New Techniques for Efficient Trapdoor Functions and Applications[*]

Sanjam Garg[†]   Romain Gay[‡]   Mohammad Hajiabadi[§]

September 23, 2018

Abstract

We develop techniques for constructing trapdoor functions (TDFs) with short image size and advanced security properties. Our approach builds on the recent framework of Garg and Hajiabadi [CRYPTO 2018]. As applications of our techniques, we obtain

- The first construction of lossy TDFs based on the Decisional Diffie-Hellman (DDH) assumption with image size linear in input size, while retaining the lossiness rate of [Peikert-Waters STOC 2008].

- The first construction of deterministic-encryption schemes for block-source inputs (both for the CPA and CCA cases) based on the Computational Diffie-Hellman (CDH) assumption. Moreover, by applying our efficiency-enhancing techniques, we obtain CDH-based schemes with ciphertext size linear in plaintext size.

Prior to our work, all DDH-based constructions of lossy TDFs had image size quadratic in input size. Moreover, all previous constructions of deterministic encryption based even on the stronger DDH assumption incurred a quadratic gap between the ciphertext and plaintext sizes. At a high level, we break the previous quadratic barriers by introducing a novel technique for encoding input bits via hardcore output bits with the use of erasure-resilient codes. All previous schemes used group elements for encoding input bits, resulting in quadratic blowup.

1 Introduction

Trapdoor functions (TDFs) are a classical primitive in cryptography and are typically used as a fundamental building block in the construction of advanced primitives such as CCA2-secure public-key encryption (PKE). Introduced in the 70's [DH76, RSA78], TDFs are a family of functions, where each individual function in the family is easy to compute, and also easy to invert if one posses an additional trapdoor key. The basic security requirement is that of one-wayness, requiring that a randomly chosen function from the family be one-way.

The usefulness of TDFs stems from the fact that the inversion algorithm recovers the entire input. This stands in sharp contrast to PKE, wherein the decryption algorithm may not recover

---

[†]University of California, Berkeley.

[‡]ENS, Paris, France. Supported in part by a Google Fellowship.

[§]University of California, Berkeley and University of Virginia.

the underlying randomness. This input recovery feature of TDFs is what makes them a useful tool, especially in applications where proofs of well-formedness are required.

On the other hand, building TDFs turns out to be much more difficult than building PKE, mostly due to the requirement of recovering the entire input, which in turn is the reason behind the lack of black-box transformations from PKE to TDFs [GMR01]. Specifically, in groups with discrete-log based hardness assumptions, this restricts the use of operations such as *exponentiation*, for which we do not have a generic trapdoor. Furthermore, in some applications we need TDFs to be *robust*, providing enhanced security properties rather than mere one-wayness (e.g., [BBO07, PW08, PVW08, BFOR08, RS09]).

Recently, Garg and Hajiabadi [GH18] introduced a new approach for building TDFs, obtaining the first construction of TDFs from the Computational Diffie-Hellman (CDH) assumption. Although their approach gives new feasibility results, their constructed TDFs are limited in certain ways: (a) Their TDFs are not robust enough — for example, it is not clear how to go beyond one-wayness, obtaining more advanced properties such as those required by deterministic encryption [BBO07, BFOR08, BFO08] or CCA2 security; and (b) The length of their TDF images grows (at least) quadratically with the length of the input.

We stress that Point (b) is not just an artifact of the construction of [GH18]. In fact, we do not know of any TDF constructions (even based on the stronger decisional Diffie-Hellman (DDH) assumption) with advanced properties, such as deterministic-encryption security, with images growing linearly in their inputs.[1] Since TDFs are typically used as building blocks in more advanced primitives, designing more efficient TDFs translates into the same features in target applications. For example, lossy TDFs [PW08] are an extremely versatile primitive with a long list of applications e.g., [BFOR08, BHY09, BBN+09, MY10, BCPT13].

1.1 Our Results

We develop techniques for constructing efficient and robust TDFs. As concrete applications of our new techniques, we obtain the first construction of deterministic encryption for block sources (in the sense of [BFO08]) under the CDH assumption. We give both CPA and CCA2 versions of our constructions. We stress that prior to our work we knew how to build (even) CPA-secure deterministic encryption only from decisional assumptions, including DDH, QR, DCR and LWE [BFO08, Wee12]. Thus, in addition, we also obtain instantiations under the hardness of factoring assumption.

Furthermore, we show how to use our efficiency techniques to obtain (1) a construction of lossy TDFs ([PW08]) from DDH with image size linear in input size and (2) CDH-based deterministic encryption schemes with ciphertext size linear in plaintext size. Our DDH-based lossy TDFs achieve the same lossiness rate as in [PW08]. All previous DDH-based lossy TDF constructions (achieving non-trivial lossiness rates) resulted in images quadratically large in their inputs. Moreover, our CDH-based deterministic-encryption schemes beat all the previous DDH-based schemes in terms of ciphertext size. The sizes of other parameters (e.g., the secret key and public key) remain the same. See Table 1 for a comparison.

---

[1]We note that building a TDF providing mere one-wayness with linear-size images is simple: if $\mathsf{TDF.F}(\mathsf{ik}, \cdot)$ maps n-bit inputs to n^c-bit outputs, define $\mathsf{TDF.F}'(\mathsf{ik}, \mathsf{x}||\mathsf{x}')$, where $|\mathsf{x}| = n$ and $|\mathsf{x}'| = n^c$, as $\mathsf{TDF.F}(\mathsf{ik}, \mathsf{x})||\mathsf{x}'$. Although this transformation results in TDFs with linear-image size, it destroy more advanced properties such as CCA2 security, deterministic-encryption security and the lossiness rate.

work	assumption	primitive	index key	trapdoor key	image
ours	CDH	CCA2 DE	$\Theta(n^2 \log p)$	$\Theta(n^2 \log p)$	$\log p + 10n$
[**BFO08**]	DDH	CCA2 DE	$\Theta(n^2 \log p)$	$\Theta(n^2 \log p)$	$\Theta(n \log p)$
ours	DDH	$(n, \log p)$-LTDF	$\Theta(n^2 \log p)$	$\Theta(n^2 \log p)$	$\log p + 13n$
[**PW08**, **FGK$^+$10**]	DDH	$(n, \log p)$-LTDF	$\Theta(n^2 \log p)$	$\Theta(n^2 \log p)$	$\Theta(n \log p)$

Table 1: Bit complexity: p is the order of the group and n is the bit size of the TDF input. Here (n, k)-LTDF means lossy TDFs where in lossy mode the image-space size is at most 2^k. We call $1 - k/n$ the lossiness rate.

1.2 Technical Overview

In this section we give an overview of our techniques for constructing robust and efficient TDFs. We will build TDFs with several abstract properties, and we will apply these techniques to the setting of deterministic encryption and lossy TDFs as concrete applications.

Our constructions rely on the same primitive of *recyclable one-way function with encryption (OWFE)* used by [GH18], so we first review this notion. An OWFE consists of a one-way function $f(pp, \cdot) \colon \{0,1\}^n \to \{0,1\}^\nu$, where pp is a public parameter, along with encapsulation/decapsulation algorithms (E, D). Specifically, E takes as input pp, an image $y \in \{0,1\}^\nu$ of $f(pp, \cdot)$, a target index $i \in [n]$ and a target bit $b \in \{0,1\}$, and produces an *encapsulated ciphertext* ct and a corresponding *key bit* $e \in \{0,1\}$. The algorithm D allows us to retrieve e from ct using any pre-image x of y, if $x_i = b$. For security, letting $y := f(pp, x)$, we require that if $(ct, e) \overset{\$}{\leftarrow} E(pp, y, (i, 1 - x_i))$, then even knowing both x and ct, one cannot distinguish e from a truly random bit. Finally, letting E_1 and E_2 refer to the first and second output pars of E, the recyclability requirement says that the output of E_1 does not depend on y, namely, we have: $ct = E_1(pp, (i, b))$ and $e = E_2(pp, y, (i, b))$. (See Definition 3.1.) The work of [GH18] gives CDH instantiations of this notion; see Section 3.1.

Approach of [GH18]. A property implied by recyclable OWFE is the following: given $x \in \{0,1\}^n$ and two fresh encapsulated ciphertexts (ct_0, ct_1) made w.r.t. $y := f(pp, x)$ and an arbitrary target index i and target bits 0 and 1 (respectively), one cannot distinguish the values of the corresponding two key bits (e_0, e_1) from a pair in which we replace e_{1-x_i} with a random bit. Exploiting this property, [GH18] set their index key to contain encapsulated ciphertexts $ct_{i,b}$ made w.r.t. each value of $i \in [n]$ and $b \in \{0,1\}$ — they put all the corresponding randomness values $[r_{i,b}]$ in the trapdoor key. The input to their TDF contains $x \in \{0,1\}^n$ and the output u consists of $y := f(pp, x)$ as well as a $2 \times n$ matrix M of bits $(e_{i,b})_{i \in [n], b \in \{0,1\}}$, where for all i, they set $e_{i,x_i} := D(pp, x, ct_{i,x_i})$ and set $e_{i,1-x_i}$ to be a random bit. Since TDFs are not allowed to make use of randomness, they draw $e_{i,1-x_i}$ for all i from an additional part of their input which they call the *blinding* part. For inverting $u := (y, M)$, the inverter may make use of its knowledge of all the randomness values underlying $ct_{i,b}$'s to form the corresponding key bits w.r.t. y. Then the inverter may check each column of the resulting matrix, M', against the corresponding column of the matrix M, and look for a matched coordinate. This would enable recovering half of the input bits (on average). The one-wayness of their scheme follows by the property alluded to above. Namely, for any $i \in [n]$, we may switch e_{1-x_i} from uniformly random to $E_2(pp, y, (i, 1 - x_i); r_{i,1-x_i})$. Consequently,

the image of the trapdoor function becomes: $(\mathsf{y}, (\mathsf{e}_{i,b})_{i,b})$, where $\mathsf{e}_{i,b} := \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (i,b); r_{i,b})$ for all $i \in [n]$ and $b \in \{0,1\}$. In other words, the entire view of a TDF adversary may be computed from y alone. At this point, the one-wayness of the TDF follows from the one-wayness of the underlying OWFE. Finally, [GH18] boosts correctness by repeating the above process in parallel. For future reference, we call the above initial TDF (which enables the recovery of half of the bits) *TDF gadget*.

Lack of perfect correctness in [GH18]. The TDF of [GH18] only achieves a weak form of correctness, under which the inversion algorithm may fail w.r.t. *any* index/trapdoor keys for a negligible fraction of the inputs. This severely restricts the applicability of CCA2-enhancing techniques, such as those of [RS09, KMO10], for obtaining CCA2 secure primitives. Even for the CPA case, the lack of perfect correctness hindered the construction of CPA-secure deterministic encryption schemes. Deterministic public-key encryption schemes [BBO07] are TDFs which hide information about plaintexts drawn from high min-entropy sources. There are various forms of this definition, e.g., [BBO07, BFO08, BFOR08, BS11, MPRS12]. Strong versions of this notion have so far been realized in the random oracle model [BBO07] and are subject to impossibility results [Wic13]. Boldyreva, Fehr and O'Neill [BFO08] formulated a relaxation of this notion (called *block-source security*), and showed how to realize this relaxed notion under standard assumptions such as DDH and LWE. Informally, block-source security requires that the (deterministic) encryptions of any two sources with high min entropy (more than a threshold k) remain computationally indistinguishable.[2] Ideally, we want $k << n$, where n is plaintext size.

The TDF of [GH18] does not achieve block-source security for the same reason that degraded their correctness property: The TDF input contains a blinding part, which in turn is copied in the clear in the output (but in hidden spots). To see how this breaks security, consider two sources, where the first one fixes the blinding part to all zeros, and the second one fixes it to all ones. Then it would be easy to distinguish between the outputs of the TDF w.r.t. these two sources, even though they may have high min entropy.

Enhancing to perfect correctness. We fix the imperfect correctness of [GH18] via a *mirroring* technique. Recall that the bits of the blinding input were previously used to form the values of $\mathsf{e}_{i,1-\mathsf{x}_i}$. Now, instead of having a blinding part in the input for making up the values of $\mathsf{e}_{i,1-\mathsf{x}_i}$, we set $\mathsf{e}_{i,1-\mathsf{x}_i} := \mathsf{a}_i - \mathsf{e}_{i,\mathsf{x}_i}$, where $\mathsf{a} := (\mathsf{a}_1, \ldots, \mathsf{a}_n) \in \{0,1\}^n$ is a random vector that comes from the index key. This way we get rid of inclusion of blinders as part of the input — the input now solely consists of a string $\mathsf{x} \in \{0,1\}^n$. We show that this method improves correctness: Our TDF is now perfectly correct for all but a negligible fraction of index/trapdoor keys; see Remark 2.2.

Lossy-like properties of our TDF toward obtaining deterministic encryption. So far, we showed how to fix the imperfect-correctness problem of [GH18], but this by itself does not guarantee deterministic-encryption security. Toward this goal, we show that the mirroring technique allows us to establish a *lossy-like* property for our TDFs, which in turn gives us block-source security.[3] Specifically, let y be an image point of $\mathsf{f}(\mathsf{pp}, \cdot)$ of the OWFE scheme, and let S be the set of all pre-images of y (which can be of exponential size under our CDH instantiation). We can now set

---

[2]This is the indistinguishability-based, single-message version of their notion, which as they show, is equivalent to the multiple-message version both for the indistinguishability- and simulation-based definitions.

[3]We note that this lossiness property is weaker than the one of [PW08], but it can be realized under CDH. We will later show efficient DDH-based instantiations of lossiness in the sense of [PW08].

the index key as $\mathsf{ik_y}$, where (a) $\mathsf{ik_y}$ loses information w.r.t. all pre-images of y: for all $\mathsf{x}, \mathsf{x}' \in \mathsf{S}$ we have $\mathsf{TDF.F(ik_y, x) = TDF.F(ik_y, x')}$ and (b) $\mathsf{ik_y}$ is computationally indistinguishable from an honestly generated ik. We exploit this property to prove block-source security for our TDFs.

Having achieved block-source CPA security, we may boost this scheme into a CCA2-secure deterministic-encryption scheme using the techniques of [RS09, KMO10].[4] Specifically, we show how to use our lossiness property to prove *k-repetition* security (introduced by [RS09]) for our TDF. Intuitively, k-repetition security requires one-wayness to hold even if the given input is evaluated under k-randomly chosen functions. The scheme of [GH18] fails to achieve k-repetition security, exactly because of the presence of blinders.

Finally, we mention that based on CDH we do not get lossiness in the sense of [PW08] as the amount of information we lose is negligible over the entire input space. Nevertheless, our weak lossiness property which can be realized under CDH suffices for our deterministic-encryption application, and may find other applications later.

Efficiency of our TDFs so far: quadratically large images. Under CDH instantiations of the above approach, for plaintexts of n bits, the bit-size of the ciphertext is $\Theta(n^2)$ in the CPA case, and $\Theta(n^2 \omega(\log n))$ in the CCA case. In contrast, the DDH-based constructions of [BFO08] give ciphertext size $\Theta(n^2)$ both for the CPA and CCA cases.

Sources of inefficiency. Recall our TDF gadget has image size $\Theta(n)$. This TDF gadget may fail to recover any given bit of the input with probability $1/2$. Thus, we ran many TDF gadgets in parallel, resulting in $\Theta(n^2)$ image size. We refer to this as *correctness repetition*. For the CCA2 case, since we relied on techniques of [RS09, KMO10] we needed to perform yet another repetition, which we call *CCA2 repetition*. This justifies the further blowup in CCA2 image size.

We develop techniques for avoiding both these repetitions, sketched below.

Erasure-resilient codes to the rescue: linear-image TDFs. We give techniques involving the use of erasure-resilient codes for making the size of our TDF images linear, while preserving other properties. Recall that under our TDF gadget, for a randomly chosen input $\mathsf{x} \in \{0, 1\}^n$ and for any index $i \in [n]$, the inversion algorithm either recovers x_i correctly, or outputs $\bot$ for this bit position (with probability $1/2$). Notice that the inversion process has a *local property*, in the sense that each bit individually may be recovered or not with probability $1/2$.

Now instead of performing parallel repetition which results in a quadratic-size blowup, we boost correctness through the use of erasure-resilient codes. Suppose (Encode, Decode) is an erasure-resilient code, where $\mathsf{Encode}\colon \{0, 1\}^n \to \{0, 1\}^m$ (for $m = cn \in O(n)$), and where Decode only needs n (noise-free) bits of a codeword $\mathsf{Encode(x)}$ in order to recover x. Such codes may be built from Reed-Solomon codes by adapting them to bitstrings; see Definition 5.1.

Now we avoid repetition as follows: on input $\mathsf{x} \in \{0, 1\}^n$, apply the TDF gadget on the encoded input $\mathsf{z} := \mathsf{Encode(x)}$. To invert, we no longer need to recover *all* the m bits of z; recovering n of them will do. By choosing c in $m = cn$ appropriately and using the mirroring idea, we can ensure perfect correctness for all but a negligible fraction of index/trapdoor keys. This way, images will grow linearly. The proof of CPA block-source security follows almost similarly as before.

---

[4]We mention that the transformation of [RS09] results in CCA-secure PKE schemes which use randomness, but this can be avoided by using the techniques of [KMO10] to get CCA2-secure TDFs.

Concretely, under CDH instantiations, CPA ciphertexts (for plaintexts of n bits) consist of one group element and (at most) $5n$ bits. This substantially improves the ciphertext size of previous DDH-based schemes under which a ciphertext consists of n group elements.

Keeping image size linear in CCA-like applications. So far, we showed how to build linear-image TDFs with additional properties (e.g., block-source CPA security). Typically, TDFs with enhanced properties (such as k-repetition security [RS09] or lossy properties [PW08]) can be boosted into CCA2 primitives, but this requires "parallel repetition" of the base object, increasing the sizes. Our linear-image TDF turns out to be k-repetition secure, but we cannot afford to use previous techniques for getting CCA2 security, because we would like to keep image size linear. Here is where our other techniques come into play: We develop a half-simulation method for proving CCA2 security for our *same* TDF scheme without any further modifications. For this, we just need to choose the constant c in $m = cn$ big enough, e.g., $c = 10$. Our CCA techniques are different from those of [PW08, RS09], which implicitly or explicitly relied on repetition.

As an application, we will get a CDH-based block-source CCA-secure deterministic encryption, beating the ciphertext size of DDH-based schemes. We now sketch our techniques.

