to LAF.LTag , the output distribution of $\mathcal{B}$ is identical to that of the real LAF.LTag oracle. We denote by abort_k the event that $\mathcal{B}$ aborts at the k -th query to the LAF.LTag oracle for $k \in [Q]$. Letting $t^* = (t_c^*, t_a^*)$ denote the lossy tag generated by $\mathcal{A}$, we parse t_c^* as $t_c^* = (\{R_i^*, S_i^*, D_i^*, E_i^*\}_{i=1}^n, r_{\text{hash}}^*)$ and compute $\tau^* = \text{CMhash}(hk_{\text{hash}}, (t_a^*, \{R_i^*, S_i^*, D_i^*, E_i^*\}_{i=1}^n), r_{\text{hash}}^*)$. In the event that $\alpha_0 + \sum_{k=1}^L \alpha_k \cdot \tau^*[k] \neq 0$, $\mathcal{B}$ aborts. We denote the latter event by abort_{ch} . If $\mathcal{B}$ did not abort (which implies $\alpha_0 + \sum_{k=1}^L \alpha_k \cdot \tau^*[k] = 0$), we have

$$\begin{aligned} D_J^* &= h^{b \cdot r_J^*} \cdot H_{\mathbb{G}}(\tau^*)^{\tilde{\rho}_J^*} \\ &= g^{ab \cdot r_J^*} \cdot g^{(a \cdot (\alpha_0 + \sum_{k=1}^L \alpha_k \cdot \tau^*[k]) + (\beta_0 + \sum_{k=1}^L \beta_k \cdot \tau^*[k])) \cdot \tilde{\rho}_J^*} \\ &= g^{ab \cdot r_J^*} \cdot E_J^{*\beta_0 + \sum_{k=1}^L \beta_k \cdot \tau^*[k]}, \end{aligned}$$

where $E_J^* = g^{\tilde{\rho}_J^*}$. Finally, $\mathcal{B}$ outputs $(R_J^*, \frac{D_J^*}{E_J^{*\beta_0 + \sum_{k=1}^L \beta_k \cdot \tau^*[k]}})$.

Clearly, if $\mathcal{B}$ did not abort, its output $(R_J, D_J / E_J^{\beta_0 + \sum_{k=1}^L \beta_k \cdot \tau[k]})$ is a valid 2-3-CDH challenge. We are left with evaluating the probability that $\mathcal{B}$ aborts.

If we define the function $\alpha : \{0, 1\}^L \rightarrow \mathbb{Z}$ that maps the string $m = m[1] \dots m[L] \in \{0, 1\}^L$ to $\alpha(m) = \alpha_0 + \sum_{k=1}^L \alpha_i \cdot m[i]$, the probability that $\mathcal{B}$ does not abort is given by

$$\begin{aligned} &\Pr[\neg \text{abort}_{ch} \wedge \neg \text{abort}_1 \wedge \dots \wedge \neg \text{abort}_Q] \\ &= \Pr[\alpha(\tau^*) = 0 \wedge \alpha(\tau_1), \dots, \alpha(\tau_Q) \neq 0], \end{aligned} \quad (28)$$

where Q is the number of queries to LAF.LTag , and τ_i denotes the output of the chameleon hash function produced at the i -th LAF.LTag query. By applying

known results on programmable hash functions [24], our choice of $\alpha_0, \alpha_1, \dots, \alpha_L$ ensures that

$$\Pr[\neg \text{abort}_{ch} \wedge \neg \text{abort}_1 \wedge \dots \wedge \neg \text{abort}_Q] \geq \delta, \quad (29)$$

where $\delta = \Omega(Q \cdot \sqrt{L})$. Putting the above altogether, we can conclude that

$$\mathbf{Adv}^{\text{eva}}(1^\lambda) \leq \mathbf{Adv}_{\text{CMH}}^{\text{CR}}(\lambda) + nQ \cdot \mathbf{Adv}^{\text{wD3DH1}}(\lambda) + \mathcal{O}(Q \cdot \sqrt{L}) \cdot \mathbf{Adv}^{2\text{-}3\text{-CDH}}(\lambda),$$

