



Trapdoor Functions from the Computational Diffie-Hellman Assumption

Sanjam Garg^{1(✉)} and Mohammad Hajiabadi^{1,2}

¹ University of California, Berkeley, USA

{sanjamg,mdhajiabadi}@berkeley.edu

² University of Virginia, Charlottesville, USA

Abstract. Trapdoor functions (TDFs) are a fundamental primitive in cryptography. Yet, the current set of assumptions known to imply TDFs is surprisingly limited, when compared to public-key encryption. We present a new general approach for constructing TDFs. Specifically, we give a generic construction of TDFs from any Chameleon Encryption (Döttling and Garg [CRYPTO'17]) satisfying a novel property which we call recyclability. By showing how to adapt current Computational Diffie-Hellman (CDH) based constructions of chameleon encryption to yield recyclability, we obtain the first construction of TDFs with security proved under the CDH assumption. While TDFs from the Decisional Diffie-Hellman (DDH) assumption were previously known, the possibility of basing them on CDH had remained open for more than 30 years.

Keywords: Trapdoor functions
Computational Diffie-Hellman assumption · Hash encryption

1 Introduction

Trapdoor functions (TDFs) are a fundamental primitive in cryptography, historically pre-dating the complexity-based development of public key encryption (PKE) [11, 28]. Informally, TDFs are a family of functions, where each function in the family is easy to compute given the function's index key, and also easy to invert given an associated trapdoor key. The security requirement is that a randomly chosen function from the family should be hard to invert without knowledge of a trapdoor key.

Research supported in part from DARPA/ARL SAFEWARE Award W911NF15C0210, AFOSR Award FA9550-15-1-0274, AFOSR YIP Award, DARPA and SPAWAR under contract N66001-15-C-4065, a Hellman Award and research grants by the Okawa Foundation, Visa Inc., and Center for Long-Term Cybersecurity (CLTC, UC Berkeley). The views expressed are those of the author and do not reflect the official policy or position of the funding agencies.

A salient difference between TDFs and PKE lies in their inversion (decryption) algorithms: while the inversion algorithm of a TDF recovers the entire pre-image in full, the decryption algorithm of a PKE only recovers the corresponding plaintext, and not necessarily the randomness. This full-input recovery feature of TDFs is useful in many applications. For example, suppose we have two image points $y_1 := F(\text{ik}_1, x_1)$ and $y_2 := F(\text{ik}_2, x_2)$ of a trapdoor function F , and we want to convince Alice — who is given both y_1 and y_2 but only a trapdoor key tk_1 for ik_1 — that $x_1 = x_2$. This will be easy for Alice to do herself: retrieve x_1 from y_1 using the trapdoor key tk_1 and check whether $y_1 = F(\text{ik}_1, x_1)$ and $y_2 = F(\text{ik}_2, x_1)$.¹ This is a very useful property, especially in the context of chosen-ciphertext (CCA2) security, and is in fact the main reason behind the success of building CCA2-secure PKE in a black-box way from various forms of TDFs [23, 27, 29]. In contrast, enabling this technique based on PKE [25] requires the use of expensive non-interactive zero knowledge proofs [4, 16], which in turn require strong assumptions and lead to non-black-box constructions.

The deterministic structure of TDFs, however, comes with a price, making the construction of TDFs more challenging than that of PKE. This belief is justified by an impossibility result of Gertner, Malkin and Reingold [18] showing that TDFs cannot be built from PKE in a black-box way. As another evidence, while it was known from the 80's how to build semantically-secure PKE from the Decisional Diffie-Hellman (DDH) assumption [15, 20, 32], it took two decades to realize TDFs based on DDH [27].

Despite the fundamental nature of TDFs and extensive research on this notion [1–3, 17, 23, 27, 29, 31] a long-standing question has remained open:

Can trapdoor functions be based on the Computational Diffie-Hellman (CDH) Assumption?

The main difficulty of the above question is that all known DDH-based constructions of TDFs, e.g., [17, 27], exploit properties of DDH, such as pseudorandomness of low rank matrices of group elements, which do not hold in the CDH setting (see Sect. 1.1).

Apart from being a natural question, it has the following theoretical motivation: since we know that TDFs are not necessary in a black-box sense for PKE [18], there may be computational assumptions that imply PKE but not TDFs. Thus, it is important to understand whether TDFs can be obtained from all existing computational assumptions that imply PKE. This provides insights into the hardness nature of TDFs as well as our computational assumptions.

1.1 Lack of CDH-Based Techniques for TDF

Diffie-Hellman related assumptions (even DDH) do not naturally lend themselves to a TDF construction. The main reason why it is more difficult to build TDFs

¹ Here we also need to assume that certifying whether a given point is in the domain of a trapdoor function can be done efficiently.

from such assumptions, compared to, say, factoring related assumptions, is that we do not know of any generic trapdoors for the *discrete log* problem. Indeed, a long standing open problem in cryptography is whether PKE can be based on the sole hardness of the discrete log problem. To see how this makes things more difficult, consider ElGamal encryption: to encrypt a group element g_m under a public key (g, g_1) , we return $(g^r, g_1^r \cdot g_m)$, where r is a random exponent. The decryption algorithm can recover g_m but not r because computing r is as hard as solving the discrete log problem.

Known DDH-based TDF constructions [17, 27] get around the above obstacle by designing their TDFs in such a way that during inversion, one will only need to solve the discrete log problem over a small space, e.g., recovering a bit b from g^b . The main idea is as follows: the index key ik of their TDF is g^M , where g is a generator of the group \mathbb{G} of order p and $M \in \mathbb{Z}_p^{n \times n}$ is a random $n \times n$ invertible matrix and g^M denotes entry-wise exponentiation. Let $tk := M^{-1}$ be the trapdoor key. Using ik , the evaluation algorithm on input $x \in \{0, 1\}^n$ may use the algebraic property of the group to compute $y := g^{Mx} \in \mathbb{G}^n$. Now using tk and y one can compute $g^x \in \mathbb{G}^n$, hence retrieving x .

To argue about one-wayness, one uses the following property implied by DDH: the matrix g^M is computationally indistinguishable from a matrix g^{M_1} , where M_1 is a random matrix of rank one. If the index key is now set to g^{M_1} and if we have $2^n \gg p$, then even an unbounded adversary cannot retrieve the original x from y . This argument is used to establish one-wayness for the TDF.

Unfortunately, the above rank indistinguishability property used to prove one-wayness is not known (and not believed) to be implied by CDH. Thus, designing TDFs based on CDH requires new techniques.

Finally, we mention that even from the Computational Bilinear Assumption [6] (i.e., pairing-based CDH) no TDF constructions are known. The closest is a result of Wee [30], showing that trapdoor relations, which are much weaker than TDFs, can be built from CDH. Roughly, trapdoor relations are a relaxed version of TDFs, in that the function might not be efficiently computable on individual points but one may sample efficiently a *random* input element together with its corresponding image.

1.2 Our Results and Techniques

We give the first construction of TDFs under the CDH assumption. Our construction is black-box and is obtained through a general construction of TDFs from a primitive we call a *recyclable one-way function with encryption (OWFE)*. Moreover, we show that an *adaptive* strengthening of our notion of recyclable OWFE yields a black-box construction of CCA2-secure PKE.

An OWFE is described by a one-way function $f_{pp}: \{0, 1\}^n \rightarrow \{0, 1\}^\nu$, where pp is a public parameter, together with encapsulation/decapsulation algorithms (E, D) . Specifically, E takes as input pp , an image $y \in \{0, 1\}^\nu$ of f_{pp} , an index $i \in [n]$ and a selector bit $b \in \{0, 1\}$, and produces an encapsulated ciphertext ct and a corresponding *key bit* $e \in \{0, 1\}$. The algorithm D allows anyone to retrieve e from ct using any pre-image x of y whose i th bit is b . For security,

letting $y := f_{pp}(x)$, we require if $(ct, e) \stackrel{\$}{\leftarrow} E(pp, y, (i, b))$ and $x_i \neq b$, then even knowing x one cannot recover e from ct with probability better than $\frac{1}{2} + \text{negl}(\lambda)$, where λ is the security parameter. That is, for any $x \in \{0, 1\}^n$, $i \in [n]$, we have $(x, ct, e) \stackrel{c}{\equiv} (x, ct, e')$, where $e' \stackrel{\$}{\leftarrow} \{0, 1\}$, $(ct, e) \stackrel{\$}{\leftarrow} E(pp, f(pp, x), (i, 1 - x_i))$ and $\stackrel{c}{\equiv}$ denotes computational indistinguishability. Our OWFE notion is a weakening of the hash encryption notion [13] in that we do not require f to be collision resistant. The following is a variant of the CDH-based construction of [13].

CDH-based instantiation of OWFE [13]. Let \mathbb{G} be a group of prime order p . The public parameter is a $2 \times n$ matrix of random group elements $pp := \begin{pmatrix} g_{1,0}, g_{2,0}, \dots, g_{n,0} \\ g_{1,1}, g_{2,1}, \dots, g_{n,1} \end{pmatrix}$ and $y := f(pp, x \in \{0, 1\}^n) = \prod_{j \in [n]} g_{j, x_j}$.

To perform $E(pp, y, (i, b))$, sample $\rho \stackrel{\$}{\leftarrow} \mathbb{Z}_p$ and return (ct, e) , where

$$ct := \begin{pmatrix} g'_{1,0}, g'_{2,0}, \dots, g'_{n,0} \\ g'_{1,1}, g'_{2,1}, \dots, g'_{n,1} \end{pmatrix}, \text{ where } g'_{i,1-b} := \perp, g'_{i,b} := g_{i,b}^{\rho}$$

and for all $j \neq i : g'_{j,0} := g_{j,0}^{\rho}$ and $g'_{j,1} := g_{j,1}^{\rho}$

and $e := HC(y^{\rho})$, where HC is a hardcore bit function. The function D is now derived easily. See the main body for the proof of security.

Recyclability. Our recyclability notion asserts that the ciphertext part output, ct , of the key encapsulation algorithm E is independent of the corresponding image input part y . That is, letting E_1 and E_2 refer to the first and second output of E , for any values of y_1 and y_2 , we always have $E_1(pp, y_1, (i, b); \rho) = E_1(pp, y_2, (i, b); \rho)$. It is easy to verify that the above CDH-based OWFE satisfies this property. Thus, we may drop y as an input to E_1 and obtain the following:

Property 1. Letting $x \in \{0, 1\}^n$, $x_i = b$, $ct := E_1(pp, (i, b); \rho)$ and $y := f(pp, x)$:

$$D(pp, x, ct) = E_2(pp, y, (i, b); \rho).$$

1.3 Sketch of Our OWFE-Based Construction and Techniques

Let (K, f, E, D) be a recyclable OWFE scheme.² Our TDF construction is based on a new technique that we call *bits planting*. Briefly, the input X to our TDF consists of a domain element $x \in \{0, 1\}^n$ of $f(pp, \cdot)$ and a *blinding* string $\mathbf{b} \in \{0, 1\}^{n \times r}$, for some r that we will specify later. The output Y is comprised of $y := f(pp, x)$, as well as a matrix of bits in which we copy all the bits of \mathbf{b} in the clear but in hidden spots determined by x ; we fill up the rest of the matrix with key bits that somehow correspond to bit-by-bit encryption of x under y . To an adversary, the matrix is “unrevealing,” with no indicative signs of what spots corresponding to the blinding part — which contain \mathbf{b} in the clear. However, using our designed trapdoor, an inverter can pull out both x and \mathbf{b} from Y with all but negligible probability.

² K is the public-parameter generation algorithm. See Definition 3.

Warm-up Construction. We first give a warm up construction in which our inversion algorithm only recovers half of the input bits (on average). Our TDF input is of the form $(x, \mathbf{b}) \in \{0, 1\}^n \times \{0, 1\}^n$. This warm-up construction contains most of the ideas behind the full-blown construction.

- **Key generation:** The trapdoor key is $\text{tk} := \begin{pmatrix} \rho_{1,0}, \dots, \rho_{n,0} \\ \rho_{1,1}, \dots, \rho_{n,1} \end{pmatrix}$ a matrix of randomness values, and the index key is $\text{ik} := \text{pp}, \begin{pmatrix} \text{ct}_{1,0}, \dots, \text{ct}_{n,0} \\ \text{ct}_{1,1}, \dots, \text{ct}_{n,1} \end{pmatrix}$, formed as:

$$\text{ik} := \text{pp}, \begin{pmatrix} \text{ct}_{1,0} := E_1(\text{pp}, (1, 0); \rho_{1,0}), \dots, \text{ct}_{n,0} := E_1(\text{pp}, (n, 0); \rho_{n,0}) \\ \text{ct}_{1,1} := E_1(\text{pp}, (1, 1); \rho_{1,1}), \dots, \text{ct}_{n,1} := E_1(\text{pp}, (n, 1); \rho_{n,1}) \end{pmatrix}.$$

- **Evaluation** $F(\text{ik}, X)$: Parse $\text{ik} := \text{pp}, \begin{pmatrix} \text{ct}_{1,0}, \dots, \text{ct}_{n,0} \\ \text{ct}_{1,1}, \dots, \text{ct}_{n,1} \end{pmatrix}$ and parse the input X as $(x \in \{0, 1\}^n, \mathbf{b} := \mathbf{b}_1 \dots \mathbf{b}_n \in \{0, 1\}^n)$. Set $y := f(\text{pp}, x)$. For $i \in [n]$ set M_i as follows:

$$\begin{aligned} M_i &:= \begin{pmatrix} D(\text{pp}, x, \text{ct}_{i,0}) \\ \mathbf{b}_i \end{pmatrix} \stackrel{*}{=} \begin{pmatrix} E_2(\text{pp}, y, (i, 0); \rho_{i,0}) \\ \mathbf{b}_i \end{pmatrix} \text{ if } x_i = 0 \\ M_i &:= \begin{pmatrix} \mathbf{b}_i \\ D(\text{pp}, x, \text{ct}_{i,1}) \end{pmatrix} \stackrel{*}{=} \begin{pmatrix} \mathbf{b}_i \\ E_2(\text{pp}, y, (i, 1); \rho_{i,1}) \end{pmatrix} \text{ if } x_i = 1 \end{aligned} \quad (1)$$

The matrix M_i is computed using the *deterministic* algorithm D and the equalities specified as $\stackrel{*}{=}$ follow by Property (1).

