Circular chosen-ciphertext security with compact ciphertexts

Dennis Hofheinz*

Karlsruhe Institute of Technlogy

Abstract. A key-dependent message (KDM) secure encryption scheme is secure even if an adversary obtains encryptions of messages that depend on the secret key. Such key-dependent encryptions naturally occur in scenarios such as harddisk encryption, formal cryptography, or in specific protocols. However, there are not many provably secure constructions of KDM-secure encryption schemes. Moreover, only one construction, due to Camenisch, Chandran, and Shoup (Eurocrypt 2009) is known to be secure against active (i.e., CCA) attacks.

In this work, we construct the first public-key encryption scheme that is KDM-secure against active adversaries and has compact ciphertexts. As usual, we allow only circular key dependencies, meaning that encryptions of arbitrary *entire* secret keys under arbitrary public keys are considered in a multi-user setting.

Technically, we follow the approach of Boneh, Halevi, Hamburg, and Ostrovsky (Crypto 2008) to KDM security, which however only achieves security against passive adversaries. We explain an inherent problem in adapting their techniques to active security, and resolve this problem using a new technical tool called "lossy algebraic filters" (LAFs). We stress that we significantly deviate from the approach of Camenisch, Chandran, and Shoup to obtain KDM security against active adversaries. This allows us to develop a scheme with compact ciphertexts that consist only of a constant number of group elements.

Keywords: key-dependent messages, chosen-ciphertext security, public-key encryption.

1 Introduction

KDM security. An encryption scheme is key-dependent message (KDM) secure if it is secure even against an adversary who has access to encryptions of messages that depend on the secret key. Such a setting arises, e.g., in harddisk encryption [10], computational soundness results in formal methods [7, 2], or specific protocols [13]. KDM security does not follow from standard security [1, 15], and there are indications [19, 5] that KDM security (at least in its most general form) cannot be proven using standard techniques; it seems that dedicated constructions and proof techniques are necessary.[1]

* Supported by DFG grant GZ HO 4534/2-1.

[1] We mention, however, that there are semi-generic transformations that *enhance* the KDM security of an already "slightly" KDM-secure scheme [5, 12, 3].

The BHHO approach to KDM-CPA security. Boneh, Halevi, Hamburg, and Ostrovsky [10] (henceforth BHHO) were the first to construct and prove a public-key encryption (PKE) scheme that is KDM secure under chosen-plaintext attacks (KDM-CPA-secure) in the standard model, under the Decisional Diffie-Hellman (DDH) assumption. While they did not prove their scheme secure under messages that *arbitrarily* depend on the secret key, their result encompasses the important case of *circular (CIRC-CPA) security*. Loosely speaking, a PKE scheme is circular secure if it is secure even in a multi-user setting where encryptions of arbitrary secret keys under arbitrary public keys are known. This notion is sufficient for certain applications [13], and can often be extended to stronger forms of KDM security [5, 12]. Inspired by BHHO, KDM-CPA-secure PKE schemes from other computational assumptions followed [4, 11, 25].

Since we will be using a similar approach, we give a high-level intuition of BHHO's approach. The crucial property of their scheme is that it is *publicly* possible to construct encryptions of the secret key (under the corresponding public key). Thus, encryptions of the secret key itself do not harm the (IND-CPA) security of that scheme. Suitable homomorphic properties of both keys and ciphertexts allow to extend this argument to circular security (for arbitrarily many users/keys), and to affine functions of all keys.

Why the BHHO approach fails to achieve KDM-CCA security. When considering an *active* adversary, we require a stronger form of KDM security. Namely, KDM-CCA, resp. CIRC-CCA security requires security against an adversary who has access to key-dependent encryptions *and* a decryption oracle. (Naturally, to avoid a trivial notion, the adversary is not allow to submit any of those given KDM encryptions to its decryption oracle.) Now if we want to extend BHHO's KDM-CPA approach to an adversary with a decryption oracle, the following problem arises: since it is publicly possible to construct (fresh) encryptions of the secret key, an adversary can generate such an encryption and then submit it to its decryption oracle, thus obtaining the full secret key. Hence, the very property that BHHO use to prove KDM-CPA security seemingly contradicts chosen-ciphertext security.

Our technical tool: lossy algebraic filters (LAFs). Before we describe our approach to KDM-CCA security, let us present the core technical tool we use. Namely, a *lossy algebraic filter* (LAF) is a family of functions, indexed by a public key and a tag. A function from that family takes a vector $X = (X_i)_{i=1}^{\mathfrak{n}}$ as input. Now if the tag is *lossy*, then the output of the function reveals only a linear combination of the X_i. If the tag is *injective*, however, then so is the function. We require that there are many lossy tags, which however require a special trapdoor to be found. On the other hand, lossy and injective tags are computationally indistinguishable. This concept is very similar to (parameterized) lossy trapdoor functions [27], and in particular to all-but-many lossy trapdoor functions (ABM-LTFs [20]). However, we do not require efficient inversion, but we do require that lossy functions always reveal *the same* linear combination about the input. In particular, evaluating the same input under many lossy tags will still leave the input (partially) undetermined.

We give a construction of LAFs under the Decision Linear (DLIN) assumption in pairing-friendly groups. Similar to ABM-LTFs, lossy tags correspond to suitably blinded signatures. (This in particular allows to release many lossy tags, while still making the generation of a fresh lossy tag hard for an adversary.) However, unlike with ABM-LTFs, functions with lossy tags always release the same information about its input. Our construction has compact tags with $\mathbf{O}(1)$ group elements, which will be crucial for our KDM-CCA secure encryption scheme.[2]

Our approach to KDM-CCA security. We can now describe our solution to the KDM-CCA dilemma explained above. We will start from a hybrid between the BHHO-like PKE schemes of Brakerski and Goldwasser [11], resp. Malkin et al. [25]. This scheme has compact ciphertexts ($\mathbf{O}(1)$ group elements), and its KDM-CPA security can be proved under the Decisional Composite Residuosity (DCR) assumption. As with the BHHO scheme, the scheme's KDM-CPA security relies on the fact that encryptions of its secret key can be publicly generated. Essentially, our modification consists of adding a suitable authentication tag to each ciphertext. This authentication tag comprises the (encrypted) image of the plaintext message under an LAF. During decryption, a ciphertext is rejected in case of a wrong authentication tag.

In our security proof, all authentication tags for the key-dependent encryptions the adversary gets are made with respect to lossy filter tags. This means that information-theoretically, little information about the secret key is released (even with many key-dependent encryptions, resp. LAF evaluations). However, any decryption query the adversary makes must refer (by the LAF properties) to an injective tag. Hence, in order to place a valid key-dependent decryption query, the adversary would have to guess the whole (hidden) secret key.[3]

Thus, adding a suitable authentication tag allows us to leverage the techniques by BHHO, resp. Brakerski and Goldwasser, Malkin et al. to chosen-ciphertext attacks. In particular, we obtain a CIRC-CCA-secure PKE scheme with compact ciphertexts (of $\mathbf{O}(1)$ group elements). We prove security under the conjunction of the following assumptions: the DCR assumption (in $\mathbb{Z}^*_{N^3}$), the DLIN assumption (in a pairing-friendly group), and the DDH assumption (somewhat curiously, in the subgroup of order $(P-1)(Q-1)/4$ of $\mathbb{Z}^*_{N^3}$, where $N = PQ$).[4]

[2] The size of the LAF public key depends on the employed signature scheme. In our main construction, we use Waters signatures, which results in very compact tags, but public keys of $\mathbf{O}(k)$ group elements, where k is the security parameter. Alternatively, at the end of Section 3.1, we sketch an LAF with constant-size (but larger than in our main construction) tags *and* constant-size public keys.

[3] We will also have to protect against a re-use of (lossy) authentication tags, and "ordinary", key-independent chosen-ciphertext attacks. This will be achieved by a combination of one-time signatures and 2-universal hash proof systems [16, 24, 22].

[4] Very roughly, we resort to the DDH assumption since we release *partial* information about our secret keys. Whereas the argument of [11, 25] relies on the fact that the secret key sk is completely hidden modulo N, where computations take place in $\mathbb{Z}_N$, we cannot avoid to leak some information about $sk \bmod N$ by releasing LAF images of sk. However, using a suitable message encoding, we *can* argue that sk is

Relation to Camenisch et al.'s CIRC-CCA-secure scheme. Camenisch, Chandran, and Shoup [14] present the only other known CIRC-CCA-secure PKE scheme in the standard model. They also build upon BHHO techniques, but instead use a Naor-Yung-style double encryption technique [26] to achieve chosen-ciphertext security. As an authentication tag, they attach to each ciphertext a non-interactive zero-knowledge proof that *either* the encryption is consistent (in the usual Naor-Yung sense), *or* that they know a signature for the ciphertext. Since they build on the original, DDH-based BHHO scheme, they can use Groth-Sahai proofs [18] to prove consistency. Compared to our scheme, their system is less efficient: they require $\mathbf{O}(k)$ group elements per ciphertext, and the secret key can only be encrypted bitwise. However, their sole computational assumption to prove circular security is the DDH (or, more generally, k-Linear) assumption in pairing-friendly groups. One thing to point out is their implicit use of a signature scheme. Their argument is conceptually not unlike our LAF argument. However, since they can apply a hybrid argument to substitute all key-dependent encryptions with random ciphertexts, they only require one-time signatures. Furthermore, the meaning of "consistent ciphertext" and "proof" in our case is very different. (Unlike Camenisch et al., we apply an argument that rests on the *information* that the adversary has about the secret key.)