which yields the statement of the theorem. $\square$

B Deferred Proofs for the MAC of Section 4.1

B.1 Proof of Lemma 6

Proof. Assuming the existence of an adversary $\mathcal{A}$ that can distinguish between $\text{Game}'_{2,(i-1)}$ and $\text{Game}'_{2,i}$, we will build a DDH distinguisher $\mathcal{B}$. Our distinguisher $\mathcal{B}$ inputs a DDH instance $(g, g^a, g^b, T) \in \mathbb{G}^4$ and decides whether $T = g^{ab}$ or $T \in_R \mathbb{G}$. To do this, $\mathcal{B}$ flips a random coin $\gamma \xleftarrow{R} \{0, 1\}$ and uses a random function $R' : \{0, 1\}^{i-1} \rightarrow \mathbb{Z}_p$, which is lazily defined as the adversary makes queries. Using R' , $\mathcal{B}$ defines another random function $R : \{0, 1\}^i \rightarrow \mathbb{Z}_p$ as

$$R(m[1] \dots m[i]) = \begin{cases} R(m[1] \dots m[i-1]) & m[i] = \gamma \\ R(m[1] \dots m[i-1]) + R'(m[1] \dots m[i-1]) & m[i] = 1 - \gamma \end{cases}.$$

We now consider the output of MAC queries. Implicitly, $\mathcal{B}$ defines $x_{i,1-\gamma}$ and $y_{i,1-\gamma}$ as $x_{i,1-\gamma} = x'_{i,1-\gamma} + \theta(1-a) \cdot y'_{i,1-\gamma}$ and $y_{i,1-\gamma} = a \cdot y'_{i,1-\gamma}$, where $x'_{i,1-\gamma}, y'_{i,1-\gamma} \xleftarrow{R} \mathbb{Z}_p$. Note that the only value in public parameter that depends on $x_{i,1-\gamma}$ and $y_{i,1-\gamma}$ is

$$\hat{Z}_{i,1-\gamma} = \hat{g}^{x_{i,1-\gamma} + \theta \cdot y_{i,1-\gamma}} = \hat{g}^{x'_{i,1-\gamma} + \theta \cdot y'_{i,1-\gamma}},$$

so that $\hat{Z}_{i,1-\gamma}$ is computable from $(x'_{i,1-\gamma}, y'_{i,1-\gamma}) \in \mathbb{Z}_p^2$. The remaining secret key components are chosen as in the real key generation algorithm, by sampling $\eta, \alpha, \beta \xleftarrow{R} \mathbb{Z}_p$, $x_{i,\gamma}, y_{i,\gamma} \xleftarrow{R} \mathbb{Z}_p$ and $x_{k,b}, y_{k,b} \xleftarrow{R} \mathbb{Z}_p$ for each $k \in [L] \setminus \{i\}$, $b \in \{0, 1\}$.

Then, $\mathcal{B}$ simulates the responses to MAC queries in the following way.

1. From $(A = g^a, B = g^b, T)$, $\mathcal{B}$ uses the random self-reducibility of DDH assumption to generate a fresh pair $(B_{m|i-1}, T_{m|i-1})$ for each value of $m|i-1 = m[1] \dots m[i-1] \in \{0, 1\}^{i-1}$ in such a way that, if (A, B, T) is a DDH tuple, so is $(A, B_{m|i-1}, T_{m|i-1})$. Otherwise, $B_{m|i-1} \in_R \mathbb{G}$ and $T_{m|i-1} \in_R \mathbb{G}$ are i.i.d. For convenience, we may associate each string $m|i-1 \in \{0, 1\}^i$ with a tuple

$$(A, B_{m|i-1}, T_{m|i-1}) = (g^a, g^{b_{m|i-1}}, g^{a \cdot b_{m|i-1} + e_{m|i-1}})$$

where either $e_{m|i-1} = 0$ or $e_{m|i-1} \in_R \mathbb{Z}_p$. Note that the pairs $(B_{m|i-1}, T_{m|i-1})$ can be sampled lazily by having $\mathcal{B}$ initially generate Q pairs since at most Q distinct prefixes $m|i-1$ can occur in all MAC queries.