Return $Y := (y, M_1 || \dots || M_n)$.

- **Inversion** $F^{-1}(\text{tk}, Y)$: Parse $Y := (y, M_1 || \dots || M_n)$ and $\text{tk} := \begin{pmatrix} \rho_{1,0}, \dots, \rho_{n,0} \\ \rho_{1,1}, \dots, \rho_{n,1} \end{pmatrix}$. Set

$$(M'_1 || \dots || M'_n) := \begin{pmatrix} E_2(\text{pp}, y, (1, 0); \rho_{1,0}), \dots, E_2(\text{pp}, y, (n, 0); \rho_{n,0}) \\ E_2(\text{pp}, y, (1, 1); \rho_{1,1}), \dots, E_2(\text{pp}, y, (n, 1); \rho_{n,1}) \end{pmatrix}. \quad (2)$$

Output $(x, \mathbf{b} := \mathbf{b}_1 \dots \mathbf{b}_n)$, where we retrieve x_i and \mathbf{b}_i as follows. If $M'_{i,1} = M_{i,1}$ and $M'_{i,2} \neq M_{i,2}$ (where $M_{i,1}$ is the first element of M_i), then set $x_i := 0$ and $\mathbf{b}_i := M_{i,2}$. If $M'_{i,1} \neq M_{i,1}$ and $M'_{i,2} = M_{i,2}$, then set $x_i := 1$ and $\mathbf{b}_i := M_{i,1}$. Else, set $x_i := \perp$ and $\mathbf{b}_i := \perp$.

One-Wayness (Sketch). We show $(\text{ik}, Y) \stackrel{c}{\equiv} (\text{ik}, Y_{\text{sim}})$, where (ik, Y) is as above and

$$Y_{\text{sim}} := y, \begin{pmatrix} E_2(\text{pp}, y, (1, 0); \rho_{1,0}), \dots, E_2(\text{pp}, y, (n, 0); \rho_{n,0}) \\ E_2(\text{pp}, y, (1, 1); \rho_{1,1}), \dots, E_2(\text{pp}, y, (n, 1); \rho_{n,1}) \end{pmatrix}.$$

Noting that we may produce $(\text{ik}, Y_{\text{sim}})$ using only pp and y , the one-wayness of $f(\text{pp}, \cdot)$ implies it is hard to recover x from $(\text{ik}, Y_{\text{sim}})$, and so also from (ik, Y) .

Why $(\text{ik}, Y) \stackrel{c}{\equiv} (\text{ik}, Y_{\text{sim}})$? Consider $Y_{\text{sim},1}$, whose first column is the same as Y_{sim} and whose subsequent columns are the same as Y . We prove $(x, \text{ik}, Y) \stackrel{c}{\equiv} (x, \text{ik}, Y_{\text{sim},1})$; the rest will follow using a hybrid argument.

Letting M_1 be formed as in Eq. 1, to prove $(x, ik, Y) \stackrel{c}{\equiv} (x, ik, Y_{\text{sim},1})$ it suffices to show

$$x, \begin{pmatrix} E_1(\text{pp}, y, (1, 0); \rho_{1,0}) \\ E_1(\text{pp}, (1, 1); \rho_{1,1}) \end{pmatrix}, M_1 \stackrel{c}{\equiv} x, \begin{pmatrix} E_1(\text{pp}, (1, 0); \rho_{1,0}) \\ E_1(\text{pp}, (1, 1); \rho_{1,1}) \end{pmatrix}, \begin{pmatrix} E_2(\text{pp}, y, (1, 0); \rho_{1,0}) \\ E_2(\text{pp}, y, (1, 1); \rho_{1,1}) \end{pmatrix}. \quad (3)$$

We prove Eq. 3 using the security property of OWFE, which says

$$(x, E_1(\text{pp}, (1, 1 - x_1); \rho), b') \stackrel{c}{\equiv} (x, E_1(\text{pp}, (1, 1 - x_1); \rho), E_2(\text{pp}, y, (1, 1 - x_1); \rho)), \quad (4)$$

where $b' \stackrel{\$}{\leftarrow} \{0, 1\}$ and ρ is random. We give an algorithm that converts a sample from either side of Eq. 4 into a sample from the same side of Eq. 3. On input (x, ct_1, b_1) , sample $(ct_2, b_2) \stackrel{\$}{\leftarrow} E(\text{pp}, y, (1, x_1))$ and

- if $x_1 = 0$, then return $x, \begin{pmatrix} ct_2 \\ ct_1 \end{pmatrix}, \begin{pmatrix} b_2 \\ b_1 \end{pmatrix}$;
- else if $x_1 = 1$, then return $x, \begin{pmatrix} ct_1 \\ ct_2 \end{pmatrix}, \begin{pmatrix} b_1 \\ b_2 \end{pmatrix}$.

The claimed property of the converter follows by inspection. Finally, we mention that the argument used to prove Eq. 3 builds on a technique used by Brakerski et al. [7] to build circularly-secure PKE.

Correctness. F^{-1} recovers on average half of the input bits: F^{-1} fails for an index $i \in [n]$ if $b_i = E_2(\text{pp}, y, (i, 1 - x_i); \rho_{i,1-x_i})$. This happens with probability $\frac{1}{2}$ because b_i is a completely random bit.

Boosting correctness. To boost correctness, we provide r blinding bits for each index i of $x \in \{0, 1\}^n$. That is, the input to the TDF is $(x, \mathbf{b}) \in \{0, 1\}^n \times \{0, 1\}^{rn}$. We will also expand ik by providing r encapsulated ciphertexts for each position $(i, b) \in [n] \times \{0, 1\}$. This extra information will bring the inversion error down to 2^{-r} . We will show that one-wayness is still preserved.

On the role of blinding. One may wonder why we need to put a blinding string \mathbf{b} in the TDF input. Why do not we simply let the TDF input be x and derive multiple key bits for every index i of x by applying D to the corresponding ciphertexts provided for that position i in the index key ik ; the inverter can still find the matching bit for every index. The reason behind our design choice is that by avoiding blinders, it seems very difficult (if not impossible) to give a reduction to the security of the OWFE scheme.

1.4 CCA2 Security

Rosen and Segev [29] show that an extended form of one-wayness for TDFs, which they term *k-repetition security*, leads to a black-box construction of CCA2-secure PKE. Informally, a TDF is *k-repetition secure* if it is hard to recover a random input X from $F(ik_1, X), \dots, F(ik_k, X)$, where ik_1, \dots, ik_k are sampled independently. They show that *k-repetition security*, for $k \in \Theta(\lambda)$, suffices for CCA2 security.

We give a CCA2 secure PKE by largely following [29], but we need to overcome two problems. The first problem is that our TDF is not k -repetition secure, due to the blinding part of the input. We overcome this by observing that a weaker notion of k -repetition suffices for us: one in which we should keep x the same across all k evaluations but may sample \mathbf{b} freshly each time. A similar weakening was also used in [26].

The second problem is that our inversion may fail with negligible probability for every choice of (ik, tk) and a bit care is needed here. In particular, the simulation strategy of [29] will fail if the adversary can create an image Y , which is a true image of a domain point X , but which the inversion algorithm fails to invert. To overcome this problem, we slightly extend the notion of security required by OWFE, calling it *adaptive OWFE*, and show that if our TDF is instantiated using this primitive, it satisfies all the properties needed to build CCA2 secure PKE.

Comparison with related CCA2 constructions. We note that CCA2-secure PKE constructions from CDH are already known, e.g., [9, 22, 30], which are more efficient than the one obtained by instantiating our construction using CDH. We presented a CCA-2 secure construction just to show the black-box utility of our base general primitive. The recent results of [7, 12–14], combined with [8], show that CCA-secure PKE can be built from a related primitive called chameleon/batch encryption, but in a non-black-box way.

1.5 Discussion

Black-box power of chameleon encryption. Our work is a contribution toward understanding the black-box power of the notion of chameleon encryption. Recent works [7, 12, 13] show that chameleon encryption (and its variants) may be used in a non-black-box way to build strong primitives such as identity-based encryption (IBE). The work of Brakerski et al. [7] shows also black-box applications of (a variant of) this notion, obtaining in turn circularly-secure and leakage-resilient PKE from CDH. Our work furthers the progress in this area, by giving a black-box construction of TDFs.

Related work. Hajiabadi and Kapron [21] show how to build TDFs from any *reproducible* circularly secure single-bit PKE. Informally, a PKE is reproducible if given a public key pk' , a public/secret key (pk, sk) and a ciphertext $c := \text{PKE.E}(pk', b'; r)$, one can recycle the randomness of c to obtain $\text{PKE.E}(pk, b; r)$ for any bit $b \in \{0, 1\}$. Brakerski et al. [7] recently built a circularly secure single-bit PKE using CDH. Their construction is not reproducible, however. (The following assumes familiarity with [7].) In their PKE, a secret key x of their PKE is an input to their hash function and the public key y is its corresponding image. To encrypt a bit b they (a) additively secret-share b into (b_1, \dots, b_n) , where $n = |x|$ and (b) form $2n$ ciphertext $ct_{i,b}$, where $ct_{i,b}$ encrypts b_i using y relative to (i, b) . Their scheme is not reproducible because the randomness used for step (a) cannot be recycled and also half of the randomness used to create hash encryption ciphertexts in step (b) cannot be recycled. (This half corresponds to

the bits of the target secret key w.r.t. which we want to recycle randomness.) It is not clear whether their scheme can be modified to yield reproducibility.

Open problems. Our work leads to several open problems. Can our TDF be improved to yield perfect correctness? Our current techniques leave us with a negligible inversion error. Can we build lossy trapdoor functions (LTDF) [27] from recyclable-OWFE/CDH? Given the utility of LTDFs, a construction based on CDH will be interesting. Can we build deterministic encryption based on CDH matching the parameters of those based on DDH [5]?

2 Preliminaries

Notation. We use λ for the security parameter. We use $\stackrel{c}{\equiv}$ to denote computational indistinguishability between two distributions and use \equiv to denote two distributions are identical. For a distribution D we use $x \stackrel{\$}{\leftarrow} D$ to mean x is sampled according to D and use $y \in D$ to mean y is in the support of D . For a set S we overload the notation to use $x \stackrel{\$}{\leftarrow} S$ to indicate that x is chosen uniformly at random from S .

Definition 1 (Trapdoor Functions (TDFs)). Let $w = w(\lambda)$ be a polynomial. A family of trapdoor functions TDF with domain $\{0, 1\}^w$ consists of three PPT algorithms TDF.K, TDF.F and TDF.F⁻¹ with the following syntax and security properties.

- TDF.K(1^λ): Takes the security parameter 1^λ and outputs a pair (ik, tk) of index/trapdoor keys.
- TDF.F(ik, X): Takes an index key ik and a domain element $X \in \{0, 1\}^w$ and outputs an image element Y.
- TDF.F⁻¹(tk, Y): Takes a trapdoor key tk and an image element Y and outputs a value $X \in \{0, 1\}^w \cup \{\perp\}$.

We require the following properties.

- **Correctness:** For any $(ik, tk) \in \text{TDF.K}(1^\lambda)$

$$\Pr[\text{TDF.F}^{-1}(\text{tk}, \text{TDF.F}(\text{ik}, X)) \neq X] = \text{negl}(\lambda), \tag{5}$$

where the probability is taken over $X \stackrel{\$}{\leftarrow} \{0, 1\}^w$.

- **One-wayness:** For any PPT adversary \mathcal{A}

$$\Pr[\mathcal{A}(\text{ik}, Y) = X] = \text{negl}(\lambda), \tag{6}$$

where $(ik, tk) \stackrel{\$}{\leftarrow} \text{TDF.K}(1^\lambda)$, $X \stackrel{\$}{\leftarrow} \{0, 1\}^w$ and $Y = \text{TDF.F}(\text{ik}, X)$.

A note about the correctness condition. Our correctness notion relaxes that of perfect correctness by allowing the inversion algorithm to fail (with respect to any trapdoor key) for a negligible fraction of evaluated elements. This relaxation

nonetheless suffices for all existing applications of perfectly-correct TDFs. Our correctness notion, however, implies a weaker notion under which the correctness probability is also taken over the choice of the index/trapdoor keys. This makes our result for constructing TDFs stronger.

