Regularly Lossy Functions and Applications

Yu Chen[1,2], Baodong Qin[3], and Haiyang Xue[1(✉)]

[1] State Key Laboratory of Information Security,
Institute of Information Engineering, Chinese Academy of Sciences,
Beijing, China
xuehaiyang@iie.ac.cn

[2] School of Cyber Security, University of Chinese Academy of Sciences,
Beijing, China

[3] National Engineering Laboratory for Wireless Security,
Xi'an University of Posts and Telecommunications, Xi'an, China

Abstract. In STOC 2008, Peikert and Waters introduced a powerful primitive called *lossy trapdoor functions* (LTFs). In a nutshell, LTFs are functions that behave in one of two modes. In the normal mode, functions are injective and invertible with a trapdoor. In the lossy mode, functions statistically lose information about their inputs. Moreover, the two modes are computationally indistinguishable. In this work, we put forward a relaxation of LTFs, namely, *regularly lossy functions* (RLFs). Compared to LTFs, the functions in the normal mode are not required to be efficiently invertible or even unnecessary to be injective. Instead, they could also be lossy, but in a regular manner. We also put forward richer abstractions of RLFs, namely *all-but-one regularly lossy functions* (ABO-RLFs).

We show that (ABO)-RLFs admit efficient constructions from both a variety of number-theoretic assumptions and hash proof system (HPS) for subset membership problems satisfying natural algebraic properties. Thanks to the relaxations on functionality, the constructions enjoy shorter key size and better computational efficiency than that of (ABO)-LTFs.

We demonstrate the applications of (ABO)-RLFs in leakage-resilient cryptography.

- As a special case of RLFs, lossy functions imply leakage-resilient injective one-way functions with optimal leakage rate $1 - o(1)$.
- ABO-RLFs immediately imply leakage-resilient message authentication code (MAC) with optimal leakage rate $1 - o(1)$, though in a weak sense.
- ABO-RLFs together with HPS give rise to leakage-resilient chosen-ciphertext (CCA) secure key encapsulation mechanisms (KEM) (this approach extends naturally to the identity-based setting). Combining the construction of ABO-RLFs from HPS, this gives the first leakage-resilient CCA-secure public-key encryption (PKE) with optimal leakage rate based solely on HPS, and thus goes beyond the barrier posed by Dodis et al. (Asiacrypt 2010).

N. P. Smart (Ed.): CT-RSA 2018, LNCS 10808, pp. 491–511, 2018.
https://doi.org/10.1007/978-3-319-76953-0_26

1 Introduction

In STOC 2008, Peikert and Waters [PW08] introduced a powerful primitive called lossy trapdoor function (LTF). Informally, LTF is a collection of functions $\mathcal{F} = \{f_{ek}\}$ whose evaluation key (i.e., function index or code) is created in one of two modes. One is injective (i.e., normal) mode: given a suitable trapdoor td for ek, the entire input x can be efficiently recovered from $f_{ek}(x)$. The other is lossy mode: f_{ek} statistically loses a significant amount of information about its input. Moreover, the two modes are computationally indistinguishable: given just ek, no efficient adversary can tell whether f_{ek} is injective or lossy. They also introduced a richer abstraction called all-but-one lossy trapdoor functions (ABO-LTFs). A collection of ABO-LTFs is associated with a set B called branches. The key generation algorithm takes a given branch $b^* \in B$ as an extra parameter, and outputs an evaluation key ek and a trapdoor td. The function $f_{ek,b}(\cdot)$ is injective and invertible with td for any branch $b \neq b^*$, while the function $f_{ek,b^*}(\cdot)$ is lossy. Moreover, the lossy branch b^* is computationally hidden by ek.

Using LTFs and ABO-LTFs, Peikert and Waters [PW08] develop new approaches for constructing several important cryptographic tools, such as injective TDFs, collision-resistant hash functions (CRHFs), oblivious transfer and CCA-secure PKE.

1.1 Related Work

Since the initial work of [PW08], there has been much additional work on LTFs and related concepts.

One direction of research is to find additional realizations of LTFs. Boyen and Waters [BW10] gave a technique to shrink the public key of matrix construction of [PW08] with the help of pairing. Rosen and Segev [RS09] and Boldyreva et al. [BFO08] independently described simple, compact constructions of LTFs and ABO-LTFs under the decisional composite residuosity (DCR) assumption. Freeman et al. [FGK+13] provided more constructions of LTFs from the quadratic residuosity (QR) and d-linear assumptions. Kiltz et al. [KOS17] and Xue et al. [XLL+13] gave constructions of LTFs based on factoring assumptions. Hemenway and Ostrovsky [HO12] gave a construction of LTFs based on the extended decisional Diffie-Hellman (eDDH) assumption, which generalizes the DDH, QR and DCR assumption. They also showed a generic construction of LTFs from homomorphic smooth HPS. Wee [Wee12] presented an alternative generic construction of LTFs from dual HPS.

Another direction of research is to explore variations and more applications. Rosen and Segev [RS09] and Kiltz et al. [KMO10] showed LTFs imply correlated-product TDFs and adaptive TDFs respectively. Boldyreva et al. [BFO08] constructed CCA-secure deterministic encryption based on LTFs and ABO-LTFs. Hemenway et al. [HLOV11] generalized ABO-LTFs to all-but-N lossy trapdoor functions (ABN-LTFs) that have N lossy branches. Hofheinz [Hof12] further generalized ABN-LTFs to all-but-many (ABM) LTFs in which the number of lossy branches is not bounded by any polynomial. Recently, Boyen and Li [BL17] realized ABM LTFs based on the learning with errors assumptions. So far, ABM-LTFs have shown their usefulness in constructing PKE with strong security properties including selective opening security [Hof12] and key-dependent message security [Hof13]. Mol and Yilek [MY10] constructed a CCA-secure PKE from any slightly lossy trapdoor functions that lose only a noticeable fraction of a bit. On the contrary, Zhandry [Zha16] introduced extremely lossy functions (whose functions in the lossy mode only have polynomial-sized image), and demonstrated extremely lossiness is useful for instantiating random oracles in several settings.

1.2 Motivations

Due to the strong requirement for the normal mode (injective and efficiently invertible with trapdoor), the concrete constructions of (ABO)-LTFs are typically not efficient in terms of the size of evaluation key and complexity of evaluation. The generic constructions of (ABO)-LTFs require advanced property for the basing primitives, such as homomorphic and invertible properties.

In all the known applications of LTFs, the normal mode is used to fulfill functionality, while the lossy mode is used to establish security. However, in many scenarios we do not require the full power of LTFs. As observed by Peikert and Waters [PW08, Sect. 3.4], some applications (such as injective OWFs, CRHFs) *do not require a trapdoor*, but only indistinguishability between normal mode and lossy mode. Thereby, they conjectured “realizing the weaker notion of lossy (non-trapdoor) functions (LFs) could be achieved more simply or efficiently than the full notion of LTFs”, and left the investigation of this question as an interesting problem.

A central goal in cryptography is to base cryptosystems on primitives that are as week as possible. With the question raised by Peikert and Waters [PW08] in mind, we ask the following questions:

> *How to realize LFs efficiently? Are there any other applications of LFs? Can we further weaken the notion of LFs while still being useful?*

1.3 Our Contributions

We answer the above questions affirmatively. An overview of our contributions is as below.

1.4 Regularly Lossy Functions and Extensions

As discussed above, when building cryptographic protocols the normal mode of LTF is used to fulfill functionality. For some applications that invertible property for the normal mode is overkilled, the injective property may also be unnecessary. This suggests that we may further relax the notion of LFs.