2. For each message M queried by $\mathcal{A}$, $\mathcal{B}$ randomly chooses $r, d \xleftarrow{R} \mathbb{Z}_p$ and computes σ in the following way.

$$\begin{aligned}\sigma_1 &= h^{(v-\theta R(m[1] \dots m[i-1])) \cdot r} \cdot (B_{m[i]}^r \cdot g^d)^{x'_{i,m[i]} + \theta y'_{i,m[i]}} \\ &\quad \cdot (T^r \cdot A^d)^{-\theta y'_{i,m[i]}} \cdot (B_{m[i-1]}^r \cdot g^d)^{\sum_{k=1 \wedge k \neq i}^L x_{k,m[k]}}, \\ \sigma_2 &= h^{R(m[1] \dots m[i-1]) \cdot r} \cdot (T^r \cdot A^d)^{y'_{i,m[i]}} \cdot (B_{m[i-1]}^r \cdot g^d)^{\sum_{k=1 \wedge k \neq i}^L y_{k,m[k]}}, \\ \sigma_3 &= B_{m[i-1]}^r \cdot g^d, \\ \sigma_4 &= g^r.\end{aligned}$$

We observe that, if we set $\rho = b_{m[i-1]} \cdot r + d$, the above equations can be written as

$$\begin{aligned}\sigma_1 &= h^{(v-\theta R(m[1] \dots m[i-1])) \cdot r} \cdot g^{\rho \cdot (x'_{i,m[i]} + \theta(1-a) \cdot y'_{i,m[i]})} \\ &\quad \cdot g^{\rho \cdot \sum_{k=1 \wedge k \neq i}^L x_{k,m[k]}} \cdot g^{-e_{m[i-1]} \cdot y_{i,m[i]} \cdot r \cdot \theta} \\ &= h^{(v-\theta \cdot R(m[1] \dots m[i-1])) \cdot r} \cdot g^{\rho \cdot x_{i,m[i]}} \cdot g^{\rho \cdot \sum_{k=1 \wedge k \neq i}^L x_{k,m[k]}} \cdot g^{-e_{m[i-1]} \cdot y'_{i,m[i]} \cdot r \cdot \theta} \\ \sigma_2 &= h^{R(m[1] \dots m[i-1]) \cdot r} \cdot g^{\rho \cdot a \cdot y'_{i,m[i]}} \cdot g^{\rho \cdot \sum_{k=1 \wedge k \neq i}^L y_{k,m[k]}} \cdot g^{e_{m[i-1]} \cdot y_{i,m[i]} \cdot r} \\ &= h^{R(m[1] \dots m[i-1]) \cdot r} \cdot g^{\rho \cdot y_{i,m[i]}} \cdot g^{\rho \cdot \sum_{k=1 \wedge k \neq i}^L y_{k,m[k]}} \cdot g^{e_{m[i-1]} \cdot y_{i,m[i]} \cdot r}\end{aligned}$$

If $(A, B_{m[i-1]}, T_{m[i-1]})$ is a Diffie-Hellman tuple (i.e., if $e_{m[i-1]} = 0$), the output distribution is the same as in $\text{Game}_{2,(i-1)}$. In contrast, if $e_{m[i-1]} \in_R \mathbb{Z}_p$, we have

$$\begin{aligned}\sigma_1 &= h^{(v-\theta \cdot R(m[1] \dots m[i])) \cdot r} \cdot g^{\rho \cdot x_{i,m[i]}} \cdot g^{\rho \cdot \sum_{k=1 \wedge k \neq i}^L x_{k,m[k]}} \\ \sigma_2 &= h^{R(m[1] \dots m[i]) \cdot r} \cdot g^{\rho \cdot y_{i,m[i]}} \cdot g^{\rho \cdot \sum_{k=1 \wedge k \neq i}^L y_{k,m[k]}}\end{aligned}$$

where the random function $R : \{0, 1\}^i \rightarrow \mathbb{Z}_p$ is defined using

$$R(m[1] \dots m[i-1]) = g^{\frac{e_{m[i-1]} \cdot y_{i,m[i]}}{\eta}}.$$

In this case, the output distribution of the MAC oracle is identical to that of $\text{Game}_{2,i}$.

If the adversary chooses to forge on a message $m^*[1] \dots m^*[L]$ such that $m^*[i] = 1 - \gamma$ (which occurs with probability $1/2$), then $\mathcal{B}$ aborts and outputs a random bit. If $m^*[i] = \gamma$, we have

$$R(m^*[1] \dots m^*[i]) = R(m^*[1] \dots m^*[i-1])$$

by the definition of R . Since $\mathcal{B}$ knows $y_{i,m^*[i]} = y_{i,\gamma}$, it can check if

$$\sigma_2^* = \sigma_4^* \cdot R(m^*[1] \dots m^*[i-1]) \cdot \sigma_3^* \cdot \sum_{k=1}^L y_{k,m^*[k]}$$

and return 1 if and only if this equality is satisfied. We thus conclude that $|\Pr[W_{2,i}] - \Pr[W_{2,(i-1)}]| \leq 2 \cdot \mathbf{Adv}^{\text{DDH}_1}(\lambda)$, as claimed. $\square$