Note about concurrent work. In a work concurrent to ours, Galindo, Herranz, and Villar [17] define and instantiate a strong notion of KDM security for identity-based encryption (IBE) schemes. Using the IBE→PKE transformation of Boneh, Canetti, Halevi, and Katz [9], they derive a KDM-CCA-secure PKE scheme. Their concrete construction is entropy-based and achieves only a bounded form of KDM security, much like the KDM-secure SKE scheme from [23]. Thus, while their ciphertexts are very compact, they can only tolerate a number of (arbitrary) KDM queries that is linear in the size of the secret key. In particular, it is not clear how to argue that the encryption of a full secret key in their scheme is secure.

2 Preliminaries

Notation. For $n \in \mathbb{N}$, let $[n] := \{1, \ldots, n\}$. Throughout the paper, $k \in \mathbb{N}$ is the security parameter. For a finite set $\mathcal{S}$, $s \leftarrow \mathcal{S}$ denotes the process of sampling s uniformly from $\mathcal{S}$. For a probabilistic algorithm A, $y \leftarrow A(x; R)$ denotes the process of running A on input x and with randomness R, and assigning y the result. We write $y \leftarrow A(x)$ for $y \leftarrow A(x; R)$ with uniformly chosen R. If A's running time is polynomial in k, A is called probabilistic polynomial-time (PPT).

Standard definitions. Due to lack of space, we postpone some standard definitions to the full version [21]. These include definitions of PKE and signature schemes, (one-time/strong) EUF-CMA security, IND-CPA security, (chameleon) hash functions, and the DCR, DDH, and DLIN assumptions.

completely hidden modulo the coprime order $(P-1)(Q-1)/4$ of quadratic residues modulo N, which enables a reduction to the DDH assumption.

Key-unique SKE schemes. A secret-key encryption (SKE) scheme (E, D) consists of two PPT algorithms. Encryption $\mathsf{E}(K, M)$ takes a key K and a message M, and outputs a ciphertext C. Decryption $\mathsf{D}(K, C)$ takes a key K and a ciphertext C, and outputs a message M. For correctness, we want $\mathsf{D}(K, C) = M$ for all M, all K, and all $C \leftarrow \mathsf{E}(K, M)$. We say that (E, D) is *key-unique* if for every ciphertext C, there is at most one key K with $\mathsf{D}(K, C) \neq \perp$. For instance, ElGamal encryption can be interpreted as a key-unique SKE scheme through $\mathsf{E}(x, M) := (g^x, g^y, g^{xy} \cdot M)$ (and the obvious D). This example assumes a publicly known group $\mathbb{G} = \langle g \rangle$ in which the DDH assumption holds.[5] If a larger message space (e.g., $\{0,1\}^*$) is desired, hybrid encryption techniques (which are easily seen to preserve key-uniqueness) can be employed.

Pairings. A (symmetric) pairing is a map $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$ between two cyclic groups $\mathbb{G}$ and $\mathbb{G}_T$ that satisfies $e(g, g) \neq 1$ and $e(g^a, g^b) = e(g, g)^{ab}$ for all generators g of $\mathbb{G}$ and all $a, b \in \mathbb{Z}$.

Waters signatures. In [28], Waters proves the following signature scheme EUF-CMA secure:[6]

- $\mathsf{Gen}(1^k)$ chooses groups $\mathbb{G}, \mathbb{G}_T$ of prime order p, along with a pairing $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$, a generator $g \in \mathbb{G}$, and uniform group elements $g^\omega, H_0, \ldots, H_k \in \mathbb{G}$. Output is $vk = (\mathbb{G}, \mathbb{G}_T, e, p, g, (H_i)_{i=0}^k, e(g, g)^\omega)$ and $sigk = (vk, g^\omega)$.
- $\mathsf{Sig}(sigk, M)$, for $M = (M_i)_{i=1}^k \in \{0,1\}^k$, picks $r \leftarrow \mathbb{Z}_p$, and outputs $\sigma := (g^r, g^\omega \cdot (H_0 \prod_{i=1}^k H_i^{M_i})^r)$.
- $\mathsf{Ver}(vk, M, (\sigma_0, \sigma_1))$, outputs 1 iff $e(g, \sigma_1) = e(g, g)^\omega \cdot e(\sigma_0, H_0 \prod_{i=1}^k H_i^{M_i})$.

KDM-CCA and CIRC-CCA security. Let $n = n(k)$ and let PKE be a PKE scheme with message space $\mathcal{M}$. PKE is chosen-ciphertext secure under key-dependent message attacks (n-KDM-CCA secure) iff

$$\mathsf{Adv}^{\mathsf{kdm\text{-}cca}}_{\mathsf{PKE},n,A}(k) \;:=\; \Pr\left[\mathsf{Exp}^{\mathsf{kdm\text{-}cca}}_{\mathsf{PKE},n,A}(k) = 1\right] - 1/2$$

is negligible for all PPT A, where experiment $\mathsf{Exp}^{\mathsf{kdm\text{-}cca}}_{\mathsf{PKE},n,A}$ is defined as follows. First, the experiment tosses a coin $b \leftarrow \{0,1\}$, and samples public parameters $pp \leftarrow \mathsf{Pars}(1^k)$ and n keypairs $(pk_i, sk_i) \leftarrow \mathsf{Gen}(pp)$. Then A is invoked with input pp and $(pk_i)_{i=1}^n$, and access to two oracles:

- a KDM oracle $\mathcal{KDM}_b(\cdot, \cdot)$ that maps $i \in [n]$ and a function $f : (\{0,1\}^*)^n \to \{0,1\}^*$ to a ciphertext $C \leftarrow \mathsf{Enc}(pp, pk_i, M)$. If $b = 0$, then $M = f((sk_i)_{i=1}^n)$; else, $M = 0^{|f((sk_i)_{i=1}^n)|}$.
- a decryption oracle $\mathcal{DEC}(\cdot, \cdot)$ that takes as input an index $i \in [n]$ and a ciphertext C, and outputs $\mathsf{Dec}(pp, sk_i, C)$.

When A finally generates an output $b' \in \{0,1\}$, the experiment outputs 1 if $b = b'$ (and 0 else). We require that (a) A never inputs a ciphertext C to $\mathcal{DEC}$ that has been produced by $\mathcal{KDM}_b$ (for the same index i), and (b) A only specifies PPT-computable functions f that always output messages of the same length. As a

[5] In our application, $\mathbb{G}$ can be made part of the public parameters.

[6] In fact, our description is a slight folklore variant of Waters's scheme. The original scheme features elements g^α, g^β in vk, so that $e(g^\alpha, g^\beta)$ takes the role of $e(g, g)^\omega$.

relevant special case, PKE is *n-CIRC-CCA-secure* if it is n-KDM-CCA secure against all A that only query $\mathcal{KDM}_b$ with functions $f \in \mathcal{F}$ for

$$\mathcal{F} := \{f_j : f_j((sk_i)_{i=1}^n) = sk_j\}_{j\in[n]} \cup \{f_M : f_M((sk_i)_{i=1}^n) = M\}_{M\in\mathcal{M}}\,.$$

(Technically, what we call "circular security" is called "clique security" in [10]. However, our notion of circular security implies that of [10].) Our main result will be a PKE scheme that is n-CIRC-CCA-secure for all polynomials $n = n(k)$.

3 Lossy algebraic filters

Informal description. An $(\ell_{\mathsf{LAF}}, \mathfrak{n})$-lossy algebraic filter (LAF) is a family of functions indexed by a public key *Fpk* and a tag t. A function $\mathsf{LAF}_{Fpk,t}$ from the family maps an input $X = (X_i)_{i=1}^{\mathfrak{n}} \in \mathbb{Z}_p^{\mathfrak{n}}$ to an output $\mathsf{LAF}_{Fpk,t}(X)$, where p is an ℓ_{LAF}-bit prime contained in the public key.

The crucial property of an LAF is its lossiness. Namely, for a given public key *Fpk*, we distinguish *injective* and *lossy* tags.[7] For an injective tag t, the function $\mathsf{LAF}_{Fpk,t}(\cdot)$ is injective, and thus has an image of size $p^{\mathfrak{n}}$. However, if t is lossy, then $\mathsf{LAF}_{Fpk,t}(\cdot)$ only depends on a linear combination $\sum_{i=1}^{\mathfrak{n}} \omega_i X_i \bmod p$ of its input. In particular, different X with the same value $\sum_{i=1}^{\mathfrak{n}} \omega_i X_i \bmod p$ are mapped to the same image. Here, the coefficients $\omega_i \in \mathbb{Z}_p$ only depend on *Fpk* (but not on t). For a lossy tag t, the image of $\mathsf{LAF}_{Fpk,t}(\cdot)$ is thus of size at most p. Note that the modulus p is public, while the coefficients ω_i may be (and in fact will have to be) computationally hidden.