Definition 2 (Computational Diffie-Hellman (CDH) Assumption). *Let \mathbb{G} be a group-generator scheme, which on input 1^λ outputs (\mathbb{G}, p, g) , where \mathbb{G} is the description of a group, p is the order of the group which is always a prime number and g is a generator of the group. We say that \mathbb{G} is CDH-hard if for any PPT adversary \mathcal{A} : $\Pr[\mathcal{A}(\mathbb{G}, p, g, g^{a_1}, g^{a_2}) = g^{a_1 a_2}] = \text{negl}(\lambda)$, where $(\mathbb{G}, p, g) \xleftarrow{\$} \mathbb{G}(1^\lambda)$ and $a_1, a_2 \xleftarrow{\$} \mathbb{Z}_p$.*

3 Recyclable One-Way Function with Encryption

We will start by defining the notion of a one-way function with encryption. This notion is similar to the chameleon encryption notion of Döttling and Garg [13]. However, it is weaker in the sense that it does not imply collision-resistant hash functions.

Next, we will define a novel *ciphertext-randomness recyclability* property for one-way function with encryption schemes. We will show that a variant of the chameleon encryption construction of Döttling and Garg [13] satisfies this ciphertext-randomness recyclability property.

3.1 Recyclable One-Way Function with Encryption

We provide the definition of a one-way function with encryption. We define the notion as a key-encapsulation mechanism with single bit keys.

Definition 3 (One-Way Function with Encryption (OWFE)). *An OWFE scheme consists of four PPT algorithms K , f , E and D with the following syntax.*

- $K(1^\lambda)$: Takes the security parameter 1^λ and outputs a public parameter pp for a function f from n bits to ν bits.
- $f(\text{pp}, x)$: Takes a public parameter pp and a preimage $x \in \{0, 1\}^n$, and outputs $y \in \{0, 1\}^\nu$.
- $E(\text{pp}, y, (i, b); \rho)$: Takes a public parameter pp , a value y , an index $i \in [n]$, a bit $b \in \{0, 1\}$ and randomness ρ , and outputs a ciphertext ct and a bit e .³
- $D(\text{pp}, x, \text{ct})$: Takes a public parameter pp , a value x and a ciphertext ct , and deterministically outputs $e' \in \{0, 1\} \cup \{\perp\}$.

We require the following properties.

- **Correctness:** For any $\text{pp} \in K(1^\lambda)$, any $i \in [n]$, any $x \in \{0, 1\}^n$ and any randomness value ρ , the following holds: letting $y := f(\text{pp}, x)$, $b := x_i$ and $(\text{ct}, e) := E(\text{pp}, y, (i, b); \rho)$, we have $e = D(\text{pp}, x, \text{ct})$.

³ ct is assumed to contain (i, b) .

– **One-wayness:** For any PPT adversary \mathcal{A} :

$$\Pr[f(\text{pp}, \mathcal{A}(\text{pp}, y)) = y] = \text{negl}(\lambda),$$

where $\text{pp} \xleftarrow{\$} \mathcal{K}(1^\lambda)$, $x \xleftarrow{\$} \{0, 1\}^n$ and $y := f(\text{pp}, x)$.

– **Security for encryption:** For any $i \in [n]$ and $x \in \{0, 1\}^n$:

$$(x, \text{pp}, \text{ct}, e) \stackrel{c}{\equiv} (x, \text{pp}, \text{ct}, e')$$

where $\text{pp} \xleftarrow{\$} \mathcal{K}(1^\lambda)$, $(\text{ct}, e) \xleftarrow{\$} \mathcal{E}(\text{pp}, f(\text{pp}, x), (i, 1 - x_i))$ and $e' \xleftarrow{\$} \{0, 1\}$.

Definition 4 (Recyclability). We say that an OWFE scheme $(f, \mathcal{K}, \mathcal{E}, \mathcal{D})$ is recyclable if the following holds. Letting E_1 and E_2 refer to the first and second output of \mathcal{E} , the value of $E_1(\text{pp}, y, (i, b); \rho)$ is always independent of y . That is, for any $\text{pp} \in \mathcal{K}(1^\lambda)$, $y_1, y_2 \in \{0, 1\}^\nu$, $i \in [n]$, $b \in \{0, 1\}$ and randomness ρ : $E_1(\text{pp}, y_1, i, b); \rho = E_1(\text{pp}, y_2, (i, b); \rho)$.

We now conclude the above definitions with two remarks.

Note 1 (Simplified Recyclability). Since under the recyclability notion the ciphertext output ct of \mathcal{E} is independent of the input value y , when referring to E_1 , we may omit the inclusion of y as an input and write $\text{ct} = E_1(\text{pp}, (i, b); \rho)$.

Note 2. If the function $f(\text{pp}, \cdot)$ is length decreasing (e.g., $f(\text{pp}, \cdot): \{0, 1\}^n \mapsto \{0, 1\}^{n-1}$), then the one-wayness condition of Definition 3 is implied by the combination of the security-for-encryption and correctness conditions. In our definition, however, we do not place any restriction on the structure of the function f , and it could be, say, a one-to-one function. As such, under our general definition, the one-wayness condition is not necessarily implied by those two other conditions.

3.2 Adaptive One-Way Function with Encryption

For our CCA application we need to work with an adaptive version of the notion of OWFE. Recall by Note 1 that a ciphertext ct does not depend on the corresponding y . The security for encryption notion (Definition 3) says if (ct, e) is formed using an image $y := f(\text{pp}, x)$ and parameters (i, b) , and if $x_i \neq b$, then even knowing x does not help an adversary in distinguishing e from a random bit. The adaptive version of this notion allows the adversary to choose x after seeing ct . This notion makes sense because ct does not depend on the image y , and so ct may be chosen first.

Definition 5 (Adaptive OWFE). We say that $\mathcal{E} = (\mathcal{K}, f, \mathcal{E}, \mathcal{D})$ is an adaptive one-way function with encryption scheme if \mathcal{E} is correct in the sense of Definition 3, f is one-way in the sense of Definition 3 and that \mathcal{E} is adaptively secure in the following sense.

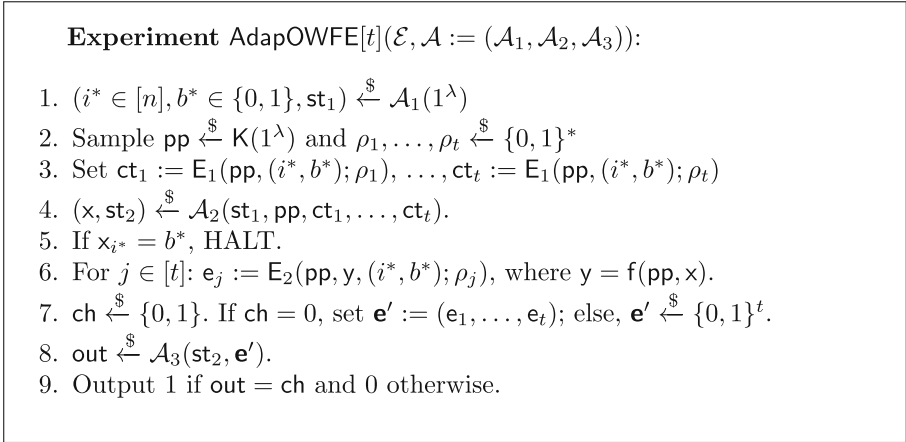


Fig. 1. The $\text{AdapOWFE}[t](\mathcal{E}, \mathcal{A})$ Experiment

– **Adaptive Security:** For any PPT adversary \mathcal{A} , we have the following: the probability that $\text{AdapOWFE}[t = 1](\mathcal{E}, \mathcal{A})$ outputs 1 is $\frac{1}{2} + \text{negl}(\lambda)$, where the experiment $\text{AdapOWFE}[t]$ is defined in Fig. 1.

We remind the reader that in Step 3 of Fig. 1 the algorithm \mathbf{E}_1 does not take any y as input because of Note 1. The following lemma is obtained using a straightforward hybrid argument, so we omit the proof.

Lemma 1. Let $\mathcal{E} = (\mathcal{K}, f, \mathbf{E}, \mathcal{D})$ be an adaptive OWFE scheme. For any polynomial $t := t(\lambda)$ and any PPT adversary \mathcal{A} , we have $\Pr[\text{AdapOWFE}[t](\mathcal{E}, \mathcal{A}) = 1] \leq \frac{1}{2} + \text{negl}(\lambda)$.

3.3 Construction from CDH

We give a CDH-based construction of a recyclable adaptive OWFE based on a group scheme \mathbb{G} (Definition 2), which is a close variant of constructions given in [10, 13].

– $\mathcal{K}(1^\lambda)$: Sample $(\mathbb{G}, p, g) \xleftarrow{\$} \mathbb{G}(1^\lambda)$. For each $j \in [n]$ and $b \in \{0, 1\}$, choose $g_{j,b} \xleftarrow{\$} \mathbb{G}$. Output

$$\text{pp} := \mathbb{G}, p, g, \left(\begin{matrix} g_{1,0}, g_{2,0}, \dots, g_{n,0} \\ g_{1,1}, g_{2,1}, \dots, g_{n,1} \end{matrix} \right). \tag{7}$$

– $f(\text{pp}, x)$: Parse pp as in Eq. 7, and output $y := \prod_{j \in [n]} g_{j, x_j}$.

– $\mathbf{E}(\text{pp}, y, (i, b))$: Parse pp as in Eq. 7. Sample $\rho \xleftarrow{\$} \mathbb{Z}_p$ and proceed as follows:

1. For every $j \in [n] \setminus \{i\}$, set $c_{j,0} := g_{j,0}^\rho$ and $c_{j,1} := g_{j,1}^\rho$.
 2. Set $c_{i,b} := g_{i,b}^\rho$ and $c_{i,1-b} := \perp$.
 3. Set $e := \text{HC}(y^\rho)$.⁴
 4. Output (ct, e) where $\text{ct} := \begin{pmatrix} c_{1,0}, c_{2,0}, \dots, c_{n,0} \\ c_{1,1}, c_{2,1}, \dots, c_{n,1} \end{pmatrix}$.
- $D(\text{pp}, x, \text{ct})$: Parse $\text{ct} := \begin{pmatrix} c_{1,0}, c_{2,0}, \dots, c_{n,0} \\ c_{1,1}, c_{2,1}, \dots, c_{n,1} \end{pmatrix}$. Output $\text{HC}(\prod_{j \in [n]} c_{j,x_j})$.

Lemma 2. *Assuming that \mathbb{G} is CDH-hard and $n \in \omega(\log p)$, the construction described above is an adaptive one-way function with encryption scheme satisfying the recyclability property.*

Proof. We start by proving one-wayness.

One-wayness. The fact that f_{pp} for a random pp is one-way follows by the discrete-log hardness (and hence CDH hardness) of \mathbb{G} . Let g^* be a random group element for which we want to find r^* such that $g^{r^*} = g^*$. Sample $i_1 \xleftarrow{\$} [n]$ and $b_1 \xleftarrow{\$} \{0, 1\}$ and set $g_{i_1, b_1} := g^*$. For all $i \in [n]$ and $b \in \{0, 1\}$ where $(i, b) \neq (i_1, b_1)$, sample $r_{i,b} \xleftarrow{\$} \mathbb{Z}_p$ and set $g_{i,b} := g^{r_{i,b}}$. Set $\text{pp} := \begin{pmatrix} g_{1,0}, \dots, g_{n,0} \\ g_{1,1}, \dots, g_{n,1} \end{pmatrix}$. Sample x' at random from $\{0, 1\}^n$ subject to the condition that $x'_{i_1} = 1 - b_1$. Set $y := \prod_{j \in [n]} g_{j,x'_j}$. Call the inverter adversary on (pp, y) to receive $x \in \{0, 1\}^n$. Now if $n \in \omega(\log p)$, then by the leftover hash lemma with probability negligibly close to $\frac{1}{2}$ we have $x_{i_1} = b_1$, allowing us to find r^* from $r_{i,b}$'s.

Recyclability. We need to show that the ciphertext output ct of E is independent of the input value y . This follows immediately by inspection.

Notation. For a matrix $\mathbf{M} := \begin{pmatrix} a_{1,0}, a_{2,0}, \dots, a_{n,0} \\ a_{1,1}, a_{2,1}, \dots, a_{n,1} \end{pmatrix}$, $i \in [n]$ and $b \in \{0, 1\}$, we define the matrix $\mathbf{M}' := \mathbf{M}|(i, b)$ to be the same as \mathbf{M} except that instead of $a_{i,b}$ we put \perp in \mathbf{M}' . If \mathbf{M} is matrix of group elements, then \mathbf{M}^r denotes element-wise exponentiation to the power of r .