We introduce a new primitive called regularly lossy functions (RLFs), which is a public function f_{ek} (the evaluation key ek serves as the function index) that is created to behave in one of two modes. In the normal mode, the function f_{ek} could be lossy, but should lose *regularly* (we will formally define this later). The intuition is that when the input x has high min-entropy, so does $f_{ek}(x)$. In the lossy mode, the function f_{ek} statistically loses a significant amount information about its input x, i.e., the average min-entropy of $x|f_{ek}(x)$ is high. Finally, the two modes are indistinguishable: no efficient adversary can tell whether f_{ek} is in normal mode or lossy mode.

In line of the above intuition, we can use image size to capture the lossy mode same as LTFs [PW08], but not for the normal mode. This is because image size is a *global* characterization for a function, which suffices to give the lower bound of the average min-entropy of $x|f_{ek}(x)$ by applying the chain rule for min-entropy [DORS08], but is insufficient to give the lower bound of the min-entropy of $f_{ek}(x)$. For instance, when the function is highly unstructured, it is possible that the image size of f_{ek} is slightly smaller the domain size, but the min-entropy $f_{ek}(x)$ is much smaller than that of x. To address this subtle issue, we choose a *local* characterization of function named regularity to capture the normal mode. In the normal mode, the function f_{ek} is ν-regular, i.e., each image has at most ν preimages under f_{ek}. With this requirement, the (average) min-entropy of $f(x)$ decreases at most $\log \nu$ compared to that of x (by applying Lemma 1 we develop in Sect. 2.2).

Clearly, our notion of RLFs differs from LFs only at the normal mode, whose functions are not required to be injective but could be flexibly lossy from injective to significantly lossy, subjected to the parameter choices of concrete applications. The only constraint is they should lose in a *regular* way.

To admit more applications, we introduce a richer abstraction called ABO-RLFs, analogously to the extension of LTFs to ABO-LTFs. Briefly, an ABO collection is associated with a branch set B. The generation algorithm of ABO-RLF takes an extra parameter $b^* \in B$, and outputs an evaluation key such that $f_{ek,b}$ is regular for any branch $b \neq b^*$ but is lossy when $b = b^*$. Moreover, the lossy branch is (computationally) hidden by ek.

1.5 Efficient Constructions of ABO-RLFs

Existing constructions of (ABO)-LTFs are less efficient due to their strong requirement for the normal mode. In contrast, RLFs require nothing but the intrinsic regularity of functions for the normal mode. Such weakening admits much more efficient constructions from both number-theoretic assumptions and HPS.

First, we mainly follow the matrix approach due to [PW08] to give a DDH-based ABO-RLFs, in which the evaluation key is specified by an $n \times m$ matrix over groups. The efficiency improvements of our construction comes from two aspects: (1) since we do not require efficiently inversion, the input x can be treated as an n-dimensional vector of elements from some large field (say $\mathbb{Z}_p$)

rather than a binary string over $\{0,1\}^n$; (2) since we even do not require injectivity, m could be set smaller than n and thus the matrix size shrinks noticeably. Our DDH-based ABP-RLFs can be naturally extended to base on the eDDH assumption.

As to generic constructions, we first give a construction of ABO-RLF from any HPS for subset membership problems (SMPs). The construction proceeds via two steps: (1) build LF from any HPS following the approach of building LTF from dual HPS [Wee12]; (2) amplify the obtained RLF to ABO-RLF with branch set $\{0,1\}^\ell$. However, this construction is inefficient in that its second step invokes ℓ individual copies of RLF and involves some degradation in lossiness. Towards a direct and efficient construction, we require the SMPs to satisfy natural algebra properties, namely L is a subgroup of X and the quotient group $H = X/L$ is a cyclic group of order p. By exploiting this properties, we manage to give an efficient ABO-RLF with branch set $B = \mathbb{Z}_p$ directly from HPS.

1.6 Applications in Leakage-Resilient Cryptography

On the surface, non-injective function without a trapdoor do not appear pretty useful, since many appealing applications of standard LTF require a trapdoor (e.g., public-key encryption) or at least injectivity (e.g., CRHFs) for the normal mode. Indeed, RLF does not suffice for most of the applications outlined above. Nevertheless, we show that this simple notion on its own or in conjunction with other tools can in fact quite useful in leakage-resilient cryptography.

Traditional security models assume complete privacy of secret keys. However, in real systems the adversary might learn partial information about secret keys by launching various "key leakage attacks" via side channels, which make this idealized assumption false in practice. This fact lead to the design of leakage-resilient cryptography, which spreads to stream ciphers, block ciphers, digital signatures, public-key encryption, identity-based encryption.

There are several models of key leakage-resilience in the literature, mainly differing in their specifications of what and how many information can be leaked to the adversary. In this work we will focus on a simple yet general model, called bounded-leakage model. In this model, the adversary can learn arbitrary information about the secret key, subjected to the restriction that the total number of leakage is bounded by some leakage bound $\ell(\lambda)$, where λ is the security parameter. The leakage rate is defined as the ratio of $\ell(\lambda)$ to the secret key size $s(\lambda)$, i.e., $\ell(\lambda)/s(\lambda)$. Clearly, $1 - o(1)$ is the optimal leakage rate in the bounded leakage model.

In this work, we demonstrate the utility of RLFs (including their special case – LFs) by exploring their applications in leakage-resilient cryptography.

Leakage-Resilient OWFs. A function is said to be ℓ-leakage-resilient one-way if one-wayness maintains even the attacker may obtain at most ℓ-bits leakage about the preimage.

It was shown in [ADW09b, DHLW10, Kom16] (and implicitly in [ADW09a, KV09]) that any weak universal one-way hash function (UOWHF)[1] from $\{0,1\}^n$ to $\{0,1\}^m$ automatically provides ℓ-leakage-resilient one-wayness, where $\ell \leq n - m - \omega(\lambda)$. The shortcoming of this construction is the resulting LR OWFs are inherently compressing, and the leakage bound is dependent on the image size. As a consequence, in some applications one has to make a trade-off between image size and leakage bound.

In this work, we give an alternative construction based on LF. The insight is that the implication of LF $\Rightarrow$ injective OWF [PW08] also holds in the leakage setting. More precisely, we show that the functions in the injective mode of LFs make up a collection of ℓ-leakage-resilient injective OWFs. The leakage bound is $\ell \leq n - \tau - \omega(\lambda)$, where n is the length of inputs and τ is the logarithm of image size for the lossy mode. Both of our construction based of LF and the construction based on UOWHF achieves optimal leakage rate with appropriate parameter choice. The advantage of our construction is that the leakage bound is independent of the image size[2], which is more applicable in practice. To the best of our knowledge, our construction appears to be the first leakage-resilient injective OWF with optimal leakage rate.

Leakage-Resilient MAC. Hazay et al. [HLAWW13] constructed a leakage-resilient MAC from standard PRF. Though their construction only requires minimum assumption (OWFs), the leakage rate $\log \lambda / s(\lambda)$ is poor. Constructing leakage-resilient MAC under general assmption with higher leakage rate was left as an open problem [HLAWW13].

In this work, we make a progress towards this problem. We construct a leakage-resilient MAC with optimal leakage rate from ABO-RLFs, though in a weaker sense. To convert a ABO-RLF to a MAC, the key generation algorithm generates an evaluation key ek as public parameter, then chooses a random x from input space as the secret key; the tag algorithm treats message m as branch and evaluate $t \leftarrow f_{ek,m}(x)$; the verification algorithm is canonical, namely recomputes the tag and checks for equality.

The resulting MAC turns out to be leakage-resilient strongly unforgeable, though in a weaker sense: the attacker only makes one tagging query and declares the query at the very beginning. The security argument leverages on the power of *lose information*. Upon the attacker submitting its target query m^*, the reduction generates ek with m^* as the lossy branch and returns $t^* \leftarrow f_{ek,m^*}(x)$. Observe that f_{ek,m^*} is a lossy function, thus the secret key x still retains sufficient min-entropy even after revealing t^* and bounded leakage. For any forge (m, t), we must have $m \neq m^*$ since the MAC is unique. Besides, $f_{ek,m}$ is a ν-regular function whenever $m \neq m^*$. In this case, the (average) min-entropy of $t = f_{ek,m}(x)$ decreases at most $\log \nu$ compared to that of x. Therefore, t is unpredictable. The leakage rate could achieve $1 - o(1)$ under proper parameter choice.