For this concept to be useful, we require that (a) lossy and injective tags are computationally indistinguishable, (b) lossy tags can be generated using a special trapdoor, but (c) new lossy (or, rather, non-injective) tags cannot be found efficiently without that trapdoor, even when having seen polynomially many lossy tags before. In view of our application, we will work with structured tags: each tag $t = (t_{\mathsf{c}}, t_{\mathsf{a}})$ consists of a *core tag* t_{c} and an *auxiliary tag* t_{a}. The auxiliary tag will be a ciphertext part that is authenticated by a filter image.

Definition 1. *An $(\ell_{\mathsf{LAF}}, \mathfrak{n})$-lossy algebraic filter (LAF) LAF consists of three PPT algorithms:*

Key generation. $\mathsf{FGen}(1^k)$ *samples a keypair (Fpk, Ftd). The public key Fpk contains an ℓ_{LAF}-bit prime p and the description of a tag space $\mathcal{T} = \mathcal{T}_{\mathsf{c}} \times \{0,1\}^*$ for efficiently samplable $\mathcal{T}_{\mathsf{c}}$. A tag $t = (t_{\mathsf{c}}, t_{\mathsf{a}})$ consists of a core tag $t_{\mathsf{c}} \in \mathcal{T}_{\mathsf{c}}$ and an auxiliary tag $t_{\mathsf{a}} \in \{0,1\}^*$. A tag may be injective, or lossy, or neither. Ftd is a trapdoor that will allow to sample lossy tags.*

Evaluation. $\mathsf{FEval}(Fpk, t, X)$*, for a public key Fpk and a tag $t = (t_{\mathsf{c}}, t_{\mathsf{a}}) \in \mathcal{T}$, maps an input $X = (X_i)_{i=1}^{\mathfrak{n}}) \in \mathbb{Z}_p^{\mathfrak{n}}$ to a unique output $\mathsf{LAF}_{Fpk,t}(X)$.*

Lossy tag generation. $\mathsf{FTag}(Ftd, t_{\mathsf{a}})$*, for a trapdoor Ftd and $t_{\mathsf{a}} \in \{0,1\}^*$, samples a core tag t_{c} such that $t = (t_{\mathsf{c}}, t_{\mathsf{a}})$ is lossy.*

We require the following:

[7] Technically, there may also be tags that are neither injective nor lossy.

Lossiness. *The function* $\mathsf{LAF}_{Fpk,t}(\cdot)$ *is injective if* t *is injective. If* t *is lossy, then* $\mathsf{LAF}_{Fpk,t}(X)$ *depends only on* $\sum_{i=1}^{\mathfrak{n}} \omega_i X_i \bmod p$ *for* $\omega_i \in \mathbb{Z}_p$ *that only depend on* Fpk*.*

Indistinguishability. *Lossy tags are indistinguishable from random tags:*

$$\mathsf{Adv}^{\mathsf{ind}}_{\mathsf{LAF},A}(k) := \Pr\left[A(1^k, Fpk)^{\mathsf{FTag}(Ftd,\cdot)} = 1\right] - \Pr\left[A(1^k, Fpk)^{\mathcal{O}_{\mathcal{T}_\mathsf{c}}(\cdot)} = 1\right]$$

is negligible for all PPT A*, where* $(Fpk, Ftd) \leftarrow \mathsf{FGen}(1^k)$*, and* $\mathcal{O}_{\mathcal{T}_\mathsf{c}}(\cdot)$ *is the oracle that ignores its input and samples a random core tag* t_c*.*

Evasiveness. *Non-injective (and in particular lossy) tags are hard to find, even given multiple lossy tags:*

$$\mathsf{Adv}^{\mathsf{eva}}_{\mathsf{LAF},A}(k) := \Pr\left[t \text{ non-injective} \;\middle|\; t \leftarrow A(1^k, Fpk)^{\mathsf{FTag}(Ftd,\cdot)}\right]$$

is negligible with $(Fpk, Ftd) \leftarrow \mathsf{FGen}(1^k)$*, and for any PPT algorithm* A *that never outputs a tag obtained through oracle queries (i.e.,* A *never outputs* $t = (t_\mathsf{c}, t_\mathsf{a})$ *when* t_c *has been obtained by an oracle query* $\mathsf{FTag}(Ftd, t_\mathsf{a})$*).*

3.1 Construction

Intuition. We present a construction based on the DLIN problem in a group $\mathbb{G}$ of order p with symmetric pairing $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$. Essentially, each tag corresponds to $\mathfrak{n}$ DLIN-encrypted Waters signatures. If the signatures are valid, the tag is lossy. The actual filter maps an input $X = (X_i)_{i=1}^{\mathfrak{n}} \in \mathbb{Z}_p^{\mathfrak{n}}$ to the tuple

$$\mathsf{LAF}_{Fpk,t}(X) \;:=\; \mathbf{M} \circ X \;:=\; \Big(\prod_{j=1}^{\mathfrak{n}} \mathbf{M}_{i,j}^{X_j}\Big)_{j=1}^{\mathfrak{n}} \;\in\; \mathbb{G}_T^{\mathfrak{n}}, \tag{1}$$

where the matrix $\mathbf{M} = (\mathbf{M}_{i,j})_{i,j\in[\mathfrak{n}]} \in \mathbb{G}_T^{\mathfrak{n}\times\mathfrak{n}}$ is computed from public key and tag. Note that this mapping is lossy if and only if the matrix

$$\widetilde{\mathbf{M}} := (\widetilde{\mathbf{M}}_{i,j}) := (\mathrm{dlog}_{e(g,g)}(\mathbf{M}_{i,j}))_{i,j} \in \mathbb{Z}_p^{\mathfrak{n}\times\mathfrak{n}} \tag{2}$$

of discrete logarithms (to some arbitrary basis $e(g,g) \in \mathbb{G}_T$) is non-invertible.

For a formal description, let $\ell_{\mathsf{LAF}}(k), \mathfrak{n}(k)$ be two functions.

Key generation. $\mathsf{FGen}(1^k)$ generates cyclic groups $\mathbb{G}, \mathbb{G}_T$ of prime order p (where p has bitlength $\lfloor \log_2(p) \rfloor = \ell_{\mathsf{LAF}}(k)$), and a symmetric pairing $e : \mathbb{G} \times \mathbb{G} \to \mathbb{G}_T$. Then FGen chooses

- a generator $g \in \mathbb{G}$ and a uniform exponent $\omega \leftarrow \mathbb{Z}_p$,
- uniform group elements $U_1, \ldots, U_{\mathfrak{n}} \leftarrow \mathbb{G}$, $H_0, \ldots, H_k \leftarrow \mathbb{G}$, and
- a keypair (Hpk, Htd) for a chameleon hash $\mathsf{CH} : \{0,1\}^* \to \{0,1\}^k$.

FGen finally outputs

$$\begin{aligned} Fpk &:= (\mathbb{G}, \mathbb{G}_T, e, p, g, (H_i)_{i=0}^k, (U_i)_{i=1}^{\mathfrak{n}}, W := e(g,g)^\omega, Hpk) \\ Ftd &:= (Fpk, g^\omega, Htd). \end{aligned}$$

For convenience, write $U_i = g^{u_i}$ for suitable (unknown) exponents u_i.

Tags. (Core) tags are of the form

$$t_{\mathsf{c}} := (R, (\widetilde{S}_i)_{i=1}^{\mathfrak{n}}, (S_{i,j})_{i,j=1}^{\mathfrak{n}}, R_{\mathsf{CH}}) \in \mathbb{G} \times \mathbb{G} \times \mathbb{G}^{\mathfrak{n}\times\mathfrak{n}} \times \mathcal{R}_{\mathsf{CH}}$$

(for CH's randomness space $\mathcal{R}_{\mathsf{CH}}$), where we require $e(U_{j'}, S_{i,j}) = e(U_j, S_{i,j'})$ whenever $i \notin \{j, j'\}$. This means we can write $R = g^r$, $\widetilde{S}_i = g^{\widetilde{s}_i}$, and $S_{i,j} = U_j^{s_i}$ (for $i \neq j$) for suitable $r, s_i, \widetilde{s}_i$. To a tag $t = (t_{\mathsf{c}}, t_{\mathsf{a}})$ (with auxiliary part $t_{\mathsf{a}} \in \{0,1\}^*$), we associate the matrix $\mathbf{M} = (\mathbf{M}_{i,j})_{i,j=1}^{\mathfrak{n}} \in \mathbb{G}_T^{\mathfrak{n}\times\mathfrak{n}}$ with

$$\begin{aligned} \mathbf{M}_{i,j} &= e(U_j, \widetilde{S}_i) \cdot e(g, S_{i,j}) = e(g,g)^{u_j(\widetilde{s}_i + s_i)} \qquad (i \neq j) \\ \mathbf{M}_{i,i} &= \frac{e(g, S_{i,i})}{W \cdot e(H_0 \prod_{i=1}^k H_i^{T_i}, R)} \end{aligned} \tag{3}$$

for $(T_i)_{i=1}^k := \mathsf{CH}_{Hpk}(R, (\widetilde{S}_i)_{i=1}^{\mathfrak{n}}, (S_{i,j})_{i,j=1}^{\mathfrak{n}}, t_{\mathsf{a}}; R_{\mathsf{CH}})$. If the matrix $\widetilde{\mathbf{M}}$ of discrete logarithms (see (2)) is invertible, we say that t is injective; if $\widetilde{\mathbf{M}}$ has rank 1, then t is lossy. Thus, for lossy tags, $\mathbf{M}_{i,j} = e(g,g)^{u_j(\widetilde{s}_i + s_i)}$ for *all* i, j.