Recall that a trapdoor key tk consists of all randomness values $r_{i,b}$'s used to form $\mathsf{ct}_{i,b}$'s (which are in turn fixed in the index key ik). On input $\mathsf{x} \in \{0,1\}^n$ we form $\mathsf{z} := \mathsf{Encode}(\mathsf{x}) \in \{0,1\}^m$ and return $\mathsf{u} := (\mathsf{y}, \mathbf{M} := \left(\begin{smallmatrix} \mathsf{e}_{1,0},\ldots,\mathsf{e}_{m,0} \\ \mathsf{e}_{1,1},\ldots,\mathsf{e}_{m,1} \end{smallmatrix}\right))$, where $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{z})$, $\mathsf{e}_{i,\mathsf{z}_i} = \mathsf{D}(\mathsf{pp}, \mathsf{z}, \mathsf{ct}_{i,\mathsf{z}_i})$ and $\mathsf{e}_{i,1-\mathsf{z}_i} = \mathsf{a}_i - \mathsf{e}_{i,\mathsf{z}_i}$, where $\mathbf{a} := (\mathsf{a}_1, \ldots, \mathsf{a}_m) \in \{0,1\}^m$ is sampled in ik. The inversion algorithm will recover the ith bit of z iff $\mathsf{e}_{i,1-\mathsf{z}_i} = 1 - \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (i, 1-\mathsf{z}_i); r_{i,1-\mathsf{z}_i})$. Say the ith column of $\mathbf{M}$ is *hung* if $\mathsf{e}_{i,1-\mathsf{z}_i} = \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (i, 1-\mathsf{z}_i); r_{i,1-\mathsf{z}_i})$ — if this happens, then the inverter cannot decide on the ith bit of z.

Let us argue CCA2 security w.r.t. two sources $\mathcal{S}_0$ and $\mathcal{S}_1$: The adversary should distinguish $(\mathsf{ik}, \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x}_0))$ from $(\mathsf{ik}, \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x}_1))$, where $\mathsf{x}_b \overset{\$}{\leftarrow} \mathcal{S}_b$. For deterministic encryption we may assume all CCA queries happen after seeing the index key and challenge ciphertext (Definition 2.3).

Our CCA2 simulation is based on a *half-trapdoor simulation* technique under which we forget one randomness value from each pair $(r_{i,0}, r_{i,1})$ in the trapdoor. Specifically, letting $\mathsf{x}^\star$ be the challenge plaintext, imagine a half-trapdoor key obtained based on $\mathsf{x}^\star$ from tk as $\mathsf{tk}_{\mathrm{rd},\mathsf{x}^\star} := (r_{1,\mathsf{z}_1^\star}, \ldots, r_{m,\mathsf{z}_m^\star})$, where $\mathsf{z}^\star = \mathsf{Encode}(\mathsf{x}^\star)$. We perform half-trapdoor inversion of a given point $\mathsf{u} := (\mathsf{y}, \left(\begin{smallmatrix} \mathsf{e}_{1,0},\ldots,\mathsf{e}_{m,0} \\ \mathsf{e}_{1,1},\ldots,\mathsf{e}_{m,1} \end{smallmatrix}\right))$ w.r.t. $\mathsf{tk}_{\mathrm{rd},\mathsf{x}^\star}$ as follows: Build (potentially) a codeword z as follows: if $\mathsf{e}_{i,\mathsf{z}_i^\star} = 1 - \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (i, \mathsf{z}_i^\star); r_{i,\mathsf{z}_i^\star})$, set $\mathsf{z}_i = 1 - \mathsf{z}_i^\star$; otherwise, set $\mathsf{z}_i = \bot$. Once z is formed, decode it to get a string and check if the re-encryption of that string gives back u. If so, return the string.

Letting $\mathsf{x}^\star$ be the challenge plaintext (and recalling that all CCA queries are post-challenge), we first show we may use $\mathsf{tk}_{\mathrm{rd},\mathsf{x}^\star}$ (instead of the full key tk) to reply to the CCA queries, without the adversary noticing any difference, using the following two facts. First, for a queried point u, if u is not a valid image (i.e., it does not have a pre-image), then both (full and half) inversions return $\bot$. This is because at the end of either inversion we re-encrypt the result to see whether we get the given image point back. So suppose for the queried $\mathsf{u} := (\mathsf{y}, \mathbf{M} := \left(\begin{smallmatrix} \mathsf{e}_{1,0},\ldots,\mathsf{e}_{m,0} \\ \mathsf{e}_{1,1},\ldots,\mathsf{e}_{m,1} \end{smallmatrix}\right))$ we have $\mathsf{u} := \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x})$ for some $\mathsf{x} \in \{0,1\}^n \setminus \{\mathsf{x}^\star\}$. (If $\mathsf{x} = \mathsf{x}^\star$, then u will be the challenge ciphertext itself and hence not a permitted query.) We show that by choosing c big enough, then for sufficiently-many indices $i \in [m]$, the ith column of $\mathbf{M}$ is *not* hung; namely, $\mathsf{e}_{i,1-\mathsf{z}_i} = 1 - \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (i, 1-\mathsf{z}_i); r_{i,1-\mathsf{z}_i})$, where $\mathsf{z} := \mathsf{Encode}(\mathsf{x})$. Concretely, letting S be the set of indices of such non-hung columns, we can ensure $|\mathsf{S}| \geq 2n$. Moreover, for any index $i \in \mathsf{S}$ such that $\mathsf{z}_i^\star = 1 - \mathsf{z}_i$, the half-inversion process (w.r.t. $\mathsf{tk}_{\mathrm{rd},\mathsf{x}^\star}$) will recover the ith bit of z (by definition). Since z and $\mathsf{z}^\star$ agree on at most n bits,[5] the

---

[5]This can be guaranteed if the code has distance at least $m - n$ (i.e., for any two codewords z and $\mathsf{z}^\star$: $\mathsf{H}_{\mathrm{dst}}(\mathsf{z}, \mathsf{z}^\star) \geq$

half-trapdoor inversion can recover at least n bits of z, hence recovering x.

Once we solely use $\mathsf{tk}_{\mathrm{rd},\mathsf{x}^\star}$ to reply to decryption queries, letting $\mathsf{u}^\star := \left(\mathsf{y}^\star, \begin{pmatrix} \mathsf{e}_{1,0}^\star, \ldots, \mathsf{e}_{m,0}^\star \\ \mathsf{e}_{1,1}^\star, \ldots, \mathsf{e}_{m,1}^\star \end{pmatrix}\right)$ be the corresponding challenge ciphertext, we may replace each $\mathsf{e}_{i,1-\mathsf{z}_i^\star}^\star$ with $\mathsf{E}_2(\mathsf{pp}, \mathsf{y}^\star, (i, 1 - \mathsf{z}_i^\star); r_{i,1-\mathsf{z}_i^\star})$, and simultaneously set the ith bit of the vector $\mathbf{a}$ of the index key as $\mathsf{a}_i := \mathsf{E}_2(\mathsf{pp}, \mathsf{y}^\star, (i, \mathsf{z}_i^\star); r_{i,\mathsf{z}_i^\star}) + \mathsf{E}_2(\mathsf{pp}, \mathsf{y}^\star, (i, 1 - \mathsf{z}_i^\star); r_{i,1-\mathsf{z}_i^\star})$. This change goes unnoticed by the security of the OWFE. At this point the challenge ciphertext and index key only depend on $\mathsf{y}^\star$ and we only use $\mathsf{z}^\star$ to decide which randomness value from each pair of tk to forget. We will now switch back to using the full trapdoor, with analysis similar to before. At this point, the entire view of the adversary may be simulated using $\mathsf{y}^\star := \mathsf{f}(\mathsf{pp}, \mathsf{z}^\star)$, and thus we have block-source security in this hybrid similar to the CPA case.

Lossy TDFs. Recall that a TDF is lossy [PW08] if one may generate index keys in a *lossy* way which is (a) indistinguishable from honestly generated index keys and (2) which results in statistical loss of information if used during the evaluation algorithm. We will show that our erasure-resilient based technique for encoding input bits via hardcore output bits allows us to obtain lossy TDFs based on DDH with image size linear in input size. All previous DDH-based constructions of lossy TDFs incur a quadratic blowup in image size [PW08, FGK$^+$10, Wee12]. See Section 6.

Other related work. OWFE is a relaxation of the notion of (chameleon) hash encryption and its variants, which in turn imply strong primitives such as laconic oblivious transfer and identity-based encryption (IBE) in a non-black-box way [CDG$^+$17, DG17b, DG17a, BLSV18, DGHM18].

Freeman et al. [FGK$^+$10] give additional constructions and simplifications to the TDF construction of [PW08]. Further constructions of (lossy) TDFs from various assumptions are given in [Wee12, HO12, HO13]. As for efficient TDFs, Boyen and Waters show that in the bilinear setting one may drastically shorten the index-key size of the Peikert-Waters lossy-TDF construction from a quadratic number of group elements to linear [BW10].

Paper organization. The rest of the paper is organized as follows. We give standard definitions and lemmas in Section 2 and OWFE-related definitions in Section 3. We give our (inefficient) construction of TDFs with deterministic-encryption security in Section 4 and give our efficient construction in Section 5. Finally, we give our DDH-based lossy TDF construction with linear image size in Section 6.

2 Preliminaries

Notation. We use λ for the security parameter. We use $\overset{c}{\equiv}$ to denote computational indistinguishability between two distributions and use $\equiv$ to denote two distributions are identical. For any $\varepsilon > 0$, we write $\approx_\varepsilon$ to denote that two distributions are statistically close, within statistical distance ε, and use $\overset{s}{\equiv}$ for statistical indistinguishability. For a distribution $\mathcal{S}$ we use $x \overset{\$}{\leftarrow} \mathcal{S}$ to mean x is sampled according to $\mathcal{S}$ and use $y \in \mathcal{S}$ to mean $y \in \sup(\mathcal{S})$, where sup denotes the support of a distribution. For a set S we overload the notation to use $x \overset{\$}{\leftarrow} \mathsf{S}$ to indicate that x is chosen uniformly at random from S. If $\mathsf{A}(x_1, \ldots, x_n)$ is a randomized algorithm, then $\mathsf{A}(a_1, \ldots, a_n)$, for

---

$m - n$, where $\mathsf{H}_{\mathrm{dst}}$ denotes hamming distance) — which can be instantiated using Reed-Solomon codes.

deterministic inputs $a_1, \ldots, a_n$, denotes the random variable obtained by sampling random coins r uniformly at random and returning $\mathsf{A}(a_1, \ldots, a_n; r)$.

The min-entropy of a distribution $\mathcal{S}$ is defined as $\mathsf{H}_\infty(\mathcal{S}) \triangleq -\log(\max_{\mathsf{x}} \Pr[\mathcal{S} = \mathsf{x}])$. We call a distribution $\mathcal{S}$ a (k, n)-source if $\mathsf{H}_\infty(\mathcal{S}) \geq k$ and $\sup(\mathcal{S}) \subseteq \{0, 1\}^n$.

2.1 Standard Definitions

Definition 2.1 (Trapdoor functions (TDFs)). *Let $n = n(\lambda)$ be a polynomial. A family of trapdoor functions* TDF *with domain $\{0, 1\}^n$ consists of three PPT algorithms* $\mathsf{TDF.KG}$, $\mathsf{TDF.F}$ *and* $\mathsf{TDF.F}^{-1}$ *with the following syntax and security properties.*

- $\mathsf{TDF.KG}(1^\lambda)$: *Takes the security parameter 1^λ and outputs a pair $(\mathsf{ik}, \mathsf{tk})$ of index/trapdoor keys.*

- $\mathsf{TDF.F}(\mathsf{ik}, \mathsf{x})$: *Takes an index key ik and a domain element $\mathsf{x} \in \{0, 1\}^n$ and deterministically outputs an image element u.*

- $\mathsf{TDF.F}^{-1}(\mathsf{tk}, \mathsf{u})$: *Takes a trapdoor key tk and an image element u and outputs a value $\mathsf{x} \in \{0, 1\}^n \cup \{\bot\}$.*

We require the following properties.

- **Correctness:**

$$\Pr_{(\mathsf{ik}, \mathsf{tk})}[\exists \mathsf{x} \in \{0, 1\}^n \ s.t. \ \mathsf{TDF.F}^{-1}(\mathsf{tk}, \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x})) \neq \mathsf{x}] = \mathsf{negl}(\lambda), \tag{1}$$

where the probability is taken over $(\mathsf{ik}, \mathsf{tk}) \xleftarrow{\$} \mathsf{TDF.KG}(1^\lambda)$.

- **One-wayness:** *For any PPT adversary $\mathcal{A}$, we have $\Pr[\mathcal{A}(\mathsf{ik}, \mathsf{u}) = \mathsf{x}] = \mathsf{negl}(\lambda)$, where $(\mathsf{ik}, \mathsf{tk}) \xleftarrow{\$} \mathsf{TDF.KG}(1^\lambda)$, $\mathsf{x} \xleftarrow{\$} \{0, 1\}^n$ and $\mathsf{u} := \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x})$.*

Remark 2.2. *The work of Garg and Hajiabadi [GH18] builds a TDF with a weaker correctness guarantee, under which for any choice of $(\mathsf{ik}, \mathsf{tk})$, we are allowed to have a negligible inversion error (over the choice of $\mathsf{x} \xleftarrow{\$} \{0, 1\}^n$). Although the correctness condition of [GH18] implies that for a randomly chosen $(\mathsf{ik}, \mathsf{tk})$ and a randomly chosen x, the probability of an inversion error is negligible, it falls short in certain applications, such as CCA2 constructions, for which a stronger correctness condition, as that given in Definition 2.1, is needed.*

We will now define a single-message-based notion of indistinguishability for deterministic encryption of block sources, which as proved in [BFO08], is equivalent to both the simulation-based and indistinguishability-based multiple-message notions.

Definition 2.3 (Deterministic-encryption security [BFO08]). *Let $\mathsf{TDF} = (\mathsf{TDF.KG}, \mathsf{TDF.F}, \mathsf{TDF.F}^{-1})$ be as in Definition 2.1. We say that TDF is (k, n)-CPA-indistinguishable if for any two (k, n)-sources $\mathcal{S}_1$ and $\mathcal{S}_2$ we have $(\mathsf{ik}, \mathsf{TDF.F}(\mathsf{ik}, \mathcal{S}_1)) \stackrel{c}{\equiv} (\mathsf{ik}, \mathsf{TDF.F}(\mathsf{ik}, \mathcal{S}_2))$, where $(\mathsf{ik}, *) \xleftarrow{\$} \mathsf{TDF.KG}(1^\lambda)$.*

We say that TDF *is* (k, n)-*CCA2-indistinguishable if for any two* (k, n)-*sources* $\mathcal{S}_0$ *and* $\mathcal{S}_1$, *and any PPT adversary* $\mathcal{A}$ *the following probability is negligible:*

$$\Pr\left[b = b' : \begin{array}{l} (\mathsf{ik}, \mathsf{tk}) \xleftarrow{\$} \mathsf{TDF.KG}(1^\lambda), b \xleftarrow{\$} \{0, 1\}, \mathsf{x}^\star \xleftarrow{\$} \mathcal{S}_b, \mathsf{u}^\star := \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x}^\star) \\ b' \leftarrow \mathcal{A}^{\mathcal{O}_{\mathsf{Dec}}}(\mathsf{ik}, \mathsf{u}^\star) \end{array} \right] - \frac{1}{2}$$

where on input u, *the decryption oracle* $\mathcal{O}_{\mathsf{Dec}}$ *returns* $\mathsf{TDF.F}^{-1}(\mathsf{tk}, \mathsf{u})$ *if* $\mathsf{u} \neq \mathsf{u}^\star$, *and* $\bot$ *otherwise.*

We remark that considering only CCA2 queries (as opposed to both CCA1 and CCA2 queries) in the CCA2-indistinguishability definition for deterministic encryption is without loss of generality, since the plaintexts are not chosen by the adversary. See [BFO08] for further explanation.

Definition 2.4 (Computational Diffie-Hellman (CDH) assumption). *Let* 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* G *is CDH-hard if for any PPT adversary* $\mathcal{A}$: $\Pr[\mathcal{A}(\mathbb{G}, p, g, g^{a_1}, g^{a_2}) = g^{a_1 a_2}] = \mathsf{negl}(\lambda)$, *where* $(\mathbb{G}, p, g) \xleftarrow{\$} \mathsf{G}(1^\lambda)$ *and* $a_1, a_2 \xleftarrow{\$} \mathbb{Z}_p$.

2.2 Standard Lemmas

Lemma 2.5 (Chernoff inequality). *Let* $\mathcal{X}_1, \ldots, \mathcal{X}_m$ *be independent Boolean variables each of expected value at least* p. *Then, for all* $\varepsilon > 0$:

$$\Pr\left[\frac{1}{m} \sum_{i=1}^{m} \mathcal{X}_i < p - \varepsilon \right] < e^{-2\varepsilon^2 m}.$$

Lemma 2.6 (Leftover hash lemma [ILL89]). *Let* $\mathcal{X}$ *be a random variable over* X *and* $h : \mathsf{S} \times \mathsf{X} \to \mathsf{Y}$ *be a 2-universal hash function, where* $|\mathsf{Y}| \leq 2^m$ *for some* $m > 0$. *If* $m \leq \mathsf{H}_\infty(\mathcal{X}) - 2 \log\left(\frac{1}{\varepsilon}\right)$, *then* $(h(\mathcal{S}, \mathcal{X}), \mathcal{S}) \approx_\varepsilon (\mathcal{U}, \mathcal{S})$, *where* $\mathcal{S}$ *is uniform over* S *and* $\mathcal{U}$ *is uniform over* Y.

3 Smooth Recyclable OWFE

We recall the definition of recyclable one-way function with encryption from [GH18]. We adapt the definition to a setting in which the underlying input distribution is not necessarily uniform. We will also define a smoothness notion, which generalizes the one-wayness notion.

Definition 3.1 (Recyclable one-way function with encryption (OWFE)). *A recyclable* (k, n)-*OWFE scheme consists of the PPT algorithms* $\mathsf{K}, \mathsf{f}, \mathsf{E}_1, \mathsf{E}_2$ *and* D *with the following syntax.*

- $\mathsf{K}(1^\lambda)$: *Takes the security parameter* 1^λ *and outputs a public parameter* pp *(by tossing coins) for a function* $\mathsf{f}(\mathsf{pp}, \cdot)$ *from* n *bits to* ν *bits.*

- $\mathsf{f}(\mathsf{pp}, \mathsf{x})$: *Takes a public parameter* pp *and a preimage* $\mathsf{x} \in \{0, 1\}^n$, *and deterministically outputs* $\mathsf{y} \in \{0, 1\}^\nu$.

- $\mathsf{E}_1(\mathsf{pp}, (i, b); \rho)$: *Takes a public parameter* pp, *an index* $i \in [n]$, *a bit* $b \in \{0, 1\}$ *and randomness* ρ, *and outputs a ciphertext* ct.[6]

---

[6] ct is assumed to contain (i, b).

- $E_2(pp, y, (i, b); \rho)$: *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 bit* e. *Notice that unlike* E_1, *which does not take* y *as input, the algorithm* E_2 *does take* y *as input.*

- $D(pp, ct, x)$: *Takes a public parameter* pp, *a ciphertext* ct *and a preimage* $x \in \{0, 1\}^n$, *and deterministically outputs a bit* e.

We require the following properties.