[1] This is sometimes called second preimage resistant functions.

[2] The leakage bound only subjects to the image size of functions in the lossy mode, which will not be used in real construction.

Leakage-Resilient PKE. A PKE is said to be ℓ-leakage-resilient if semantic security maintains even if the attacker can obtain at most ℓ-bits leakage about the secret key.

Akavia et al. [AGV09] first formalized the notion of leakage-resilient chosen-plaintext security (LR CPA) in the bounded-leakage model. Since then, many existing PKE schemes [Reg05, GPV08, BHHO08] have been proved secure in the bounded-leakage model. Later Naor and Segev [NS09] generalized the main ideas behind these constructions to by giving a generic construction of LR CPA-secure PKE schemes from universal$_1$ hash proof system (HPS) [CS02]. Moreover, they also show how to achieve LR CCA security by either: (1) applying the Naor-Yung paradigm to obtain impractical PKE schemes with leakage-rate $1 - o(1)$ or (2) combining universal$_2$ HPS to obtain practical PKE schemes (variants of the Cramer-Shoup cryptosystems) with leakage-rate $1/6 - o(1)$. Later, Liu et al. [LWZ13] proposed a new variant of the Cramer-Shoup cryptosystems which is LR CCA-secure with leakage-rate $1/4 - o(1)$. Dodis et al. [DHLW10] realized that the HPS approach to building LR CCA-secure PKE seems to be inherently limited to leakage-rates below 1/2: because the secret-key consists of two components (sk_1 of universal$_1$ HPS for decrypting ciphertext and sk_2 of universal$_2$ HPS for verifying the well-formedness of the ciphertext) and the proofs break down if either of the components is individually leaked in its entirety.[3] Later, Qin and Liu [QL13, QL14] bypassed the bound by replacing the universal$_2$ HPS in the HPS approach [NS09] with a new primitive called one-time lossy filters (OT-LFs). By delicate instantiations of universal$_1$ HPS and OT-LF, they obtained LR CCA-secure PKE schemes with leakage rate $1 - o(1)$. However, if OT-LF is implied by HPS is unknown. The problem of whether we can build LR CCA-secure PKE with optimal leakage-rate based on solely HPS is still open.

In this work, we resolve this problem by building LR CCA-secure PKE with leakage rate $1 - o(1)$ based solely on HPS. This goes beyond previous believed bound conjectured by Dodis et al. [DHLW10]. Our starting point is the work of Qin and Liu [QL13]. It is well-known that key encapsulation mechanism (KEM) is more preferable than PKE from both theoretic and practice interest, thus we focus on the construction of leakage-resilient KEM.

Observe that in the setting of PKE the challenge ciphertext depends on attacker's choice of target messages, whereas in the setting of KEM the challenge ciphertext is entirely determined by the challenger in the setting of KEM. Such feature allows us to replace OT-LFs with all-but-one lossy functions (ABO-LFs), which saves at least a chameleon hash for the KEM construction.[4] Moreover, we

[3] Kiltz et al. [KPSY09] showed that CCA-secure PKE can be constructed from a universal$_2$ HPS with an authenticated one-time secure symmetric encryption, while universal$_2$ HPS can be generically obtained from universal$_1$ HPS via 4-wise independent hash function. At a first glance, their construction can be easily augmented to be leakage-resilient CCA-secure by applying randomness extractor to the projective hash. However, such augment could be very subtle in that the adding of a random seed may render the overall ciphertext easily malleable, and thus cannot be CCA-secure.

[4] As shown in [QL13], OT-LFs can be build from ABO-LFs and chameleon hash.

show that ABO-LFs can be relaxed to ABO-RLFs. As we show in Sect. 5, ABO-RLFs can be efficiently constructed from any HPS for subgroup membership problem with natural algebraic properties. Taken together, the secret key in our approach consists of just one component for verifying the well-formedness of the ciphertext and for decrypting it simultaneously. Therefore, the leakage rate of our construction can go beyond the limitation of 1/2, being subject to the leakage tolerance of the underlying universal$_1$ HPS. For instance, applying the DDH-based universal$_1$ HPS from [QL13], we obtain a LR CCA-secure KEM with leakage rate $1/2 - o(1)$; applying the universal$_1$ HPS from refined subgroup indistinguishability problem [QL14], we obtain a LR CCA-secure KEM with leakage rate $1 - o(1)$.

Note that a KEM can be bootstrapped to a PKE by combining a data encapsulation mechanism (DEM) with appropriate security properties [CS02,KD04, HK07], and the composition applies well in the leakage-resilient setting (without requiring DEM to be leakage-resilient). Taken together, our KEM construction indicates that LR-CCA secure PKE with optimal leakage ratio are achievable based on solely HPS.

2 Preliminaries

2.1 Basic Notations

For a distribution or random variable X, we write $x \xleftarrow{\text{R}} X$ to denote the operation of sampling a random x according to X. For a set X, we use $x \xleftarrow{\text{R}} X$ to denote the operation of sampling x uniformly at random from X, and use $|X|$ to denote its size. We use U_X to denote the uniform distribution over X.

We denote $\lambda \in \mathbb{N}$ as the security parameter. Unless described otherwise, all quantities are implicit functions of λ, and all cryptographic algorithms (including the adversary) take λ as an input. We say that a quantity is negligible, written $\mathsf{negl}(\lambda)$, if it vanishes faster than the inverse of any polynomial in λ. A probabilistic polynomial time (PPT) algorithm is a randomized algorithm that runs in time $\mathsf{poly}(\lambda)$. If $\mathcal{A}$ is a randomized algorithm, we write $z \leftarrow \mathcal{A}(x_1, \ldots, x_n; r)$ to indicate that $\mathcal{A}$ outputs z on inputs $(x_1, \ldots, x_n)$ and random coins r. For notational clarity we usually omit r and write $z \leftarrow \mathcal{A}(x_1, \ldots, x_n)$.

Due to space limit, we defer the definition of standard cryptographic primitives and information background to the full version.

2.2 Regular Functions

A function f is injective (akin, 1-to-1) if every image has one and only one preimage. Following [BHSV98], we measure the amount of "non-injectivity" by looking at the maximum preimage size. Let ν be a quantity of security parameter λ. We say that f is ν-to-1 (or ν-approximately-regular) if ν bounds the maximum preimage size of f: any image has at most ν preimages under f. Particularly, if every image has the same number (say ν) of preimages, we say f is ν-regular.

We develop the following useful lemma which establishes the relation between the min-entropy of X and $f(X)$.

Lemma 1. *Let $f : D \to R$ is a ν-to-1 function and X is a random variable over domain D. Then we have:*

$$\mathsf{H}_\infty(f(X)) \geq \mathsf{H}_\infty(X) - \log \nu$$

Proof. Let x^* be the value in the domain that maximizes $\Pr[X = x]$ and y^* be the value in the range that maximizes $\Pr[f(X) = y]$. Since every image has at most ν preimages, it follows that $\Pr[f(X) = y^*] = \sum_{x \in f^{-1}(y^*)} \Pr[X = x] \leq \nu \cdot \Pr[X = x^*]$. According to the definition of min-entropy, the lemma immediately follows. The equality achieves when f is ν-regular and X follows the uniform distribution. Moreover, the above relation applies to average min-entropy as well. Suppose X is correlated to another random variable Y, we have $\tilde{\mathsf{H}}_\infty(f(X)|Y) \geq \tilde{\mathsf{H}}_\infty(X|Y) - \log \nu$. □

Hereafter, we do not distinguish ν-approximately-regular and ν-regular. For ease of presentation, we refer to them collectively as ν-regular.