Security for encryption. We show if \mathbb{G} is CDH-hard, then the scheme is adaptively secure. Suppose that there exists an adversary \mathcal{A} for which we have $\Pr[\text{AdapOWFE}[t = 1](\mathcal{E}, \mathcal{A})] = \frac{1}{2} + \frac{1}{q} > \frac{1}{2} + \text{negl}(\lambda)$. Using standard techniques we may transform \mathcal{A} into a predictor \mathcal{B} who wins with probability at least $\frac{1}{2} + \frac{1}{q}$ in the following experiment:

1. $(i^*, b^*) \xleftarrow{\$} \mathcal{B}(1^\lambda)$.
2. Sample

$$\text{pp} := \begin{pmatrix} g_{1,0}, g_{2,0}, \dots, g_{n,0} \\ g_{1,1}, g_{2,1}, \dots, g_{n,1} \end{pmatrix} \xleftarrow{\$} \mathbb{G}^{2 \times n}. \tag{8}$$

⁴ We assume that the $\text{HC}(\cdot)$ is a hardcore bit function. If a deterministic hard-core bit for the specific function is not known then we can use the Goldreich-Levin [19] construction. We skip the details of that with the goal of keeping exposition simple.

3. Sample $\rho \xleftarrow{\$} \mathbb{Z}_p$ and set $\text{ct} := \text{pp}^\rho|(i^*, b^*)$.
4. $(x, b) \xleftarrow{\$} \mathcal{B}(\text{pp}, \text{ct})$.
5. \mathcal{B} wins if $x_{i^*} = b^*$ and $b = \text{HC}(y^\rho)$, where $y := \prod_{j \in [n]} g_{j, x_j}$.

Using the Goldreich-Levin theorem we know that there should be an adversary \mathcal{B}_1 that wins with non-negligible probability in the following:

1. $(i^*, b^*) \xleftarrow{\$} \mathcal{B}_1(1^\lambda)$.
2. Sample

$$\text{pp} := \left(g_{1,0}, g_{2,0} \dots, g_{n,0} \right) \xleftarrow{\$} \mathbb{G}^{2 \times n}. \tag{9}$$

3. Sample $\rho \xleftarrow{\$} \mathbb{Z}_p$ and set $\text{ct} := \text{pp}^\rho|(i^*, b^*)$.
4. $(x, g^*) \xleftarrow{\$} \mathcal{B}_1(\text{pp}, \text{ct})$.
5. \mathcal{B}_1 wins if $x_{i^*} = b^*$ and $g^* = y^\rho$, where $y := \prod_{j \in [n]} g_{j, x_j}$.

We now show how to use \mathcal{B}_1 to solve the CDH problem.

CDH Adversary $\mathcal{A}_1(g, g_1, g_2)$:

- Run $\mathcal{B}_1(1^\lambda)$ to get (i^*, b^*) .
- For any $j \in [n] \setminus \{i^*\}$ and $b \in \{0, 1\}$ sample $\alpha_{j,b} \xleftarrow{\$} \mathbb{Z}_p$ and set $g_{j,b} = g^{\alpha_{j,b}}$. Set $g_{i^*, b^*} := g_1$ and $g_{i^*, 1-b^*} = g^\alpha$, where $\alpha \xleftarrow{\$} \mathbb{Z}_p$. Set

$$\text{pp} := \left(g_{1,0}, g_{2,0} \dots, g_{n,0} \right) \left(g_{1,1}, g_{2,1}, \dots, g_{n,1} \right).$$

- Set $g'_{i^*, 1-b^*} = g_2$ and $g'_{i^*, b^*} = \perp$. For any $j \in [n] \setminus \{i^*\}$ and $b \in \{0, 1\}$ set $g'_{j,b} = g_2^{(\alpha^{-1} \cdot \alpha_{j,b})}$. Set

$$\text{ct} := \left(g'_{1,0}, g'_{2,0} \dots, g'_{n,0} \right) \left(g'_{1,1}, g'_{2,1}, \dots, g'_{n,1} \right). \tag{10}$$

- Run $\mathcal{B}_1(\text{pp}, \text{ct})$ to get (x, g^*) . If $x_i \neq b_i^*$ then return \perp . Otherwise
 - Set

$$g_u := \frac{g^*}{\prod_{j=1}^{i^*-1} g'_{j, x_j} \cdot \prod_{j=i^*+1}^n g'_{j, x_j}}.$$

- Return g_u^α .

By inspection one may easily verify that whenever \mathcal{B}_1 wins, \mathcal{A}_1 also wins. The proof is now complete. □

4 TDF Construction

In this section we describe our TDF construction. We first give the following notation.

Extending the notation for D. For a given pp , a sequence $\text{ct} := (\text{ct}_1, \dots, \text{ct}_r)$ of encapsulated ciphertexts and a value x , we define $D(\text{pp}, x, \text{ct})$ to be the concatenation of $D(\text{pp}, x, \text{ct}_i)$ for $i \in [r]$.

Algorithm Perm. For two lists \mathbf{u}_1 and \mathbf{u}_2 and a bit b we define $\text{Perm}(\mathbf{u}_1, \mathbf{u}_2, b)$ to output $(\mathbf{u}_1, \mathbf{u}_2)$ if $b = 0$, and $(\mathbf{u}_2, \mathbf{u}_1)$ otherwise.

Construction 3 (TDF Construction).

Base Primitive. A recyclable OWFE scheme $\mathcal{E} = (\text{K}, \text{f}, \text{E}, \text{D})$. Let Rand be the randomness space of the encapsulation algorithm E .

Construction. The construction is parameterized over two parameters $n = n(\lambda)$ and $r = r(\lambda)$, where n is the input length to the function f , and r will be instantiated in the correctness proof. The input space of each TDF is $\{0, 1\}^{n+nr}$. We will make use of the fact explained in Note 1.

– TDF.K(1^λ):

- Sample $\text{pp} \leftarrow \text{K}(1^\lambda)$.
- For each $i \in [n]$ and selector bit $b \in \{0, 1\}$:

$$\rho_{i,b} := (\rho_{i,b}^{(1)}, \dots, \rho_{i,b}^{(r)}) \xleftarrow{\$} \text{Rand}^r$$

$$\text{ct}_{i,b} := (\text{E}_1(\text{pp}, (i, b); \rho_{i,b}^{(1)}), \dots, \text{E}_1(\text{pp}, (i, b); \rho_{i,b}^{(r)})).$$

- Form the index key ik and the trapdoor key tk as follows:

$$\text{ik} := (\text{pp}, \text{ct}_{1,0}, \text{ct}_{1,1}, \dots, \text{ct}_{n,0}, \text{ct}_{n,1}) \tag{11}$$

$$\text{tk} := (\text{pp}, \rho_{1,0}, \rho_{1,1}, \dots, \rho_{n,0}, \rho_{n,1}). \tag{12}$$

– TDF.F(ik, X):

- Parse ik as in Eq. 11 and parse

$$X := (x \in \{0, 1\}^n, \mathbf{b}_1 \in \{0, 1\}^r, \dots, \mathbf{b}_n \in \{0, 1\}^r).$$

- Set $y := \text{f}(\text{pp}, x)$.
- For all $i \in [n]$ set

$$\mathbf{e}_i := D(\text{pp}, x, \text{ct}_{i,x_i}).$$

- Return

$$Y := (y, \text{Perm}(\mathbf{e}_1, \mathbf{b}_1, x_1), \dots, \text{Perm}(\mathbf{e}_n, \mathbf{b}_n, x_n)).$$

– TDF.F⁻¹(tk, Y):

- Parse tk as in Eq. 12 and $Y := (y, \widetilde{\mathbf{b}}_{1,0}, \widetilde{\mathbf{b}}_{1,1}, \dots, \widetilde{\mathbf{b}}_{n,0}, \widetilde{\mathbf{b}}_{n,1})$.
- Reconstruct $x := x_1 \cdots x_n$ bit-by-bit and $\mathbf{b} := (\mathbf{b}_1, \dots, \mathbf{b}_n)$ vector-by-vector as follows. For $i \in [n]$:

* Parse $\rho_{i,0} := (\rho_{i,0}^{(1)}, \dots, \rho_{i,0}^{(r)})$ and $\rho_{i,1} := (\rho_{i,1}^{(1)}, \dots, \rho_{i,1}^{(r)})$.
 * If

$$\begin{aligned} \widetilde{\mathbf{b}}_{i,0} &= \left(E_2(\text{pp}, y, (i, 0); \rho_{i,0}^{(1)}), \dots, E_2(\text{pp}, y, (i, 0); \rho_{i,0}^{(r)}) \right) \text{ and} \\ \widetilde{\mathbf{b}}_{i,1} &\neq \left(E_2(\text{pp}, y, (i, 1); \rho_{i,1}^{(1)}), \dots, E_2(\text{pp}, y, (i, 1); \rho_{i,1}^{(r)}) \right), \end{aligned} \quad (13)$$

then set $x_i = 0$ and $\mathbf{b}_i = \widetilde{\mathbf{b}}_{i,1}$.
 * Else, if

$$\begin{aligned} \widetilde{\mathbf{b}}_{i,0} &\neq \left(E_2(\text{pp}, y, (i, 0); \rho_{i,0}^{(1)}), \dots, E_2(\text{pp}, y, (i, 0); \rho_{i,0}^{(r)}) \right) \text{ and} \\ \widetilde{\mathbf{b}}_{i,1} &= \left(E_2(\text{pp}, y, (i, 1); \rho_{i,1}^{(1)}), \dots, E_2(\text{pp}, y, (i, 1); \rho_{i,1}^{(r)}) \right), \end{aligned} \quad (14)$$

then set $x_i = 1$ and $\mathbf{b}_i = \widetilde{\mathbf{b}}_{i,0}$.
 * Else, halt and return \perp .

- If $y \neq f(\text{pp}, x)$, then return \perp . Otherwise, return (x, \mathbf{b}) .

We will now give the correctness and one-wayness statements about our TDF, and will prove them in subsequent subsections.

Lemma 3 (TDF Correctness). *The inversion error of our constructed TDF is at most $\frac{n}{2^r}$. That is, for any $(ik, tk) \in \text{TDF.K}(1^\lambda)$ we have*

$$\beta := \Pr[\text{TDF.F}^{-1}(tk, (\text{TDF.F}(ik, X))) \neq X] \leq \frac{n}{2^r}, \quad (15)$$

where the probability is taken over $X := (x, \mathbf{b}_1, \dots, \mathbf{b}_n) \xleftarrow{\$} \{0, 1\}^{n+nr}$. By choosing $r \in \omega(\log \lambda)$ we will have a negligible inversion error.

For one-wayness we will prove something stronger: parsing $X := (x, \dots)$, then recovering any x' satisfying $f(\text{pp}, x) = f(\text{pp}, x')$ from $(ik, \text{TDF.F}(ik, X))$ is infeasible.

Lemma 4 (One-Wayness). *The TDF $(\text{TDF.K}, \text{TDF.F}, \text{TDF.F}^{-1})$ given in Construction 3 is one-way. That is, for any PPT adversary \mathcal{A}*

$$\Pr[\mathcal{A}(ik, Y) = x' \text{ and } f(\text{pp}, x') = y] = \text{negl}(\lambda), \quad (16)$$

where $(ik := (\text{pp}, \dots), tk) \xleftarrow{\$} \text{TDF.K}(1^\lambda)$, $X := (x, \dots) \xleftarrow{\$} \{0, 1\}^{n+nr}$ and $Y := (y, \dots) := \text{TDF.F}(ik, X)$.

By combining Lemmas 2, 3 and 4 we will obtain our main result below.

Theorem 4 (CDH Implies TDF). *There is a black-box construction of TDFs from CDH-hard groups.*

4.1 Proof of Correctness: Lemma 3

Proof Let $X := (x, \mathbf{b}_1, \dots, \mathbf{b}_n) \stackrel{\$}{\leftarrow} \{0, 1\}^{n+nr}$ be as in the lemma and

$$Y := \text{TDF.F}(\text{ik}, X) := (y, \widetilde{\mathbf{b}}_{1,0}, \widetilde{\mathbf{b}}_{1,1}, \dots, \widetilde{\mathbf{b}}_{n,0}, \widetilde{\mathbf{b}}_{n,1}). \quad (17)$$

By design, for all $i \in [n]$: $\widetilde{\mathbf{b}}_{i,1-x_i} = \mathbf{b}_i$. Parse

$$\begin{aligned} \text{tk} &:= (\rho_{1,0}, \rho_{1,1}, \dots, \rho_{n,0}, \rho_{n,1}), \\ \rho_{i,b} &:= (\rho_{i,b}^{(1)}, \dots, \rho_{i,b}^{(r)}), \text{ for } i \in [n] \text{ and } b \in \{0, 1\}. \end{aligned}$$

Consider the execution of $\text{TDF.F}^{-1}(\text{tk}, Y)$. By the correctness of our recyclable OWFE \mathcal{E} we have the following: the probability that $\text{TDF.F}^{-1}(\text{tk}, Y) \neq X$ is the probability that for some $i \in [n]$:

$$\mathbf{b}_i = \left(\mathbf{E}_2(\text{pp}, y, (i, 1 - x_i); \rho_{i,1-x_i}^{(1)}), \dots, \mathbf{E}_2(\text{pp}, y, (i, 1 - x_i); \rho_{i,1-x_i}^{(r)}) \right). \quad (18)$$

Now since \mathbf{b}_i , for all i , is chosen uniformly at random and independently of x , the probability of the event in Eq. 18 is $\frac{1}{2^r}$. A union bound over $i \in [n]$ gives us the claimed error bound. \square

4.2 Proof of One-wayness: Lemma 4

We will prove Lemma 4 through a couple of hybrids, corresponding to the real and a simulated view. We first give the following definition which will help us describe the two hybrids in a compact way.