C Proof of Theorem 4

Proof. Let us assume that a PPT adversary $\mathcal{A}$ can break the evasiveness property with noticeable advantage. We show that this contradicts either: (i) The indistinguishability of the scheme; (ii) The collision-resistance of the chameleon hash function; (iii) The SXDH assumption. We will prove this claim via a sequence of hybrid games:

In Game_0 , the challenger interacts with the adversary $\mathcal{A}$ as in the real evasiveness experiment. In the final game, we show that, if the adversary can output lossy tag with non-negligible probability, we can create an PPT algorithm breaks SXDH assumption with noticeable advantage.

For each i , we denote by bad_i , the event $\mathcal{A}$ manages to output a non-trivial lossy tag in Game_i .

Game₀: In this game, the adversary $\mathcal{A}$ has access to two oracles: (i) the lossy tag generation oracle $\text{LAF.LTag}(tk, \cdot)$ that always outputs lossy tags; (ii) the lossy tag verification oracle $\text{LAF.IsInjective}(\cdot)$ that uses a trapdoor to decide if a tag is lossy or injective. By definition.

$$\Pr[\text{bad}_0] = \Pr[\mathcal{A}(1^\lambda, ek)^{\text{LAF.LTag}(tk, \cdot), \text{LAF.IsInjective}(\cdot)}]. \quad (30)$$

Game₁: In this game, we define bad_{hash} to be the event that the adversary $\mathcal{A}$ manages to output a tag $t = (t_a, t_c = (\{R_i, S_i, D_i, E_i, F_i\}_{i=1}^n, r_{hash}))$ for which the corresponding chameleon hash collides with that of some tags produced by the oracle $\text{LAF.LTag}(tk, \cdot)$. The only difference between Game_0 and Game_1 is that the latter aborts when bad_{hash} occurs. It is straightforward that

$$|\Pr[\text{bad}_1] - \Pr[\text{bad}_0]| = \Pr[\text{bad}_{hash} \text{ in Game}_1]. \quad (31)$$

As in the proof of Theorem 2, we need to use the collision resistance property of the CMH to bound the probability $\Pr[\text{bad}_{hash} \text{ in Game}_1]$. Since the $\text{LAF.LTag}(tk, \cdot)$ oracle uses the trapdoor td_{CMH} to create lossy tags, we cannot immediately rely on the collision-resistance of CMH. Instead, we need to consider Game'_1 , where the $\text{LAF.LTag}(tk, \cdot)$ oracle always outputs injective tags instead of lossy ones. Using the result of Theorem 3, we have

$$\begin{aligned} & |\Pr[\text{bad}_{hash} \text{ in Game}_1] - \Pr[\text{bad}_{hash} \text{ in Game}'_1]| \\ & \leq n \cdot (\text{Adv}_{\mathcal{B}_1}^{\text{R-wD3DH1}}(\lambda) + \text{Adv}_{\mathcal{B}_2}^{\text{DDH1}}(\lambda)). \end{aligned}$$

Since Game_0 and Game_1 only differ when bad_{hash} occurs in Game_1 , we can bound to probability (31) as

$$\begin{aligned} |\Pr[\text{bad}_1] - \Pr[\text{bad}_0]| &= \Pr[\text{bad}_{hash} \text{ in Game}_1] \\ &\leq \Pr[\text{bad}_{hash} \text{ in Game}'_1] \\ &\quad + n \cdot (\text{Adv}_{\mathcal{B}_1}^{\text{R-wD3DH1}}(\lambda) + \text{Adv}_{\mathcal{B}_2}^{\text{DDH1}}(\lambda)). \end{aligned} \quad (32)$$

Since the trapdoor td_{CMH} is never been used in Game'_1 , a straightforward reduction shows that $\Pr[\text{bad}_{hash} \text{ in Game}_1] \leq \text{Adv}_{\text{CMH}}^{\text{CR}}(\lambda)$.

We now prove that, if event bad_1 occurs with noticeable probability, we can construct a PPT adversary $\mathcal{B}$ that breaks the unforgeability of the MAC of Section 4.1. As this property is proven by Theorem 4 under SXDH assumption, this will conclude the proof.