Evaluation. $\mathsf{FEval}(Fpk, t, X)$, for $t = (t_{\mathsf{c}}, t_{\mathsf{a}})$, $t_{\mathsf{a}} \in \{0,1\}^*$, $X = (X_i)_{i=1}^{\mathfrak{n}} \in \mathbb{Z}_p^{\mathfrak{n}}$, and Fpk and t_{c} as above, computes $\mathbf{M}$ as in (3) and then $(Y_i)_{i=1}^{\mathfrak{n}} := \mathsf{LAF}_{Fpk,t}(X) \in \mathbb{G}_T^{\mathfrak{n}}$ as in (1).

Lossiness. If we write $Y_i = e(g,g)^{y_i}$, the definition of FEval implies $(y_i)_{i=1}^{\mathfrak{n}} = \widetilde{\mathbf{M}} \cdot X$. Since injective tags satisfy that $\widetilde{\mathbf{M}}$ is invertible, they lead to injective functions $\mathsf{LAF}_{Fpk,t}(\cdot)$. But for a lossy tag, $\widetilde{\mathbf{M}}_{i,j} = u_j(\widetilde{s}_i + s_i)$, so that

$$y_i = \sum_{j=1}^{\mathfrak{n}} u_j(\widetilde{s}_i + s_i) X_j = (\widetilde{s}_i + s_i) \cdot \sum_{j=1}^{\mathfrak{n}} u_j X_j \mod p.$$

Specifically, $\mathsf{LAF}_{Fpk,t}(X)$ depends only on $\sum_i \omega_i X_i \bmod p$ for $\omega_i := u_i$.

Lossy tag generation. $\mathsf{FTag}(Ftd, t_{\mathsf{a}})$, for Ftd as above and $t_{\mathsf{a}} \in \{0,1\}^*$, first chooses a random CH-image $T = (T_i)_{i=1}^k \in \{0,1\}^k$ that can later be explained, using Htd, as the CH-image of an arbitrary preimage. FTag then chooses uniform $r, s_i, \widetilde{s}_i \leftarrow \mathbb{Z}_p$ and sets (for $i \neq j$)

$$\begin{aligned} R &:= g^r, \qquad \widetilde{S}_i := g^{\widetilde{s}_i}, \\ S_{i,j} &:= U_j^{s_i}, \quad S_{i,i} := U_i^{\widetilde{s}_i + s_i} \cdot g^{\omega} \cdot \left(H_0 \prod_{i=1}^k H_i^{T_i} \right)^r. \end{aligned} \tag{4}$$

Finally, FTag chooses R_{CH} with $\mathsf{CH}_{Hpk}(R, (\widetilde{S}_i)_{i=1}^{\mathfrak{n}}, (S_{i,j})_{i,j=1}^{\mathfrak{n}}, t_{\mathsf{a}}; R_{\mathsf{CH}}) = T$ and outputs $t_{\mathsf{c}} = (R, (\widetilde{S}_i)_{i=1}^{\mathfrak{n}}, (S_{i,j})_{i,j=1}^{\mathfrak{n}}, R_{\mathsf{CH}})$. Intuitively, t_{c} consists of $\mathfrak{n}$ DLIN encryptions (with correlated randomness $s_i, \widetilde{s}_i$) of Waters signatures $(g^r, g^{\omega} \cdot (H_0 \prod_{i=1}^k H_i^{T_i})^r)$ for message T. Indeed, substituting into (3) yields

$$\mathbf{M}_{i,i} := \frac{e(g,g)^{u_i(\widetilde{s}_i + s_i)} \cdot W \cdot e(g, (H_0 \prod_{i=1}^k H_i^{T_i})^r)}{W \cdot e(g^r, H_0 \prod_{i=1}^k H_i^{T_i})} = e(g,g)^{u_i(\widetilde{s}_i + s_i)}.$$

Hence, $\widetilde{\mathbf{M}}_{i,j} = u_j(\widetilde{s}_i + s_i)$ for *all* i, j, and thus the resulting tag $t = (t_{\mathsf{c}}, t_{\mathsf{a}})$ is lossy.

A generalization. In the full paper [21], we also show how to generalize the above construction to achieve constant-size tags *and* evaluation keys.

Other instances and further applications of LAFs. Since LAFs can be seen as "disguised signature schemes", it seems interesting to try to convert other signature schemes (and in particular schemes that do not require pairing-friendly groups) to LAFs. Besides, LAFs would seem potentially useful in other settings, specifically in settings with inherently many challenges (e.g., the selective-opening setting [6]).

Theorem 1. *If the DLIN assumption holds in* $\mathbb{G}$*, and* CH *is a chameleon hash function, then the LAF construction* LAF *from Section 3.1 satisfies Definition 1.*

The lossiness of LAF has already been discussed in Section 3.1. We prove indistinguishability and evasiveness separately.

Lemma 1. *For every adversary* A *on* LAF*'s indistinguishability, there exists a DLIN distinguisher* B *such that* $\mathsf{Adv}^{\mathsf{ind}}_{\mathsf{LAF},A}(k) = \mathfrak{n} \cdot \mathsf{Adv}^{\mathsf{dlin}}_{B}(k)$.

Intuitively, to see Lemma 1, observe that lossy tags differ from random tags only in their $S_{i,i}$ components, and in how the CH randomness R_{CH} is generated. For lossy tags, the $S_{i,i}$ are (parts of) DLIN ciphertexts, which are pseudorandom under the DLIN assumption. Furthermore, the uniformity property of CH guarantees that the distribution of R_{CH} is the same for lossy and random tags. We formally prove Lemma 1 in the full version [21].

Lemma 2. *For every adversary* A *on* LAF*'s evasiveness, there exist adversaries* B*,* C*, and* F *with* $\mathsf{Adv}^{\mathsf{eva}}_{\mathsf{LAF},A}(k) \leq \left|\mathsf{Adv}^{\mathsf{ind}}_{\mathsf{LAF},B}(k)\right| + \mathsf{Adv}^{\mathsf{cr}}_{\mathsf{CH},C}(k) + \mathsf{Adv}^{\mathsf{euf\text{-}cma}}_{\mathsf{Sig}_{\mathsf{Wat}},F}(k)$.

Intuitively, Lemma 2 holds because lossy (or, rather, non-injective) tags correspond to DLIN-encrypted Waters signatures. Hence, even after seeing many lossy tags (i.e., encrypted signatures), an adversary cannot produce a fresh encrypted signature. We note that the original Waters signatures from [28] are re-randomizable and thus not *strongly* unforgeable. To achieve evasiveness, we have thus used a chameleon hash function, much like Boneh et al. [8] did to make Waters signatures strongly unforgeable. We give a formal proof in [21].

Combining Lemma 1, Lemma 2, and the fact that Waters signatures are EUF-CMA secure already under the CDH assumption, we obtain Theorem 1.

4 CIRC-CCA-secure encryption scheme

Setting and ingredients. First, we assume an algorithm GenN that outputs ℓ_N-bit Blum integers $N = PQ$ along with their prime factors P and Q. If N is clear from the context, we write $\mathbb{G}_{\mathsf{rnd}}$ and $\mathbb{G}_{\mathsf{msg}}$ for the unique subgroups of $\mathbb{Z}^*_{N^3}$ of order $(P-1)(Q-1)/4$, resp. N^2. We also write $h := 1 + N \bmod N^3$, so $\langle h \rangle = \mathbb{G}_{\mathsf{msg}}$. Note that it is efficiently possible to compute $\mathrm{dlog}_h(X) := x$ for $X := h^x \in \mathbb{G}_{\mathsf{msg}}$ and $x \in \mathbb{Z}_{N^2}$. Specifically, it is efficiently possible to test for membership in $\mathbb{G}_{\mathsf{msg}}$. In our scheme, $\mathbb{G}_{\mathsf{msg}}$ will be used to embed a suitably encoded message, and $\mathbb{G}_{\mathsf{rnd}}$ will be used for blinding. We require that

– P and Q are safe primes of bitlength between $\ell_N/2 - k$ and $\ell_N/2 + k$,
– $\gcd((P-1)(Q-1)/4, N) = 1$ (as, e.g., for uniform P, Q of a certain length),
– $\ell_N \geq 25k + 8$ (e.g., $k = 80$ and $\ell_N = 2048$)
– the DCR assumption holds in $\mathbb{Z}^*_{N^3}$, and the DDH assumption holds in $\mathbb{G}_{\mathsf{rnd}}$.