- ***Correctness:*** *For any choice of* $pp \in K(1^\lambda)$, *any index* $i \in [n]$, *any preimage* $x \in \{0, 1\}^n$ *and any randomness value* ρ, *the following holds: letting* $y := f(pp, x)$, $b := x_i$ *and* $ct = E_1(pp, (i, x_i); \rho)$, *we have* $E_2(pp, y, (i, x_i); \rho) = D(pp, ct, x)$.

- (k, n)-***One-wayness:*** *For any* (k, n) *source* $\mathcal{S}$ *and any PPT adversary* $\mathcal{A}$:

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

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

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

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

  where $pp \xleftarrow{\$} K(1^\lambda)$, $\rho \xleftarrow{\$} \{0, 1\}^*$, $ct \xleftarrow{\$} E_1(pp, (i, 1 - x_i); \rho)$, $e \xleftarrow{\$} E_2(pp, f(pp, x), (i, 1 - x_i); \rho)$ *and* $e' \xleftarrow{\$} \{0, 1\}$.

Whenever we say an OWFE scheme (without specifying the parameters), we mean $k = n$.

Notation 3.2. *We define* $E(pp, y, (i, b); \rho) \overset{\triangle}{=} (E_1(pp, (i, b); \rho), E_2(pp, y, (i, b); \rho))$.

We will now define the notion of smoothness which extends the one-wayness property to an indistinguishability-based property.

Definition 3.3 (Smoothness). *Let* (K, f, E, D) *be as in Definition 3.1. We say that* (K, f, E, D) *is* (k, n)-*smooth if for any two* (k, n)-*sources* $\mathcal{S}_1$ *and* $\mathcal{S}_2$ *we have* $(pp, f(pp, x_1)) \overset{c}{\equiv} (pp, f(pp, x_2))$, *where* $pp \xleftarrow{\$} K(1^\lambda)$, $x_1 \xleftarrow{\$} \mathcal{S}_1$ *and* $x_2 \xleftarrow{\$} \mathcal{S}_2$.

3.1 Smooth Recyclable OWFE from CDH

Here we show that the recyclable OWFE from [GH18] is (k, n)-smooth, for any $k \geq \log p + \omega(\log \lambda)$, where p is the order of the underlying CDH-hard group. We first present the construction.

Construction 3.4 (Smooth recyclable OWFE from CDH [GH18]). *Let* G *be a CDH-hard group-generator scheme.*

- $K(1^\lambda)$: *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*

$$pp := \begin{pmatrix} g_{1,0}, g_{2,0}, \dots, g_{n,0} \\ g_{1,1}, g_{2,1}, \dots, g_{n,1} \end{pmatrix}. \tag{2}$$

- $\mathsf{f}(\mathsf{pp}, \mathsf{x})$: *Parse* pp *as in Equation 2, and output* $\mathsf{y} := \prod_{j \in [n]} g_{j, \mathsf{x}_j}$.

- $\mathsf{E}_1(\mathsf{pp}, (i, b); \rho)$: *Parse* pp *as in Equation 2. Given the randomness* $\rho \xleftarrow{\$} \mathbb{Z}_p$, *proceed as follows:*

    - *For every* $j \in [n] \setminus \{i\}$, *set* $c_{j,0} := g_{j,0}^\rho$, *and* $c_{j,1} := g_{j,1}^\rho$.
    - *Set* $c_{i,b} := g_{i,b}^\rho$ *and* $c_{i,1-b} := \bot$.
    - *Output*

$$\mathsf{ct} := \begin{pmatrix} c_{1,0}, c_{2,0}, \ldots, c_{n,0} \\ c_{1,1}, c_{2,1}, \ldots, c_{n,1} \end{pmatrix}. \tag{3}$$

- $\mathsf{E}_2(\mathsf{pp}, (\mathsf{y}, i, b); \rho)$: *Given the randomness* $\rho \xleftarrow{\$} \mathbb{Z}_p$, *output* $\mathsf{HC}(\mathsf{y}^\rho)$, *where* HC *denotes a hardcore function (e.g., the Goldreich-Levin hardcore function).*

- $\mathsf{D}(\mathsf{pp}, \mathsf{ct}, \mathsf{x})$: *Parse* ct *as in Equation 3, and output* $\mathsf{HC}\left(\prod_{j \in [n]} c_{j, \mathsf{x}_j}\right)$.

[GH18] showed that this OWFE satisfies the correctness, one-wayness, and security for encryption properties. We show that for any $k \geq \log p + \omega(\log \lambda)$, where p is the order of the underlying group, we also have the smoothness property.

Lemma 3.5 (CDH implies (k, n)-smooth recyclable OWFE). *Assuming that* G *is CDH-hard and* $k \geq \log p + \omega(\log \lambda)$, *then the recyclable OWFE of Construction 3.4 is* (k, n)-*smooth.*

Proof. Let $\mathcal{S}_1$ and $\mathcal{S}_2$ be two (k, n) sources. The smoothness follows directly from the left over hash lemma. Namely, assuming $\mathsf{pp} \xleftarrow{\$} \mathsf{K}(1^\lambda)$, $\mathsf{x}_1 \xleftarrow{\$} \mathcal{S}_1$ and $\mathsf{x}_2 \xleftarrow{\$} \mathcal{S}_2$, since f is 2-universal hash function, by Lemma 2.6 we know that the outputs of both $\mathsf{f}(\mathsf{pp}, \mathsf{x}_1)$ and $\mathsf{f}(\mathsf{pp}, \mathsf{x}_2)$ are statistically $\frac{1}{2^{\omega(\log \lambda)}}$ close to the uniform over $\mathbb{G}$, and hence negligibly close (statistically) to each other. $\square$

4 Strong TDFs from Smooth Recyclable OWFE

In this section we show that recyclable OWFE implies the existence of TDFs with almost-perfect correctness in the sense of Definition 2.1. This improves the correctness property of [GH18]; see Remark 2.2. Moreover, we show that if the base recyclable OWFE scheme is smooth (Definition 3.3), then the resulting TDF satisfies the notions of security for deterministic encryption (Definition 2.3). We will then use this statement along with Lemma 3.5 to obtain the first deterministic-encryption scheme based on CDH. In particular, the existence of deterministic encryption (even) with CPA security from CDH has been open until now.

A central new tool developed in this work is a *mirroring* technique, which we will describe below. As notation, for a matrix $\mathbf{M} \in \mathbb{Z}_2^{k \times n}$, we define $\mathsf{RSum}(\mathbf{M}) \overset{\triangle}{=} \mathbf{M}_1 + \cdots + \mathbf{M}_k \in \mathbb{Z}_2^n$, where $\mathbf{M}_i$ for $i \in [k]$ denotes the ith row of $\mathbf{M}$.

Definition 4.1. *(The mirror Function* Mir*) Let* $(\mathsf{K}, \mathsf{f}, \mathsf{E}_1, \mathsf{E}_2, \mathsf{D})$ *be a recyclable OWFE scheme. For a public parameter* pp, *a value* $\mathsf{x} \in \{0,1\}^n$, *a matrix* $\mathbf{CT} := \begin{pmatrix} \mathsf{ct}_{1,0}, \mathsf{ct}_{2,0}, \ldots, \mathsf{ct}_{n,0} \\ \mathsf{ct}_{1,1}, \mathsf{ct}_{2,1}, \ldots, \mathsf{ct}_{n,1} \end{pmatrix}$ *of ciphertexts outputted by* E_1, *and a vector* $\mathbf{a} \in \{0,1\}^n$, *the function* $\mathsf{Mir}(\mathsf{pp}, \mathsf{x}, \mathbf{CT}, \mathbf{a})$ *outputs a matrix* $\mathbf{M} \in \mathbb{Z}_2^{2 \times n}$, *where* $\mathbf{M} := \begin{pmatrix} \mathsf{b}_{1,0}, \mathsf{b}_{2,0}, \ldots, \mathsf{b}_{n,0} \\ \mathsf{b}_{1,1}, \mathsf{b}_{2,1}, \ldots, \mathsf{b}_{n,1} \end{pmatrix}$ *is formed deterministically and uniquely according to the following two rules:*

1. *for all $i \in [n]$: $\mathsf{b}_{i,\mathsf{x}_i} = \mathsf{D}(\mathsf{pp}, \mathsf{ct}_{i,\mathsf{x}_i}, \mathsf{x})$; and*

2. $\mathsf{RSum}(\mathbf{M}) = \mathbf{a}$.

Note that the above computation is deterministic and can be done efficiently.

Construction 4.2 (TDF construction). *We now present our TDF construction.*

Base primitive. *A recyclable OWFE scheme $\mathcal{E} = (\mathsf{K}, \mathsf{f}, \mathsf{E}, \mathsf{D})$. Let Rand be the randomness space of the 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 $\mathsf{f}(\mathsf{pp}, \cdot)$, and r will be instantiated in the correctness proof. The input space of the TDF is $\{0,1\}^n$.*

- TDF.KG(1^λ):

    1. *Sample $\mathsf{pp} \leftarrow \mathsf{K}(1^\lambda)$.*
    2. *For each $h \in [r]$:*

    $$\mathbf{P}_h := \begin{pmatrix} \rho_{1,0}^{(h)}, \rho_{2,0}^{(h)}, \dots, \rho_{n,0}^{(h)} \\ \rho_{1,1}^{(h)}, \rho_{2,1}^{(h)}, \dots, \rho_{n,1}^{(h)} \end{pmatrix} \xleftarrow{\$} \mathsf{Rand}^{2 \times n}, \tag{4}$$

    $$\mathbf{CT}_h := \begin{pmatrix} \mathsf{E}_1(\mathsf{pp}, (1,0); \rho_{1,0}^{(h)}), \mathsf{E}_1(\mathsf{pp}, (2,0); \rho_{2,0}^{(h)}), \dots, \mathsf{E}_1(\mathsf{pp}, (n,0); \rho_{n,0}^{(h)}) \\ \mathsf{E}_1(\mathsf{pp}, (1,1); \rho_{1,1}^{(h)}), \mathsf{E}_1(\mathsf{pp}, (2,1); \rho_{2,1}^{(h)}), \dots, \mathsf{E}_1(\mathsf{pp}, (n,1); \rho_{n,1}^{(h)}) \end{pmatrix}. \tag{5}$$

    3. *For $h \in [r]$ sample $\mathbf{a}_h \xleftarrow{\$} \{0,1\}^n$.*
    4. *Form the index key ik and the trapdoor key tk as follows:*

    $$\mathsf{ik} := (\mathsf{pp}, \mathbf{CT}_1, \dots, \mathbf{CT}_r, \mathbf{a}_1, \dots, \mathbf{a}_r), \tag{6}$$
    $$\mathsf{tk} := (\mathsf{pp}, \mathbf{P}_1, \dots, \mathbf{P}_r). \tag{7}$$

- TDF.F(ik, x): *Parse ik as in Equation 6. Set $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{x})$. Return*

    $$\mathsf{u} := (\mathsf{y}, \mathsf{Mir}(\mathsf{pp}, \mathsf{x}, \mathbf{CT}_1, \mathbf{a}_1), \dots, \mathsf{Mir}(\mathsf{pp}, \mathsf{x}, \mathbf{CT}_r, \mathbf{a}_r)). \tag{8}$$

- TDF.F^{-1}(tk, u):

    1. *Parse $\mathsf{tk} := (\mathsf{pp}, \mathbf{P}_1, \dots, \mathbf{P}_r)$ and parse $\mathbf{P}_h$ for $h \in [r]$ as in Equation 4.*
    2. *Parse $\mathsf{u} := (\mathsf{y}, \mathbf{M}_1, \dots, \mathbf{M}_r)$, where for all $h \in [r]$, $\mathbf{M}_h \in \mathbb{Z}_2^{2 \times n}$.*
    3. *Reconstruct $\mathsf{x} := \mathsf{x}_1 \cdots \mathsf{x}_n \in \{0,1\}^n$ bit-by-bit as follows. To recover the ith bit of x:*

        (a) *If for some $h \in [r]$, $\mathbf{M}_h[i] = \begin{pmatrix} \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (i,0); \rho_{i,0}^{(h)}) \\ 1 - \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (i,1); \rho_{i,1}^{(h)}) \end{pmatrix}$, set $\mathsf{x}_i = 0$. Here $\mathbf{M}_h[i]$ denotes the ith column of $\mathbf{M}_h$.*

        (b) *Else, if for some $h \in [r]$, $\mathbf{M}_h[i] = \begin{pmatrix} 1 - \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (i,0); \rho_{i,0}^{(h)}) \\ \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (i,1); \rho_{i,1}^{(h)}) \end{pmatrix}$, set $\mathsf{x}_i = 1$.*

        (c) *Otherwise, halt and return $\perp$.*

4. Return x.

We first give a remark about our TDF construction and will then give the correctness and one-wayness statements about our TDF.

Remark 4.3. *We mention that for an image* $u := (y, M_1, \ldots, M_r)$*, we do not need* u *to contain the entire part of matrices* $M_1, \ldots, M_r$*. Just outputting the first row of each matrix will be sufficient because the other row can be computed from the first row and the vectors* $a_1, \ldots, a_r$ *of the index key* ik*. For simplicity, we keep the entire matrices in the image* u *for our security and correctness proofs, but we will make use of this observation in our efficiency estimates.*

Lemma 4.4 (TDF correctness)**.** *We have*

$$\Pr_{(ik,tk)}[\exists x \in \{0,1\}^n \text{ s.t. } \mathsf{TDF.F}^{-1}(tk, (\mathsf{TDF.F}(ik,x))) \neq x] \leq \frac{n2^n}{2^r}, \tag{9}$$

where the probability is taken over $(ik, tk) \xleftarrow{\$} \mathsf{TDF.KG}(1^\lambda)$*. For instance, setting:* $r = n + \omega(\log \lambda)$ *gives a negligible inversion error.*

Lemma 4.5 (TDF one-wayness and CPA-indistinguishability security)**.** *Assuming* $\mathcal{E}$ *is an OWFE scheme (i.e., an* (n, n)*-OWFE scheme), the TDF* $(\mathsf{TDF.KG}, \mathsf{TDF.F}, \mathsf{TDF.F}^{-1})$ *given in Construction 4.2 is one-way. That is, for any PPT adversary* $\mathcal{A}$

$$\Pr[\mathcal{A}(ik, \mathsf{TDF.F}(ik, x)) = x] = \mathsf{negl}(\lambda), \tag{10}$$

where $(ik, tk) \xleftarrow{\$} \mathsf{TDF.KG}(1^\lambda)$ *and* $x \xleftarrow{\$} \{0,1\}^n$*. Moreover, if* $\mathcal{E}$ *is* (k, n)*-smooth (Definition 3.3), the constructed TDF is* (k, n)*-CPA-indistinguishable (Definition 2.3).*

We may now combine Lemmas 3.5, 4.4 and 4.5 to get the first CPA-secure deterministic encryption scheme from CDH.

Corollary 4.6 (CDH implies deterministic encryption)**.** *Let* G *be a CDH-hard group scheme. For any* $k \geq \log p + \omega(\log \lambda)$ *and any* $n \geq k$ *(where* p *is the order of the underlying group), there exists a* (k, n)*-CPA-indistinguishable deterministic encryption scheme with plaintext size* n *(in bits) and ciphertext size* $\Theta(n^2)$*.*

4.1 Proof of Correctness: Lemma 4.4

Proof. We will use notation given in Construction 4.2. Note that for a given $x \in \{0,1\}^n$, the inversion succeeds unless there exists an index $i \in [n]$ for which the following bad event happens.

- $\mathsf{Bad}_{x,i}$: for all $h \in [r]$, $\mathbf{a}_h[i] = \mathsf{E}_2(\mathsf{pp}, y, (i, 0); \rho_{i,0}^{(h)}) + \mathsf{E}_2(\mathsf{pp}, y, (i, 1); \rho_{i,0}^{(h)}) \in \mathbb{Z}_2$, where $\mathbf{a}_h[i]$ denotes the i'th coordinate of $\mathbf{a}_h \in \{0,1\}^n$.

Since the bits $\mathbf{a}_h[i]$ for all $h \in [r]$ are chosen uniformly at random (independently of pp and $\rho_{i,b}$'s), we have: $\Pr[\mathsf{Bad}_{x,i}] = 2^{-r}$. Doing a union bound over all column $i \in [n]$ gives the probability $n \cdot 2^{-r}$ of an inversion error for a given x. We conclude using a union bound over all $x \in \{0,1\}^n$. $\square$

4.2 Proof of One-Wayness and CPA Security: Lemma 4.5

To prove Lemma 4.5 we first give a simulated way of sampling an index key together with an image point for a target input value.

Definition 4.7 (Simulated distribution Sim). *Let $\mathcal{E} = (\mathsf{K}, \mathsf{f}, \mathsf{E}, \mathsf{D})$ be the underlying recyclable OWFE scheme. Fix $\mathsf{x} \in \{0,1\}^n$ and let $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{x})$. We define a simulator $\mathsf{Sim}(\mathsf{pp}, n, \mathsf{y})$, which samples a simulated index key $\mathsf{ik}_{\mathrm{sim}}$ with a corresponding simulated TDF image $\mathsf{u}_{\mathrm{sim}}$ for x, as follows. For $h \in [r]$ sample*

$$(\mathbf{CT}_h, \mathbf{M}_h) \overset{\$}{\leftarrow} \begin{pmatrix} \mathsf{E}_1(\mathsf{pp}, (1,0)), \ldots, \mathsf{E}_1(\mathsf{pp}, (n,0)) \\ \mathsf{E}_1(\mathsf{pp}, (1,1)), \ldots, \mathsf{E}_1(\mathsf{pp}, (n,1)) \end{pmatrix}, \begin{pmatrix} \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (1,0)), \ldots, \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (n,0)) \\ \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (1,1)), \ldots, \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (n,1)) \end{pmatrix}.$$

Let

$$\mathsf{ik}_{\mathrm{sim}} := (\mathsf{pp}, \mathbf{CT}_1, \ldots, \mathbf{CT}_r, \mathsf{RSum}(\mathbf{M}_1), \ldots, \mathsf{RSum}(\mathbf{M}_r))$$
$$\mathsf{u}_{\mathrm{sim}} := (\mathsf{y}, \mathbf{M}_1, \ldots, \mathbf{M}_r).$$

Equipped with the above definition, we now give of the proof of Lemma 4.5.

Proof of Lemma 4.5. For any distribution $\mathcal{S}$ over $\{0,1\}^n$, we show that the sole security-for-encryption requirement of the recyclable OWFE implies that

$$(\mathsf{x}, \mathsf{ik}, \mathsf{TDF}.\mathsf{F}(\mathsf{ik}, \mathsf{x})) \overset{c}{\equiv} (\mathsf{x}, \mathsf{Sim}(\mathsf{pp}, n, \mathsf{y})), \tag{11}$$

where $\mathsf{x} \overset{\$}{\leftarrow} \mathcal{S}$, $\mathsf{pp} \overset{\$}{\leftarrow} \mathsf{K}(1^\lambda)$, $(\mathsf{ik}, *) \overset{\$}{\leftarrow} \mathsf{TDF}.\mathsf{KG}(1^\lambda)$ and $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{x})$.