3 Regularly Lossy Functions and Extensions

3.1 Regularly Lossy Functions

Now, we define the notion of RLFs. Suppose the size of domain is $2^{n(\lambda)}$ where $n(\lambda) = \mathsf{poly}(\lambda)$. Define $\nu(\lambda) \leq 2^{n(\lambda)}$ to represent the *non-injectivity* of the collection, and $2^\tau(\lambda) \leq 2^{n(\lambda)}$ to represent the *image size* of the collection. For all these quantities, we often omit the dependence on the security parameter λ.

A collection of (ν, τ)-RLFs is given by four polynomial time algorithms satisfying the following properties:

- $\mathsf{Setup}(\lambda)$: on input λ, output public parameter pp which includes the descriptions of evaluation key space EK, domain X and range Y.
- $\mathsf{GenNormal}(pp)$: on input pp, output an evaluation key ek. $f_{ek}(\cdot)$ is a ν-regular function from X to Y.
- $\mathsf{GenLossy}(pp)$: on input pp, output an evaluation key ek. $f_{ek}(\cdot)$ is a lossy function from X to Y whose image has size at most 2^τ. The *lossiness* is defined as $n - \tau$.
- $\mathsf{Eval}(ek, x)$: on input ek and an element $x \in X$, output $y \leftarrow f_{ek}(x)$.

Hard to distinguish normal from lossy. For all $pp \leftarrow \mathsf{Setup}(\lambda)$, the outputs of $\mathsf{GenNormal}(pp)$ and $\mathsf{GenLossy}(pp)$ are computationally indistinguishable.

Remark 1. Our notion of RLFs is a generalization of LFs. In the case $\nu = 1$, RLFs obviously boil down to LFs.

3.2 All-But-One Regularly Lossy Functions

To admit more applications, it is convenient to work with a richer notion named ABO-RLFs. The extension is an analog of LTFs to ABO-LTFs in [PW08]. In an ABO collection, each function has an extra input called its *branch*. All of the branches are regular functions, except for one branch is lossy. The lossy branch is an auxiliary input to the evaluation key generation algorithm, and its value is hidden (computationally) by the resulting evaluation key.

We retain the same notation for n, ν, τ as above, and let B be the set of branches. A collection of (ν, τ)-ABO-RLFs consists of three polynomial time algorithms satisfying the following properties:

- $\mathsf{Setup}(\lambda)$: on input λ, output public parameter pp which specifies of evaluation key space EK, branch set B, domain X and range Y.
- $\mathsf{Gen}(pp, b^*)$: on input pp and any $b^* \in B$, output an evaluation key ek. For any $b \neq b^*$, $f_{ek,b}(\cdot)$ is a ν-regular function from X to Y, while $f_{ek,b^*}(\cdot)$ is a lossy function from X to Y whose image has size at most 2^τ.
- $\mathsf{Eval}(ek, b, x)$: on input an evaluation key ek and a branch $b \in B$ and an element $x \in X$, output $y \leftarrow f_{ek,b}(x)$.

Hidden lossy branch. For any $b_0^*, b_1^* \in B \times B$, the output ek_0 of $\mathsf{Gen}(pp, b_0^*)$ and the output ek_1 of $\mathsf{Gen}(pp, b_1^*)$ are computationally indistinguishable.

Peikert and Waters [PW08] showed that LTFs and ABO-LTFs are equivalent for appropriate choices of parameters and degree of lossiness. It is straightforward to verify the equivalence also holds in our regularly lossy setting. We list the results as below for completeness. The security proofs are omitted here since they follow readily from [PW08].

Lemma 2. *There exists a collection of (ν, τ)-ABO-RLFs having exactly two branches if and only if there exists a collection of (ν, τ)-RLFs.*

4 Concrete Construction of ABO-RLFs

In this section, we build ABO-RLFs from the DDH assumption. Our construction mainly follow the matrix approach due to [PW08], but with important refinement for better efficiency.

We first recall the algorithm named $\mathsf{GenConceal}$ for generating a pseudorandom concealer matrix that enjoys certain useful linearity properties from [PW08]. In a nutshell, $\mathsf{GenConceal}$ takes as input positive integers n and m (where $n \geq m$), outputs a $n \times m$ matrix $\mathbb{G}^{n \times m}$, in which the matrix is pseudorandom and all the columns lie in a one-dimensional subspace. More precisely, it works as follows:

- Choose $\mathbf{r} = (r_1, \ldots, r_n) \leftarrow \mathbb{Z}_p^n$ and $\mathbf{s} = (s_1, \ldots, s_m) \leftarrow \mathbb{Z}_p^m$ uniformly at random.
- Let $\mathbf{V} = \mathbf{r} \otimes \mathbf{s} = \mathbf{r}^t \mathbf{s} \in \mathbb{Z}_p^{n \times m}$ be the outer product of $\mathbf{r}$ and $\mathbf{s}$.
- Output $\mathbf{C} = g^{\mathbf{V}} \in \mathbb{G}^{n \times m}$ as the concealer matrix.

Lemma 3 **([PW08]).** *Let $n, m = \mathsf{poly}(\lambda)$. Under the DDH assumption, the conceal matrix $\mathbf{C} = g^{\mathbf{V}} \leftarrow \mathsf{GenConceal}(n, m)$ is pseudorandom over $\mathbb{G}^{n \times m}$.*

Our construction of ABO-RLFs from the DDH assumption is as below.

- $\mathsf{Setup}(\lambda)$: run $(\mathbb{G}, g, p) \leftarrow \mathsf{GroupGen}(\lambda)$, output $pp = (\mathbb{G}, g, p)$ and $B = \mathbb{Z}_p$.
- $\mathsf{Gen}(pp, b^*)$: on input pp and $b^* \in \mathbb{Z}_p$, invoke $\mathsf{GenConceal}(n, m)$ to generate $\mathbf{C} = g^{\mathbf{V}} \in \mathbb{G}^{n \times m}$, output $ek = g^{\mathbf{Y}} = g^{\mathbf{V} - b^* \mathbf{I}'}$, where $\mathbf{I}' \in \mathbb{Z}_p^{n \times m}$, i.e., the ith column is the standard basis vector $\mathbf{e}_i \in \mathbb{Z}_p^n$ for $i \leq n$, and the rest columns are zero vectors.
- $\mathsf{Eval}(ek, b, \mathbf{x})$: on input evaluation key $ek = g^{\mathbf{Y}}$, a branch $b \in \mathbb{Z}_p$ and an element $\mathbf{x} \in \mathbb{Z}_p^n$, output $\mathbf{y} = g^{\mathbf{x}(\mathbf{Y} + b\mathbf{I}')} = g^{\mathbf{x}(\mathbf{V} + (b - b^*)\mathbf{I}')} \in \mathbb{G}^m$.

Lemma 4. *Under the DDH assumption, the above construction is a collection of $(p^{n-m}, \log p)$-ABO-RLFs for $n > 1$.*

Proof. For any $b \neq b^*$, $(\mathbf{V}, b)$ determines p^{n-m}-to-1 function because the rank of $(\mathbf{Y} + b\mathbf{I}')$ is m and the size of the solution space for every $y \in \mathbb{G}^m$ is p^{n-m}. For $b = b^*$, every output $\mathbf{y}$ is of the form $g^{r' \mathbf{s}}$, where $r' = \mathbf{x}\mathbf{r}^t \in \mathbb{Z}_p$. Because $\mathbf{s}$ is fixed by the function index $\mathbf{V}$, there are at most p distinct outputs of any particular function determined by $(\mathbf{V}, b^*)$. The lossiness is $(n-1)\log p$.

The hidden lossy branch property (under the DDH assumption) follows by an elementary reduction: for any branch $b^* \in \mathbb{Z}_p$ the output of $\mathsf{Gen}(\lambda, b^*)$ is computationally indistinguishable from uniform over $\mathbb{G}^{n \times m}$.