Definition 6 Fix $\text{pp}, x \in \{0, 1\}^n$ and $y := f(\text{pp}, x)$. We define two PPT algorithms *Real* and *Sim*, where *Real* takes as input (pp, x) and *Sim* takes as input (pp, y) . We stress that *Sim* does not take x as input.

The algorithm *Real* (pp, x) outputs $(\mathbf{CT}, \mathbf{E})$ and the algorithm *Sim* (pp, y) outputs $(\mathbf{CT}, \mathbf{E}_{\text{sim}})$, sampled in the following way.

- Sample $\begin{pmatrix} \rho_{1,0}, \dots, \rho_{n,0} \\ \rho_{1,1}, \dots, \rho_{n,1} \end{pmatrix} \stackrel{\$}{\leftarrow} \text{Rand}^{2 \times n}$.
- Set

$$\mathbf{CT} := \begin{pmatrix} \text{ct}_{1,0}, \dots, \text{ct}_{n,0} \\ \text{ct}_{1,1}, \dots, \text{ct}_{n,1} \end{pmatrix} := \begin{pmatrix} \mathbf{E}_1(\text{pp}, (1, 0); \rho_{1,0}), \dots, \mathbf{E}_1(\text{pp}, (n, 0); \rho_{n,0}) \\ \mathbf{E}_1(\text{pp}, (1, 1); \rho_{1,1}), \dots, \mathbf{E}_1(\text{pp}, (n, 1); \rho_{n,1}) \end{pmatrix}.$$

- Set

$$\mathbf{E} := \begin{pmatrix} \mathbf{b}_{1,0}, \dots, \mathbf{b}_{n,0} \\ \mathbf{b}_{1,1}, \dots, \mathbf{b}_{n,1} \end{pmatrix},$$

where, for all $i \in [n]$:

- if $x_i = 0$, then $\mathbf{b}_{i,0} := \mathbf{D}(\text{pp}, x, \text{ct}_{i,0})$ and $\mathbf{b}_{i,1} \stackrel{\$}{\leftarrow} \{0, 1\}$.
- if $x_i = 1$, then $\mathbf{b}_{i,0} \stackrel{\$}{\leftarrow} \{0, 1\}$ and $\mathbf{b}_{i,1} := \mathbf{D}(\text{pp}, x, \text{ct}_{i,1})$.

– Set

$$\mathbf{E}_{\text{sim}} := \begin{pmatrix} E_2(\text{pp}, y, (1, 0); \rho_{1,0}), \dots, E_2(\text{pp}, y, (n, 0); \rho_{n,0}) \\ E_2(\text{pp}, y, (1, 1); \rho_{1,1}), \dots, E_2(\text{pp}, y, (n, 1); \rho_{n,1}) \end{pmatrix}$$

We now prove the following lemma which will help us to prove the indistinguishability of the two hybrids in our main proof.

Lemma 5. Fix polynomial $r := r(\lambda)$ and let $x \in \{0, 1\}^n$. We have

$$(\text{pp}, x, \mathbf{CT}_1, \mathbf{E}_1, \dots, \mathbf{CT}_r, \mathbf{E}_r) \stackrel{c}{\equiv} (\text{pp}, x, \mathbf{CT}_1, \mathbf{E}_{\text{sim},1}, \dots, \mathbf{CT}_r, \mathbf{E}_{\text{sim},r}), \quad (19)$$

where $\text{pp} \stackrel{\$}{\leftarrow} K(1^\lambda)$, and for all $i \in [r]$, we sample $(\mathbf{CT}_i, \mathbf{E}_i) \stackrel{\$}{\leftarrow} \text{Real}(\text{pp}, x)$ and $(\mathbf{CT}_i, \mathbf{E}_{\text{sim},i}) \stackrel{\$}{\leftarrow} \text{Sim}(\text{pp}, f(\text{pp}, x))$.

Proof. Fix $x \in \{0, 1\}^n$ and let $y := f(\text{pp}, x)$. For the purpose of doing a hybrid argument we define two algorithms SReal and SSim below.

– $\text{SReal}(i, \text{pp}, x)$: sample $\rho_0, \rho_1 \stackrel{\$}{\leftarrow} \text{Rand}$ and return $(\mathbf{ct}, \mathbf{e})$, where

$$\mathbf{ct} := \begin{pmatrix} \text{ct}_0 \\ \text{ct}_1 \end{pmatrix} := \begin{pmatrix} E_1(\text{pp}, (i, 0); \rho_0) \\ E_1(\text{pp}, (i, 1); \rho_1) \end{pmatrix} \quad (20)$$

and \mathbf{e} is defined as follows:

- if $x_i = 0$, then $\mathbf{e} := \begin{pmatrix} D(\text{pp}, x, \text{ct}_0) \\ b \end{pmatrix}$, where $b \stackrel{\$}{\leftarrow} \{0, 1\}$;
- if $x_i = 1$, then $\mathbf{e} := \begin{pmatrix} b \\ D(\text{pp}, x, \text{ct}_1) \end{pmatrix}$, where $b \stackrel{\$}{\leftarrow} \{0, 1\}$.

– $\text{SSim}(i, \text{pp}, y)$: Return $(\mathbf{ct}, \mathbf{e}_{\text{sim}})$, where \mathbf{ct} is sampled as in Eq. 20 and \mathbf{e}_{sim} is sampled as

$$\mathbf{e}_{\text{sim}} := \begin{pmatrix} E_2(\text{pp}, y, (i, 0); \rho_0) \\ E_2(\text{pp}, y, (i, 1); \rho_1) \end{pmatrix}.$$

We will show that for all $i \in [n]$ and $x \in \{0, 1\}^n$

$$(\text{pp}, x, \mathbf{ct}, \mathbf{e}) \stackrel{c}{\equiv} (\text{pp}, x, \mathbf{ct}, \mathbf{e}_{\text{sim}}), \quad (21)$$

where

$$\text{pp} \stackrel{\$}{\leftarrow} K(1^\lambda), (\mathbf{ct}, \mathbf{e}) \stackrel{\$}{\leftarrow} \text{SReal}(i, \text{pp}, x) \text{ and } (\mathbf{ct}, \mathbf{e}_{\text{sim}}) \stackrel{\$}{\leftarrow} \text{SSim}(i, \text{pp}, f(\text{pp}, x)).$$

From Eq. 21 using a simple hybrid argument the indistinguishability claimed in the lemma (Eq. 19) is obtained. Note that for the hybrid argument we need to make use of the fact that x is provided in both sides of Eq. 21, because we need to know x to be able to build the intermediate hybrids between those of Eq. 19. Thus, in what follows we will focus on proving Eq. 21.

To prove Eq. 21, first note that by the correctness of the OWFE scheme \mathcal{E} , we have

$$(\text{pp}, x, \mathbf{ct}, \mathbf{e}) \equiv (\text{pp}, x, \mathbf{ct}, \mathbf{e}'),$$

where \mathbf{ct} and \mathbf{e} are sampled according to $\text{SReal}(i, \text{pp}, x)$ as above (using randomness values ρ_0 and ρ_1), and \mathbf{e}' is sampled as:

- if $x_i = 0$, then $\mathbf{e}' := \begin{pmatrix} E_2(\mathbf{pp}, y, (i, 0); \rho_0) \\ b \end{pmatrix}$, where $b \stackrel{\$}{\leftarrow} \{0, 1\}$;
- if $x_i = 1$, then $\mathbf{e}' := \begin{pmatrix} b \\ E_2(\mathbf{pp}, y, (i, 1); \rho_1) \end{pmatrix}$, where $b \stackrel{\$}{\leftarrow} \{0, 1\}$.

Thus, we will prove

$$(\mathbf{pp}, x, \mathbf{ct}, \mathbf{e}') \stackrel{c}{\equiv} (\mathbf{pp}, x, \mathbf{ct}, \mathbf{e}_{\text{sim}}). \tag{22}$$

We derive Eq. 22 from the security-for-encryption requirement of the scheme (K, f, E, D) .

Recall that the security for encryption requirement asserts that no PPT adversary can distinguish between $(x, \mathbf{ct}_1, \mathbf{e}_1)$ and $(x, \mathbf{ct}_1, \mathbf{e}_2)$, where $\mathbf{pp} \stackrel{\$}{\leftarrow} K(1^\lambda)$, $(\mathbf{ct}_1, \mathbf{e}_1) \stackrel{\$}{\leftarrow} E(\mathbf{pp}, f(\mathbf{pp}, x), (i, 1 - x_i))$ and $\mathbf{e}_2 \stackrel{\$}{\leftarrow} \{0, 1\}$. Let us call $(x, \mathbf{ct}_1, \mathbf{e}_1)$ the simulated challenge and $(x, \mathbf{ct}_1, \mathbf{e}_2)$ the random challenge.

To build the reduction we show the existence of a procedure Turn that generically turns a simulated challenge into a sample of $\text{SSim}(i, \mathbf{pp}, x)$ and turns a random challenge into a sample of $\text{SReal}(i, \mathbf{pp}, y)$.

The algorithm $\text{Turn}(x, \mathbf{ct}, \mathbf{e})$ returns $(\mathbf{ct}_1, \mathbf{e}_1)$, formed as follows:

- Sample $\rho \stackrel{\$}{\leftarrow} \text{Rand}$. Then
 - if $x_i = 0$, then return

$$\mathbf{ct}_1 = \begin{pmatrix} E_1(\mathbf{pp}, (i, 0); \rho) \\ \mathbf{ct} \end{pmatrix} \quad \mathbf{e}_1 = \begin{pmatrix} E_2(\mathbf{pp}, y, (i, 0); \rho) \\ \mathbf{e} \end{pmatrix}$$

- if $x_i = 1$, then return

$$\mathbf{ct}_1 = \begin{pmatrix} \mathbf{ct} \\ E_1(\mathbf{pp}, (i, 0); \rho) \end{pmatrix} \quad \mathbf{e}_1 = \begin{pmatrix} \mathbf{e} \\ E_2(\mathbf{pp}, y, (i, 0); \rho) \end{pmatrix}$$

It should be clear by inspection that the output of $\text{Turn}(x, \mathbf{ct}, \mathbf{e})$ is identically distributed to $\text{SReal}(i, \mathbf{pp}, x)$ if $(x, \mathbf{ct}, \mathbf{e})$ is a random challenge (defined above), and identically distributed to $\text{SSim}(i, \mathbf{pp}, x)$ if $(x, \mathbf{ct}, \mathbf{e})$ is a simulated challenge. The proof is now complete. \square

Proof (of Lemma 4). To prove Lemma 4 we define two hybrids and will use the notation view_i to refer to the view sampled in **Hybrid** i .

Hybrid 0. The view (ik, Y) is produced honestly as in the real executions of the scheme TDF. That is,

- Sample $\mathbf{pp} \stackrel{\$}{\leftarrow} K(1^\lambda)$, $x \stackrel{\$}{\leftarrow} \{0, 1\}^n$ and let $y := f(\mathbf{pp}, x)$.
- For all $j \in [r]$ sample $(\mathbf{CT}^{(j)}, \mathbf{E}^{(j)}) \stackrel{\$}{\leftarrow} \text{Real}(\mathbf{pp}, x)$. Parse

$$\mathbf{CT}^{(j)} := \begin{pmatrix} \mathbf{ct}_{1,0}^{(j)}, \dots, \mathbf{ct}_{n,0}^{(j)} \\ \mathbf{ct}_{1,1}^{(j)}, \dots, \mathbf{ct}_{n,1}^{(j)} \end{pmatrix} \quad \mathbf{E}^{(j)} := \begin{pmatrix} \mathbf{b}_{1,0}^{(j)}, \dots, \mathbf{b}_{n,0}^{(j)} \\ \mathbf{b}_{1,1}^{(j)}, \dots, \mathbf{b}_{n,1}^{(j)} \end{pmatrix}.$$

– For all $i \in [n]$ and $d \in \{0, 1\}$ set

$$\begin{aligned} \mathbf{ct}_{i,d} &:= (\mathbf{ct}_{i,d}^{(1)}, \dots, \mathbf{ct}_{i,d}^{(r)}) \\ \mathbf{b}_{i,d} &:= (\mathbf{b}_{i,d}^{(1)}, \dots, \mathbf{b}_{i,d}^{(r)}). \end{aligned}$$

– Form the view $(\mathbf{ik}, \mathbf{Y})$ as follows:

$$\underbrace{((\mathbf{pp}, \mathbf{ct}_{1,0}, \mathbf{ct}_{1,1}, \dots, \mathbf{ct}_{n,0}, \mathbf{ct}_{n,1}))}_{\mathbf{ik}}, \underbrace{(y, \mathbf{b}_{1,0}, \mathbf{b}_{1,1}, \dots, \mathbf{b}_{n,0}, \mathbf{b}_{n,1})}_{\mathbf{Y}} \quad (23)$$

Hybrid 1. The view $(\mathbf{ik}, \mathbf{Y})$ is produced the same as **Hybrid 0** except that for all $j \in [r]$ we sample $(\mathbf{CT}^{(j)}, \mathbf{E}^{(j)})$ now as $(\mathbf{CT}^{(j)}, \mathbf{E}^{(j)}) \stackrel{\$}{\leftarrow} \text{Sim}(\mathbf{pp}, y)$.