We first show how to use Equation 11 to derive the one-wayness and indistinguishability security of the resulting TDF from the corresponding one-wayness and smoothness of the underlying OWFE scheme, and will then prove Equation 11.

For one-wayness, if there exists an inverter $\mathcal{A}$ that with non-negligible probability can compute x from $(\mathsf{ik}, \mathsf{TDF}.\mathsf{F}(\mathsf{ik}, \mathsf{x})$ — where $(\mathsf{ik}, *) \overset{\$}{\leftarrow} \mathsf{TDF}.\mathsf{KG}(1^\lambda)$ and $\mathsf{x} \overset{\$}{\leftarrow} \{0,1\}^n$ — then Equation 11 implies that with non-negligible probability the adversary $\mathcal{A}$ can compute x from $\mathsf{Sim}(\mathsf{pp}, n, \mathsf{y})$, where $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{x})$. However, this latter violates the one-wayness of f, because the computation of $\mathsf{Sim}(\mathsf{pp}, n, \mathsf{y})$ may be done efficiently with knowledge of pp, n and y.

For indistinguishability security (Definition 2.3) let $\mathcal{S}_0$ and $\mathcal{S}_1$ be two (k, n) sources and assume that the recyclable OWFE scheme is k-smooth (Definition 3.3).

Letting $(\mathsf{ik}, *) \overset{\$}{\leftarrow} \mathsf{TDF}.\mathsf{KG}(1^\lambda)$, $\mathsf{x}_0 \overset{\$}{\leftarrow} \mathcal{S}_0$, $\mathsf{x}_1 \overset{\$}{\leftarrow} \mathcal{S}_1$, $\mathsf{y}_0 := \mathsf{f}(\mathsf{pp}, \mathsf{x}_0)$ and $\mathsf{y}_1 := \mathsf{f}(\mathsf{pp}, \mathsf{x}_1)$, by Equation 11 we have

$$(\mathsf{ik}, \mathsf{TDF}.\mathsf{F}(\mathsf{ik}, \mathsf{x}_0)) \overset{c}{\equiv} \mathsf{Sim}(\mathsf{pp}, n, \mathsf{y}_0) \overset{c}{\equiv} \mathsf{Sim}(\mathsf{pp}, n, \mathsf{y}_1) \overset{c}{\equiv} (\mathsf{ik}, \mathsf{TDF}.\mathsf{F}(\mathsf{ik}, \mathsf{x}_1)),$$

where the second indistinguishability follows from the k-smoothness of the recyclable OWFE scheme, which states $(\mathsf{pp}, \mathsf{y}_0) \overset{c}{\equiv} (\mathsf{pp}, \mathsf{y}_1)$.

We are left to prove Equation 11. Fix the distribution $\mathcal{S}$ for which we want to prove Equation 11. To this end, we change the simulator Sim given in Definition 4.7 to define a new simulator Sim' which on input $\mathsf{Sim}'(\mathsf{pp}, \mathsf{x})$ samples a pair $(\mathsf{ik}'_{\mathrm{sim}}, \mathsf{u}'_{\mathrm{sim}})$ as follows. Let $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{x})$. For all $h \in [r]$, let $\mathbf{CT}_h$ be sampled as in $\mathsf{Sim}(\mathsf{pp}, n, \mathsf{y})$, but with the following modification to $\mathbf{M}_h$:

- Letting $\mathbf{M}_h := \begin{pmatrix} e_{1,0}^{(h)},...,e_{n,0}^{(h)} \\ e_{1,1}^{(h)},...,e_{n,1}^{(h)} \end{pmatrix}$ be formed as in $\mathsf{Sim}(\mathsf{pp},\mathsf{y})$, for any $i \in [n]$ change $e_{i,1-\mathsf{x}_i}^{(h)}$ to a random bit (fresh for each index).

Having defined how $\mathbf{CT}_h$ and $\mathbf{M}_h$ are sampled for $h \in [r]$ during $\mathsf{Sim}'(\mathsf{pp},\mathsf{x})$, form $(\mathsf{ik}_{\mathrm{sim}}',\mathsf{u}_{\mathrm{sim}}')$ exactly as how $(\mathsf{ik}_{\mathrm{sim}},\mathsf{u}_{\mathrm{sim}})$ is formed during $\mathsf{Sim}(\mathsf{pp},n,\mathsf{y})$.

The security-for-encryption requirement of the OWFE scheme implies that $(\mathsf{x},\mathsf{ik}_{\mathrm{sim}},\mathsf{u}_{\mathrm{sim}}) \overset{c}{\equiv} (\mathsf{x},\mathsf{ik}_{\mathrm{sim}}',\mathsf{u}_{\mathrm{sim}}')$, where $\mathsf{x} \overset{\$}{\leftarrow} \mathcal{S}$, $\mathsf{y} := \mathsf{f}(\mathsf{pp},\mathsf{x})$, $(\mathsf{ik}_{\mathrm{sim}},\mathsf{u}_{\mathrm{sim}}) \overset{\$}{\leftarrow} \mathsf{Sim}(\mathsf{pp},n,\mathsf{y})$ and $(\mathsf{ik}_{\mathrm{sim}}',\mathsf{u}_{\mathrm{sim}}') \overset{\$}{\leftarrow} \mathsf{Sim}'(\mathsf{pp},\mathsf{x})$. Moreover, it is easy to verify that $(\mathsf{x},\mathsf{ik}_{\mathrm{sim}}',\mathsf{u}_{\mathrm{sim}}')$ is identically distributed to $(\mathsf{x},\mathsf{ik},\mathsf{TDF.F}(\mathsf{ik},\mathsf{x}))$, where $(\mathsf{ik},\mathsf{tk}) \overset{\$}{\leftarrow} \mathsf{TDF.KG}(1^\lambda)$. The proof is now complete.

$\square$

4.3 CCA2 Security

The TDF given in Construction 4.2 is CPA secure (in a deterministic-encryption sense), but it is not hard to show that the construction is not CCA2 secure. However, using techniques of [RS09, KMO10] one may use the TDF of Construction 4.2 to build another TDF which is CCA2 secure. This upgrading further increases the ciphertext size, resulting in ciphertext size $\Theta(n^3)$ (for the CDH-based instantiation), where n is the plaintext size. We skip the details of this construction, because later in Section 5 we will give constructions of efficient CCA2-secure deterministic encryption schemes with linear-size ciphertexts. In what comes below we explain about a feature of the TDF of Construction 4.2 which enables the upgrading into CCA2-secure deterministic encryption.

Specifically, Rosen and Segev [RS09] show that *t-repetition security*, an extension of the basic notion of one-wayness for TDFs, suffices for the construction of CCA2-secure PKE. We first review this definition, and then show that the TDF of Construction 4.2 provides *t*-repetition security for any *t*. This observation allows us to obtain the first *t*-repetition secure TDFs from CDH.

Definition 4.8 (*t*-repetition-secure TDFs [RS09]). *Let* $\mathsf{TDF} := (\mathsf{TDF.KG},\mathsf{TDF.F},\mathsf{TDF.F}^{-1})$ *be as in Definition 2.1. We say that* TDF *is t-repetition secure if for any PPT adversary* $\mathcal{A}$ *we have* $\Pr[\mathcal{A}((\mathsf{ik}_1,\mathsf{u}_1),\ldots,(\mathsf{ik}_t,\mathsf{u}_t)) = \mathsf{x}] = \mathsf{negl}(\lambda)$, *where* $\mathsf{x} \overset{\$}{\leftarrow} \{0,1\}^n$ *and for all* $i \in [t]$, $(\mathsf{ik}_i,*) \overset{\$}{\leftarrow} \mathsf{TDF.KG}(1^\lambda)$ *and* $\mathsf{u}_i := \mathsf{TDF.F}(\mathsf{ik}_i,\mathsf{x})$.

We will now show that our TDF given in Construction 4.2 provides *t*-repetition security, for all t.[7] This gives us the first CDH-based construction of *t*-repetition-secure TDFs.

Lemma 4.9. *Assuming* $\mathcal{E}$ *is a recyclable* (n,n)-*OWFE scheme, the TDF* $(\mathsf{TDF.KG},\mathsf{TDF.F},\mathsf{TDF.F}^{-1})$ *given in Construction 4.2 is t-repetition secure, for all t.*

Proof. The proof is an immediate generalization of the proof of Lemma 4.5. Specifically, using a hybrid argument, we may extend Equation 11 (Page 14) to obtain

$$(\mathsf{x},\mathsf{ik}_1,\mathsf{TDF.F}(\mathsf{ik}_1,\mathsf{x}),\ldots,\mathsf{ik}_t,\mathsf{TDF.F}(\mathsf{ik}_t,\mathsf{x})) \overset{c}{\equiv} (\mathsf{x},\mathsf{Sim}(\mathsf{pp},n,\mathsf{y};r_1),\ldots,\mathsf{Sim}(\mathsf{pp},n,\mathsf{y};r_t)), \qquad (12)$$

where $\mathsf{x} \overset{\$}{\leftarrow} \{0,1\}^n$, $\mathsf{pp} \overset{\$}{\leftarrow} \mathsf{K}(1^\lambda)$, $\mathsf{y} := \mathsf{f}(\mathsf{pp},\mathsf{x})$ and for $i \in [t]$, $(\mathsf{ik}_i,*) \overset{\$}{\leftarrow} \mathsf{TDF.KG}(1^\lambda)$ and $r_i \overset{\$}{\leftarrow} \{0,1\}^*$. Now the *t*-repetition security of the TDF follows immediately from Equation 12, by noting the whole view of the adversary may be formed in a computationally-close way using (pp,y) only. $\square$

---

[7] For *t*-repetition security, we assume that pp is a public parameter (the same across all index keys).

5   Efficient Strong TDFs from Smooth OWFE

The TDF and deterministic encryption presented in Section 4 have the drawback that the output size grows at least quadratically with the input size. The reason for this blowup is that we need to do "repetitions," resulting in $\Theta(n/2)$ output bits for every single bit of the input. In this section we show how to do away with the repetition and to obtain TDFs (and deterministic encryption) whose image/ciphertext size grows linearly with input size. Our main idea involves the use of erasure codes, taking advantage of the local inversion property of our basic TDF. As a result, we will obtain the first CPA-secure deterministic encryption scheme with linear ciphertext size based on CDH or DDH. We stress that, even relying on DDH, previous DDH-based deterministic-encryption and TDF schemes resulted in quadratically large ciphertexts.

Definition 5.1 (Erasure codes)**.** *We will work with an optimal erasure-resilient code* $(\mathsf{Encode}, \mathsf{Decode})$, *where* $\mathsf{Encode} : \{0,1\}^n \to \{0,1\}^m$ *and* $m = c \cdot n$ *for some constant* c. *We denote such a code with* (n, c). *For optimal erasure resilience, we require that*

- *For any two distinct* $\mathsf{x}_1, \mathsf{x}_2 \in \{0,1\}^n$ *we have* $\mathsf{H}_{\mathrm{dst}}(\mathsf{Encode}(\mathsf{x}_1), \mathsf{Encode}(\mathsf{x}_2)) \geq m - n$, *where* $\mathsf{H}_{\mathrm{dst}}$ *denotes the Hamming distance.*

- *For any* $x \in \{0,1\}^n$, *letting* $\mathsf{z} := \mathsf{Encode}(\mathsf{x})$, *given any string* $\mathsf{z}' \in \{0,1,\bot\}^m$, *which has at least* n *non-*$\bot$ *symbols, and whose all non-*$\bot$ *symbols agree with* z, *we have* $\mathsf{Decode}(\mathsf{z}') = \mathsf{x}$.

Such codes can be built from, say, the Reed-Solomon codes for strings over alphabets of large fields. For binary strings, there are efficient transformations that preserve the linear-output-size property, e.g., [BKK+95].[8]

Construction 5.2 (TDF construction)**.** *We now describe our TDF construction.*

Base primitive. *An optimal erasure code* $(\mathsf{Encode}, \mathsf{Decode})$, *where* $\mathsf{Encode}: \{0,1\}^n \to \{0,1\}^m$, *and a recyclable OWFE scheme* $\mathcal{E} = (\mathsf{K}, \mathsf{f}, \mathsf{E}, \mathsf{D})$, *where* f*'s input space is* $\{0,1\}^m$. *Let* Rand *be the randomness space of the encapsulation algorithm* E.

Construction.

- $\mathsf{TDF}.\mathsf{KG}(1^\lambda)$*:*

  1. *Sample* $\mathsf{pp} \leftarrow \mathsf{K}(1^\lambda)$ *and*

  $$\mathbf{P} := \begin{pmatrix} \rho_{1,0}, \rho_{2,0}, \ldots, \rho_{m,0} \\ \rho_{1,1}, \rho_{2,1}, \ldots, \rho_{m,1} \end{pmatrix} \xleftarrow{\$} \mathsf{Rand}^{2 \times m}, \tag{13}$$

  $$\mathbf{CT} := \begin{pmatrix} \mathsf{ct}_{1,0}, \mathsf{ct}_{2,0}, \ldots, \mathsf{ct}_{m,0} \\ \mathsf{ct}_{1,1}, \mathsf{ct}_{2,1}, \ldots, \mathsf{ct}_{m,1} \end{pmatrix}, \tag{14}$$

  where for all $i \in [m]$ *and* $b \in \{0,1\}$, $\mathsf{ct}_{i,b} := \mathsf{E}_1(\mathsf{pp}, (i,b); \rho_{i,b})$.

  2. *Sample* $\mathbf{a} \xleftarrow{\$} \{0,1\}^m$.

  3. *Form the index key* ik *and the trapdoor key* tk *as follows:*

  $$\mathsf{ik} := (\mathsf{pp}, \mathbf{a}, \mathbf{CT}) \qquad \mathsf{tk} := (\mathsf{pp}, \mathbf{a}, \mathbf{P}). \tag{15}$$

---

[8]The naive transformation of encoding a large field element with a binary string will destroy the optimal erasure-resilience property.

- TDF.F(ik, x): *Parse* ik := (pp, **a**, **CT**). *Let* z := Encode(x) *and* y := f(pp, z). *Return*

$$u := (y, \text{Mir}(pp, z, \mathbf{CT}, \mathbf{a})). \tag{16}$$

- TDF.F^{-1}(tk, u):

    1. *Parse* tk := (pp, **a**, **P**) *and parse* **P** *as in Equation* (13). *Parse* u := (y, **M**), *where* $\mathbf{M} \in \mathbb{Z}_2^{2 \times m}$. *If* RSum(**M**) $\neq$ **a**, *then return* $\perp$.

    2. *Construct* $z' := z'_1 \cdots z'_m$ *bit-by-bit as follows. To recover the ith bit of* z':

        (a) *If* $\mathbf{M}[i] = \begin{pmatrix} \mathsf{E}_2(pp,y,(i,0);\rho_{i,0}) \\ 1-\mathsf{E}_2(pp,y,(i,1);\rho_{i,1}) \end{pmatrix}$, *set* $z'_i = 0$. *Here* $\mathbf{M}[i]$ *denotes the ith column of* **M**.

        (b) *Else if* $\mathbf{M}[i] = \begin{pmatrix} 1-\mathsf{E}_2(pp,y,(i,0);\rho_{i,0}) \\ \mathsf{E}_2(pp,y,(i,1);\rho_{i,1}) \end{pmatrix}$, *set* $z'_i = 1$.

        (c) *Else, set* $z'_i = \perp$.

    3. *Letting* x := Decode(z'), *if* TDF.F(ik, x) = u, *then return* x. *Otherwise, return* $\perp$.

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

Lemma 5.3 (Correctness). *Let* (n, c) *be the parameters of the underlying optimal erasure code* (Encode, Decode) *(Definition 5.1). For any* $c > 2$, *we have*

$$\Pr_{(ik,tk)}[\exists x \in \{0,1\}^n \ s.t. \ \mathsf{TDF.F}^{-1}(tk, (\mathsf{TDF.F}(ik, x))) \neq x] \leq 2^n \cdot e^{\frac{-n(c-2)^2}{2c}}. \tag{17}$$

In particular, if $c \geq 5$, *the probability in Equation* (19) *will be negligible.*

Lemma 5.4 (TDF one-wayness and CPA-indistinguishability security). *Assuming* $\mathcal{E}$ *is an* (n, m)-*OWFE scheme, the TDF* (TDF.KG, TDF.F, TDF.F^{-1}) *given in Construction 4.2 is one-way. That is, for any PPT adversary* $\mathcal{A}$

$$\Pr[\mathcal{A}(ik, \mathsf{TDF.F}(ik, x)) = x] = \mathsf{negl}(\lambda), \tag{18}$$

where (ik, tk) $\xleftarrow{\$}$ TDF.KG(1^λ) *and* x $\xleftarrow{\$}$ $\{0, 1\}^n$. *Moreover, assuming that the underlying OWFE scheme is* (k, m)-*smooth (Definition 3.3), the constructed TDF is* (k, n)-*indistinguishable (Definition 2.3).*

Theorem 5.5 (CCA2-indistinguishability security). *Assuming that the underlying OWFE scheme is* (k, m)-*smooth and that* $m \geq 10n$, *the constructed TDF is* (k, n)-*CCA2-indistinguishable.*

We may now combine Lemmas 3.5, 5.3 with Theorem 5.5 and also Remark 4.3 to get the following corollary.

Corollary 5.6 (CDH implies efficient deterministic encryption). *Let* G *be a CDH-hard group scheme. For any* $k \geq \log p + \omega(\log \lambda)$ *and any* $n \geq k$ *(where* p *is the order of the underlying group), there exists a* (k, n)-*CCA2-indistinguishable deterministic encryption scheme with plaintext size* n *(in bits) and ciphertext size* $\log p + 10n$.

5.1 Proof of Correctness: Lemma 5.3

Proof. We will use notation and variables used in construction 5.2. For any $x \in \{0,1\}^n$, the inversion of $u := \mathsf{TDF.F}(ik, x)$ w.r.t. tk succeeds if the inversion algorithm can recover at least n coordinates $i \in [m]$ of $z := \mathsf{Encode}(x)$. That is, for at least n coordinates i we should have $z_i' \neq \perp$, where recall that z' is constructed during $\mathsf{TDF.F}^{-1}(tk, u)$.

Using the Boolean random variables $X_i = 1$ if $z_i' \neq \perp$, and $X_i = 0$ otherwise, we may write the probability of an inversion error for a given x as: $\Pr[\sum_{i=1}^m X_i < n]$.