Remark 2. The parameter n controls the size of domain, while the parameter m allows us to manipulate the regularity for the ABO branches in a flexible manner. When $m = n$ the above construction becomes the standard ABO lossy functions because the ABO branches are injective.

In the DDH-based ABO-LTF construction [PW08], the input space is restricted to $\{0,1\}^n$ and m must be larger than n to ensure invertible property. In our construction, we do not require invertible property. Therefore, the input space dramatically extends from $\{0,1\}^n$ to $\mathbb{Z}_p^n$ without expanding the conceal matrix. Moreover, when injective property is not necessary, we could further shrink the matrix by setting m smaller than n. In the matrix-based construction, both the size of evaluation key and the computation cost of evaluation are dominated by n and m. Therefore, compared to the DDH-based ABO-LTFs, our DDH-based ABO-RLFs allows much larger inputs and much better efficiency. The flexible choice of m gives rise to more compact evaluation key.

Following a similar approach due to Hemenway and Ostrovsky [HO12], the above DDH-based construction naturally extends to construction based on the eDDH assumption [HO12], which generalized the DDH, QR and DQR assumptions.

5 Generic Construction of ABO-RLFs

In this section, we focus on generic construction of ABO-RLFs.

5.1 Construction from HPS for Subset Membership Problem

Lemma 2 indicates that ABO-RLF is implied by RLF. Thus, the task of constructing ABO-RLF can be reduced to seeking generic construction of RLF.

Wee [Wee12] introduced the notion of dual HPS. As with universal HPS, dual HPS also centers around a family of hash function $\{\Lambda_{sk}\}$ indexed by secret key sk and whose input x comes from some "hard" language. As before, dual HPS requires that for $x \in L$ (YES instance), the hash value $\Lambda_{sk}(x)$ is completely determined by x and $pk = \alpha(sk)$. On the other hand, for $x \notin L$ (NO instance), dual HPS requires *invertibility* – that $\alpha(sk)$ and $\Lambda_{sk}(x)$ jointly determine sk, and there exists an inversion trapdoor td that enables us to efficiently recover sk given $(\alpha(sk), \Lambda_{sk}(x))$[5] along with x. Wee showed an elegant construction of LTF from dual HPS, which is depicted in Eq. (1) as below.

$$f_x(sk) = \alpha(sk)||\Lambda_x(sk) \tag{1}$$

In Wee's construction, instance x serves as the evaluation key and secret key sk acts as input. The injective mode (when $x \notin L$) follows from the invertible property of dual HPS, whereas the lossy mode (when $x \in L$) follows from the projective property of $\Lambda_{sk}(\cdot)$. Moreover, the indistinguishability of injective and lossy mode follows from the hardness of subset membership problem.

Interestingly, we can build RLF from any HPS via the same construction shown as above. Since RLF is much weaker then LTF, we only need the projective property of HPS; any additional properties such as smooth, universal or invertible properties are unnecessary. Formally, let $(X, L, W, \mathsf{R}, PK, SK, \alpha, \Pi, \Lambda)$ be public parameter of HPS. Assume $f_x(sk) = \alpha(sk)||\Lambda_x(sk)$ is a ν-to-1 function from SK to Π for any $x \notin L$.[6] We have the following lemma.

Lemma 5. *Under the subset membership assumption, Eq. (1) yields a collection of* $(\nu, \log|\mathrm{Img}(\alpha)|)$*-RLFs.*

Proof. Correctness for the normal mode follows readily from the fact that $f_x(\cdot)$ is a ν-to-1 function. Lossiness for the lossy mode follows readily from the projective property, which implies that for any $x \in L$, $\mathrm{Img}(f_x) = \mathrm{Img}(\alpha)$. The indistinguishability between normal mode and lossy mode can be directly reduced to the subset membership assumption. □

Putting all the above together, we can generically construct ABO-RLF from any HPS. The construction proceeds via two steps: (1) build RLF from any HPS; (2) amplify the obtained RLF to ABO-RLF with branch set $\{0,1\}^\ell$. However, this generic construction is not efficient in that its second step invokes ℓ individual copies of RLF and involves some degradation in lossiness.

[5] Following the treatment of [Wee12], we will write $\Lambda_{sk}(x)$ as $\Lambda_x(sk)$ occasionally.

[6] The regularity of α gives an upper bound of ν.

5.2 Efficient Construction from HPS for Algebraic Subset Membership Problem

The above construction serves as a proof of concept that one can generically build ABO-RLF from any HPS. It is intriguing to know if there exists more efficient construction.

Our idea is to exploit more algebra property of the associated subset membership problem. More precisely, we choose to work with group-oriented SMPs, which we call algebraic subgroup membership problem.

Algebraic subset membership problems. We first formally introduce a new class of cryptographic indistinguishability problem called algebraic subset membership problems (ASMPs), which is a special type of SMPs (cf. definition in Sect. 8) with the following requirements.

1. X forms a finite Abelian group, L forms a subgroup of X.
2. The quotient group $H = X/L$ is cyclic with order $p = |X|/|L|$.

With the above algebraic properties, we have the following two useful facts:

- Let $\overline{a} = aL$ for some $a \in X \backslash L$ be a generator of H, then the co-sets $(aL, 2aL, \dots, (p-1)aL, paL = L)$ constitute a partition of X.
- For each $x \in L$, $ia + x \in X \backslash L$ for $1 \leq i < p$.

The hardness of ASMPs is same as that of SMPs, which stipulates the uniform distributions over L and $X \backslash L$ are computationally indistinguishable. Define the density of L as $\rho = |L|/|X|$. When ρ is negligible, $U_L \approx_c U_{X \backslash L}$ is equivalent to $U_L \approx_c U_X$ in that $U_{X \backslash L}$ and U_X are statistically close. When ρ is known, $U_L \approx_c U_{X \backslash L}$ implies $U_L \approx_c U_X$ since one can efficiently reconstruct U_X from U_L, $U_{X \backslash L}$ and ρ.

To demonstrate the generality of ASMP, we instantiate it based the DDH, d-linear, QR and DCR assumptions respectively. Due to space limit, we defer the instantiations to the full version.

Remark 3. ASMP could also be thought as an enhancement of subgroup membership problems with requirement (2). For our application in this work, requirement (2) could be further relaxed to H contains a cyclic subgroup.

Comparison to (refined) subgroup indistinguishability problems. Brakerski and Goldwasser [BG10] introduced the so called subgroup indistinguishability problems (SIPs). SIPs is also defined w.r.t. a finite Abelian group X and a subgroup L. In addition, SIPs require X is isomorphic to direct product of two groups: $X \simeq L \times M$ and $\mathsf{gcd}(\mathrm{ord}(L), \mathrm{ord}(M)) = 1$. Qin and Liu [QL14] introduced refined SIPs, which further requires M to be cyclic. Compared to (refined) SIPs, ASMPs only require the quotient group X/L to be cyclic. Therefore, ASMP is strictly stronger than RSIP, and also arguably stronger than SIP because SIP is unlikely to be implied by the DDH and d-linear problems. Correspondingly, our algebraic subset membership assumption is potentially weaker.

Now we are ready to construct ABO-RLF from HPS for ASMP.

- $\mathsf{Setup}(\lambda)$: run HPS.$\mathsf{Setup}(\lambda)$ to generate $pp = (X, L, W, \mathsf{R}, PK, SK, \alpha, \Pi, \Lambda)$, pick a random generator aL for the quotient group H, output $\hat{pp} = (pp, a)$.
- $\mathsf{Gen}(\hat{pp}, b^*)$: on input $\hat{pp} = (pp, a)$ and a given lossy branch $b^* \in \mathbb{Z}_p$, run $(x, w) \leftarrow$ HPS.$\mathsf{SampYes}(pp)$ to sample a random element from L, compute the evaluation key $ek = -b^*a + x \in X$.
- $\mathsf{Eval}(ek, b, sk)$: on input an evaluation key $ek = -b^*a + x$, a branch b and an input sk, compute $\alpha(sk)||\Lambda_{sk}(ek + ba)$. This algorithm defines $f_{ek,b}(sk) := \alpha(sk)||\Lambda_{sk}(ek + ba)$.