We also assume an $(\ell_{\mathsf{LAF}}, \mathfrak{n})$-lossy algebraic filter LAF for $\mathfrak{n} = 6$ and $\ell_{\mathsf{LAF}} = (\ell_N + k + 1)/(\mathfrak{n} - 2)$. Our scheme will encrypt messages from the domain

$$\mathcal{M} := \mathbb{Z}_{2^{3k}} \times \mathbb{Z}_{p \cdot 2^k} \times \mathbb{Z}_{N \cdot 2^{k-2}},$$

where p is the modulus of the used LAF. (The reason for this weird-looking message space will become clearer in the proof.) During encryption, we will have to treat a message $M = (a, b, c) \in \mathcal{M}$ both as an element of $\mathbb{Z}_{N^2}$ and as an LAF-input from $\mathbb{Z}^{\mathfrak{n}}_p$. In these cases, we can encode

$$\begin{aligned} \mathbb{Z} &:= a + 2^{3k} \cdot b + p \cdot 2^{4k} \cdot c \in \mathbb{Z}, \\ [M]_{\mathbb{Z}^{\mathfrak{n}}_p} &:= (a, b \bmod p, c_0, \ldots, c_{\mathfrak{n}-3}) \in \mathbb{Z}^{\mathfrak{n}}_p \end{aligned} \tag{5}$$

for the natural interpretation of $\mathbb{Z}_i$-elements as integers between 0 and $i-1$, and c's p-adic representation $(c_i)_{i=0}^{\mathfrak{n}-3} \in \mathbb{Z}_p^{\mathfrak{n}-2}$ with $c = \sum_{i=0}^{\mathfrak{n}-3} c_i \cdot p^i$. By our choice of ℓ_N and ℓ_{LAF}, we have $0 \leq [M]_{\mathbb{Z}} < N^2 - 2^k$. However, the encoding $[M]_{\mathbb{Z}^{\mathfrak{n}}_p}$ is *not* injective, since it only depends on $b \bmod p$ (while $0 \leq b < p \cdot 2^k$).

Finally, we assume a strongly OT-EUF-CMA secure signature scheme $\mathsf{Sig} = (\mathsf{SGen}, \mathsf{Sig}, \mathsf{Ver})$ with k-bit verification keys, and a key-unique IND-CPA secure symmetric encryption scheme (E, D) (see Section 2) with k-bit symmetric keys K and message space $\{0,1\}^*$.

Now consider the following PKE scheme PKE:

Public parameters. $\mathsf{Pars}(1^k)$ first runs $(N, P, Q) \leftarrow \mathsf{GenN}(1^k)$. Recall that this fixes the groups $\mathbb{G}_{\mathsf{rnd}}$ and $\mathbb{G}_{\mathsf{msg}}$. Then, Pars selects two generators g_1, g_2 of $\mathbb{G}_{\mathsf{rnd}}$. Finally, Pars runs $(Fpk, Ftd) \leftarrow \mathsf{FGen}(1^k)$ and outputs $pp = (N, g_1, g_2, Fpk)$. In the following, we denote with p the LAF modulus contained in Fpk.

Key generation. $\mathsf{Gen}(pp)$ uniformly selects four messages $s_j = (a_j, b_j, c_j) \in \mathcal{M}$ (for $1 \leq j \leq 4$) as secret key, and sets $pk := \left(u := g_1^{[s_1]_{\mathbb{Z}}} g_2^{[s_2]_{\mathbb{Z}}}, v := g_1^{[s_3]_{\mathbb{Z}}} g_2^{[s_4]_{\mathbb{Z}}}\right)$ and $sk := (s_j)_{j=1}^4$.

Encryption. $\mathsf{Enc}(pp, pk, M)$, for pp and pk as above, and $M \in \mathcal{M}$, uniformly selects exponents $r, \widetilde{r} \leftarrow \mathbb{Z}_{N/4}$, a random filter core tag t_{c}, a Sig-keypair $(vk, sigk) \leftarrow \mathsf{SGen}(1^k)$, and a random symmetric key $K \in \{0,1\}^k$ for (E, D), and computes

$$\begin{aligned} (G_1, G_2) &:= (g_1^r, g_2^r) & C_{\mathsf{E}} &\leftarrow \mathsf{E}(K, \mathsf{LAF}_{Fpk,t}([M]_{\mathbb{Z}^{\mathfrak{n}}_p})) \\ (\widetilde{G}_1, \widetilde{G}_2) &:= (g_1^{\widetilde{r}}, g_2^{\widetilde{r}}) & \sigma &\leftarrow \mathsf{Sig}(sigk, ((G_j, \widetilde{G}_j)_{j=1}^2, Z, \widetilde{Z}, C_{\mathsf{E}}, t_{\mathsf{c}})) \\ Z &:= (u^{vk} v)^{r \cdot N^2} & C &:= ((G_j, \widetilde{G}_j)_{j=1}^2, Z, \widetilde{Z}, C_{\mathsf{E}}, t_{\mathsf{c}}, vk, \sigma) \\ \widetilde{Z} &:= (u^{vk} v)^r u^{\widetilde{r}} h^{K + 2^k \cdot [M]_{\mathbb{Z}}} \end{aligned}$$

for the auxiliary tag $t_{\mathsf{a}} := vk$, and the resulting filter tag $t := (t_{\mathsf{c}}, t_{\mathsf{a}})$.

Decryption. $\mathsf{Dec}(pp, sk, C)$, for pp, sk and C as above, first checks the signature σ and rejects with $\perp$ if $\mathsf{Ver}(vk, ((G_j, \widetilde{G}_j)_{j=1}^2, Z, \widetilde{Z}, C_{\mathsf{E}}, t_{\mathsf{c}}), \sigma) = 0$, or if

$$Z \neq \left(G_1^{[s_1]_{\mathbb{Z}} \cdot vk + [s_3]_{\mathbb{Z}}} G_2^{[s_2]_{\mathbb{Z}} \cdot vk + [s_4]_{\mathbb{Z}}}\right)^{N^2}.$$

Then Dec computes

$$Z' := G_1^{[s_1]_{\mathbb{Z}} \cdot vk + [s_3]_{\mathbb{Z}}} G_2^{[s_2]_{\mathbb{Z}} \cdot vk + [s_4]_{\mathbb{Z}}} \widetilde{G}_1^{[s_1]_{\mathbb{Z}}} \widetilde{G}_2^{[s_2]_{\mathbb{Z}}}$$

and then $K \in \{0,1\}^k, M \in \mathcal{M}$ with $K + 2^k \cdot [M]_{\mathbb{Z}} := \mathrm{dlog}_h(\widetilde{Z}/Z')$. If $\widetilde{Z}/Z' \notin \mathbb{G}_{\mathsf{msg}}$, or no such M exists, or $\mathsf{D}(K, C_{\mathsf{E}}) \neq \mathsf{LAF}_{Fpk,t}([M]_{\mathbb{Z}_p^{\mathfrak{n}}})$ (for $t = (t_{\mathsf{c}}, t_{\mathsf{a}})$ computed as during encryption), then Dec rejects with $\perp$. Else, Dec outputs M.

Secret keys as messages. Our scheme has secret keys $s = (s_j)_{j=1}^4 \in \mathcal{M}^4$; hence, we can only encrypt one quarter s_j of a secret key at a time. In the security proof below, we will thus only consider KDM queries that ask to encrypt a specific secret key *part*. Alternatively, we can change our scheme, so that 4-tuples of $\mathcal{M}$-elements are encrypted. To avoid malleability (which would destroy CCA security), we of course have to use only one LAF tag for this. Our CIRC-CCA proof below applies to such a changed scheme with minor syntactic changes.

Efficiency. When instantiated with our DLIN-based LAF construction from Section 3, and taking $\mathfrak{n} = 6$ as above, our scheme has ciphertexts with 43 $\mathbb{G}$-elements, 6 $\mathbb{Z}_{N^3}$-elements, plus chameleon hash randomness, a one-time signature and verification key, and a symmetric ciphertext (whose size could be in the range of one $\mathbb{Z}_{N^2}$-element plus some encryption randomness). The number of group elements in the ciphertext is constant, and does not grow in the security parameter. The public parameters contain $\mathbf{O}(k)$ group elements[8] (most of them from $\mathbb{G}$), and public keys contain two $\mathbb{Z}_{N^3}$-elements; secret keys consist of four $\mathbb{Z}_{N^2}$-elements. While these parameters are not competitive with current non-KDM-secure schemes, they are significantly better than those from the circular-secure scheme of Camenisch et al. [14].[9]

Security proof (single-user user). It is instructive to first treat the single-user case. Here, we essentially only require that PKE is IND-CCA secure, even if encryptions of its secret key are made public. For the multi-user case (see [21]), we can then proceed like [10, 11] and re-randomize keys and ciphertexts of a single PKE instance. This enables an analysis analogous to the single-user case.

Theorem 2. *Assume the DCR assumption holds in $\mathbb{Z}_{N^3}$, the DDH assumption holds in $\mathbb{G}_{\mathsf{rnd}}$, LAF is an LAF, Sig is a strongly OT-EUF-CMA secure signature scheme, H is collision-resistant, and (E, D) is a key-unique IND-CPA secure SKE scheme. Then PKE is 1-CIRC-CCA-secure.*

[8] Using the generalized LAF mentioned at the end of Section 3.1, public parameters with $\mathbf{O}(1)$ group elements are possible, at the cost of a (constant) number of extra group elements per tag.

[9] For instance, Section 7 of the full version of [14] implies that their scheme has a public key, resp. ciphertext of about 500, resp. 1000 $\mathbb{G}$-elements (for $\log_2(|\mathbb{G}|) = 160$).