Let $y := f(pp, x)$. Note that $X_i = 1$ iff $\mathbf{a}_i \neq \mathsf{E}_2(pp, y, (i,0); \rho_{i,0}) + \mathsf{E}_2(pp, y, (i,0); \rho_{i,1})$, where $\mathbf{a}_i \in \mathbb{Z}_2$ denotes the i'th bit of $\mathbf{a} \in \mathbb{Z}_2^m$. (Recall that $\mathbf{a}$ is given as part of the index key ik.) This is because, parsing $u := (y, \mathbf{M})$, for any ith column $\mathbf{M}[i]$, by design we have $\mathbf{a}_i = \mathsf{RSum}(\mathbf{M}[i])$, and that we have $\mathbf{M}[i] = \begin{pmatrix} \mathsf{E}_2(pp,y,(i,0);\rho_{i,0}) \\ 1-\mathsf{E}_2(pp,y,(i,1);\rho_{i,1}) \end{pmatrix}$ or $\mathbf{M}[i] = \begin{pmatrix} 1-\mathsf{E}_2(pp,y,(i,0);\rho_{i,0}) \\ \mathsf{E}_2(pp,y,(i,1);\rho_{i,1}) \end{pmatrix}$ or $\mathbf{M}[i] = \begin{pmatrix} \mathsf{E}_2(pp,y,(i,0);\rho_{i,0}) \\ \mathsf{E}_2(pp,y,(i,1);\rho_{i,1}) \end{pmatrix}$. The first two cases lead to $z_i' \neq \perp$, and the last case causes $z_i' = \perp$.

Since the bits $\mathbf{a}_i$ are picked uniformly at random (independently of pp, $\rho_{i,0}$ and $\rho_{i,1}$), the Boolean variables X_i are independent and of same expected value $1/2$. Thus, using the Chernoff inequality (Theorem 2.5 with $p = 1/2$ and $\varepsilon = 1/2 - 1/c$, where $m = c \cdot n$), we have: $\Pr[\sum_{i=1}^m X_i < n] \leq e^{-\frac{n(c-2)^2}{2c}}$. We conclude using a union bound over all $x \in \{0,1\}^n$. $\square$

We will now prove a more general result than that of Lemma 5.3. Informally, in the proof of Lemma 5.3, we argued that for any given x, the inversion algorithm can recover at least n bits of z (with high probability), which in turn was sufficient for recovering the original input x. In our later proofs, we will need to argue that the inversion algorithm is capable of retrieving a much larger number of bits, say $2n$ bits of z. To formalize this statement, we give the following definition, which formally defines the notion of a "correct" column and a "hung" column; the former leads to the recovery of the corresponding bit of z, while the latter leads to a failure of discovering the corresponding bit.

Notation 5.7 (NCor Function). *Fix a trapdoor key $tk := (pp, \mathbf{a}, \mathbf{P})$ and let ik be the corresponding index key. Let $u := (y, \mathbf{M})$ be an (alleged) image point. We let $\mathsf{NCor}(tk, u)$ denote the number of "correct" columns in $\mathbf{M}$ with respect to y and tk, in the following sense. We say that the ith column $\mathbf{M}[i]$ of $\mathbf{M}$ is correct if either $\mathbf{M}[i] = \begin{pmatrix} \mathsf{E}_2(pp,y,(i,0);\rho_{i,0}) \\ 1-\mathsf{E}_2(pp,y,(i,1);\rho_{i,1}) \end{pmatrix}$ or $\mathbf{M}[i] = \begin{pmatrix} 1-\mathsf{E}_2(pp,y,(i,0);\rho_{i,0}) \\ \mathsf{E}_2(pp,y,(i,1);\rho_{i,1}) \end{pmatrix}$. We say that the ith column is hung if $\mathbf{M}[i] = \begin{pmatrix} \mathsf{E}_2(pp,y,(i,0);\rho_{i,0}) \\ \mathsf{E}_2(pp,y,(i,1);\rho_{i,1}) \end{pmatrix}$. Note that if u is a true image (i.e., $u := \mathsf{TDF}(ik, x)$ for some $x \in \{0,1\}^n$), then each column is either correct or hung.*

Following the notation of Construction 5.2, an ith correct column means that the inversion algorithm $\mathsf{TDF.F}^{-1}(tk, u)$ can recover the ith bit of z.

Having defined the function NCor, we may easily verify that the proof of Lemma 5.3 gives us the following corollary. We will use this corollary in our CCA proofs.

Corollary 5.8. *Let (n, c) be the parameters of the underlying optimal erasure code $(\mathsf{Encode}, \mathsf{Decode})$. For any $c > 4$ we have*

$$\Pr_{(ik,tk)}[\exists x' \in \{0,1\}^n \ s.t. \ \mathsf{NCor}(tk, \mathsf{TDF.F}(ik, x')) < 2n] \leq 2^n \cdot e^{\frac{-n(c-4)^2}{2c}}, \tag{19}$$

where $(ik, tk) \leftarrow \mathsf{TDF.KG}(1^\lambda)$. As a result, by taking $c \geq 8$, we have

$$\Pr_{(ik,tk)}[\exists x' \in \{0,1\}^n \ s.t. \ \mathsf{NCor}(tk, \mathsf{TDF.F}(ik, x')) < 2n] = \mathsf{negl}(\lambda), \tag{20}$$

Proof. Exactly the same as in the proof of Lemma 5.3, except that now we want to bound $\Pr[\sum_{i=1}^{m} X_i < 2n]$. Using the Chernoff inequality (Theorem 2.5 with $p = 1/2$ and $\varepsilon = 1/2 - 2/c$, where $m = c \cdot n$), we have: $\Pr[\sum_{i=1}^{m} X_i < n] \le e^{-\frac{n(c-4)^2}{2c}}$. We conclude using a union bound over all $x \in \{0,1\}^n$. $\square$

5.2 Proof of CPA Security: Lemma 5.4

Proof. The proof of this lemma is similar to the proof of Lemma 5.4. More specifically, for any distribution $\mathcal{S}$ over $\{0,1\}^n$, in the same way as in the proof of Lemma 5.4, we may show that the sole security-for-encryption requirement of the OWFE implies that

$$(\mathsf{x}, \mathsf{ik}, \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x})) \stackrel{c}{\equiv} (\mathsf{x}, \mathsf{Sim}(\mathsf{pp}, m, \mathsf{y})), \tag{21}$$

where $\mathsf{x} \stackrel{\$}{\leftarrow} \mathcal{S}$, $(\mathsf{ik}, *) \stackrel{\$}{\leftarrow} \mathsf{TDF.KG}(1^\lambda)$, $\mathsf{z} := \mathsf{Encode}(\mathsf{x})$ and $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{z})$.

Given Equation 21 we show how to derive both the one-wayness and indistinguishability claims of the lemma.

For one-wayness, if there exists an inverter $\mathcal{A}$ that with non-negligible probability can compute x from $(\mathsf{ik}, \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x})$ — where $(\mathsf{ik}, *) \stackrel{\$}{\leftarrow} \mathsf{TDF.KG}(1^\lambda)$ and $\mathsf{x} \stackrel{\$}{\leftarrow} \{0,1\}^n$ — then Equation 21 implies that with non-negligible probability the adversary $\mathcal{A}$ can compute x (and hence $\mathsf{z} := \mathsf{Encode}(\mathsf{x})$) from $\mathsf{Sim}(\mathsf{pp}, m, \mathsf{y})$, where $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{z})$. Now since $\mathsf{z} \in \{0,1\}^m$ and z has min entropy n, this latter condition violates the (n, m)-one-wayness of f, because the computation of $\mathsf{Sim}(\mathsf{pp}, m, \mathsf{y})$ may be done efficiently with knowledge of pp, m and y.

For indistinguishability security (Definition 2.3) let $\mathcal{S}_0$ and $\mathcal{S}_1$ be two (k, n) sources on the TDF input and assume that the underlying OWFE scheme is (k, m)-smooth (Definition 3.3).

Letting $(\mathsf{ik}, *) \stackrel{\$}{\leftarrow} \mathsf{TDF.KG}(1^\lambda)$, $\mathsf{x}_0 \stackrel{\$}{\leftarrow} \mathcal{S}_0$, $\mathsf{x}_1 \stackrel{\$}{\leftarrow} \mathcal{S}_1$, $\mathsf{z}_0 := \mathsf{Encode}(\mathsf{x}_0)$, $\mathsf{z}_1 := \mathsf{Encode}(\mathsf{x}_1)$, $\mathsf{y}_0 := \mathsf{f}(\mathsf{pp}, \mathsf{z}_0)$ and $\mathsf{y}_1 := \mathsf{f}(\mathsf{pp}, \mathsf{z}_1)$, by Equation 21 we have

$$(\mathsf{ik}, \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x}_0)) \stackrel{c}{\equiv} \mathsf{Sim}(\mathsf{pp}, m, \mathsf{y}_0) \stackrel{c}{\equiv} \mathsf{Sim}(\mathsf{pp}, m, \mathsf{y}_1) \stackrel{c}{\equiv} (\mathsf{ik}, \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x}_1)),$$

where the second indistinguishability follows from the (k, m)-smoothness of the OWFE scheme, which states $(\mathsf{pp}, \mathsf{y}_0) \stackrel{c}{\equiv} (\mathsf{pp}, \mathsf{y}_1)$. The proof is now complete. $\square$

5.3 Proof of CCA2 Security: Theorem 5.5

In this section we show that our TDF is CCA2 indistinguishable.

Notation 5.9. *For an OWFE scheme* $(\mathsf{K}, \mathsf{f}, \mathsf{E}_1, \mathsf{E}_2, \mathsf{D})$, *letting* $\mathbf{P} := \left(\begin{smallmatrix} \rho_{1,0}, \rho_{2,0}, \ldots, \rho_{m,0} \\ \rho_{1,1}, \rho_{2,1}, \ldots, \rho_{m,1} \end{smallmatrix} \right)$ *we define*

$\mathsf{E}(\mathsf{pp}, \mathsf{y}, \mathbf{P}) \stackrel{\triangle}{=}$
$$\begin{pmatrix} \mathsf{E}_1(\mathsf{pp}, (1,0); \rho_{1,0}), \ldots, \mathsf{E}_1(\mathsf{pp}, (m,0); \rho_{m,0}) \\ \mathsf{E}_1(\mathsf{pp}, (1,1); \rho_{1,1}), \ldots, \mathsf{E}_1(\mathsf{pp}, (m,1); \rho_{m,1}) \end{pmatrix}, \begin{pmatrix} \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (1,0); \rho_{1,0}), \ldots, \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (m,0); \rho_{m,0}) \\ \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (1,1); \rho_{1,1}), \ldots, \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (m,1); \rho_{m,1}) \end{pmatrix}.$$

Half-trapdoor keys. In the proof of Theorem 5.5 we will make use of an alternative way of inversion which works with respect to knowledge of half of all the randomness values that were

19

fixed in the trapdoor key. We refer to such trapdoor keys as *half* trapdoor keys (or simulated trapdoor keys). Recall that a real trapdoor key is of the form

$$(\mathsf{pp}, \mathbf{a}, (\rho_{1,0}, \rho_{1,1}), \ldots, (\rho_{m,0}, \rho_{m,1})). \tag{22}$$

A half-trapdoor key is a reduced version of a full trapdoor key in that we forget one randomness value from each pair, while remembering whether we chose to keep the first or the second coordinate of that pair. Formally, given a full trapdoor key as in Equation (22), a half trapdoor key is obtained based on a string $\mathsf{s} \in \{0,1\}^m$ as $\mathsf{tk}_{\mathrm{rd}} := (\mathsf{pp}, \mathbf{a}, \mathsf{s}, (\rho_1, \ldots, \rho_m))$, where $\rho_i = \rho_{i,\mathsf{s}_i}$. (The subscript rd stands for "reduced.")

We will now define how to perform inversion w.r.t. half-trapdoor keys.

Definition 5.10 (Half-trapdoor inversion $\mathsf{TDF.F}_{\mathrm{rd}}^{-1}$). *For an image* $\mathsf{u} := (\mathsf{y}, \mathbf{M})$ *of our constructed TDF and a half-trapdoor key* $\mathsf{tk}_{\mathrm{rd}} := (\mathsf{pp}, \mathbf{a}, \mathsf{s}, (\rho_1, \ldots, \rho_m))$ *we define* $\mathsf{TDF.F}_{\mathrm{rd}}^{-1}(\mathsf{tk}_{\mathrm{rd}}, \mathsf{u})$ *as follows:*

1. *If* $\mathsf{RSum}(\mathbf{M}) \neq \mathbf{a} \in \mathbb{Z}_2^m$, *then return* $\perp$.

2. *Construct* $\mathsf{z}' \in \{0,1\}^m$ *bit by bit as follows. For all* $i \in [m]$, *we denote by* $\mathbf{M}[i] = \binom{\mathsf{e}_{i,0}}{\mathsf{e}_{i,1}}$ *the* i'*th column of* $\mathbf{M}$. *If* $\mathsf{e}_{i,\mathsf{s}_i} = 1 - \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (i, \mathsf{s}_i); \rho_i)$, *then set* $\mathsf{z}'_i = 1 - \mathsf{s}_i$; *otherwise set* $\mathsf{z}'_i = \perp$.

3. *Letting* $\mathsf{x} := \mathsf{Decode}(\mathsf{z})$, *if* $\mathsf{TDF.F}(\mathsf{ik}, \mathsf{x}) = \mathsf{u}$, *then return* x. *Otherwise, return* $\perp$.

As terminology, we say that $\mathsf{TDF.F}_{\mathrm{rd}}^{-1}(\mathsf{tk}_{\mathrm{rd}}, \mathsf{u})$ *is able to open the ith column of* $\mathbf{M}$ *if* $\mathsf{z}'_i \neq \perp$ *(i.e., if* $\mathsf{e}_{i,\mathsf{s}_i} = 1 - \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (i, \mathsf{s}_i); \rho_i))$.

We first fix some notation and will then prove a useful property about half-inversion simulation, which in turn will be used in the CCA2 proof.

Notation 5.11 (Half trapdoor keys). *For a given* $\mathsf{tk} := (\mathsf{pp}, \mathbf{a}, (\rho_{1,0}, \rho_{1,1}), \ldots, (\rho_{m,0}, \rho_{m,1}))$ *and* $\mathsf{z} \in \{0,1\}^m$ *we define* $\mathsf{tk}/\mathsf{z} \triangleq (\mathsf{pp}, \mathbf{a}, \mathsf{z}, \rho_{1,\mathsf{z}_1}, \ldots, \rho_{m,\mathsf{z}_m})$.

Let $(\mathsf{ik}, \mathsf{tk}) \overset{\$}{\leftarrow} \mathsf{TDF.KG}(1^\lambda)$. It is not hard to verify that for any $\mathsf{x} \in \{0,1\}^n$, letting $\mathsf{z} := \mathsf{Encode}(\mathsf{x})$ and $\mathsf{u} := \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x})$, we have $\mathsf{TDF.F}_{\mathrm{rd}}^{-1}(\mathsf{tk}/\mathsf{z}, \mathsf{u}) = \perp$. The following lemma shows that this is the only point where half-inversion simulation fails. Namely, for any input $\mathsf{x}' \in \{0,1\}^n \setminus \{\mathsf{x}\}$, the half-trapdoor inversion (w.r.t. tk/z) of the resulting image will be as good as inversion w.r.t. the full trapdoor.

Lemma 5.12 (Half-trapdoor inversion). *Fix* $\mathsf{x} \in \{0,1\}^n$ *and let* $\mathsf{z} := \mathsf{Encode}(\mathsf{x})$. *Assume* $c \geq 8$, *where recall that* (n, c) *denotes the parameters of the underlying erasure code* $(\mathsf{Encode}, \mathsf{Decode})$. *We have*

$$\Pr_{(\mathsf{ik}, \mathsf{tk})}[\exists \mathsf{x}' \in \{0,1\}^n \setminus \{\mathsf{x}\} \text{ s.t. } \mathsf{TDF.F}^{-1}(\mathsf{tk}, \mathsf{u}') \neq \mathsf{TDF.F}_{\mathrm{rd}}^{-1}(\mathsf{tk}_{\mathrm{rd}}, \mathsf{u}')] = \mathsf{negl}(\lambda),$$

where $(\mathsf{ik}, \mathsf{tk}) \overset{\$}{\leftarrow} \mathsf{TDF.KG}(1^\lambda)$, $\mathsf{u}' := \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x}')$ *and* $\mathsf{tk}_{\mathrm{rd}} := \mathsf{tk}/\mathsf{z}$.

Proof. Fix $\mathsf{x} \in \{0,1\}^n$ and let $\mathsf{z} := \mathsf{Encode}(\mathsf{x})$. For a sampled $(\mathsf{ik}, \mathsf{tk})$ we define the event Bad as

$$\mathrm{Bad} := \exists \mathsf{x}' \in (\{0,1\}^n \setminus \{\mathsf{x}\}) \text{ s.t. } \mathsf{TDF.F}^{-1}(\mathsf{tk}, \mathsf{u}') \neq \mathsf{TDF.F}_{\mathrm{rd}}^{-1}(\mathsf{tk}_{\mathrm{rd}}, \mathsf{u}'),$$

where $\mathsf{u}' := \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x}')$ and $\mathsf{tk}_{\mathrm{rd}} := \mathsf{tk}/\mathsf{z}$. We need to show $\Pr[\mathrm{Bad}] = \mathsf{negl}(\lambda)$.

First, note that if $\mathsf{TDF.F}^{-1}(\mathsf{tk}, \mathsf{TDF.F}(\mathsf{ik}, x')) = \perp$, then $\mathsf{TDF.F}_{\mathrm{rd}}^{-1}(\mathsf{tk}_{\mathrm{rd}}, \mathsf{TDF.F}(\mathsf{ik}, x')) = \perp$, and if $\mathsf{TDF.F}_{\mathrm{rd}}^{-1}(\mathsf{tk}_{\mathrm{rd}}, \mathsf{TDF.F}(\mathsf{ik}, x')) \neq \perp$, then $\mathsf{TDF.F}^{-1}(\mathsf{tk}, \mathsf{TDF.F}(\mathsf{ik}, x')) = x' = \mathsf{TDF.F}_{\mathrm{rd}}^{-1}(\mathsf{tk}_{\mathrm{rd}}, \mathsf{TDF.F}(\mathsf{ik}, x'))$. This follows from the descriptions of $\mathsf{TDF.F}_{\mathrm{rd}}^{-1}$ and D (see Definition 5.10).

Thus, defining

$$\mathrm{Bad}' : \exists x' \in (\{0,1\}^n \setminus \{x\}) \text{ s.t. } \left(\mathsf{TDF.F}_{\mathrm{rd}}^{-1}(\mathsf{tk}_{\mathrm{rd}}, (\mathsf{TDF}(\mathsf{ik}, x'))) = \perp \right),$$

we have $\Pr[\mathrm{Bad}] \leq \Pr[\mathrm{Bad}']$. In what follows, we will show $\Pr[\mathrm{Bad}'] = \mathsf{negl}(\lambda)$, and this will complete the proof.

Since $c \geq 8$, by Corollary 5.8 (Equation (20)) we have

$$\Pr_{(\mathsf{ik},\mathsf{tk})}[\forall x' \in \{0,1\}^n : \mathsf{NCor}(\mathsf{tk}, \mathsf{TDF.F}(\mathsf{ik}, x')) \geq 2n] = 1 - \mathsf{negl}(\lambda). \tag{23}$$

We will use Equation (23) to show that $\Pr[\mathrm{Bad}'] = \mathsf{negl}(\lambda)$. First, notice that if $x \neq x'$, then letting $z' := \mathsf{Encode}(x')$, we have $\mathsf{H}_{\mathrm{dst}}(z, z') \geq m - n$.