Theorem 1. *Assume $X = \{0,1\}^n$ and the function $f_x(sk) = \alpha(sk)||\Lambda_x(sk)$ is a ν-regular for any $x \notin L$. The above construction yields a collection of $(\nu, \log|\mathrm{Img}\alpha|)$-ABO-RLFs under the algebraic subset membership problem.*

Proof. By the group property of the ASMP, $ek + ba = x + (b - b^*)a \notin L$ as long as $b \neq b^*$. In this case, $f_{ek,b}(\cdot)$ is a ν-regular function. When $b = b^*$, $ek + ba = x + (b - b^*)a = x \in L$. In this case, $f_{ek,b}(\cdot)$ is a lossy function by the projective property. For the security, the hidden lossy branch property follows readily from the subgroup membership problem. For any $b_0^*, b_1^* \in \mathbb{Z}_p$, $(-b_0^*a + x) \equiv_c (-b_0^*a + u) \equiv u \equiv (-b_1^*a + u) \equiv_c (-b_1^*a + x)$, where $u \xleftarrow{\mathrm{R}} X$. This proves the theorem. □

6 Leakage-Resilient One-Way Functions

We now show LFs implies a family of leakage-resilient OWFs. The construction and security proof are in the same spirit of the implication LTFs ⇒ injective TDFs given in [PW08]. We prove the implication also holds in the leakage setting.

Theorem 2. *Suppose* $(\mathsf{Setup}, \mathsf{GenInj}, \mathsf{GenLossy}, \mathsf{Eval})$ *give a collection of lossy functions over $\{0,1\}^n$ whose the image size of functions in the lossy mode is at most 2^τ. Then* $(\mathsf{Setup}, \mathsf{GenInj}, \mathsf{Eval})$ *is a collection of ℓ-leakage-resilient injective OWFs over $\{0,1\}^n$ for any $\ell \leq n - \tau - \omega(\lambda)$.*

Due to space limit, we defer to proof to the full version.

7 Leakage-Resilient Message Authentication Code

In this section, we construct leakage-resilient MAC from ABO-RLFs and OT-RLFs, respectively.

7.1 Construction from ABO Regularly Lossy Functions

We show how to convert an ABO-RLF to a MAC. The high-level idea is treating input as secret key and branch as message, outputting the function value as tag.

- $\mathsf{Setup}(\lambda)$: run ABORLF.$\mathsf{Setup}(\lambda)$ to generate $pp = (EK, B, X, Y)$ where $|X| = 2^n$ and $B = \{0,1\}^b$, generate $ek \leftarrow$ ABORLF.$\mathsf{Gen}(pp, 0^b)$, output $\hat{pp} = (pp, ek)$. The key space $K = X$, the message space $M = B$ and the tag space $T = Y$.
- $\mathsf{Gen}(\hat{pp})$: pick $k \xleftarrow{\text{R}} X$ as the secret key.
- $\mathsf{Tag}(k, m)$: compute $t \leftarrow f_{ek,m}(k)$, output (m, t).
- $\mathsf{Vefy}(k, m, t)$: output 1 if $t = f_{ek,m}(k)$ and 0 otherwise.

Theorem 3. *If ABORLF is a collection of (ν, τ)-ABO-RLFs, the above construction is ℓ-leakage-resilient selectively one-time sUF as long as $\omega(\log \lambda) \leq n - \tau - \ell - \log \nu$.*

Due to space limit, we defer the proof to the full version.

8 Leakage-Resilient CCA-secure KEM

Our starting point is the work of Qin and Liu [QL13]. By combining a universal HPS and an OT-LF in a clever manner, they obtained a simple and efficient leakage-resilient CCA-secure PKE scheme with higher leakage rate than previous constructions based on HPS [NS09,LWZ13].

To better illustrate our idea, we first briefly review their construction and security proof. Their construction can be divided in two steps. In the first step, they followed the approach of [NS09] to build a LR CPA-secure PKE from a universal$_1$-HPS. The first part ciphertext is $(x, s, z = \mathsf{ext}(\pi, s) + m)$, where x is a random element in L with witness w, s is a random seed for randomness extractor ext, m is the message, and $\pi =$ HPS.$\mathsf{Pub}(pk, x, w)$. In the second step, they employed an OT-LF $f_{ek,\cdot}(\cdot)$ to generate a randomized tag to authenticate the first part ciphertext. The second part ciphertext is (b_c, t), where b_c is randomly chosen core branch, $x||s||z$ serves as the auxiliary branch b_a, and $t = f_{ek,b_c||b_a}(k)$. This differs from previous (leakage-resilient) CCA-secure PKE constructions which use an independent universal$_2$ HPS to authenticate the first part ciphertext, and eventually allows high leakage ratio.

To establish security, the challenge ciphertext $c^* = (x^*, s^*, z^*, b_c^*, t^*)$ evolves via a sequence of hybrids. In the last hybrid, x^* is sampled from $X \backslash L$ and t^* is evaluated via a lossy core branch $b_c^* \leftarrow$ OTLF.$\mathsf{SampLossy}(td, b_a^* = x^*||s^*||z^*)$. No PPT adversary can tell the changes due to the hardness of subset membership problem and the indistinguishability of lossy branches and injective ones. Conditioned on c^*, it is possible that $\pi^* =$ HPS.$\mathsf{Priv}(sk, x^*)$ maintains high min-entropy by proper parameter choice of ext and the fact that t^* is evaluated under a lossy branch. On one hand, when a PPT adversary makes decryption queries, $f_{ek,(b_c,b_a)}(\cdot)$ is an injective function with overwhelming probability due to the evasiveness of OT-LF, and thus the resulting t maintains the min-entropy of its input. According to the universal property of HPS and the fact that t^* is evaluated under a lossy branch, $\Lambda_{sk}(x)$ has high average min-entropy when $x \notin L$ even after exposing c^*. Thereby, the reduction can safely reject all invalid decryption queries with $x \notin L$. On the other hand, due to the projection of

$\mathit{\Lambda}_{sk}$, the responses to all valid decryption queries do not reveal more information about sk other than pk and c^*. In summary, the decryption oracle does not reveal more information of π^* to the adversary. Upon the this point, ext can be used to distill the leftover entropy from π^* as the session key to mask m.

From both theoretic and practical interest, KEM is more preferable than PKE. In Qin-Liu's PKE, the auxiliary branch b_a is of the from (x, s, z). During the security proof, $z^* = m^* + \mathsf{ext}(\pi^*, s^*)$ cannot be determined by the reduction in advance, in that m^* is one of the two messages outputted by the adversary in the challenge stage. Thereby, the reduction is unable to decide the lossy branch at the very beginning, but has to generate it with the help of trapdoor on-the-fly. In contrast, in the KEM setting the reduction has fully control of the challenge ciphertext $c^* = (x^*, s^*)$, which could be programmed as the lossy branch before the generation of evaluation key. Thereby, the agility of OT-LF is overkilled and its static version – ABO-LF suffices. Moreover, we note that both OT-LF and ABO-LF act as a leakage-resilient MAC in the construction. Combining this observation with the implication we have shown in Sect. 7, a HPS and an ABO-RLF suffice for the construction of leakage-resilient CCA-secure KEM.

Next, we formally show how to construct leakage-resilient CCA-secure KEM from HPS and ABO-RLF. We first recall the notion of HPS [CS02] as below.