We prove that the two views are indistinguishable and then we will show that inverting the image under view_1 is computationally infeasible.

Indistinguishability of the views: By Lemma 5 we have $\text{view}_0 \stackrel{c}{\equiv} \text{view}_1$. The reason is that the view in either hybrid is produced entirely based on $(\mathbf{CT}^{(1)}, \mathbf{E}^{(1)}, \dots, \mathbf{CT}^{(r)}, \mathbf{E}^{(r)})$ and that this tuple is sampled from the distribution $\text{Real}(\mathbf{pp}, x)$ in one hybrid and from $\text{Sim}(\mathbf{pp}, y)$ in the other.

One-wayness in Hybrid 1: We claim that for any PPT adversary \mathcal{A}

$$\Pr[\mathcal{A}(\text{view}_1) = x' \text{ and } f(\mathbf{pp}, x') = y] = \text{negl}(\lambda). \quad (24)$$

Recall that $\text{view}_1 := (\mathbf{ik}, \mathbf{Y})$ is the view in **Hybrid 1** and that the variables \mathbf{pp} and y are part of $\mathbf{ik} := (\mathbf{pp}, \dots)$ and $\mathbf{Y} := (y, \dots)$. The proof of Eq. 24 follows from the one-wayness of f , taking into account the fact that view_1 in its entirety is produced solely based on \mathbf{pp} and $y := f(\mathbf{pp}, x)$ (and especially without knowing x). This is because all the underlying variables $(\mathbf{CT}^{(j)}, \mathbf{E}^{(j)})$ — for all j — are produced as $(\mathbf{CT}^{(j)}, \mathbf{E}^{(j)}) \stackrel{\$}{\leftarrow} \text{Sim}(\mathbf{pp}, y)$, which can be formed without knowledge of x .

Completing the Proof of Lemma 4. Let $\text{view}_0 := (\mathbf{ik}, \mathbf{Y})$ and parse $\mathbf{ik} := (\mathbf{pp}, \dots)$ and $\mathbf{Y} := (y, \dots)$. For any PPT adversary \mathcal{A} we need to show that the probability that \mathcal{A} on input view_0 outputs $x' \in \{0, 1\}^n$ such that $f(\mathbf{pp}, x') = y$ is negligible. We know that \mathcal{B} fails to compute such a string x' with non-negligible probability if the view $((\mathbf{pp}, \dots), (y, \dots))$ is sampled according to view_1 . Since $\text{view}_0 \stackrel{c}{\equiv} \text{view}_1$, the claim follows. \square

4.3 Extended One-Wayness

For our CCA2 application we need to prove a stronger property than the standard one-wayness for our constructed TDF. This extension requires that if we evaluate m correlated inputs under m independent functions from the TDF family, the result still cannot be inverted.

Lemma 6 (Extended One-Wayness). *Let $\text{TDF} = (\text{TDF.K}, \text{TDF.F}, \text{TDF.F}^{-1})$ be the TDF built in Construction 3 based on an arbitrary parameter $r = r(\lambda)$. Let $m := m(\lambda)$. For any PPT adversary \mathcal{A}*

$$\Pr[\mathcal{A}(\text{view} := (\text{ik}_1, \dots, \text{ik}_m, \text{Y}_1, \dots, \text{Y}_m)) = \mathbf{x}] = \text{negl}(\lambda),$$

where $\mathbf{x} \xleftarrow{\$} \{0, 1\}^n$ and for $i \in [m]$, $(\text{ik}_i, \text{tk}_i) \xleftarrow{\$} \text{TDF.K}(1^\lambda)$, $\mathbf{b}_i \xleftarrow{\$} \{0, 1\}^{nr}$ and $\text{Y}_i := \text{TDF.F}(\text{ik}_i, \mathbf{x} \parallel \mathbf{b}_i)$. Thus, there exists a hardcore function HC such that $\text{HC}(\mathbf{x})$ remains pseudorandom in the presence of view.

Proof. For any PPT adversary \mathcal{A} we need to show that the probability that $\mathcal{A}(\text{view})$ outputs \mathbf{x} is negligible. It is easy to verify by inspection that the distribution of view can be perfectly formed based on the view $(\text{ik}^*, \text{Y}^*)$ of an inverter against the one-wayness of the trapdoor function $(\text{TDF.K}, \text{TDF.F}, \text{TDF.F}^{-1})$ of Construction 3 but under the new parameter $r' = m \times r$. Invoking Lemma 4 our claimed one-wayness extension follows. \square

5 CCA2-Secure Public-Key Encryption

In this section we show how to use our constructed TDF to build a CCA2 secure PKE. For the proof of CCA2 security we need to assume that the OWFE scheme underlying the TDF is adaptively secure (Definition 5).

Notation. Let $\text{TDF} := (\text{TDF.K}, \text{TDF.F}, \text{TDF.F}^{-1})$ be as in Sect. 4. We will interpret the input X to the TDF as (\mathbf{x}, \mathbf{s}) , where $\mathbf{x} \in \{0, 1\}^n$ corresponds to f 's pre-image part and $\mathbf{s} \in \{0, 1\}^{n_1}$ corresponds to the *blinding* part. In particular, if r is the underlying parameter of the constructed TDF as in Construction 3, then $n_1 = n \times r$.

Ingredients of our CCA2-secure PKE. Apart from a TDF with the above syntax, our CCA2 secure construction also makes use of a one-time signature scheme $\text{SIG} = (\text{SIG.K}, \text{SIG.Sign}, \text{SIG.Ver})$ with perfect correctness, which in turn can be obtained from any one-way function. A one-time signature scheme SIG with message space $\{0, 1\}^\eta$ is given by three PPT algorithms SIG.K , SIG.Sign and SIG.Ver satisfying the following syntax. The algorithm SIG.K on input a security parameter 1^λ outputs a pair (vk, sgk) consisting of a verification key vk and a signing key sgk . The signing algorithm SIG.Sign on input a signing key sgk and a message $\mathbf{m} \in \{0, 1\}^\eta$ outputs a signature σ . For correctness, we require that for any $(\text{vk}, \text{sgk}) \in \text{SIG.K}(1^\lambda)$, any message $\mathbf{m} \in \{0, 1\}^\eta$ and any signature $\sigma \in \text{SIG.Sign}(\text{sgk}, \mathbf{m})$: $\text{SIG.Ver}(\text{vk}, \mathbf{m}, \sigma) = \top$. The one-time unforgeability property requires that the success probability of any PPT adversary \mathcal{A} in the following game be at most negligible. Sample $(\text{vk}, \text{sgk}) \xleftarrow{\$} \text{SIG.K}(1^\lambda)$ and give vk to \mathcal{A} . Now, $\mathcal{A}(\text{vk})$ may call a signing oracle $\text{SgnOracle}[\text{sgk}](\cdot)$ only once, where the oracle $\text{SgnOracle}[\text{sgk}](\cdot)$ on input \mathbf{m} returns $\sigma \xleftarrow{\$} \text{SIG.Sign}(\text{sgk}, \mathbf{m})$. Finally, $\mathcal{A}(\text{vk})$ should return a pair (\mathbf{m}', σ') of message/signature and will win if $(\mathbf{m}, \sigma) \neq (\mathbf{m}', \sigma')$ and that $\text{SIG.Ver}(\text{vk}, \mathbf{m}', \sigma') = \top$.

Our CCA2 primitive. We will build a CCA2 secure single-bit PKE, which by the result of [24] can be boosted into many-bit CCA2 secure PKE. Since we deal with single-bit CCA2 PKE, we may assume without loss of generality that the CCA adversary issues all her CCA oracles after seeing the challenge ciphertext.

We will now describe our CCA2-secure PKE scheme.

Construction 5 (CCA2 Secure PKE). *The construction is parameterized over a parameter $m := m(\lambda)$, which denotes the size of the verification key of the underlying signature scheme SIG. Let HC be a bit-valued hardcore function whose existence was proved in Lemma 6.*

- PKE.K(1^λ): For $i \in [m]$ and $b \in \{0, 1\}$, sample $(ik_i^b, tk_i^b) \xleftarrow{\$} \text{TDF.K}(1^\lambda)$. Form (pk, sk) the public/secret key as follows:

$$pk := (ik_1^0, ik_1^1, \dots, ik_m^0, ik_m^1), \quad sk := (tk_1^0, tk_1^1, \dots, tk_m^0, tk_m^1). \quad (25)$$

- PKE.E(pk, b): Parse pk as in Eq. 25. Sample $(vk, sgk) \xleftarrow{\$} \text{SIG.K}(1^\lambda)$, $x \xleftarrow{\$} \{0, 1\}^n$ and set

$$X_1 := (x, s_1 \xleftarrow{\$} \{0, 1\}^{n_1}), \quad \dots, \quad X_m := (x, s_m \xleftarrow{\$} \{0, 1\}^{n_1}). \quad (26)$$

Let $b' = b \oplus \text{HC}(x)$ and for $i \in [m]$ let $Y_i = \text{TDF.F}(ik_i^{b'}, X_i)$. Return

$$c := (vk, Y_1, \dots, Y_m, b', \text{Sign}(sgk, Y_1 || \dots || Y_m || b')). \quad (27)$$

- PKE.D(sk, c): Parse sk as in Eq. 25 and parse

$$c := (vk, Y_1, \dots, Y_m, b', \sigma). \quad (28)$$

- Set $msg := Y_1 || \dots || Y_m || b'$. If $\text{SIG.Ver}(vk, msg, \sigma) = \perp$, then return \perp .
- Otherwise, for $i \in [m]$ set $X_i := \text{TDF.F}^{-1}(tk_i^{b'}, Y_i)$. Check that for all $i \in [m]$: $Y_i = \text{TDF.F}(ik_i^{b'}, X_i)$. If not, return \perp .
- If there exists $x \in \{0, 1\}^n$ and $s_1, \dots, s_m \in \{0, 1\}^{n_1}$ such that for all $i \in [m]$, $X_i = (x, s_i)$, then return $b' \oplus \text{HC}(x)$. Otherwise, return \perp .

Correctness. If the underlying signature scheme $\text{SIG} = (\text{SIG.K}, \text{SIG.Sign}, \text{SIG.Ver})$ is correct and also that the underlying TDF $(\text{TDF.K}, \text{TDF.F}, \text{TDF.F}^{-1})$ is correct in the sense of Definition 1, the above constructed PKE is correct in a similar sense: for any $(pk, sk) \in \text{PKE.K}(1^\lambda)$ and plaintext bit $b \in \{0, 1\}$ we have $\Pr[\text{PKE.D}(sk, \text{PKE.E}(pk, b))] = \text{negl}(\lambda)$. The proof of this is straightforward.

6 Proof of CCA2 Security

We will prove the following theorem.

Theorem 6 (CCA2 security). *Let $(\text{TDF.K}, \text{TDF.F}, \text{TDF.F}^{-1})$ be the TDF that results from Construction 3 based on a recyclable OWFE (K, f, E, D) . Assuming (K, f, E, D) is adaptively secure, the PKE given in Construction 5 is CCA2 secure.*

We need to show that the probability of success of any CCA2 adversary is the CCA2 game is at most $\frac{1}{2} + \text{negl}(\lambda)$. Fix the adversary \mathcal{A} in the remainder of this section. We give the following event that describes exactly the success of \mathcal{A} .

Event Success. Let $(pk, sk) \xleftarrow{\$} \text{PKE.K}(1^\lambda)$, $b_{\text{plain}} \xleftarrow{\$} \{0, 1\}$, $c \xleftarrow{\$} \text{PKE.E}(pk, b_{\text{plain}})$. Run the adversary \mathcal{A} on (pk, c) and reply to any query $c' \neq c$ of \mathcal{A} with $\text{PKE.D}(sk, c')$. We say that the event **Success** holds if \mathcal{A} outputs b_{plain} .

Road Map. To prove Theorem 6, in Sect. 6.1 we define a simulated experiment *Sim* and we show that the probability of success of any CCA2 adversary in this experiment is $\frac{1}{2} + \text{negl}(\lambda)$. Next, in Sect. 6.2 we will show that the probabilities of success of any CCA2 adversary in the real and simulated experiments are negligibly close, establishing Theorem 6.

6.1 Simulated CCA2 Experiment

We now define a simulated way of doing the CCA2 experiment. Roughly, our simulator does not have the full secret key (needed to reply to CCA2 queries of the adversary), but some part of it. Our simulation is enabled a syntactic property of our constructed TDF. We first state the property and then prove that it is satisfied by our TDF. We require the existence of an efficient algorithm *Recover* for our constructed TDF ($\text{TDF.K}, \text{TDF.F}, \text{TDF.F}^{-1}$) that satisfies the following property.