Our MAC adversary $\mathcal{B}$ will simulate Game_1 using the access to MAC oracle and MAC verification oracle. Algorithm $\mathcal{B}$ receives as input MAC public parameters

$$\text{pp} = ((\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T), g, \hat{g}, h, \hat{g}^\theta, (\hat{V}, \hat{\mathbf{Z}}_0, \hat{\mathbf{Z}}_1)).$$

As mentioned in Game_0 of the proof for Lemma 4, $\mathcal{B}$ additionally obtains the discrete logarithm $\eta = \log_g(h)$ from the MAC challenger, which will allow it to simulate the LAF.IsInjective oracle in the evasiveness experiment. Then, $\mathcal{B}$ extends these public parameters into public key of a LAF public key. To this end, $\mathcal{B}$ randomly guesses the position $i^* \xleftarrow{R} [n]$ of the MAC forgery in the non-injective tag it is expected to produce and sets

$$\text{pp}_{i^*} = ((\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T), g, \hat{g}, h, \hat{g}^\theta, (\hat{V}, \hat{\mathbf{Z}}_0, \hat{\mathbf{Z}}_1))$$

using the public parameters pp obtained from its challenger. Next, for all $i \in [n] \setminus \{i^*\}$, $\mathcal{B}$ sets $\hat{h} = \hat{g}^\eta$ and generates $n - 1$ keys $(\{\text{pp}_i, \text{sk}_{\text{mac}, i}\})_{i \in [n] \setminus \{i^*\}}$ for the MAC of Section 4.1 which all share the same $g, h \in \mathbb{G}$ and $\hat{g} \in \hat{\mathbb{G}}$. For each $i \in [n] \setminus \{i^*\}$, the i -th set of MAC public parameters thus consist of

$$\text{pp}_i = ((\mathbb{G}, \hat{\mathbb{G}}, \mathbb{G}_T), g, \hat{g}, h, \hat{g}^{\theta_i}, (\hat{V}_i, \hat{\mathbf{Z}}_{i,0}, \hat{\mathbf{Z}}_{i,1})).$$

To complete the generation the LAF evaluation key, $\mathcal{B}$ chooses $a \xleftarrow{R} \mathbb{Z}_p$, which is used to set $u = g^a$ and $\hat{u} = \hat{g}^a$. In addition, $\mathcal{B}$ computes $\hat{H}_i = (\hat{V}_i^{i \cdot \eta} \cdot \hat{V}_i^a)$ for all $i \in [n]$. It also chooses a key pair $(hk_{\text{CMH}}, td_{\text{CMH}})$ for the chameleon hash function CMH and includes hk_{CMH} in the LAF evaluation key.

To simulate the generation of lossy tags at each LAF.LTag -query t_a made by $\mathcal{A}$, $\mathcal{B}$ forwards t_a to its challenger and obtains a MAC $(R_{i^*}, D_{i^*}, E_{i^*}, F_{i^*})$ of the message t_a . Then, $\mathcal{B}$ computes on its own $n - 1$ MACs $\{(R_i, D_i, E_i, F_i)\}_{i \in [n] \setminus \{i^*\}}$ using the secret MAC keys $\{\text{sk}_{\text{mac}, i}\}_{i \in [n] \setminus \{i^*\}}$. For all $i \in [n]$, $\mathcal{B}$ also computes $S_i = R_i^{\eta \cdot i + a}$. It finally uses the trapdoor of the chameleon hash function to generate r_{hash} such that

$$\tau = \text{CMhash}(hk_{\text{CMH}}, (t_a, \{R_i, S_i, D_i, E_i, F_i\}_{i=1}^n), r_{\text{hash}}) \in \{0, 1\}^L$$

and returns $t = (t_a, t_c = (\{R_i, S_i, D_i, E_i, F_i\}_{i=1}^n, r_{\text{hash}}))$ to $\mathcal{A}$.