For all x', letting $u' := \mathsf{TDF.F}(\mathsf{ik}, x')$ and parsing $u' := (y', \mathbf{M}')$, by Equation (23) we know that $\mathbf{M}'$ has at least $2n$ correct columns (except with negligible probability). Recall that we say that the ith column $\mathbf{M}'[i]$ of $\mathbf{M}'$ is correct if

- $\mathbf{M}'[i] = \begin{pmatrix} \mathsf{E}_2(\mathsf{pp}, y', (i, 0); \rho_{i,0}) \\ 1 - \mathsf{E}_2(\mathsf{pp}, y', (i, 1); \rho_{i,1}) \end{pmatrix}$ if $z'_i = 0$; and

- $\mathbf{M}'[i] = \begin{pmatrix} 1 - \mathsf{E}_2(\mathsf{pp}, y', (i, 0); \rho_{i,0}) \\ \mathsf{E}_2(\mathsf{pp}, y', (i, 1); \rho_{i,1}) \end{pmatrix}$ if $z'_i = 1$,

where $z' := \mathsf{Encode}(x')$ and $\mathsf{tk} := (\mathsf{pp}, \mathbf{a}, (\rho_{1,0}, \rho_{1,1}), \ldots, (\rho_{m,0}, \rho_{m,1}))$.

We claim that except with negligible probability the following holds: for all $x' \in \{0,1\}^n \setminus \{x\}$, the algorithm $\mathsf{TDF.F}_{\mathrm{rd}}^{-1}(\mathsf{tk}_{\mathrm{rd}}, u')$ can "open" at least n columns of $\mathbf{M}'$, hence recovering at least n bits of z'. (See Definition 5.10 for the meaning of "open".) Toward proving this claim, for any given $x' \in \{0,1\}^n \setminus \{x\}$, let $\Gamma_{x'}$ denote the set of indices of the correct columns of $\mathbf{M}'$, and note that, except with negligible probability, for all x' we have $|\Gamma_{x'}| \geq 2n$. For any index $i \in \Gamma_{x'}$, the algorithm $\mathsf{TDF.F}_{\mathrm{rd}}^{-1}(\mathsf{tk}_{\mathrm{rd}}, u')$ fails to open the ith column of $\mathbf{M}'_i$ iff $z_i = z'_i$. Since $\mathsf{H}_{\mathrm{dst}}(z, z') \geq m - n$, we conclude that the execution of $\mathsf{TDF.F}_{\mathrm{rd}}^{-1}(\mathsf{tk}_{\mathrm{rd}}, u')$ can open at least n columns of $\mathbf{M}'$. By the optimal erasure resilience of $(\mathsf{Encode}, \mathsf{Decode})$, we have $\mathsf{TDF.F}_{\mathrm{rd}}^{-1}(\mathsf{tk}_{\mathrm{rd}}, u') = x'$, as desired. The proof is now complete. $\square$

Our CCA2 hybrids will also make use of a simulated way of producing index/trapdoor keys. This procedure is described below.

Definition 5.13 (Simulated TDF key generation). *We define a simulated key-generation algorithm for the TDF given in Construction 5.2. Let $(\mathsf{K}, \mathsf{f}, \mathsf{E}_1, \mathsf{E}_2, \mathsf{D})$ be the underlying OWFE scheme. The simulated key generation algorithm $\mathsf{TDF.KG}_{\mathrm{sim}}(\mathsf{pp}, y)$ takes pp and an image y of the function f as input, and outputs $(\mathsf{ik}, \mathsf{tk})$ formed as follows. Sample $\mathbf{P} \xleftarrow{\$} \mathsf{Rand}^{2 \times m}$ and set $(\mathbf{CT}, \mathbf{M}) := \mathsf{E}(\mathsf{pp}, y, \mathbf{P})$. (See Notation 5.9.) Set $\mathsf{ik} := (\mathsf{pp}, \mathsf{RSum}(\mathbf{M}), \mathbf{CT})$ and $\mathsf{tk} := (\mathsf{pp}, \mathsf{RSum}(\mathbf{M}), \mathbf{P})$.*

We will now describe the hybrids for proving CCA2 security of the deterministic encryption scheme. We define the hybrids with respect to a distribution $\mathcal{D}$ and will then instantiate the distribution in the subsequent lemmas.

Hybrid $H_0[\mathcal{D}]$: real game.

- **Index/trapdoor keys.** Sample $(\mathsf{ik}, \mathsf{tk}) \xleftarrow{\$} \mathsf{TDF.KG}(1^\lambda)$.

- **Challenge ciphertext.** Set $\mathsf{u} := \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x})$, where $\mathsf{x} \leftarrow \mathcal{D}$.

- **CCA2 inversion queries.** Reply to each inversion query $\mathsf{u}' \neq \mathsf{u}$ with $\mathsf{TDF.F}^{-1}(\mathsf{tk}, \mathsf{u}')$.

Hybrid $H_1[\mathcal{D}]$: half-trapdoor inversion. Same as H_0 except we reply to inversion queries using a half trapdoor and by using the algorithm $\mathsf{TDF.F}_{\mathrm{rd}}^{-1}$.

- **Index/trapdoor keys.** Sample $(\mathsf{ik}, \mathsf{tk}) \xleftarrow{\$} \mathsf{TDF.KG}(1^\lambda)$. Set the index key to be ik and form the trapdoor key as follows: sample $\mathsf{x} \leftarrow \mathcal{D}$, let $\mathsf{z} := \mathsf{Encode}(\mathsf{x})$ and set the trapdoor key to be $\mathsf{tk}_{\mathrm{rd}} := \mathsf{tk}/\mathsf{z}$ (Notation 5.11).

- **Challenge ciphertext.** Return $\mathsf{u} := \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x})$, where recall that x was sampled in the previous step.

- **CCA2 inversion queries.** Reply to each inversion query $\mathsf{u}' \neq \mathsf{u}$ with $\mathsf{TDF.F}_{\mathrm{rd}}^{-1}(\mathsf{tk}_{\mathrm{rd}}, \mathsf{u}')$.

Hybrid $H_2[\mathcal{D}]$: half-trapdoor inversion with a simulated index key. Same as $H_1[\mathcal{D}]$ except that we sample the index key and the challenge ciphertext jointly in a simulated way.

- **Index/trapdoor keys:**

  1. Sample $\mathsf{x} \leftarrow \mathcal{D}$, and let $\mathsf{z} := \mathsf{Encode}(\mathsf{x})$. Set $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{z})$.

  2. Sample $(\mathsf{ik}, \mathsf{tk}) \xleftarrow{\$} \mathsf{TDF.KG}_{\mathrm{sim}}(\mathsf{pp}, \mathsf{y})$.

  3. Set the index key to be ik and the trapdoor key to be $\mathsf{tk}_{\mathrm{rd}} := \mathsf{tk}/\mathsf{z}$.

- **Challenge ciphertext.** Return $\mathsf{u} := \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x})$, where recall that x was sampled above.

- **CCA2 inversion queries.** Reply to each inversion query $\mathsf{u}' \neq \mathsf{u}$ with $\mathsf{TDF.F}_{\mathrm{rd}}^{-1}(\mathsf{tk}_{\mathrm{rd}}, \mathsf{u}')$.

Hybrid $H_3[\mathcal{D}]$: Full trapdoor inversion with a simulated index key. Same as $H_2[\mathcal{D}]$ except we use tk as the trapdoor key (instead of $\mathsf{tk}_{\mathrm{rd}}$) and will reply to each CCA2 inversion query $\mathsf{u}' \neq \mathsf{u}$ with $\mathsf{TDF.F}^{-1}(\mathsf{tk}, \mathsf{u}')$. That is:

- **Index/trapdoor keys:**

  1. Sample $\mathsf{x} \leftarrow \mathcal{D}$, and let $\mathsf{z} := \mathsf{Encode}(\mathsf{x})$. Set $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{z})$.

  2. Let the index/trapdoor key be $(\mathsf{ik}, \mathsf{tk}) \xleftarrow{\$} \mathsf{TDF.KG}_{\mathrm{sim}}(\mathsf{pp}, \mathsf{y})$.

- **Challenge ciphertext.** Return $\mathsf{u} := \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x})$.

- **CCA2 inversion queries.** Reply to each inversion query $\mathsf{u}' \neq \mathsf{u}$ with $\mathsf{D}(\mathsf{tk}, \mathsf{u}')$.

The above concludes the description of the hybrids. We now define some notation and will then prove some lemmas.

Notation. For $i \in \{0, 1, 2, 3\}$ we use $\mathsf{out}_i[\mathcal{D}]$ to denote the output bit of an underlying adversary in hybrid $\mathsf{H}_i[\mathcal{D}]$. For $i, j \in \{0, 1, 2, 3\}$ and two distributions $\mathcal{S}_0$ and $\mathcal{S}_1$, we write $\mathsf{H}_i[\mathcal{S}_0] \stackrel{c}{\equiv} \mathsf{H}_j[\mathcal{S}_1]$ to mean that for all PPT adversaries $\mathcal{A}$ we have $|\Pr[\mathsf{out}_i[\mathcal{S}_0] = 1] - \Pr[\mathsf{out}_j[\mathcal{S}_1] = 1]| = \mathsf{negl}(\lambda)$.

The proof of Theorem 5.5 follows from the following lemmas.

Lemma 5.14 (Indistinguishability of Hybrids H_0 and H_1)**.** *Assuming $c \geq 8$ (i.e., $m \geq 8n$), for any distribution $\mathcal{D}$ and any PPT adversary $\mathcal{A}$ we have $|\Pr[\mathsf{out}_0[\mathcal{D}] = 1] - \Pr[\mathsf{out}_1[\mathcal{D}] = 1]| = \mathsf{negl}(\lambda)$.*

Lemma 5.15 (Indistinguishability of Hybrids H_1 and H_2)**.** *If the OWFE satisfies the security-for-encryption property, then for any distribution $\mathcal{D}$ and any PPT adversary $\mathcal{A}$, we have $|\Pr[\mathsf{out}_1[\mathcal{D}] = 1] - \Pr[\mathsf{out}_2[\mathcal{D}] = 1]| = \mathsf{negl}(\lambda)$.*

Lemma 5.16 (Indistinguishability of Hybrids H_2 and H_3)**.** *If the OWFE satisfies the security-for-encryption property and $c \geq 10$, then for any distribution $\mathcal{D}$ and any PPT adversary $\mathcal{A}$, we have $|\Pr[\mathsf{out}_2[\mathcal{D}] = 1] - \Pr[\mathsf{out}_3[\mathcal{D}] = 1]| = \mathsf{negl}(\lambda)$.*

Lemma 5.17 (CCA2 Security in H_3)**.** *If the OWFE is (k, m)-smooth, then for any two (k, n) sources $\mathcal{S}_0$ and $\mathcal{S}_1$ and any PPT adversary $\mathcal{A}$, we have $|\Pr[\mathsf{out}_3[\mathcal{S}_0] = 1] - \Pr[\mathsf{out}_3[\mathcal{S}_1] = 1]| = \mathsf{negl}(\lambda)$.*

Proof of Theorem 5.5. By applying the above lemmas, for any (k, n)-sources $\mathcal{S}_0$ and $\mathcal{S}_1$, we have:
$$\mathsf{H}_0[\mathcal{S}_0] \stackrel{c}{\equiv} \mathsf{H}_1[\mathcal{S}_0] \stackrel{c}{\equiv} \mathsf{H}_2[\mathcal{S}_0] \stackrel{c}{\equiv} \mathsf{H}_3[\mathcal{S}_0] \stackrel{c}{\equiv} \mathsf{H}_3[\mathcal{S}_1] \stackrel{c}{\equiv} \mathsf{H}_2[\mathcal{S}_1] \stackrel{c}{\equiv} \mathsf{H}_1[\mathcal{S}_1] \stackrel{c}{\equiv} \mathsf{H}_0[\mathcal{S}_1].$$

$\square$

We will first prove Lemma 5.17 and will then prove the other lemmas.

Proof of Lemma 5.17

Proof. The proof of this lemma follows from the fact that in Hybrid H_3 the whole view of a CCA2 adversary with a challenge ciphertext $\mathsf{u} := \mathsf{TDF}(\mathsf{ik}, \mathsf{x})$ may be simulated via knowledge of $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{Encode}(\mathsf{x}))$, and specially without knowing x.

Concretely, suppose for two (k, n) sources $\mathcal{S}_0$ and $\mathcal{S}_1$, and for a PPT adversary $\mathcal{A}$ we have $|\Pr[\mathsf{out}_3[\mathcal{S}_0] = 1] - \Pr[\mathsf{out}_3[\mathcal{S}_1] = 1]| > \mathsf{negl}(\lambda)$. Let $\mathcal{S}_0' := \mathsf{Encode}(\mathcal{S}_0)$ and $\mathcal{S}_1' := \mathsf{Encode}(\mathcal{S}_1)$, and note that both $\mathcal{S}_0'$ and $\mathcal{S}_1'$ are (k, m)-sources. We build an adversary $\mathcal{B}$ that breaks the (k, m)-smoothness of the underlying OWFE scheme $\mathcal{E} = (\mathsf{K}, \mathsf{f}, \mathsf{E}, \mathsf{D})$ w.r.t. the distributions $\mathcal{S}_0'$ and $\mathcal{S}_1'$.

Letting $(\mathsf{pp}, \mathsf{y})$ be the challenge input of $\mathcal{B}$, $\mathcal{B}$ does the following:

1. Sample $\mathbf{P} \stackrel{\$}{\leftarrow} \mathsf{Rand}^{2 \times m}$ and set $(\mathbf{CT}, \mathbf{M}) := \mathsf{E}(\mathsf{pp}, \mathsf{y}, \mathbf{P})$ (Notation 5.9). Set $\mathsf{ik} := (\mathsf{pp}, \mathsf{RSum}(\mathbf{M}), \mathbf{CT})$ and $\mathsf{tk} := (\mathsf{pp}, \mathsf{RSum}(\mathbf{M}), \mathbf{P})$.

2. Set $\mathsf{u} := (\mathsf{y}, \mathbf{M})$.

3. Run the CCA2 adversary $\mathcal{A}$ on input $(\mathsf{ik}, \mathsf{ct})$ and reply to to each inversion query $\mathsf{u}' \neq \mathsf{u}$ with $\mathsf{D}(\mathsf{tk}, \mathsf{u}')$. Return the bit output by $\mathcal{A}$.

It is now easy to verify that if y corresponds to $\mathsf{f}(\mathsf{pp}, \mathcal{S}_b')$ for $b \in \{0, 1\}$, then the output of $\mathcal{A}$ will be equally distributed to $\mathsf{out}_3[\mathcal{S}_b]$. The proof is now complete. $\square$

Proof of Lemma 5.14

Proof. Notice that the only difference between H_0 and H_1 lies in the way CCA2 decryption queries are handled: we reply to a permitted (i.e different from the challenge ciphertext) CCA2 query $\mathsf{u}' := (\mathsf{y}, \mathbf{M}')$ with $\mathsf{TDF.F}^{-1}(\mathsf{tk}, \mathsf{u}')$ in H_0 and reply to it with $\mathsf{TDF.F}_{\mathrm{rd}}^{-1}(\mathsf{tk}_{\mathrm{rd}}, \mathsf{u}')$ in H_1.

We prove that this difference is negligible, using Lemma 5.12. Notice that for a given CCA2 query u', if $\mathsf{u}' = \mathsf{u}$ (i.e., u' is the challenge image) or if u' is not a valid image (i.e., if there does not exist any x such that $\mathsf{TDF.F}(\mathsf{ik}, \mathsf{x}) = \mathsf{u}'$), then in both H_0 and H_1 we reply to u' with $\perp$. The reason for the latter is that at the end of both $\mathsf{TDF.F}^{-1}$ and $\mathsf{TDF.F}_{\mathrm{rd}}^{-1}$ we will check whether the evaluation of $\mathsf{TDF.F}(\mathsf{ik}, \cdot)$ on the recovered string gives back u'. Also, if u' is a valid image and $\mathsf{u}' \neq \mathsf{u}$, then Lemma 5.12 implies that this query is replied to with the same answer in both hybrids. $\square$

Proof of Lemma 5.15

For this proof, we need the following lemma.

Lemma 5.18. *Let* $(\mathsf{K}, \mathsf{f}, \mathsf{E}_1, \mathsf{E}_2, \mathsf{D})$ *be an OWFE scheme. For any* $\mathsf{z} \in \{0,1\}^m$ *we have*

$$(\mathsf{pp}, \mathbf{a}_1, \rho_{1,\mathsf{z}_1}, \ldots, \rho_{m,\mathsf{z}_m}, \mathsf{z}, \mathbf{CT}) \stackrel{c}{\equiv} (\mathsf{pp}, \mathbf{a}_2, \rho_{1,\mathsf{z}_1}, \ldots, \rho_{m,\mathsf{z}_m}, \mathsf{z}, \mathbf{CT}), \tag{24}$$

where $\mathsf{pp} \stackrel{\$}{\leftarrow} \mathsf{K}(1^\lambda)$, $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{z})$, $\mathbf{P} := \left(\begin{smallmatrix} \rho_{1,0}, \rho_{2,0}, \ldots, \rho_{m,0} \\ \rho_{1,1}, \rho_{2,1}, \ldots, \rho_{m,1} \end{smallmatrix} \right) \stackrel{\$}{\leftarrow} \mathsf{Rand}^{2 \times m}$, $(\mathbf{CT}, \mathbf{M}) \stackrel{\$}{\leftarrow} \mathsf{E}(\mathsf{pp}, \mathsf{y}, \mathbf{P})$, $\mathbf{a}_1 \stackrel{\$}{\leftarrow} \mathbb{Z}_2^m$ *and* $\mathbf{a}_2 := \mathsf{RSum}(\mathbf{M})$.

Proof. First, note that the security-for-encryption requirement of the OWFE implies that

$$\left(\mathsf{z}, \mathsf{pp}, \{\mathsf{E}_1(\mathsf{pp}, (i, 1 - \mathsf{z}_i)), \mathsf{e}_i\}_{i \in [m]} \right) \stackrel{c}{\equiv} \left(\mathsf{z}, \mathsf{pp}, \{\mathsf{E}_1(\mathsf{pp}, (i, 1 - \mathsf{z}_i)), \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (i, 1 - \mathsf{z}_i))\}_{i \in [m]} \right), \tag{25}$$

where $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{z})$ *and* $\mathsf{e}_1, \ldots, \mathsf{e}_m \stackrel{\$}{\leftarrow} \{0,1\}$.

Given a sample $(\mathsf{z}, \mathsf{pp}, \{\mathsf{ct}_i, \mathsf{u}_i\}_{i \in [m]})$ from either side of Equation (25), we show how to turn it into a sample from the same side of Equation (24). As notation, for two elements w_1 and w_2 and a bit b we define $\mathsf{Perm}(\mathsf{w}_1, \mathsf{w}_2, b)$ to output $(\mathsf{w}_1, \mathsf{w}_2)$ if $b = 0$, and $(\mathsf{w}_2, \mathsf{w}_1)$ otherwise.