Proof. Assume a PPT adversary A on PKE's 1-CIRC-CCA security. Say that A always makes $q = q(k)$ KDM queries. We proceed in games. Let out_i denote the output of Game i.

Game 1 is the 1-KDM-CCA experiment with PKE and A. By definition, $\Pr[out_1 = 1] - 1/2 = \mathsf{Adv}^{\mathsf{kdm\text{-}cca}}_{\mathsf{PKE},A}(k)$.

In **Game** 2, we modify the way KDM queries are answered. Namely, in each ciphertext prepared for A, we set up Z and $\widetilde{Z}$ up as

$$\begin{aligned} Z &:= \left(G_1^{[s_1]_{\mathbb{Z}} \cdot vk + [s_3]_{\mathbb{Z}}} G_2^{[s_2]_{\mathbb{Z}} \cdot vk + [s_4]_{\mathbb{Z}}}\right)^{N^2} \\ \widetilde{Z} &:= G_1^{[s_1]_{\mathbb{Z}} \cdot vk + [s_3]_{\mathbb{Z}}} G_2^{[s_2]_{\mathbb{Z}} \cdot vk + [s_4]_{\mathbb{Z}}} \widetilde{G}_1^{[s_1]_{\mathbb{Z}}} \widetilde{G}_2^{[s_2]_{\mathbb{Z}}} \cdot h^{K + 2^k \cdot [M]_{\mathbb{Z}}}. \end{aligned} \tag{6}$$

for the already prepared $(G_j, \widetilde{G}_j) = (g_j^r, g_j^{\widetilde{r}})$. This change is only conceptual by our setup of u, v, so $\Pr[out_2 = 1] = \Pr[out_1 = 1]$.

In **Game** 3, we again change how KDM ciphertexts are prepared. Intuitively, our goal is now to prepare the G_j and $\widetilde{G}_j$ with additional $\mathbb{G}_{\mathsf{msg}}$-components, such that $\widetilde{Z}$, as computed in (6), is of the form $g \cdot h^K$ for some $g \in \mathbb{G}_{\mathsf{rnd}}$. (That is, we want the $\mathbb{G}_{\mathsf{msg}}$-components of the $G_j, \widetilde{G}_j$ to cancel out the $h^{2^k \cdot [M]_{\mathbb{Z}}}$ term in (6).) To do so, we prepare $G_j = g_j^r / h^{\alpha_j \cdot 2^k}$ and $\widetilde{G}_j = g_j^r / h^{\widetilde{\alpha}_j \cdot 2^k}$ for $j \in \{1, 2\}$ and suitable $\alpha_j, \widetilde{\alpha}_j$ to be determined. $\widetilde{Z}$ is still computed as in (6), so

$$\widetilde{Z} = g \cdot h^{K + 2^k \cdot [M]_{\mathbb{Z}} - 2^k(\alpha_1([s_1]_{\mathbb{Z}} \cdot vk + [s_3]_{\mathbb{Z}}) + \alpha_2([s_2]_{\mathbb{Z}} \cdot vk + [s_4]_{\mathbb{Z}}) + \widetilde{\alpha}_1[s_1]_{\mathbb{Z}} + \widetilde{\alpha}_2[s_2]_{\mathbb{Z}})}$$

for $g = g_1^{r \cdot ([s_1]_{\mathbb{Z}} \cdot vk + [s_3]_{\mathbb{Z}}) + \widetilde{r}[s_1]_{\mathbb{Z}}} g_2^{r \cdot ([s_2]_{\mathbb{Z}} \cdot vk + [s_4]_{\mathbb{Z}}) + \widetilde{r}[s_2]_{\mathbb{Z}}} = \left(u^{vk} v\right)^r u^{\widetilde{r}} \in \mathbb{G}_{\mathsf{rnd}}$. So to prepare a KDM encryption of s_{j^*} with $\widetilde{Z} = g \cdot h^K$, we can set $(\alpha_1, \alpha_2, \widetilde{\alpha}_1, \widetilde{\alpha}_2)$ to $(0, 0, 1, 0)$ for $j^* = 1$, to $(1, 0, -vk, 0)$ for $j^* = 2$, to $(0, 0, 1, 0)$ for $j^* = 3$, and to $(0, 1, 0, -vk)$ for $j^* = 4$. (vk can be chosen independently in advance.) The remaining parts of C are prepared as in Game 2. We claim

$$\Pr[out_3 = 1] - \Pr[out_2 = 1] \leq 4 \cdot \mathsf{Adv}^{\mathsf{dcr}}_{\mathbb{Z}^*_{N^3}, B}(k) + \mathbf{O}(2^{-k}) \tag{7}$$

for a suitable DCR distinguisher B that simulates Game 2, resp. Game 3. Concretely, B gets as input a value $\widetilde{W} \in \mathbb{Z}^*_{N^3}$ of the form $\widetilde{W} = \tilde{g}^{N^2} \cdot h^b$ for $b \in \{0, 1\}$. Note that if we set $W := \widetilde{W}^{-2^k}$, we have $W = g^{\widehat{r}} / h^{b \cdot 2^k} \in \mathbb{Z}^*_{N^3}$, with uniform $g^{\widehat{r}} \in \mathbb{G}_{\mathsf{rnd}}$. First, B guesses a value of $j^* \in [4]$. (This gives a very small hybrid argument, in which in the j^*-th step, only encryptions of s_{j^*} are changed.) We only detail B's behavior for the case $j^* = 3$; the other cases are easier or analogous. First, B sets up $g_1 := W^{N^2}$ and $g_2 := W^{\gamma N^4}$ for uniform $\gamma \in \mathbb{Z}_{N/4}$. To prepare an encryption of s_3, B chooses uniform $\rho, \widetilde{\rho} \in \mathbb{Z}_{N^2/4}$ and sets

$$\begin{aligned} G_1 &:= W^{\rho \cdot (\rho^{-1})} & G_2 &:= W^{\gamma \cdot \rho \cdot (\rho^{-1}) \cdot N^2} \\ \widetilde{G}_1 &:= W^{vk \cdot \widetilde{\rho} \cdot (\widetilde{\rho}^{-1})} & \widetilde{G}_2 &:= W^{\gamma \cdot vk \cdot \widetilde{\rho} \cdot (\widetilde{\rho}^{-1}) \cdot N^2}, \end{aligned}$$

where the values $\rho^{-1}, \widetilde{\rho}^{-1}$ are computed modulo N^2. This implicitly sets $r = \rho \cdot (\rho^{-1})/N^2 \bmod |\mathbb{G}_{\mathsf{rnd}}|$ and $\widetilde{r} = vk \cdot \widetilde{\rho} \cdot (\widetilde{\rho}^{-1})/N^2 \bmod |\mathbb{G}_{\mathsf{rnd}}|$, both of which are statistically close to uniform. Furthermore, $G_j = g_j^r / h^{b \cdot \alpha_j \cdot 2^k}$ and $\widetilde{G}_j = g_j^{\widetilde{r}} / h^{b \cdot \widetilde{\alpha_j} \cdot 2^k}$;

so, depending on B's challenge, encryptions of s_3 are prepared as in Game 2 or Game 3. Similar arguments work for $j^* = 1, 2, 4$, and (7) follows. (The $\mathbf{O}(2^{-k})$ term in (7) accounts for the statistical defect caused by choosing $\mathbb{G}_{\mathsf{rnd}}$-exponents from $\mathbb{Z}_{N/4}$, resp. $\mathbb{Z}_{N^2/4}$.)

Using the definition of u and v, our change in Game 3 implies $\widetilde{Z} = (u^{vk}v)^r \cdot u^{\widetilde{r}} \cdot h^K$ when a key part s_j is to be encrypted. (However, note that we still have $Z = (u^{vk}v)^{r \cdot N^2}$ in any case.) This means that A still obtains information about the s_j (beyond what is public from pk) from its KDM queries, but this information is limited to values $\mathsf{LAF}_{Fpk,t}([s_j]_{\mathbb{Z}_p^n})$. We will now further cap this leaked information by making $\mathsf{LAF}_{Fpk,t}(\cdot)$ lossy. Namely, in **Game** 4, we use the LAF trapdoor Ftd initially sampled along with Fpk. Concretely, when preparing a ciphertext C for A, we sample t_{c} using $t_{\mathsf{c}} \leftarrow \mathsf{FTag}(Ftd, t_{\mathsf{a}})$ for the corresponding auxiliary tag $t_{\mathsf{a}} = vk$. A simple reduction shows

$$\Pr[out_4 = 1] - \Pr[out_3 = 1] \;=\; \mathsf{Adv}^{\mathsf{ind}}_{\mathsf{LAF},C_2}(k)$$

for a suitable adversary C_2 on LAF's indistinguishability.