Algorithm Recover: The input to the algorithm is an index key ik , a pre-fix input $x \in \{0, 1\}^n$ and a possible image Y . The output of the algorithm is $X \in \{0, 1\}^{n+n_1} \cup \{\perp\}$. As for correctness we require the following. For any $(ik, *) \in \text{TDF.K}(1^\lambda)$, $x \in \{0, 1\}^n$ and Y both the following two properties hold:

- if for no $s \in \{0, 1\}^{n_1}$ $\text{TDF.F}(ik, x||s) = Y$, then $\text{Recover}(ik, x, Y) = \perp$
- if for some s , $\text{TDF.F}(ik, x||s) = Y$, then $\text{Recover}(ik, x, Y)$ returns (x, s) .

Lemma 7 (Existence of Recover). *There exists an efficient algorithm Recover with the above properties for our constructed TDF .*

Proof. To build *Recover*, first parse the given inputs as follows: $ik = (pp, \dots)$, $x \in \{0, 1\}^n$ and $Y := (y, \widetilde{\mathbf{b}}_{1,0}, \widetilde{\mathbf{b}}_{1,1}, \dots, \widetilde{\mathbf{b}}_{n,0}, \widetilde{\mathbf{b}}_{n,1})$. Do the following steps:

1. For all $i \in [n]$ set $\mathbf{b}_i := \widetilde{\mathbf{b}}_{i,1-x_i}$.
2. Let $s = \mathbf{b}_1 || \dots || \mathbf{b}_n$
3. Check if $Y = \text{TDF.F}(ik, x||s)$. If the check holds, return (x, s) . Otherwise, return \perp .

The correctness of the algorithm *Recover* follows by inspection. □

Sim($ik_1, \dots, ik_m, Y_1, \dots, Y_m, \mathbf{b}$): The simulated experiment differs from the real experiment in that the challenger does not know the trapdoor keys of half of the index keys that are given to the CCA adversary as part of the public key. The challenger, however, tries to produce accurate answers based on her partial knowledge.

Formally, do the following steps.

1. **Initializing the CCA Adversary:**

- (a) Sample $(vk, \text{sgk}) \xleftarrow{\$} \text{SIG.K}(1^\lambda)$.
- (b) For all $i \in [m]$ set $ik_i^{\text{vk}_i} := ik_i$ and sample $(ik_i^{1-\text{vk}_i}, tk_i^{1-\text{vk}_i}) \xleftarrow{\$} \text{TDF.K}(1^\lambda)$.
- (c) Sample a challenge bit $b_{\text{plain}} \xleftarrow{\$} \{0, 1\}$ and let $b_1 := b_{\text{plain}} \oplus b$. Set

$$\text{msg} := Y_1 || \dots || Y_m || b_1.$$

Sample $\sigma \xleftarrow{\$} \text{SIG.Sign}(\text{sgk}, \text{msg})$ and set

$$\text{pk} := (ik_1^0, ik_1^1, \dots, ik_m^0, ik_m^1), c := (vk, Y_1, \dots, Y_m, b_1, \sigma).$$

Run the CCA2 adversary \mathcal{A} on (pk, c) .

2. **Simulating the CCA responses:**

Respond to a CCA2-oracle query $c' := (vk', Y'_1, \dots, Y'_m, b', \sigma')$ as follows:

- (a) Letting $\text{msg}' := Y'_1 || \dots || Y'_m || b'$ if $\text{Ver}(vk', \text{msg}', \sigma') = \perp$, then return \perp . Otherwise, if $vk' = vk$, then halt and return \perp .
- (b) Otherwise, let Q consist all of all indices $i \in [m]$ for which we have $vk'_i \neq vk_i$. For $i \in Q$ set $X'_i := \text{TDF.F}^{-1}(tk_i^{\text{vk}'_i}, Y'_i)$ and check if $Y'_i = \text{TDF.F}(ik_i^{\text{vk}'_i}, X'_i)$; if this fails for any $i \in Q$, then return \perp . Now if there exists $x' \in \{0, 1\}^n$ such that for all $i \in Q$ we have $X'_i = (x', *)$ then continue with the following steps and otherwise return \perp .
 - (I) for all $j \in [m] \setminus Q$, let $X'_j := \text{Recover}(ik_j^{\text{vk}'_j}, x', Y'_j)$. Reply to the query c' with $\text{HC}(x')$ if for all j we have $X'_j \neq \perp$; otherwise, reply to the query with \perp .

3. **Forming the output of the experiment:** The experiment outputs 1 if \mathcal{A} outputs b_{plain} ; otherwise, the experiments outputs 0.

Event $\text{Success}_{\text{sim}}$. The event that $\text{Sim}(ik_1, \dots, ik_m, Y_1, \dots, Y_m, b)$ outputs 1 where $x \xleftarrow{\$} \{0, 1\}^n$, $b := \text{HC}(x)$ and for $i \in [m]$, $(ik_i, tk_i) \xleftarrow{\$} \text{TDF.K}(1^\lambda)$, $s_i \xleftarrow{\$} \{0, 1\}^{nr}$ and $Y_i := \text{TDF.K}(ik_i, x || s_i)$.

We now show that the probability of the event $\text{Success}_{\text{sim}}$ is $\frac{1}{2} + \text{negl}(\lambda)$. We will then show in the next section that the probability of the event Success is close to that of $\text{Success}_{\text{sim}}$, hence obtaining our main result.

Lemma 8.

$$\alpha := \Pr[\text{Success}_{\text{sim}}] \leq \frac{1}{2} + \text{negl}(\lambda). \tag{29}$$

Proof. This lemma follows by Lemma 6. Suppose the input to Sim is sampled exactly as done in the event $\text{Success}_{\text{sim}}$, except that we sample $b \xleftarrow{\$} \{0, 1\}$ (instead of setting $b := \text{HC}(x)$). In this case the output of the simulation is 1 with probability 1/2. Now Lemma 6 implies that $\alpha = \frac{1}{2} + \text{negl}(\lambda)$ (Eq. 29), and the proof is complete. \square

6.2 Relating the Simulated and Real Experiments

We will now show that the probabilities of the events **Success** and $\text{Success}_{\text{sim}}$ are negligibly close, hence obtaining Theorem 6 from Lemma 8. To this end, we define below two events **Forge** and **Spoof**, and show that the difference of the probabilities of **Success** and $\text{Success}_{\text{sim}}$ is at most the sum of the probabilities of **Forge** and **Spoof**. We will then show that both these events happen with negligible probability.

Event Forge: In the experiment $\text{Sim}(\text{ik}_1, \dots, \text{ik}_m, Y_1, \dots, Y_m, \mathbf{b})$ we let **Forge** be the event that \mathcal{A} issues a CCA2 query

$$c' := (\text{vk}', Y'_1, \dots, Y'_m, \mathbf{b}', \sigma')$$

such that $\text{vk} = \text{vk}'$ and $\text{Ver}(\text{vk}, Y'_1 || \dots || Y'_m || \mathbf{b}', \sigma') = \top$. Recall that vk is part of the challenge ciphertext $c := (\text{vk}, \dots)$ given to the adversary.

Informally, the event **Spoof** below describes a situation in which a decryption query during the execution of Sim is answered with a non \perp string, but the same query may be replied to with \perp under the real decryption oracle.

Event Spoof: Let $c := (\text{vk}, \dots)$ be the challenge ciphertext formed during $\text{Sim}(\text{ik}_1, \dots, \text{ik}_m, Y_1, \dots, Y_m, \mathbf{b})$. We let **Spoof** be the event that \mathcal{A} issues a CCA2 query

$$c' := (\text{vk}', Y'_1, \dots, Y'_m, \mathbf{b}', \sigma')$$

for which the following holds. Let Q be the set of indices $i \in [m]$ for which we have $\text{vk}_i \neq \text{vk}'_i$. For some $h \in Q$ and for some $w \in [m] \setminus Q$ we have

- $\text{TDF.F}^{-1}(\text{tk}_h^{\text{vk}'_h}, Y'_h) = (x', *) \neq \perp$; and
- $s'_j := \text{Recover}(\text{ik}_w^{\text{vk}'_w}, x', Y'_w) \neq \perp$ but $\text{TDF.F}^{-1}(\text{tk}_w^{\text{vk}'_w}, Y'_w) = \perp$.

We will now prove the following three lemmas.

Lemma 9. *We have*

$$|\Pr[\text{Success}] - \Pr[\text{Success}_{\text{sim}}]| \leq \Pr[\text{Forge}] + \Pr[\text{Spoof}].$$

Lemma 10.

$$\Pr[\text{Forge}] \leq \text{negl}(\lambda).$$

Lemma 11.

$$\Pr[\text{Spoof}] \leq \text{negl}(\lambda).$$

Let us first derive the proof of Theorem 6 and then prove each of the lemmas.

Proof of Theorem 6. Follows from Lemmas 8 and 9, taking into account the fact that the experiment $\text{Real}(1^\lambda)$ is the real CCA2 experiment. \square

We prove Lemmas 9 and 10 and will prove Lemma 11 in Sect. 6.3.

Proof (of Lemma 9). First of all, note that the input (pk, c) given to the CCA2 adversary under the simulated experiment is identically distributed to that under the real CCA2 experiment. Thus, any possible difference between the simulated and real experiments must be due to the ways in which decryption queries are answered. We will now show that if at any point a decryption query is answered to differently under *Sim* and *Real*, then either *Forge* or *Spoof* must happen. First, in order for a query c' to be answered differently, either (1) the query c' is replied to with \perp under *Sim* and with some $b' \in \{0, 1\}$ under *Real*; or (2) c' is replied to with some $b' \in \{0, 1\}$ under *Sim* and with \perp under *Real*. In particular, we cannot have a situation in which c' is replied to with some $b' \in \{0, 1\}$ under *Sim* and with $1 - b'$ under *Real*. The reason for this is that if both experiments reply to a query with something other than \perp , then the underlying recovered pre-image x' must be the same, hence both will end up replying with the same bit.

Let $c := (vk, \dots)$ be the challenge ciphertext of the underlying CCA2 adversary and $c' := (vk', msg', \sigma')$ be an issued query. We now consider all possible cases:

- If *Sim* replies to $c' := (vk', msg', \sigma')$ with \perp , then one of the following must hold.
 - $SIG.Ver(vk', msg', \sigma') = \perp$: in this case *Real* also replies to with \perp ;
 - $SIG.Ver(vk', msg', \sigma') = \top$ and $vk = vk'$: in this case the event *Forge* happens;
 - *Sim* returns \perp as a result of Step 2b of the execution of *Sim*: that is, $TDF.F^{-1}(tk_i^{vk'_i}, Y'_i) = \perp$: in this case *Real* also replies with \perp
 - *Sim* replies with \perp as a result of Step (I): in this case by correctness of *Recover* we will know *Real* will also reply with \perp .
- If *Real* replies to c' with \perp and *Sim* replies to c' with some $b' \in \{0, 1\}$, then we may easily verify that the event *Spoof* must necessarily hold. We omit the details. □

Proof (of Lemma 10). Suppose $\Pr[\text{Forge}] > \text{negl}(\lambda)$. We show how to build an adversary \mathcal{B} against the one-time unforgeability of $SIG = (SIG.K, SIG.Sign, SIG.Ver)$. Build $\mathcal{B}^{SgnOracle[sgk](\cdot)}(vk)$ as follows. Sample the input $(ik_1, \dots, ik_m, Y_1, \dots, Y_m, b)$ to *Sim* and form the tuple msg as in the execution of *Sim* on this input. Then request a signature σ for the message msg by calling $SgnOracle[sgk](\cdot)$ on msg . Form (pk, c) as in *Sim* and run the CCA2 adversary \mathcal{A} on (pk, c) . Let q be the number queries that \mathcal{A} asks. Choose $i \stackrel{\$}{\leftarrow} [q]$ to be a guess for the index of the first query for which the event *Forge* occurs and output the pair of message/signature contained in that query. Note that \mathcal{B} can perfectly reply to all the previous $i - 1$ queries of \mathcal{A} , because all of those can be replied to without knowing sgk . If $\alpha := \Pr[\text{Forge}]$, then \mathcal{B} will win with probability at least $\frac{\alpha}{q}$. □

6.3 Proof of Lemma 11

The proof of Lemma 11 is based on a property of our TDF that we now state and prove. Informally, if the OWFE scheme used in our TDF construction is

adaptively secure, then the constructed TDF has the property that given a random index key ik , it is infeasible to produce an image element which is in the range of the trapdoor function $\text{TDF.F}(\text{ik}, \cdot)$, but which “inverts” to \perp .

Lemma 12. *Let $\text{TDF} = (\text{TDF.K}, \text{TDF.F}, \text{TDF.F}^{-1})$ be the TDF built in Sect. 4, with the underlying parameter $r := r(\lambda) \in \omega(\log \lambda)$, based on a recyclable OWFE scheme $\text{OWFE} = (\text{K}, \text{f}, \text{E}, \text{D})$. Assuming OWFE is adaptive (Definition 5), for any PPT adversary \mathcal{A} :*

$$\Pr[(\text{X}, \text{Y}) \stackrel{\$}{\leftarrow} \mathcal{A}(\text{ik}) \text{ s.t. } \text{Y} = \text{TDF.F}(\text{ik}, \text{X}), \text{TDF.F}^{-1}(\text{tk}, \text{Y}) = \perp] = \text{negl}(\lambda), \quad (30)$$

where $(\text{ik}, \text{tk}) \stackrel{\$}{\leftarrow} \text{TDF.K}(1^\lambda)$.