To this end,

1. Sample $\rho_1, \ldots, \rho_m \stackrel{\$}{\leftarrow} \mathsf{Rand}$. For $i \in [m]$ let $(\mathsf{ct}_i', \mathsf{e}_i') := (\mathsf{E}_1(\mathsf{pp}, (i, \mathsf{z}_i)), \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (i, \mathsf{z}_i)))$.

2. Let $\mathbf{a} := (\mathsf{u}_1, \ldots, \mathsf{u}_m) + (\mathsf{e}_1', \ldots, \mathsf{e}_m')$.

3. Return $(\mathsf{pp}, \mathbf{a}, \rho_1, \ldots, \rho_m, \mathsf{z}, \mathsf{Perm}(\mathsf{ct}_1, \mathsf{ct}_1', \mathsf{z}_1), \ldots, \mathsf{Perm}(\mathsf{ct}_m, \mathsf{ct}_m', \mathsf{z}_m))$.

Using simple inspection we may verify that starting from a sample from either side of Equation (25), the above procedure produces a sample from the same side of Equation (24). The proof is now complete. $\square$

We now proceed to prove Lemma 5.15 (indistinguishability of $\mathsf{H}_1[\mathcal{D}]$ and $\mathsf{H}_2[\mathcal{D}]$).

Proof. We will prove this by showing that $(\mathsf{ik}, \mathsf{tk}_{\mathrm{rd}}, \mathsf{u})$ — which is the joint distribution of the index key, the simulated half-trapdoor key and the challenge ciphertext — is computationally indistinguishable between $\mathsf{H}_1[\mathcal{D}]$ and $\mathsf{H}_2[\mathcal{D}]$. This will imply the statement of the claim, because the entire view in either hybrid can be constructed via knowledge of the corresponding tuple $(\mathsf{ik}, \mathsf{tk}_{\mathrm{rd}}, \mathsf{u})$.

By Lemma 5.18, we have:

$$(\mathsf{pp}, \mathbf{a}_1, \rho_{1,\mathsf{z}_1}, \ldots, \rho_{m,\mathsf{z}_m}, \mathsf{z}, \mathbf{CT}) \stackrel{c}{\equiv} (\mathsf{pp}, \mathbf{a}_2, \rho_{1,\mathsf{z}_1}, \ldots, \rho_{m,\mathsf{z}_m}, \mathsf{z}, \mathbf{CT}), \tag{26}$$

where $\mathsf{pp} \stackrel{\$}{\leftarrow} \mathsf{K}(1^\lambda)$, $\mathsf{x} \stackrel{\$}{\leftarrow} \mathcal{D}$, $\mathsf{z} := \mathsf{Encode}(\mathsf{x})$, $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{z})$, $\mathbf{P} \stackrel{\$}{\leftarrow} \mathsf{Rand}^{2 \times m}$, $(\mathbf{CT}, \mathbf{M}) \stackrel{\$}{\leftarrow} \mathsf{E}(\mathsf{pp}, \mathsf{y}, \mathbf{P})$, $\mathbf{a}_1 \stackrel{\$}{\leftarrow} \mathbb{Z}_2^m$ and $\mathbf{a}_2 := \mathsf{RSum}(\mathbf{M})$.

Given a sample $\mathbf{v} := (\mathsf{pp}, \mathbf{a}, \rho_{1,\mathsf{z}_1}, \ldots, \rho_{m,\mathsf{z}_m}, \mathsf{z}, \mathbf{CT})$ from an unknown side of Equation (26) we may form $(\mathsf{ik}, \mathsf{tk}_{\mathrm{rd}})$ as follows. Let $\mathsf{ik} := (\mathsf{pp}, \mathbf{a}, \mathbf{CT})$, $\mathsf{tk}_{\mathrm{rd}} := (\mathsf{pp}, \mathbf{a}, \mathsf{z}, \rho_{1,\mathsf{z}_1}, \ldots, \rho_{m,\mathsf{z}_m})$ and $\mathsf{u} := (\mathsf{y}, \mathbf{M})$, where $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{z})$, and $\mathbf{M}$ is formed as follows. Parse $\mathbf{a} := (\mathsf{a}_1, \ldots, \mathsf{a}_m)$ and for $i \in [m]$ let $\mathsf{e}_i := \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (1, \mathsf{z}_1); \rho_{1,\mathsf{z}_1})$ and set

$$\mathbf{M} := (\mathsf{Perm}(\mathsf{e}_1, \mathsf{a}_1 - \mathsf{e}_1, \mathsf{z}_1), \ldots, \mathsf{Perm}(\mathsf{e}_m, \mathsf{a}_m - \mathsf{e}_m, \mathsf{z}_m)).$$

It is now easy to verify that if $\mathbf{v}$ is a sample from the left-side (resp., right-side) part of Equation (26), then the produced $(\mathsf{ik}, \mathsf{tk}_{\mathrm{rd}}, \mathsf{u})$ corresponds to $\mathsf{H}_1[\mathcal{D}]$ (resp., $\mathsf{H}_2[\mathcal{D}]$). The proof is now complete. $\square$

Proof of Lemma 5.16

Indistinguishability of $\mathsf{H}_2[\mathcal{D}]$ and $\mathsf{H}_3[\mathcal{D}]$. The proof of the final hop follows the same way of reasoning as the proof of the first hop. Namely, the main idea is to show that any honestly generated ciphertext has sufficiently many correct columns, leading to the same decryption w.r.t. full trapdoors versus half trapdoors. However, arguing about the number of correct claims requires a new proof since index keys are now generated in a simulated way (Definition 5.13) and thus the counting strategy given in Corollary 5.8 does not apply anymore. We now give the following lemma aimed at proving this required statement.

Lemma 5.19. *Let $(\mathsf{K}, \mathsf{f}, \mathsf{E}_1, \mathsf{E}_2, \mathsf{D})$ be an OWFE scheme. . Fix $\mathsf{x}, \mathsf{x}' \in \{0,1\}^n$ and let $\mathsf{z} := \mathsf{Encode}(\mathsf{x})$ and $\mathsf{z}' := \mathsf{Encode}(\mathsf{x}')$. Fix $i \in [m]$ to be an index for which $\mathsf{z}_i \neq \mathsf{z}'_i$. We have*

$$(\mathsf{pp}, \mathsf{z}, \mathsf{z}', \mathsf{e}_{i,0} + \mathsf{e}_{i,1}, \mathsf{e}'_{i,0} + \mathsf{e}'_{i,1}) \stackrel{c}{\equiv} (\mathsf{pp}, \mathsf{z}, \mathsf{z}', \mathsf{b}, \mathsf{b}'), \tag{27}$$

where $\mathsf{pp} \stackrel{\$}{\leftarrow} \mathsf{K}(1^\lambda)$, $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{z})$, $\mathsf{y}' := \mathsf{f}(\mathsf{pp}, \mathsf{z}')$, $\rho_0, \rho_1 \stackrel{\$}{\leftarrow} \{0,1\}^*$, $\mathsf{e}_{i,0} \stackrel{\$}{\leftarrow} \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (i,0); \rho_0)$, $\mathsf{e}_{i,1} \stackrel{\$}{\leftarrow} \mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (i,1); \rho_1)$, $\mathsf{e}'_{i,0} \stackrel{\$}{\leftarrow} \mathsf{E}_2(\mathsf{pp}, \mathsf{y}', (i,0); \rho_0)$, $\mathsf{e}'_{i,1} \stackrel{\$}{\leftarrow} \mathsf{E}_2(\mathsf{pp}, \mathsf{y}', (i,1); \rho_1)$ and $\mathsf{b}, \mathsf{b}' \stackrel{\$}{\leftarrow} \{0,1\}$.
In particular, letting $p_i := \Pr[\mathsf{e}_{i,0} + \mathsf{e}_{i,1} = \mathsf{e}'_{i,0} + \mathsf{e}'_{i,1}]$ we have $|p_i - 1/2| = \mathsf{negl}(\lambda)$.

Proof. The "in particular" part follows immediately from the first part, and so we focus on proving the first part of the lemma.

For the first part, by the security-for-encryption requirement of the OWFE we have

$$(\mathsf{pp}, \mathsf{z}, \mathsf{z}', \mathsf{ct}, \mathsf{ct}', \mathsf{e}, \mathsf{e}') \equiv (\mathsf{pp}, \mathsf{z}, \mathsf{z}', \mathsf{ct}, \mathsf{ct}', \mathsf{b}, \mathsf{b}'), \tag{28}$$

where $\rho, \rho' \xleftarrow{\$} \{0,1\}^*$, $(\mathsf{ct}, \mathsf{e}) \xleftarrow{\$} \mathsf{E}(\mathsf{pp}, \mathsf{y}, (i, 1 - \mathsf{z}_i); \rho)$, $(\mathsf{ct}', \mathsf{e}') \xleftarrow{\$} \mathsf{E}(\mathsf{pp}, \mathsf{y}', (i, 1 - \mathsf{z}'_i); \rho)$ and $\mathsf{b}, \mathsf{b}' \xleftarrow{\$} \{0,1\}$. Recall that we have $1 - \mathsf{z}_i = \mathsf{z}'_i$. Thus, by the correctness of the OWFE we have $\mathsf{E}_2(\mathsf{pp}, \mathsf{y}, (i, \mathsf{z}_i); \rho') = \mathsf{D}(\mathsf{pp}, \mathsf{z}, \mathsf{ct}')$ and $\mathsf{E}_2(\mathsf{pp}, \mathsf{y}', (i, \mathsf{z}'_i); \rho) = \mathsf{D}(\mathsf{pp}, \mathsf{z}', \mathsf{ct})$. With this intuition in mind, we show how to generically turn a sample from either side of Equation (28) into a sample from the same side of Equation (27).

Given $(\mathsf{pp}, \mathsf{z}, \mathsf{z}', \mathsf{ct}, \mathsf{ct}', \mathsf{e}_1, \mathsf{e}'_1)$ return

$$(\mathsf{pp}, \mathsf{z}, \mathsf{z}', \mathsf{e}_1 + \mathsf{D}(\mathsf{pp}, \mathsf{z}, \mathsf{ct}'), \mathsf{e}'_1 + \mathsf{D}(\mathsf{pp}, \mathsf{z}', \mathsf{ct})).$$

The proof is now complete. $\qquad\square$

Lemma 5.20. *Let (n, c) be the parameters of the underlying optimal erasure code $(\mathsf{Encode}, \mathsf{Decode})$. Fix $\mathsf{x} \in \{0,1\}^n$. Assuming $c \geq 10$, we have*

$$\Pr_{(\mathsf{ik},\mathsf{tk})}[\exists \mathsf{x}' \in \{0,1\}^n \ s.t. \ \mathsf{NCor}(\mathsf{tk}, \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x}')) < 2n] = \mathsf{negl}(\lambda), \tag{29}$$

where $\mathsf{pp} \xleftarrow{\$} \mathsf{K}(1^\lambda)$, $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{Encode}(\mathsf{x}))$ and $(\mathsf{ik}, \mathsf{tk}) \leftarrow \mathsf{TDF.KG}_{\mathrm{sim}}(\mathsf{pp}, \mathsf{y})$. Recall that NCor was defined in Notation 5.7.

Proof. As notation, for a TDF image $\mathsf{u} := (\mathsf{y}, \mathbf{M})$, where $\mathbf{M} \in \mathbb{Z}_2^{2 \times m}$, we define $\mathsf{vec}(\mathsf{u}) := \mathsf{RSum}(\mathbf{M})$. Fix $\mathsf{x}, \mathsf{x}' \in \{0,1\}^n$.

We will prove that

$$\Pr[\mathsf{H}_{\mathrm{dst}}(\mathbf{a}, \mathbf{a}') \leq 2n] \leq \frac{1}{\mathsf{e}^{\frac{18n}{16}}}, \tag{30}$$

where $\mathsf{pp} \xleftarrow{\$} \mathsf{K}(1^\lambda)$, $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{Encode}(\mathsf{x}))$, $(\mathsf{ik}, \mathsf{tk}) \xleftarrow{\$} \mathsf{TDF.KG}_{\mathrm{sim}}(\mathsf{pp}, \mathsf{y})$, $\mathsf{ik} := (\mathsf{pp}, *, \mathbf{a})$, $\mathsf{tk} := (\mathsf{pp}, \mathbf{P})$, $\mathsf{y}' := \mathsf{f}(\mathsf{pp}, \mathsf{Encode}(\mathsf{x}'))$, $(\mathbf{CT}', \mathbf{M}') := \mathsf{E}(\mathsf{pp}, \mathsf{y}', \mathbf{P})$ and $\mathbf{a}' := \mathsf{RSum}(\mathbf{M}')$. (See Notation 5.9 for the definition of $\mathsf{E}(\mathsf{pp}, \mathsf{y}, \mathbf{P})$.)

Using notation above, notice that we have $\mathsf{NCor}(\mathsf{tk}, \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x}')) = \mathsf{H}_{\mathrm{dst}}(\mathbf{a}, \mathbf{a}')$. Thus, letting α be the value of the probability in Equation (29), using a union bound we have $\alpha \leq \frac{2^n}{\mathsf{e}^{\frac{18n}{16}}} = \mathsf{negl}(\lambda)$, as desired.

It remains to prove Equation (30). Letting $\mathbf{b} := \mathbf{a} + \mathbf{a}'$, we need to show $\Pr[\sum_{i=1}^{10n} \mathbf{b}_i \leq 2n] \leq \frac{1}{\mathsf{e}^{\frac{18n}{16}}}$. Let $\mathsf{z} := \mathsf{Encode}(\mathsf{x})$ and $\mathsf{z}' := \mathsf{Encode}(\mathsf{x}')$, and let $\mathsf{S} \subseteq [m]$ be the set of indices i on which $\mathsf{z}_i \neq \mathsf{z}'_i$. Noting that $|\mathsf{S}| = m - n = 9n$, let $\mathbf{b}' \in \{0,1\}^{9n}$ denote the induced sub-string of $\mathbf{b}$ on the index set S. By Lemma 5.19, for any $i \in [9n]$, we have $|\Pr[\mathbf{b}'_i = 1] - 1/2| = \mathsf{negl}(\lambda)$, and in particular $\Pr[\mathbf{b}'_i] \geq \frac{17}{36}$. We now have

$$\Pr[\sum_{i=1}^{10n} \mathbf{b}_i \leq 2n] \leq \Pr[\sum_{i=1}^{9n} \mathbf{b}'_i \leq 2n] = \Pr[\frac{1}{9n} \sum_{i=1}^{9n} \mathbf{b}'_i \leq \frac{2}{9}]$$

$$= \Pr[\frac{1}{9n} \sum_{i=1}^{9n} \mathbf{b}'_i \leq \frac{17}{36} - \frac{1}{4}] \leq \frac{1}{\mathsf{e}^{2\frac{9n}{16}}} = \frac{1}{\mathsf{e}^{\frac{18n}{16}}},$$

where the last inequality follows from the Chernoff bound (Theorem 2.5). The proof is now complete.

$\qquad\square$

Equipped with Lemma 5.20 we now give the proof for the last hybrid hop.

Claim 5.21. *For any distribution $\mathcal{D}$ and any PPT adversary $\mathcal{A}$, we have $|\Pr[\mathsf{out}_2[\mathcal{D}] = 1] - \Pr[\mathsf{out}_3[\mathcal{D}] = 1]| = \mathsf{negl}(\lambda)$.*

Proof. The only difference between H_2 and H_3 lies in the way CCA2 queries are handled: we reply to a permitted CCA2 query $\mathsf{u}' := (\mathsf{y}, \mathbf{M}')$ with $\mathsf{TDF.F}^{-1}(\mathsf{tk}_{\mathrm{rd}}, \mathsf{u}')$ in H_2 and reply to it with $\mathsf{TDF.F}^{-1}_{\mathrm{rd}}(\mathsf{tk}, \mathsf{u}')$ in H_3, where $\mathsf{x} \xleftarrow{\$} \mathcal{D}$, $\mathsf{z} := \mathsf{Encode}(\mathsf{x})$, $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{z})$, $(\mathsf{ik}, \mathsf{tk}) \xleftarrow{\$} \mathsf{TDF.KG}_{\mathrm{sim}}(\mathsf{pp}, \mathsf{y})$ and $\mathsf{tk}_{\mathrm{rd}} := \mathsf{tk}/\mathsf{z}$. (See Notation 5.11.)

To prove that this difference is negligible, we will set to proceed as in the proof of Lemma 5.12: namely, we show that for any $\mathsf{x} \in \{0, 1\}^n$, letting $\mathsf{z} := \mathsf{Encode}(\mathsf{x})$, then

$$\Pr_{(\mathsf{ik},\mathsf{tk})}[\exists \mathsf{x}' \in \{0, 1\}^n \setminus \{\mathsf{x}\} \text{ s.t. } \mathsf{TDF.F}^{-1}(\mathsf{tk}, \mathsf{u}') \neq \mathsf{TDF.F}^{-1}_{\mathrm{rd}}(\mathsf{tk}_{\mathrm{rd}}, \mathsf{u}')] = \mathsf{negl}(\lambda), \qquad (31)$$

where $\mathsf{y} := \mathsf{f}(\mathsf{pp}, \mathsf{Encode}(\mathsf{x}))$, $(\mathsf{ik}, \mathsf{tk}) \xleftarrow{\$} \mathsf{TDF.KG}_{\mathrm{sim}}(\mathsf{pp}, \mathsf{y})$, $\mathsf{u}' := \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x}')$ and $\mathsf{tk}_{\mathrm{rd}} := \mathsf{tk}/\mathsf{z}$. In order to prove Equation (31), note that by Lemma 5.20 we have

$$\Pr_{(\mathsf{ik},\mathsf{tk})}[\exists \mathsf{x}' \in \{0, 1\}^n \text{ s.t. } \mathsf{NCor}(\mathsf{tk}, \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x}')) < 2n] = \mathsf{negl}(\lambda). \qquad (32)$$

Given Equation (32), exactly in the same way as in the proof of Lemma 5.12, we can show that decryptions with the half trapdoor $\mathsf{tk}_{\mathrm{rd}}$ will be identical to decryptions using the full trapdoor tk, hence obtaining Equation (31). The proof is now complete. $\square$

6 DDH-Based Lossy TDFs with Linear-Image Size

In this section we show how to use our erasure-resilient code techniques in order to give lossy trapdoor functions with images growing linearly in their inputs. This gives us the first linear-image size construction of lossy TDFs from DDH. We first review the notion of lossy TDFs from [PW08].