In **Game** 5, we reject all decryption queries of A that re-use a verification key vk from one of the KDM ciphertexts. To show that this change does not significantly affect A's view, assume a decryption query C that re-uses a key $vk = vk^*$ from a KDM ciphertext C^*. Recall that C contains a signature σ of $X := ((G_j, \widetilde{G}_j)_{j=1}^2, Z, \widetilde{Z}, C_{\mathsf{E}}, t_{\mathsf{c}})$ under an honestly generated Sig-verification-key $vk = t_{\mathsf{a}} = t^*_{\mathsf{a}} = vk^*$. Since we assumed $t = (t_{\mathsf{c}}, t_{\mathsf{a}}) = (t^*_{\mathsf{c}}, t^*_{\mathsf{a}}) = t^*$, and A is not allowed to query unchanged challenge ciphertexts for decryption, we must have $(X, \sigma) \neq (X^*, \sigma^*)$ for the corresponding message X^* and signature σ^* from C^*. Hence, Game 4 and Game 5 only differ when A manages to forge a signature. A straightforward reduction to the strong OT-EUF-CMA security of Sig yields

$$\Pr[out_5 = 1] - \Pr[out_4 = 1] \;=\; q(k) \cdot \mathsf{Adv}^{\mathsf{seuf\text{-}cma}}_{\mathsf{LAF},F}(k)$$

for a forger F against Sig that makes at most one signature query.

In **Game** 6.i (for $0 \leq i \leq q$), the first i challenge ciphertexts are prepared using $Z = \widehat{g}^{N^2}$ and $\widetilde{Z} = \widehat{g} \cdot u^{\widetilde{r}} \cdot h^K$ (if a key component s_j is to be encrypted), resp. $\widetilde{Z} = \widehat{g} \cdot u^{\widetilde{r}} \cdot h^{K + 2^k[M]_{\mathbb{Z}}}$ (if a constant $M \in \mathcal{M}$ is to be encrypted) for an independently uniform $\widehat{g} \leftarrow \mathbb{G}_{\mathsf{rnd}}$ drawn freshly for each ciphertext. Obviously, Game 6.0 is identical to Game 5: $\Pr[out_{6.0} = 1] \;=\; \Pr[out_5 = 1]$. We will move from Game 6.i to Game 6.$(i+1)$ in several steps. During these steps, we denote with $C = ((G_j, \widetilde{G}_j)_{j=1}^2, Z, \widetilde{Z}, C_{\mathsf{E}}, t_{\mathsf{c}}, vk, \sigma)$ the $(i+1)$-st KDM ciphertext.

In **Game** 6.i.1, we change the $\mathbb{G}_{\mathsf{rnd}}$ parts of G_1, G_2 from a Diffie-Hellman tuple (with respect to g_1, g_2) to a random tuple. Concretely, if an s_j is to be encrypted, we set $(G_1, G_2) = (g_1^{r_1}/h^{\alpha_1 \cdot 2^k}, g_2^{r_2}/h^{\alpha_2 \cdot 2^k})$; if a constant M is encrypted, we set $(C_1, C_2) = (g_1^{r_1}, g_2^{r_2})$, in both cases for independently uniform $r_1, r_2 \leftarrow \mathbb{Z}_{N/4}$. The $\mathbb{G}_{\mathsf{msg}}$ parts of G_1, G_2 are thus unchanged compared to Game 6.i. Note that the $\widetilde{G}_j$ are still prepared as $\widetilde{G}_j = g_j^{\widetilde{r}}/h^{\widetilde{\alpha} \cdot 2^k}$, resp. $\widetilde{G}_j = g_j^{\widetilde{r}}$. A

straightforward reduction to the DDH assumption in $\mathbb{G}_{\mathsf{rnd}}$ yields

$$\sum_{i=1}^{q(k)} \left(\Pr\left[out_{6.i} = 1\right] - \Pr\left[out_{6.i.1} = 1\right]\right) \;=\; q(k) \cdot \mathsf{Adv}^{\mathsf{ddh}}_{\mathbb{G}_{\mathsf{rnd}},D_1}(k) + \mathbf{O}(2^{-k})$$

for a suitable D_1. The $\mathbf{O}(2^{-k})$ error term accounts for the statistical difference caused by the choice of exponents $r_j \leftarrow \mathbb{Z}_{N/4}$, which induces an only almost-uniform distribution on group elements g^{r_j}. Note that at this point, Z and $\widetilde{Z}$ are still computed as in (6), even if an s_j is to be encrypted.

In **Game** $6.i.2$, we compute Z and $\widetilde{Z}$ as $Z = \widehat{g}^{N^2}$ and $\widetilde{Z} = \widehat{g} \cdot u^{\widetilde{r}} \cdot h^K$, resp. $\widetilde{Z} = \widehat{g} \cdot u^{\widetilde{r}} \cdot h^{K+2^k[M]_{\mathbb{Z}}}$ for a fresh $\widehat{g} \leftarrow \mathbb{G}_{\mathsf{rnd}}$. Thus, the difference to Game $6.i.1$ is that we substitute a $\mathbb{G}_{\mathsf{rnd}}$-element computed as $G_1^{[s_1]_{\mathbb{Z}} \cdot vk + [s_3]_{\mathbb{Z}}} G_2^{[s_2]_{\mathbb{Z}} \cdot vk + [s_4]_{\mathbb{Z}}}$ with a fresh random $\widehat{g}$. To show that this change affects A's view only negligibly, it suffices to show that A's statistical information about

$$\begin{aligned} X \;:=\; \mathrm{dlog}_g \left(G_1^{[s_1]_{\mathbb{Z}} \cdot vk + [s_3]_{\mathbb{Z}}} G_2^{[s_2]_{\mathbb{Z}} \cdot vk + [s_4]_{\mathbb{Z}}} \right) \\ = \; \gamma_1 r_1([s_1]_{\mathbb{Z}} \cdot vk + [s_3]_{\mathbb{Z}}) + \gamma_2 r_2([s_2]_{\mathbb{Z}} \cdot vk + [s_4]_{\mathbb{Z}}) \bmod |\mathbb{G}_{\mathsf{rnd}}| \end{aligned}$$

(for some arbitrary generator g of $\mathbb{G}_{\mathsf{rnd}}$ and $\gamma_j = \mathrm{dlog}_g(g_j)$) is negligible. This part will be rather delicate, since we will have to argue that both A's KDM queries and A's decryption queries yield (almost) no information about X.

First, observe that A gets the following information about the s_j:

- pk reveals (through u and v) precisely the two linear equations $\gamma_1[s_1]_{\mathbb{Z}} + \gamma_2[s_2]_{\mathbb{Z}} \bmod |\mathbb{G}_{\mathsf{rnd}}|$ and $\gamma_1[s_3]_{\mathbb{Z}} + \gamma_2[s_4]_{\mathbb{Z}} \bmod |\mathbb{G}_{\mathsf{rnd}}|$ about the s_j, where the γ_j are as above. For $r_1 \neq r_2$, these equations are linearly independent of the equation that defines X. Hence, for uniform r_1, r_2, X is (almost) independent of pk.
- By LAF's lossiness, KDM queries yield (through $C_{\mathsf{E}} = \mathsf{E}(K, \mathsf{LAF}_{Fpk,t}([s_j]_{\mathbb{Z}_p^{\mathtt{n}}}))$ in total at most one equation $\omega_1 a_j + \omega_2 b_j + \sum_{i=0}^{\mathtt{n}-2} \omega_{3+i} c_{j,i} \bmod p$ for each j, where $(a_j, b_j, c_{j,0}, \ldots, c_{j,\mathtt{n}-3}) := [s_j]_{\mathbb{Z}_p^{\mathtt{n}}}$, and the ω_i are the (fixed) coefficients from LAF's lossiness property. (Recall the encodings $[s_j]_{\mathbb{Z}}, [s_j]_{\mathbb{Z}_p^{\mathtt{n}}}$ of the $s_j = (a_j, b_j, c_j)$ from (5).) Hence, the $b_j \in \mathbb{Z}_{p \cdot 2^k}$ fully blind the information released about the $c_j \in \mathbb{Z}_{2^{k-2}N}$ through the KDM ciphertexts. Thus, KDM ciphertexts reveal no information about $c_j \bmod |\mathbb{G}_{\mathsf{rnd}}|$ and hence also about $[s_j]_{\mathbb{Z}} \bmod |\mathbb{G}_{\mathsf{rnd}}|$.

Consequently, even given pk and the KDM ciphertexts, X is statistically close to independently uniform. This already shows that our change from Game $6.i.2$ affects A's view only negligibly if A makes no decryption queries. It remains to show that decryption queries yield no additional information about the s_j.