In order to simulate the LAF.IsInjective oracle for an input tag $t = (t_a, t_c)$ containing $t_c = (\{R_i, S_i, D_i, E_i, F_i\}_{i=1}^n, r_{\text{hash}})$, $\mathcal{B}$ computes the chameleon hash

$$\tau = \text{CMhash}(hk_{\text{CMH}}, (t_a, \{R_i, S_i, D_i, E_i, F_i\}_{i=1}^n), r_{\text{hash}}).$$

It then uses η to run the MAC verification oracle and check that $(R_{i^*}, D_{i^*}, E_{i^*}, F_{i^*})$ is a valid MAC by testing equation (12), which becomes

$$\frac{e(D_{i^*}, \hat{g}) \cdot e(E_{i^*}, \hat{g}^\theta)}{e(F_{i^*}, \prod_{k=1}^L Z_{i^*, k, \tau[k]})} = e(R_{i^*}, \hat{V}_{i^*})^\eta$$

here. Since $\mathcal{B}$ also knows $\text{sk}_{\text{mac},i}$ for all $i \in [n] \setminus \{i^*\}$, it can efficiently check that $\{(R_i, D_i, E_i, F_i)\}_{i \in [n] \setminus \{i^*\}}$ are all valid MACs of t_a . If all the MACs are valid, $\mathcal{B}$ also checks that for all $i \in [n]$, $S_i = R_i^{\eta \cdot i + a}$. If there exists $i \in [n]$ such that $S_i \neq R_i^{\eta \cdot i + a}$, $\mathcal{B}$ returns $\perp$ to indicate that t does not belong to the space $\mathcal{T}$ of valid tags. Finally, $\mathcal{B}$ outputs 0 (meaning that the tag t is non-injective) if it contains at least one valid MAC. If the n MACs contained in t_c are all invalid and $S_i = R_i^{\eta \cdot i + a}$ for all $i \in [n]$, it outputs 1 meaning that the tag is injective.

It remains to show how $\mathcal{B}$ can extract a MAC forgery when bad_1 occurs in Game_1 . Namely, we let $\mathcal{B}$ halt as soon as $\mathcal{A}$ queries LAF.IsInjective on a non-injective tag $t = (t_a, t_c = (\{R_i, S_i, D_i, E_i, F_i\}_{i=1}^n, r_{\text{hash}}))$ (recall that $\mathcal{B}$ can always simulate LAF.IsInjective itself without invoking its challenger) for which $\tau = \text{CMhash}(hk_{\text{CMH}}, (t_a, \{R_i, S_i, D_i, E_i, F_i\}_{i=1}^n, r_{\text{hash}})) \in \{0, 1\}^L$ has never been queried to LAF.LTag . Since $\mathcal{B}$ only queries the MAC oracle in the simulation of LAF.LTag , we know that τ has never been queried to MAC oracle. Since t is non-injective, the tuples $\{(R_i, S_i, D_i, E_i, F_i)\}_{i=1}^n$ must contain at least one valid MAC. If $(R_{i^*}, D_{i^*}, E_{i^*}, F_{i^*})$ is a valid MAC (which occurs with probability $1/n$ since i^* was chosen uniformly and independently of the adversary's view), $\mathcal{B}$ can successfully break the unforgeability of MAC by outputting $(R_{i^*}, D_{i^*}, E_{i^*}, F_{i^*})$. We thus have

$$\Pr[\text{bad}_1] \leq n \cdot \text{Adv}_{\mathcal{A}}^{\text{uf-mac}}(\lambda) \leq n \cdot (\text{Adv}_{\mathcal{B}_1}^{\text{DDH}_2}(\lambda) + 2L \cdot \text{Adv}_{\mathcal{B}_2}^{\text{DDH}_1}(\lambda)).$$

Putting the above arguments altogether, we obtain

$$\begin{aligned} \text{Adv}_{\mathcal{Q}}^{\mathcal{A}, \text{eva}} &\leq \text{Adv}_{\mathcal{B}_0}^{\text{CMH-CR}}(\lambda) + n \cdot (\text{Adv}_{\mathcal{B}_1}^{\text{R-wD3DH1}}(\lambda) + \text{Adv}_{\mathcal{B}_3}^{\text{DDH}_1}(\lambda) \\ &\quad + \text{Adv}_{\mathcal{B}_2}^{\text{DDH}_2}(\lambda) + 2L \cdot \text{Adv}_{\mathcal{B}_3}^{\text{DDH}_1}(\lambda)) \\ &= \text{Adv}_{\mathcal{B}_0}^{\text{CMH-CR}}(\lambda) + n \cdot \text{Adv}_{\mathcal{B}_1}^{\text{R-wD3DH1}}(\lambda) \\ &\quad + n \cdot \text{Adv}_{\mathcal{B}_2}^{\text{DDH}_2}(\lambda) + n \cdot (1 + 2L) \cdot \text{Adv}_{\mathcal{B}_3}^{\text{DDH}_1}(\lambda). \end{aligned} \tag{33}$$

By applying Lemma 1, we obtained the stated upper bound. $\square$