Hash Proof System. A HPS consists of the following algorithms:

- $\mathsf{Setup}(\lambda)$: on input a security parameter λ, output public parameter $pp = (X, L, W, \mathsf{R}, PK, SK, \alpha, \mathit{\Pi}, \mathit{\Lambda})$. Here X is a finite non-empty set, L is a proper subset of X defined by binary relation $\mathsf{R} \subset X \times W$ such that $x \in L$ if and only if $(x, w) \in \mathsf{R}$ for some witness $w \in W$. Here PK is the public key space, SK is the secret key space, $\alpha : SK \rightarrow PK$ is a projective map, $\mathit{\Pi}$ is the proof space, $\mathit{\Lambda} = \{\mathit{\Lambda}_{sk} : X \rightarrow \mathit{\Pi}\}_{sk \in SK}$ is a family of hash functions indexed by SK.
- $\mathsf{SampYes}(pp)$: on input pp, outputs a random element $x \in L$, together with a witness $w \in W$ for x. We refer to elements belong to L as Yes instances.
- $\mathsf{SampNo}(pp)$: on input pp, output a random element $x \in X \backslash L$. We refer to elements belong to $X \backslash L$ as No instances.
- $\mathsf{KeyGen}(pp)$: on input pp, pick $sk \xleftarrow{\mathrm{R}} SK$, compute $pk \leftarrow \alpha(sk)$, output a key pair (pk, sk).
- $\mathsf{Priv}(sk, x)$: on input sk and $x \in X$, output its hash proof $\pi \leftarrow \mathit{\Lambda}_{sk}(x)$.
- $\mathsf{Pub}(pk, x, w)$: on input pk, $x \in L$ together with a witness w, output $\pi \in \mathit{\Pi}$.

Subset membership problem. Cramer and Shoup [CS02] introduced the subset membership problems (SMP) to abstract natural cryptographic indistinguishability problems such as the DDH and QR problems as well as others.

SMP w.r.t. (X, L, W, R) requires the random distributions over L and $X \backslash L$ are computationally indistinguishable, i.e., for any PPT adversary $\mathcal{A}$, we have:

$$\mathsf{Adv}_{\mathcal{A}}^{\mathrm{smp}}(\lambda) = |\Pr[\mathcal{A}(pp, x_0)] - \Pr[\mathcal{A}(pp, x_1)]| \leq \mathsf{negl}(\lambda)$$

where $pp \leftarrow \mathsf{Gen}(\lambda)$, $(x_0, w) \leftarrow \mathsf{SampYes}(pp)$, and $x_1 \leftarrow \mathsf{SampNo}(pp)$.

Projection. Λ is projective if the action of Λ_{sk} on L is determined by $pk = \alpha(sk)$, i.e., for all $(pk, sk) \leftarrow \mathsf{KeyGen}(pp)$ and all $x \in L$ with witness w, we have:

$$\Lambda_{sk}(x) = \mathsf{Pub}(pk, x, w)$$

Universal$_1$. Λ is ϵ_1-universal$_1$ if for all $pk \in PK$, all $x \in X \backslash L$ and all $\pi \in \Pi$, we have:

$$\Pr[\Lambda_{sk}(x) = \pi | (pk, x)] \leq \epsilon_1$$

where the probability is over all possible sk with $\alpha(sk) = pk$.

The lemma below follows directly from the definition of min-entropy.

Lemma 6. *If Λ is ϵ_1-universal$_1$, then for all $pk \in PK$ and $x \in X \backslash L$, it holds that $\mathsf{H}_\infty(\Lambda_{sk}(x)|(pk, x)) \geq \log 1/\epsilon_1$, where $sk \leftarrow SK$ with $pk = \alpha(sk)$.*

8.1 Construction from HPS and ABO-RLF

Now, we show how to construct LR CCA-secure KEM from a universal$_1$ HPS, an ABO-RLF and randomness extractor. An overview of our construction is depicted in Fig. 1

- $\mathsf{Setup}(\lambda)$:
  run HPS.$\mathsf{Setup}(\lambda)$ to generate $pp_1 = (X, L, W, \mathsf{R}, PK, SK, \alpha, \Pi, \Lambda)$[7], where Λ is ϵ_1-universal$_1$ for $n = \log 1/\epsilon_1$; run ABORLF.$\mathsf{Setup}(\lambda)$ to generate $pp_2 = (EK, B = X \times \{0,1\}^d, \Pi, T)$; pick an average-case $(n - \tau - \ell, k, \epsilon_2)$-extractor $\mathsf{ext} : \Pi \times \{0,1\}^d \rightarrow K$ where $k = \log |K|$; output $pp = (pp_1, pp_2)$.
- $\mathsf{KeyGen}(pp)$: parse $pp = (pp_1, pp_2)$, then run $(pk, sk) \leftarrow$ HPS.$\mathsf{KeyGen}(pp_1)$ and $ek \leftarrow$ ABORLF.$\mathsf{Gen}(pp_2, 0^{m+d})$, output public key $\hat{pk} = (pk, ek)$ and secret key sk.
- $\mathsf{Encaps}(\hat{pk})$: on input $\hat{pk} = (pk, ek)$, sample $(x, w) \leftarrow$ HPS.$\mathsf{SampYes}(pp_1)$, compute $\pi \leftarrow$ HPS.$\mathsf{Pub}(pk, x, w)$, pick a random seed $s \xleftarrow{\mathrm{R}} \{0,1\}^d$, compute $t \leftarrow f_{ek, x||s}(\pi)$, output $c = (x, s, t)$ and $k \leftarrow \mathsf{ext}(\pi, s)$.
- $\mathsf{Decaps}(sk, c)$: on input sk and $c = (x, s, t)$, compute $\pi \leftarrow$ HPS.$\mathsf{Priv}(sk, x)$, output $k \leftarrow \mathsf{ext}(\pi, s)$ if $t = f_{ek, x||s}(\pi)$ and $\perp$ otherwise.

Theorem 4. *Assuming SMP is hard, HPS is an ϵ_1-universal$_1$ hash proof system, ABORLF is a collection of (ν, τ)-ABO-RLFs and ext be an average-case $(n - \tau - \ell, k, \epsilon_2)$-strong extractor, the above construction is ℓ-leakage-resilient CCA-secure as long as $\omega(\log \lambda) \leq n - \tau - \ell - k - \log \nu$.*

Due to space limit, we defer the proof to the full version.

Comparison. Compared to Qin-Liu's PKE [QL13, QL14], our construction is more efficient and conceptually simpler. Note that Qin-Liu's PKE requires a universal HPS and an OT-LF, while our construction requires a universal HPS

[7] Assume each element in X can be uniquely encoded as a binary string in $\{0,1\}^m$.

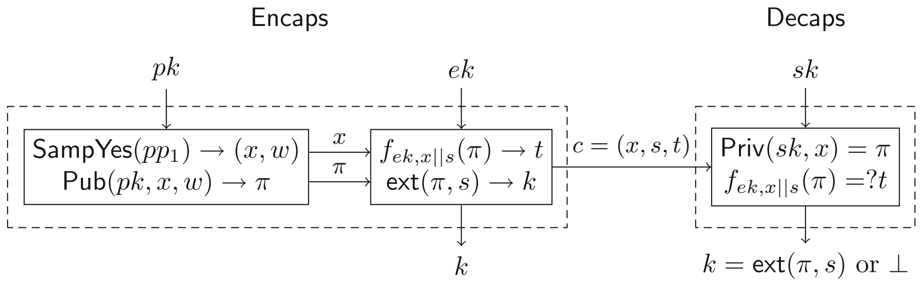
Encaps
Decaps
pk
ek
sk
$\mathsf{SampYes}(pp_1) \rightarrow (x, w)$
$\mathsf{Pub}(pk, x, w) \rightarrow \pi$
x
π
$f_{ek,x||s}(\pi) \rightarrow t$
$\mathsf{ext}(\pi, s) \rightarrow k$
$c = (x, s, t)$
$\mathsf{Priv}(sk, x) = \pi$
$f_{ek,x||s}(\pi) =?t$
k
$k = \mathsf{ext}(\pi, s)$ or $\perp$

Fig. 1. Our approach of KEM construction from HPS and ABORLF.

and an ABO-RLF. To date, the only known construction of OT-LF is from ABO-LF and chameleon hash function. As we have shown in Sect. 4, ABO-RLFs admit more efficient realizations than ABO-LFs. Moreover, as we have show in Sect. 5, ABO-RLFs can be generically build from any HPS. This implication indicates that our construction can be based solely on HPS, and help us to further reduce the footprint of cryptographic code.