Definition 6.1 (Lossy TDFs [PW08])**.** *An (n, k)-lossy TDF $((n, k)$-LTDF$)$ is given by four PPT algorithms $\mathsf{TDF.KG}$, $\mathsf{TDF.KG}_{\mathrm{ls}}$, $\mathsf{TDF.F}$, $\mathsf{TDF.F}^{-1}$, where $\mathsf{TDF.KG}_{\mathrm{ls}}(1^\lambda)$ only outputs a single key (as opposed to a pair of keys), and where the following properties hold:*

- ***Correctness in real mode.*** *The TDF $(\mathsf{TDF.KG}, \mathsf{TDF.F}, \mathsf{TDF.F}^{-1})$ satisfies correctness in the sense of Definition 2.1.*

- ***k-Lossiness.*** *For any $\mathsf{ik}_{\mathrm{ls}} \in \mathsf{TDF.KG}_{\mathrm{ls}}(1^\lambda)$ we have $|\mathsf{TDF.F}(\mathsf{ik}_{\mathrm{ls}}, \{0, 1\}^n)| \leq 2^k$, where we use $\mathsf{TDF.F}(\mathsf{ik}_{\mathrm{ls}}, \{0, 1\}^n)$ to denote the set of all images of $\mathsf{TDF.F}(\mathsf{ik}_{\mathrm{ls}}, \cdot)$.*

- ***Indistinguishability of real and lossy modes.*** *We have $\mathsf{ik} \stackrel{c}{\equiv} \mathsf{ik}_{\mathrm{ls}}$, where $(\mathsf{ik}, *) \xleftarrow{\$} \mathsf{TDF.KG}(1^\lambda)$ and $\mathsf{ik}_{\mathrm{ls}} \xleftarrow{\$} \mathsf{TDF.KG}_{\mathrm{ls}}(1^\lambda)$.*

Lossiness rate. In the definition above, we refer to the fraction $1 - k/n$ as the *lossiness rate*, describing the fraction of the bits lost. Ideally, we want this fraction to be as close to 1 as possible, e.g., $1 - o(1)$.

Our LTDF construction makes use of the following notation.

Notation. Letting $x \in \{0,1\}^n$ and $\mathbf{M} := \begin{pmatrix} g_{1,0}, g_{2,0} \cdots, g_{n,0} \\ g_{1,1}, g_{2,1}, \ldots, g_{n,1} \end{pmatrix}$ we define $x \odot \mathbf{M} = \prod_{j \in [n]} g_{j,x_j}$. For $i \in [n]$ and $b \in \{0,1\}$, we define the matrix $\mathbf{M}' := (\mathbf{M} \xrightarrow[(i,b)]{} g')$ to be the same as $\mathbf{M}$ except that instead of $g_{i,b}$ we put g' in $\mathbf{M}'$. If $\mathbf{M}$ is matrix of group elements, then $\mathbf{M}^r$ denotes entry-wise exponentiation to the power of r.

Construction 6.2 (Linear-image lossy TDF)**.** *Let* G *be a group scheme and let* $(\mathsf{Encode}, \mathsf{Decode})$ *be an optimal erasure code, where* $\mathsf{Encode}\colon \{0,1\}^n \to \{0,1\}^m$ *(Definition 5.1). We define our LTDF construction* $(\mathsf{TDF.KG}, \mathsf{TDF.KG_{ls}}, \mathsf{TDF.F}, \mathsf{TDF.F}^{-1})$ *as follows.*

- $\mathsf{TDF.KG}(1^\lambda)$*:*

  1. *Sample* $(\mathbb{G}, p, g) \xleftarrow{\$} \mathsf{G}(1^\lambda)$*, and*

  $$\mathbf{M} := \begin{pmatrix} g_{1,0}, g_{2,0}, \ldots, g_{m,0} \\ g_{1,1}, g_{2,1}, \ldots, g_{m,1} \end{pmatrix} \xleftarrow{\$} \mathbb{G}^{2 \times m}. \tag{33}$$

  2. *For all* $i \in [m]$*, sample* $g_i \xleftarrow{\$} \mathbb{G}$*,* $r_i \xleftarrow{\$} \mathbb{Z}_p$ *and* $b_i \xleftarrow{\$} \{0,1\}$*.*

  3. *Set the index and trapdoor keys as*

  $$\mathsf{ik} := (\mathbf{M}, (\mathbf{M}^{r_1} \xrightarrow[(1,b_1)]{} g_1), \ldots, (\mathbf{M}^{r_n} \xrightarrow[(m,b_m)]{} g_n)) \tag{34}$$

  $$\mathsf{tk} := (\mathbf{M}, (r_1, b_1, g_1), \ldots, (r_m, b_m, g_m)). \tag{35}$$

- $\mathsf{TDF.KG_{ls}}(1^\lambda)$*: Return* $\mathsf{ik_{ls}} := (\mathbf{M}, \mathbf{M}^{r_1}, \ldots, \mathbf{M}^{r_m})$*, where* $\mathbf{M}$ *and* r_i *for* $i \in [m]$ *are sampled as above.*

- $\mathsf{TDF.F}(\mathsf{ik}, x \in \{0,1\}^n)$*: Parse* $\mathsf{ik} := (\mathbf{M}, \mathbf{M}_1, \ldots, \mathbf{M}_m)$*. Set* $z := \mathsf{Encode}(x)$ *and return*

  $$u := (z \odot \mathbf{M}, \mathsf{HC}(z \odot \mathbf{M}_1), \ldots, \mathsf{HC}(z \odot \mathbf{M}_m)). \tag{36}$$

- $\mathsf{TDF.F}^{-1}(\mathsf{tk}, u)$*:*

  1. *Parse* $\mathsf{tk} := (\mathbf{M}, (r_1, b_1, g_1), \ldots, (r_m, b_m, g_m))$ *and* $u := (g_c, b_1', \ldots, b_m')$*.*

  2. *Construct* $z' := z_1' \cdots z_m' \in \{0, 1, \bot\}^m$ *as follows. For* $i \in [m]$*: if* $b_i' \neq \mathsf{HC}(g_c^{r_i})$*, then set* $z_i' = b_i$*; otherwise, set* $z_i' = \bot$*.*

  3. *Return* $\mathsf{Decode}(z')$*.*

The following theorem gives the lossiness property of the scheme.

Theorem 6.3 (Linear-image LTDF from DDH)**.** *Let* (n, c) *be the parameters of the underlying optimal erasure code* $(\mathsf{Encode}, \mathsf{Decode})$ *(Definition 5.1). Assuming* $c \geq 13$*, the LTDF of Construction 6.2 is* $(n, \log p)$*-lossy with image size* $\log p + cn \in \Theta(n)$*. By setting* $n \in \omega(\log p)$ *we obtain* $1 - o(1)$ *lossiness rate.*

We prove all the required properties below.

Lemma 6.4 (log p-Lossiness). *For any* $\mathsf{ik}_{\mathrm{ls}} \in \mathsf{TDF.KG}_{\mathrm{ls}}(1^\lambda)$ *we have* $|\mathsf{TDF.F}(\mathsf{ik}_{\mathrm{ls}}, \{0, 1\}^n)| \leq p$, *where recall that p is the order of the underlying group.*

Proof. Parse $\mathsf{ik}_{\mathrm{ls}} := (\mathbf{M}, \mathbf{M}_1, \ldots, \mathbf{M}_m)$. It is easy to verify that for any two messages $\mathsf{x}, \mathsf{x}' \in \{0, 1\}^n$ we have

$$\mathsf{TDF.F}(\mathsf{ik}_{\mathrm{ls}}, \mathsf{x}) \neq \mathsf{TDF.F}(\mathsf{ik}_{\mathrm{ls}}, \mathsf{x}') \Longleftrightarrow \mathsf{x} \odot \mathbf{M} \neq \mathsf{x}' \odot \mathbf{M}. \tag{37}$$

The statement of the lemma now follows, since $\{\mathsf{x} \odot \mathbf{M} \mid \mathsf{x} \in \{0, 1\}^n\} \subseteq \mathbb{G}$, and thus we have $|\{\mathsf{x} \odot \mathbf{M} \mid \mathsf{x} \in \{0, 1\}^n\}| \leq p$. □

Lemma 6.5 (Indistinguishability of real and lossy modes). *We have* $\mathsf{ik} \stackrel{c}{\equiv} \mathsf{ik}_{\mathrm{ls}}$, *where* $(\mathsf{ik}, *) \stackrel{\$}{\leftarrow}$ $\mathsf{TDF.KG}(1^\lambda)$ *and* $\mathsf{ik}_{\mathrm{ls}} \stackrel{\$}{\leftarrow} \mathsf{TDF.KG}_{\mathrm{ls}}(1^\lambda)$.

Proof. Immediate by the DDH assumption. □

Lemma 6.6 (Correctness in real mode). *Let (n, c) be the parameters of the underlying optimal erasure code* $(\mathsf{Encode}, \mathsf{Decode})$. *For any $c \geq 4$, we have*

$$\Pr_{(\mathsf{ik},\mathsf{tk})}[\exists \mathsf{x} \in \{0, 1\}^n \ s.t. \ \mathsf{TDF.F}^{-1}(\mathsf{tk}, (\mathsf{TDF.F}(\mathsf{ik}, \mathsf{x}))) \neq \mathsf{x}] \leq 2^n \cdot e^{\frac{-n(c-4)^2}{8c}}. \tag{38}$$

In particular, if $c \geq 13$, the probability in Equation (38) will be negligible.

Proof. Fix $\mathsf{x} \in \{0, 1\}^n$ and let $\mathsf{z} := \mathsf{Encode}(\mathsf{x})$. All probabilities below are taken over the random choice of $(\mathsf{ik}, \mathsf{tk})$. Parse

$$\mathsf{tk} := \left(\mathbf{M} := \begin{pmatrix} g_{1,0}, g_{2,0}, \ldots, g_{m,0} \\ g_{1,1}, g_{2,1}, \ldots, g_{m,1} \end{pmatrix}, (r_1, b_1, g_1), \ldots, (r_m, b_m, g_m) \right).$$

For input $\mathsf{x} \in \{0, 1\}^n$, let Fail_x be the event that $\mathsf{TDF.F}^{-1}(\mathsf{tk}, \mathsf{TDF.F}(\mathsf{ik}, \mathsf{x})) \neq \mathsf{x}$. Fix $\mathsf{x} \in \{0, 1\}^n$ and let $\mathsf{z} := \mathsf{Encode}(\mathsf{x}) \in \{0, 1\}^m$. As notation, for $i \in [m]$ define

$$\mathsf{w}_i := \prod_{j \in [m] \setminus \{i\}} g_{j, \mathsf{z}_j} \tag{39}$$

For $i \in [m]$ we define an event Good_i which corresponds to the event that the inversion algorithm manages to recover z_i. Formally, define Good_i as follows:

- Good_i: The event that $b_i = \mathsf{z}_i$ and $\underbrace{\mathsf{HC}(g_i \cdot \mathsf{w}_i^{r_i})}_{\text{evaluation}} \neq \underbrace{\mathsf{HC}(g_{i,b_i}^{r_i} \cdot \mathsf{w}_i^{r_i})}_{\text{inversion}}$.

Note that all Good_i are i.i.d. events and we have $\Pr[\mathsf{Good}_i] = 1/4$. The reason for this is that all of $(g_1, \ldots, g_n)$ are sampled uniformly at random independently of all other values.

We now have $\Pr[\mathsf{Fail}_\mathsf{x}] = \Pr[\sum_{i \in [m]} \mathsf{Good}_i < n]$.

Thus, using the Chernoff inequality (Theorem 2.5 with $p = 1/4$ and $\varepsilon = 1/4 - 1/c$, where $m = c \cdot n$), we have: $\Pr[\sum_{i=1}^m \mathsf{Good}_i < n] \leq e^{-\frac{n(c-4)^2}{8c}}$. We conclude using a union bound over all $\mathsf{x} \in \{0, 1\}^n$. □

7 Acknowledgements

We would like to thank Xiao Liang for suggesting a simplification to Construction 4.2.

References

[BBN+09] Mihir Bellare, Zvika Brakerski, Moni Naor, Thomas Ristenpart, Gil Segev, Hovav Shacham, and Scott Yilek. Hedged public-key encryption: How to protect against bad randomness. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 232–249, Tokyo, Japan, December 6–10, 2009. Springer, Heidelberg, Germany. 2

[BBO07] Mihir Bellare, Alexandra Boldyreva, and Adam O'Neill. Deterministic and efficiently searchable encryption. In Alfred Menezes, editor, *CRYPTO 2007*, volume 4622 of *LNCS*, pages 535–552, Santa Barbara, CA, USA, August 19–23, 2007. Springer, Heidelberg, Germany. 2, 4

[BCPT13] Eleanor Birrell, Kai-Min Chung, Rafael Pass, and Sidharth Telang. Randomness-dependent message security. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 700–720, Tokyo, Japan, March 3–6, 2013. Springer, Heidelberg, Germany. 2

[BFO08] Alexandra Boldyreva, Serge Fehr, and Adam O'Neill. On notions of security for deterministic encryption, and efficient constructions without random oracles. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 335–359, Santa Barbara, CA, USA, August 17–21, 2008. Springer, Heidelberg, Germany. 2, 3, 4, 5, 8, 9

[BFOR08] Mihir Bellare, Marc Fischlin, Adam O'Neill, and Thomas Ristenpart. Deterministic encryption: Definitional equivalences and constructions without random oracles. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 360–378, Santa Barbara, CA, USA, August 17–21, 2008. Springer, Heidelberg, Germany. 2, 4

[BHY09] Mihir Bellare, Dennis Hofheinz, and Scott Yilek. Possibility and impossibility results for encryption and commitment secure under selective opening. In Antoine Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 1–35, Cologne, Germany, April 26–30, 2009. Springer, Heidelberg, Germany. 2

[BKK+95] Johannes Bloemer, Malik Kalfane, Richard Karp, Marek Karpinski, Michael Luby, and David Zuckerman. An XOR-based erasure-resilient coding scheme. 1995. 16

[BLSV18] Zvika Brakerski, Alex Lombardi, Gil Segev, and Vinod Vaikuntanathan. Anonymous IBE, leakage resilience and circular security from new assumptions. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 535–564, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Heidelberg, Germany. 7

[BS11] Zvika Brakerski and Gil Segev. Better security for deterministic public-key encryption: The auxiliary-input setting. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 543–560, Santa Barbara, CA, USA, August 14–18, 2011. Springer, Heidelberg, Germany. 4

[BW10]   Xavier Boyen and Brent Waters. Shrinking the keys of discrete-log-type lossy trapdoor functions. In Jianying Zhou and Moti Yung, editors, *ACNS 10*, volume 6123 of *LNCS*, pages 35–52, Beijing, China, June 22–25, 2010. Springer, Heidelberg, Germany. 7

[CDG+17]   Chongwon Cho, Nico Döttling, Sanjam Garg, Divya Gupta, Peihan Miao, and Antigoni Polychroniadou. Laconic oblivious transfer and its applications. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 33–65, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Heidelberg, Germany. 7

[DG17a]   Nico Döttling and Sanjam Garg. From selective IBE to full IBE and selective HIBE. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 372–408, Baltimore, MD, USA, November 12–15, 2017. Springer, Heidelberg, Germany. 7

[DG17b]   Nico Döttling and Sanjam Garg. Identity-based encryption from the Diffie-Hellman assumption. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 537–569, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Heidelberg, Germany. 7

[DGHM18]   Nico Döttling, Sanjam Garg, Mohammad Hajiabadi, and Daniel Masny. New constructions of identity-based and key-dependent message secure encryption schemes. In Michel Abdalla and Ricardo Dahab, editors, *PKC 2018, Part I*, volume 10769 of *LNCS*, pages 3–31, Rio de Janeiro, Brazil, March 25–29, 2018. Springer, Heidelberg, Germany. 7

[DH76]   Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976. 1

[FGK+10]   David Mandell Freeman, Oded Goldreich, Eike Kiltz, Alon Rosen, and Gil Segev. More constructions of lossy and correlation-secure trapdoor functions. In Phong Q. Nguyen and David Pointcheval, editors, *PKC 2010*, volume 6056 of *LNCS*, pages 279–295, Paris, France, May 26–28, 2010. Springer, Heidelberg, Germany. 3, 7

[GH18]   Sanjam Garg and Mohammad Hajiabadi. Trapdoor functions from the computational Diffie-Hellman assumption. LNCS, pages 362–391, Santa Barbara, CA, USA, 2018. Springer, Heidelberg, Germany. 2, 3, 4, 5, 8, 9, 10, 11

[GMR01]   Yael Gertner, Tal Malkin, and Omer Reingold. On the impossibility of basing trapdoor functions on trapdoor predicates. In *42nd FOCS*, pages 126–135, Las Vegas, NV, USA, October 14–17, 2001. IEEE Computer Society Press. 2

[HO12]   Brett Hemenway and Rafail Ostrovsky. Extended-DDH and lossy trapdoor functions. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 627–643, Darmstadt, Germany, May 21–23, 2012. Springer, Heidelberg, Germany. 7

[HO13]   Brett Hemenway and Rafail Ostrovsky. Building lossy trapdoor functions from lossy encryption. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 241–260, Bengalore, India, December 1–5, 2013. Springer, Heidelberg, Germany. 7

[ILL89]   Russell Impagliazzo, Leonid A. Levin, and Michael Luby. Pseudo-random generation from one-way functions (extended abstracts). In *21st ACM STOC*, pages 12–24, Seattle, WA, USA, May 15–17, 1989. ACM Press. 9

[KMO10]   Eike Kiltz, Payman Mohassel, and Adam O'Neill. Adaptive trapdoor functions and chosen-ciphertext security. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 673–692, French Riviera, May 30 – June 3, 2010. Springer, Heidelberg, Germany. 4, 5, 15

[MPRS12] Ilya Mironov, Omkant Pandey, Omer Reingold, and Gil Segev. Incremental deterministic public-key encryption. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 628–644, Cambridge, UK, April 15–19, 2012. Springer, Heidelberg, Germany. 4

[MY10]   Petros Mol and Scott Yilek. Chosen-ciphertext security from slightly lossy trapdoor functions. In Phong Q. Nguyen and David Pointcheval, editors, *PKC 2010*, volume 6056 of *LNCS*, pages 296–311, Paris, France, May 26–28, 2010. Springer, Heidelberg, Germany. 2

[PVW08]   Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 554–571, Santa Barbara, CA, USA, August 17–21, 2008. Springer, Heidelberg, Germany. 2

[PW08]   Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 187–196, Victoria, British Columbia, Canada, May 17–20, 2008. ACM Press. 2, 3, 4, 5, 6, 7, 27

[RS09]   Alon Rosen and Gil Segev. Chosen-ciphertext security via correlated products. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 419–436. Springer, Heidelberg, Germany, March 15–17, 2009. 2, 4, 5, 6, 15

[RSA78]   Ronald L. Rivest, Adi Shamir, and Leonard M. Adleman. A method for obtaining digital signature and public-key cryptosystems. *Communications of the Association for Computing Machinery*, 21(2):120–126, 1978. 1

[Wee12]   Hoeteck Wee. Dual projective hashing and its applications - lossy trapdoor functions and more. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 246–262, Cambridge, UK, April 15–19, 2012. Springer, Heidelberg, Germany. 2, 7

[Wic13]   Daniel Wichs. Barriers in cryptography with weak, correlated and leaky sources. In Robert D. Kleinberg, editor, *ITCS 2013*, pages 111–126, Berkeley, CA, USA, January 9–12, 2013. ACM. 4