Proof. Let Surprise be the event of the lemma. Parse

$$\text{ik} := (\text{pp}, \mathbf{ct}_{1,0}, \mathbf{ct}_{1,1}, \dots, \mathbf{ct}_{n,0}, \mathbf{ct}_{n,1}), \text{tk} := (\rho_{1,0}, \rho_{1,1}, \dots, \rho_{n,0}, \rho_{n,1}), \quad (31)$$

and for all $i' \in [n]$ and $b' \in \{0, 1\}$ parse

$$\begin{aligned} \mathbf{ct}_{i',b'} &:= (\mathbf{ct}_{i',b'}^{(1)}, \dots, \mathbf{ct}_{i',b'}^{(r)}) \\ \rho_{i',b'} &:= (\rho_{i',b'}^{(1)}, \dots, \rho_{i',b'}^{(r)}). \end{aligned} \quad (32)$$

Recall that for all $i' \in [n]$, $b' \in \{0, 1\}$ and $j \in [r]$ we have

$$\mathbf{ct}_{i',b'}^{(j)} = \text{E}_1(\text{pp}, (i', b'); \rho_{i',b'}^{(j)}). \quad (33)$$

Also, parse (X, Y) , the output of $\mathcal{A}(\text{ik})$, as

$$\begin{aligned} \text{X} &:= (\mathbf{x} \in \{0, 1\}^n, \mathbf{b}_1 \in \{0, 1\}^r, \dots, \mathbf{b}_n \in \{0, 1\}^r) \\ \text{Y} &:= (\mathbf{y}, \mathbf{b}_{1,0} \in \{0, 1\}^r, \mathbf{b}_{1,1} \in \{0, 1\}^r, \dots, \mathbf{b}_{n,0} \in \{0, 1\}^r, \mathbf{b}_{n,1} \in \{0, 1\}^r). \end{aligned} \quad (34)$$

If the event Surprise happens, then by definition we have $\text{Y} = \text{TDF.F}(\text{ik}, \text{X})$ and $\text{TDF.F}^{-1}(\text{tk}, \text{Y}) = \perp$. Thus, by definition of TDF.F^{-1} , for some $i \in [n]$ we must have

$$\mathbf{b}_i = \left(\text{E}_2(\text{pp}, \mathbf{y}, (i, 1 - x_i); \rho_{i,1-x_i}^{(1)}), \dots, \text{E}_2(\text{pp}, \mathbf{y}, (i, 1 - x_i); \rho_{i,1-x_i}^{(r)}) \right). \quad (35)$$

We show how to use Eq. 35 to break the adaptive security of OWFE.

We show how to build an adversary against the adaptive security of OWFE in the sense of Lemma 1. Sample $i \stackrel{\$}{\leftarrow} [n]$ and $b \stackrel{\$}{\leftarrow} \{0, 1\}$ — The value of \mathbf{b} will serve as a guess bit for $1 - x_i$ (see Eq. 35). Give the pair (i, b) to the challenger to receive $(\text{pp}, \mathbf{ct}_1, \dots, \mathbf{ct}_r)$. Set $\mathbf{ct}_{i,b} := (\mathbf{ct}_1, \dots, \mathbf{ct}_r)$ and sample all other $\mathbf{ct}_{i',b'}$, for $(i', b') \neq (i, b)$, as in Eq. 32 and form the index ik as in Eq. 31. Call $\mathcal{A}(\text{ik})$ to receive (X, Y) and parse them as in Eq. 34. If $x_i = b$ then return \perp . Otherwise, give \mathbf{x} to the challenger to receive some $\mathbf{e}' \in \{0, 1\}^r$. If $\mathbf{e}' = \mathbf{b}_i$ then return 0, and otherwise return 1.

If the probability of the event *Surprise* is non-negligible, then \mathcal{A} wins with a probability non-negligibly greater than $\frac{1}{2}$. The reason is if \mathbf{e}' was generated uniformly at random from $\{0, 1\}^r$, then the probability that $\mathbf{e}' = \mathbf{b}_i$ is $\frac{1}{2^r} = \text{negl}(\lambda)$. On the other hand, if \mathbf{e}' was generated as a result of true encapsulation encryptions (see the description of the game in Lemma 1), then the probability that $\mathbf{e}' = \mathbf{b}_i$ is the probability of the event in Eq. 35, which is non-negligible. Thus, we break the adaptive security of OWFE, a contradiction to Lemma 1. \square

Proof (of Lemma 11). The proof of this lemma follows easily from Lemma 12, so we will give a sketch of the proof. Let $\beta := \Pr[\text{Spoof}]$. We show how to build an adversary \mathcal{B} in the sense of Lemma 12 that wins with probability $\frac{\beta}{\text{poly}(\lambda)}$.

Recall h and w from the event *Spoof*. The adversary $\mathcal{B}(\text{ik})$ acts as follows.

- Sample $(\text{vk}, \text{sgk}) \xleftarrow{\$} \text{SIG.K}(1^\lambda)$.
- Guess $w \xleftarrow{\$} [m]$
- Set $\text{ik}_w^{\text{vk}_w} := \text{ik}$. Also, sample $(\text{ik}_w^{1-\text{vk}_w}, \text{tk}_w^{1-\text{vk}_w}) \xleftarrow{\$} \text{TDF.K}(1^\lambda)$
- For all $i \in [m] \setminus \{w\}$ and $b \in \{0, 1\}$, sample $(\text{ik}_i^b, \text{tk}_w^b) \xleftarrow{\$} \text{TDF.K}(1^\lambda)$.
- Sample $\mathbf{x} \xleftarrow{\$} \{0, 1\}^n$, $\mathbf{b} := \text{HC}(\mathbf{x})$ and for $i \in [m]$, $Y_i := \text{TDF.K}(\text{ik}_i^{\text{vk}_i}, \mathbf{x} \parallel \mathbf{s}_i)$.
- Sample $\mathbf{b}_{\text{plain}} \xleftarrow{\$} \{0, 1\}$ and set $\mathbf{b}_1 := \mathbf{b}_{\text{plain}} \oplus \mathbf{b}$.
- Set the challenge public key and ciphertext (pk, \mathbf{c}) as in *Sim* and run the CCA2 adversary \mathcal{A} on (pk, \mathbf{c}) .

Now guess η to be the index of the first query of \mathcal{A} that causes *Spoof* to happen and guess $h \xleftarrow{\$} [m]$ be the underlying index defined in the event *Forge*. Note that \mathcal{B} can perfectly simulate the response all the first $\eta - 1$ queries of \mathcal{A} as in *Sim*. The reason is that \mathcal{B} has the trapdoor key for all $\text{ik}_i^{1-\text{vk}_i}$, and so it can perform as in *Sim*.

The η th query. Letting the η th query be

$$\mathbf{c}' := (\text{vk}', Y'_1, \dots, Y'_m, \mathbf{b}'_{\text{ciph}}, \sigma')$$

\mathcal{B} acts as follows: set $(\mathbf{x}', \mathbf{s}_h) := \text{TDF.F}^{-1}(\text{tk}_h^{\text{vk}'_h}, Y'_h)$ and $X'_w := \text{Recover}(\text{ik}, \mathbf{x}', Y'_w)$. Finally, \mathcal{B} returns (X'_w, Y'_w) .

It is now easy to verify if *Spoof* occurs and that all the guesses of the adversary \mathcal{B} were correct (i.e., the guessed values for h , w and vk'_w) — which happens with probability $\frac{1}{\text{poly}(\lambda)}$ — the adversary \mathcal{B} wins in the sense of Lemma 12. \square

Acknowledgments. We would like to thank the anonymous reviewers for their useful comments, and thank Mohammad Mahmoody and Adam O’Neill for useful discussions.

References

1. Bellare, M., Boldyreva, A., O’Neill, A.: Deterministic and efficiently searchable encryption. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 535–552. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74143-5_30
2. Bellare, M., Fischlin, M., O’Neill, A., Ristenpart, T.: Deterministic encryption: definitional equivalences and constructions without random oracles. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 360–378. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85174-5_20
3. Bellare, M., Halevi, S., Sahai, A., Vadhan, S.P.: 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>
4. Blum, M., Feldman, P., Micali, S.: Proving security against chosen ciphertext attacks. In: Goldwasser, S. (ed.) CRYPTO 1988. LNCS, vol. 403, pp. 256–268. Springer, New York (1990). https://doi.org/10.1007/0-387-34799-2_20
5. 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
6. Boneh, D., Boyen, X.: Secure identity based encryption without random oracles. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 443–459. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-28628-8_27
7. Brakerski, Z., Lombardi, A., Segev, G., Vaikuntanathan, V.: Anonymous IBE, leakage resilience and circular security from new assumptions. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part I. LNCS, vol. 10820, pp. 535–564. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78381-9_20
8. Canetti, R., Halevi, S., Katz, J.: Chosen-ciphertext security from identity-based encryption. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 207–222. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24676-3_13
9. Cash, D., Kiltz, E., Shoup, V.: The twin Diffie-Hellman problem and applications. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 127–145. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78967-3_8
10. Cho, C., Döttling, N., Garg, S., Gupta, D., Miao, P., Polychroniadou, A.: Laconic oblivious transfer and its applications. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part II. LNCS, vol. 10402, pp. 33–65. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63715-0_2
11. Diffie, W., Hellman, M.E.: New directions in cryptography. *IEEE Trans. Inf. Theory* **22**(6), 644–654 (1976)
12. Döttling, N., Garg, S.: From selective IBE to full IBE and selective HIBE. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017, Part I. LNCS, vol. 10677, pp. 372–408. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70500-2_13
13. Döttling, N., Garg, S.: Identity-based encryption from the Diffie-Hellman assumption. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part I. LNCS, vol. 10401, pp. 537–569. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63688-7_18
14. Döttling, N., Garg, S., Hajiabadi, M., Masny, D.: New constructions of identity-based and key-dependent message secure encryption schemes. In: Abdalla, M., Dahab, R. (eds.) PKC 2018, Part I. LNCS, vol. 10769, pp. 3–31. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-76578-5_1

15. ElGamal, T.: A public key cryptosystem and a signature scheme based on discrete logarithms. In: Blakley, G.R., Chaum, D. (eds.) CRYPTO 1984. LNCS, vol. 196, pp. 10–18. Springer, Heidelberg (1985). https://doi.org/10.1007/3-540-39568-7_2
16. Feige, U., Lapidot, D., Shamir, A.: Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In: 31st FOCS, St. Louis, Missouri, 22–24 October 1990, pp. 308–317. IEEE Computer Society Press (1990)
17. Freeman, D.M., Goldreich, O., Kiltz, E., Rosen, A., Segev, G.: More constructions of lossy and correlation-secure trapdoor functions. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 279–295. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13013-7_17
18. Gertner, Y., Malkin, T., Reingold, O.: On the impossibility of basing trapdoor functions on trapdoor predicates. In: 42nd FOCS, Las Vegas, NV, USA, 14–17 October 2001, pp. 126–135. IEEE Computer Society Press (2001)
19. Goldreich, O., Levin, L.A.: A hard-core predicate for all one-way functions. In: 21st ACM STOC, Seattle, WA, USA, 15–17 May 1989, pp. 25–32. ACM Press (1989)
20. Goldwasser, S., Micali, S.: Probabilistic encryption and how to play mental poker keeping secret all partial information. In: 14th ACM STOC, San Francisco, CA, USA, 5–7 May 1982, pp. 365–377. ACM Press (1982)
21. Hajiabadi, M., Kapron, B.M.: Reproducible circularly-secure bit encryption: applications and realizations. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015, Part I. LNCS, vol. 9215, pp. 224–243. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-47989-6_11
22. Haralambiev, K., Jager, T., Kiltz, E., Shoup, V.: Simple and efficient public-key encryption from computational Diffie-Hellman in the standard model. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 1–18. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13013-7_1
23. 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
24. Myers, S., Shelat, A.: Bit encryption is complete. In: 50th FOCS, Atlanta, GA, USA, 25–27 October 2009, pp. 607–616. IEEE Computer Society Press (2009)
25. Naor, M., Yung, M.: Public-key cryptosystems provably secure against chosen ciphertext attacks. In: 22nd ACM STOC, Baltimore, MD, USA, 14–16 May 1990, pp. 427–437. ACM Press (1990)
26. Peikert, C.: Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In: Mitzenmacher, M. (ed.) 41st ACM STOC, Bethesda, MD, USA, 31 May–2 June 2009, pp. 333–342. ACM Press (2009)
27. Peikert, C., Waters, B.: Lossy trapdoor functions and their applications. In: Ladner, R.E., Dwork, C. (eds.) 40th ACM STOC, Victoria, British Columbia, Canada, 17–20 May 2008, pp. 187–196. ACM Press (2008)
28. Rivest, R.L., Shamir, A., Adleman, L.M.: A method for obtaining digital signature and public-key cryptosystems. *Commun. Assoc. Comput. Mach.* **21**(2), 120–126 (1978)
29. 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
30. Wee, H.: Efficient chosen-ciphertext security via extractable hash proofs. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 314–332. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14623-7_17

31. 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
32. Yao, A.C.-C: Theory and applications of trapdoor functions (extended abstract). In: 23rd FOCS, pp. 80–91. IEEE Computer Society Press, Chicago, 3–5 November 1982