Acknowledgement. We thank the anonymous reviewers of CT-RSA 2018 for their useful comments.

The first author is supported by the National Key Research and Development Plan (Grant No. 2016YFB0800403), the National Natural Science Foundation of China (Grant No. 61772522), Youth Innovation Promotion Association CAS and Key Research Program of Frontier Sciences, CAS (Grant No. QYZDB-SSW-SYS035). The second author is supported by the National Natural Science Foundation of China (Grant No. 61502400). The third author is supported by the National Natural Science Foundation of China (Grant No. 61602473) and the National Cryptography Development Fund (Grant No. MMJJ20170116).

References

[ADW09a] Alwen, J., Dodis, Y., Wichs, D.: Leakage-resilient public-key cryptography in the bounded-retrieval model. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 36–54. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-03356-8_3

[ADW09b] Alwen, J., Dodis, Y., Wichs, D.: Survey: leakage resilience and the bounded retrieval model. In: Kurosawa, K. (ed.) ICITS 2009. LNCS, vol. 5973, pp. 1–18. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14496-7_1

[AGV09] Akavia, A., Goldwasser, S., Vaikuntanathan, V.: Simultaneous hardcore bits and cryptography against memory attacks. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 474–495. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-00457-5_28

[BFO08] Boldyreva, A., Fehr, S., O'Neill, A.: On notions of security for deterministic encryption, and efficient constructions without random Oracles. In:

Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 335–359. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85174-5_19

[BG10] Brakerski, Z., Goldwasser, S.: Circular and leakage resilient public-key encryption under subgroup indistinguishability. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 1–20. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14623-7_1

[BHHO08] Boneh, D., Halevi, S., Hamburg, M., Ostrovsky, R.: Circular-secure encryption from decision Diffie-Hellman. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 108–125. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85174-5_7

[BHSV98] Bellare, M., Halevi, S., Sahai, A., Vadhan, S.: Many-to-one trapdoor functions and their relation to public-key cryptosystems. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 283–298. Springer, Heidelberg (1998). https://doi.org/10.1007/BFb0055735

[BL17] Boyen, X., Li, Q.: All-but-many lossy trapdoor functions from lattices and applications. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017. LNCS, vol. 10403, pp. 298–331. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63697-9_11

[BW10] Boyen, X., Waters, B.: Shrinking the keys of discrete-log-type lossy trapdoor functions. In: Zhou, J., Yung, M. (eds.) ACNS 2010. LNCS, vol. 6123, pp. 35–52. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13708-2_3

[CS02] Cramer, R., Shoup, V.: Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 45–64. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-46035-7_4

[DHLW10] Dodis, Y., Haralambiev, K., López-Alt, A., Wichs, D.: Efficient public-key cryptography in the presence of key leakage. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 613–631. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-17373-8_35

[DORS08] Dodis, Y., Ostrovsky, R., Reyzin, L., Smith, A.: Fuzzy extractors: how to generate strong keys from biometrics and other noisy data. SIAM J. Comput. **38**(1), 97–139 (2008)

[FGK+13] Freeman, D.M., Goldreich, O., Kiltz, E., Rosen, A., Segev, G.: More constructions of lossy and correlation-secure trapdoor functions. J. Cryptol. **26**(1), 39–74 (2013)

[GPV08] Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: STOC, pp. 197–206 (2008)

[HK07] Hofheinz, D., Kiltz, E.: Secure hybrid encryption from weakened key encapsulation. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 553–571. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74143-5_31

[HLAWW13] Hazay, C., López-Alt, A., Wee, H., Wichs, D.: Leakage-resilient cryptography from minimal assumptions. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 160–176. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38348-9_10

[HLOV11] Hemenway, B., Libert, B., Ostrovsky, R., Vergnaud, D.: Lossy encryption: constructions from general assumptions and efficient selective opening chosen ciphertext security. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 70–88. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25385-0_4

[HO12] Hemenway, B., Ostrovsky, R.: Extended-DDH and lossy trapdoor functions. In: Fischlin, M., Buchmann, J., Manulis, M. (eds.) PKC 2012. LNCS, vol. 7293, pp. 627–643. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-30057-8_37

[Hof12] Hofheinz, D.: All-but-many lossy trapdoor functions. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 209–227. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29011-4_14

[Hof13] Hofheinz, D.: Circular chosen-ciphertext security with compact ciphertexts. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 520–536. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38348-9_31

[KD04] Kurosawa, K., Desmedt, Y.: A new paradigm of hybrid encryption scheme. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 426–442. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-28628-8_26

[KMO10] Kiltz, E., Mohassel, P., O'Neill, A.: Adaptive trapdoor functions and chosen-ciphertext security. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 673–692. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13190-5_34

[Kom16] Komargodski, I.: Leakage resilient one-way functions: the auxiliary-input setting. In: Hirt, M., Smith, A. (eds.) TCC 2016. LNCS, vol. 9985, pp. 139–158. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53641-4_6

[KOS17] Kiltz, E., O'Neill, A., Smith, A.D.: Instantiability of RSA-OAEP under chosen-plaintext attack. J. Cryptol. **30**(3), 889–919 (2017)

[KPSY09] Kiltz, E., Pietrzak, K., Stam, M., Yung, M.: A new randomness extraction paradigm for hybrid encryption. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 590–609. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-01001-9_34

[KV09] Katz, J., Vaikuntanathan, V.: Signature schemes with bounded leakage resilience. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 703–720. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-10366-7_41

[LWZ13] Liu, S., Weng, J., Zhao, Y.: Efficient public key cryptosystem resilient to key leakage chosen ciphertext attacks. In: Dawson, E. (ed.) CT-RSA 2013. LNCS, vol. 7779, pp. 84–100. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36095-4_6

[MY10] Mol, P., Yilek, S.: Chosen-Ciphertext Security from Slightly Lossy Trapdoor Functions. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 296–311. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13013-7_18

[NS09] Naor, M., Segev, G.: Public-key cryptosystems resilient to key leakage. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 18–35. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-03356-8_2

[PW08] Peikert, C., Waters, B.: Lossy trapdoor functions and their applications. In: STOC, pp. 187–196 (2008)

[QL13] Qin, B., Liu, S.: Leakage-resilient chosen-ciphertext secure public-key encryption from hash proof system and one-time lossy filter. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013. LNCS, vol. 8270, pp. 381–400. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-42045-0_20

[QL14] Qin, B., Liu, S.: Leakage-flexible CCA-secure public-key encryption: simple construction and free of pairing. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 19–36. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54631-0_2

[Reg05] Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: STOC, pp. 84–93 (2005)

[RS09] Rosen, A., Segev, G.: Chosen-ciphertext security via correlated products. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 419–436. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-00457-5_25

[Wee12] Wee, H.: Dual projective hashing and its applications — lossy trapdoor functions and more. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 246–262. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29011-4_16

[XLL+13] Xue, H., Li, B., Lu, X., Jia, D., Liu, Y.: Efficient lossy trapdoor functions based on subgroup membership assumptions. In: Abdalla, M., Nita-Rotaru, C., Dahab, R. (eds.) CANS 2013. LNCS, vol. 8257, pp. 235–250. Springer, Cham (2013). https://doi.org/10.1007/978-3-319-02937-5_13

[Zha16] Zhandry, M.: The magic of ELFs. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9814, pp. 479–508. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53018-4_18