To do so, let us say that a ciphertext C is *consistent* iff there exist $r, \widetilde{r}$ with $(G_j, \widetilde{G}_j) = (g_j^r, g_j^{\widetilde{r}})$ for both $j \in \{1, 2\}$. Note that the decryption of a consistent ciphertext yields no information about the s_j beyond pk. (pk and $r, \widetilde{r}$ determine the values Z, Z' computed during decryption; everything else follows from Z' and C.) So it suffices to prove the following lemma (which we do in [21]):

Lemma 3. *In the situation of Game* $6.i.\ell$ *(for* $\ell \in \{1,2\}$*), let* $\mathsf{bad}_{\mathsf{query.i}.\ell}$ *be the event that* A *places an inconsistent decryption query that is not rejected. Then*

$$\sum_{i=1}^{q(k)} \left(\Pr\left[\mathsf{bad}_{\mathsf{query.i.1}}\right] + \Pr\left[\mathsf{bad}_{\mathsf{query.i.2}}\right]\right) \;\leq\; 2 \cdot q(k) \cdot \mathsf{Adv}^{\mathsf{eva}}_{\mathsf{LAF},F}(k) + \mathbf{O}(2^{-3k}).$$

for a suitable evasiveness adversary F *on* LAF*.*

By our discussion above and Lemma 3, we obtain that

$$\sum_{i=1}^{q(k)} \left|\Pr\left[out_{6.i.2} = 1\right] - \Pr\left[out_{6.i.1} = 1\right]\right| \;\leq\; 2 \cdot q(k) \cdot \mathsf{Adv}^{\mathsf{eva}}_{\mathsf{LAF},F}(k) + \mathbf{O}(2^{-3k}).$$

In **Game** $6.i.3$, we reverse the change from Game $6.i.1$. Concretely, we prepare the G_j as $G_j = g_j^r / h^{\alpha_j \cdot 2^k}$, resp. $G_j = g_j^r$ for a single $r \leftarrow \mathbb{Z}_{N/4}$. Another straightforward reduction to the DDH assumption in $\mathbb{G}_{\mathsf{rnd}}$ yields that

$$\sum_{i=1}^{q(k)} \left(\Pr\left[out_{6.i.3} = 1\right] - \Pr\left[out_{6.i.2} = 1\right]\right) \;=\; q(k) \cdot \mathsf{Adv}^{\mathsf{ddh}}_{\mathbb{G}_{\mathsf{rnd}},D_2}(k) + \mathbf{O}(2^{-k})$$

for a suitable D_2. To close the hybrid argument, note that Games $6.i.3$ and $6.(\mathrm{i}{+}1)$ are identical.

In **Game** 7, we clear the $\mathbb{G}_{\mathsf{msg}}$-component of $\widetilde{Z}$ in all ciphertexts prepared for A. That is, instead of computing $\widetilde{Z} = \widehat{g} \cdot u^{\tilde{r}} \cdot h^K$, resp. $\widetilde{Z} = \widehat{g} \cdot u^{\tilde{r}} \cdot h^{K+[M]_{\mathbb{Z}}}$ for a freshly uniform $\widehat{g} \leftarrow \mathbb{G}_{\mathsf{rnd}}$, we set $\widetilde{Z} = \widehat{g} \cdot u^{\tilde{r}}$. (We stress that we still compute $Z = \widehat{g}^{N^2}$.) Since all $\widetilde{Z}$ already have an independently uniform $\mathbb{G}_{\mathsf{rnd}}$-component, a straightforward reduction to the DCR assumption yields

$$\Pr\left[out_{6.q} = 1\right] - \Pr\left[out_7 = 1\right] \;=\; \mathsf{Adv}^{\mathsf{dcr}}_{\mathbb{Z}^*_{N^3},E}(k) + \mathbf{O}(2^{-k})$$

for a DCR distinguisher E. Note that because of the re-randomizability of DCR, there is no factor of $q(k)$, even though we substitute many group elements at once. However, since the precise order of $\mathbb{G}_{\mathsf{rnd}}$ is not known, this re-randomization costs us an error term of $\mathbf{O}(2^{-k})$.

In **Game** 8, we substitute the symmetric ciphertexts C_{E} in all KDM ciphertexts by encryptions of random messages. By our change in Game 7, we do not use the symmetric keys K used to produce C_{E} anywhere else. Thus, a reduction to the IND-CPA security of (E,D) gives

$$\Pr\left[out_7 = 1\right] - \Pr\left[out_8 = 1\right] \;=\; q(k) \cdot \mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{(\mathsf{E},\mathsf{D}),G}(k)$$

for an IND-CPA adversary G. Note that in Game 8, A's view is independent of the challenge bit b initially selected by the KDM challenger. Hence, we have $\Pr\left[out_8 = 1\right] \;=\; 1/2$. Taking things together yields the theorem.

Acknowledgements. I would like to thank Hoeteck Wee for pointing out to me the connection between hash proof systems and KDM security. His interpretation of the schemes of [10] and [11] was the starting point for this work. I am grateful to Tibor Jager for very useful feedback and discussions. I am indebted to Baodong Qin, who pointed out two flaws (along with very helpful suggestions on how to fix them). Finally, I am thankful to the Crypto 2012 referees who found an annoying cyclic dependency in the PKE construction, and to the Eurocrypt 2013 referees for helpful comments.

References

[1] Tolga Acar, Mira Belenkiy, Mihir Bellare, and David Cash. Cryptographic agility and its relation to circular encryption. In *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 403–422. Springer, May 2010.

[2] Pedro Adão, Gergei Bana, Jonathan Herzog, and Andre Scedrov. Soundness of formal encryption in the presence of key-cycles. In Sabrina De Capitani di Vimercati, Paul F. Syverson, and Dieter Gollmann, editors, *ESORICS 2005*, volume 3679 of *LNCS*, pages 374–396. Springer, September 2005.

[3] Benny Applebaum. Key-dependent message security: Generic amplification and completeness. In *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 527–546. Springer, May 2011.

[4] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *CRYPTO 2009*, volume 5677 of *LNCS*, pages 595–618. Springer, August 2009.

[5] Boaz Barak, Iftach Haitner, Dennis Hofheinz, and Yuval Ishai. Bounded key-dependent message security. In *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 423–444. Springer, May 2010.

[6] Mihir Bellare, Dennis Hofheinz, and Scott Yilek. Possibility and impossibility results for encryption and commitment secure under selective opening. In *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 1–35. Springer, April 2009.

[7] John Black, Phillip Rogaway, and Thomas Shrimpton. Encryption-scheme security in the presence of key-dependent messages. In Kaisa Nyberg and Howard M. Heys, editors, *SAC 2002*, volume 2595 of *LNCS*, pages 62–75. Springer, August 2003.

[8] Dan Boneh, Emily Shen, and Brent Waters. Strongly unforgeable signatures based on computational Diffie-Hellman. In Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin, editors, *PKC 2006*, volume 3958 of *LNCS*, pages 229–240. Springer, April 2006.

[9] Dan Boneh, Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-ciphertext security from identity-based encryption. *SIAM Journal on Computing*, 36(5): 1301–1328, 2007.

[10] Dan Boneh, Shai Halevi, Michael Hamburg, and Rafail Ostrovsky. Circular-secure encryption from decision diffie-hellman. In *CRYPTO 2008*, volume 5157 of *LNCS*, pages 108–125. Springer, August 2008.

[11] Zvika Brakerski and Shafi Goldwasser. Circular and leakage resilient public-key encryption under subgroup indistinguishability - (or: Quadratic residuosity strikes back). In *CRYPTO 2010*, volume 6223 of *LNCS*, pages 1–20. Springer, August 2010.

[12] Zvika Brakerski, Shafi Goldwasser, and Yael Tauman Kalai. Black-box circular-secure encryption beyond affine functions. In *TCC 2011*, volume 6597 of *LNCS*, pages 201–218. Springer, March 2011.

[13] Jan Camenisch and Anna Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 93–118. Springer, May 2001.

[14] Jan Camenisch, Nishanth Chandran, and Victor Shoup. A public key encryption scheme secure against key dependent chosen plaintext and adaptive chosen ciphertext attacks. In *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 351–368. Springer, April 2009.

[15] David Cash, Matthew Green, and Susan Hohenberger. New definitions and separations for circular security. Cryptology ePrint Archive, Report 2010/144, 2010. `http://eprint.iacr.org/`.

[16] Ronald Cramer and Victor Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 45–64. Springer, April / May 2002.

[17] David Galindo, Javier Herranz, and Jorge L. Villar. Identity-based encryption with master key-dependent message security and leakage-resilience. In Sara Foresti, Moti Yung, and Fabio Martinelli, editors, *ESORICS 2012*, volume 7459 of *LNCS*, pages 627–642. Springer, September 2012.

[18] Jens Groth and Amit Sahai. Efficient non-interactive proof systems for bilinear groups. In *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 415–432. Springer, April 2008.

[19] Iftach Haitner and Thomas Holenstein. On the (im)possibility of key dependent encryption. In *TCC 2009*, volume 5444 of *LNCS*, pages 202–219. Springer, March 2009.

[20] Dennis Hofheinz. All-but-many lossy trapdoor functions. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 209–227. Springer, April 2012.

[21] Dennis Hofheinz. Circular chosen-ciphertext security with compact ciphertexts. Cryptology ePrint Archive, Report 2012/150, 2012. `http://eprint.iacr.org/`.

[22] Dennis Hofheinz and Eike Kiltz. Secure hybrid encryption from weakened key encapsulation. In *CRYPTO 2007*, volume 4622 of *LNCS*, pages 553–571. Springer, August 2007.

[23] Dennis Hofheinz and Dominique Unruh. Towards key-dependent message security in the standard model. In *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 108–126. Springer, April 2008.

[24] Kaoru Kurosawa and Yvo Desmedt. A new paradigm of hybrid encryption scheme. In *CRYPTO 2004*, volume 3152 of *LNCS*, pages 426–442. Springer, August 2004.

[25] Tal Malkin, Isamu Teranishi, and Moti Yung. Efficient circuit-size independent public key encryption with KDM security. In *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 507–526. Springer, May 2011.

[26] Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *22nd ACM STOC*. ACM Press, May 1990.

[27] Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 187–196. ACM Press, May 2008.

[28] Brent R. Waters. Efficient identity-based encryption without random oracles. In *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 114–127. Springer, May 2005.