

More Constructions of Lossy and Correlation-Secure Trapdoor Functions

David Mandell Freeman¹, Oded Goldreich², Eike Kiltz³,
Alon Rosen⁴, and Gil Segev²

¹ Stanford University, USA

dfreeman@cs.stanford.edu

² Weizmann Institute of Science, Rehovot 76100, Israel

{oded.goldreich,gil.segev}@weizmann.ac.il

³ CWI, Netherlands

kiltz@cwi.nl

⁴ Efi Arazi School of Computer Science, Herzliya Interdisciplinary Center (IDC),
Herzliya 46150, Israel
alon.rosen@idc.ac.il

Abstract. We propose new and improved instantiations of lossy trapdoor functions (Peikert and Waters, STOC '08), and correlation-secure trapdoor functions (Rosen and Segev, TCC '09). Our constructions widen the set of number-theoretic assumptions upon which these primitives can be based, and are summarized as follows:

- Lossy trapdoor functions based on the quadratic residuosity assumption. Our construction relies on modular squaring, and whereas previous such constructions were based on seemingly stronger assumptions, we present the first construction that is based solely on the quadratic residuosity assumption.
- Lossy trapdoor functions based on the composite residuosity assumption. Our construction guarantees essentially any required amount of lossiness, where at the same time the functions are more efficient than the matrix-based approach of Peikert and Waters.
- Lossy trapdoor functions based on the d -Linear assumption. Our construction both simplifies the DDH-based construction of Peikert and Waters, and admits a generalization to the whole family of d -Linear assumptions without any loss of efficiency.
- Correlation-secure trapdoor functions related to the hardness of syndrome decoding.

Keywords: Public-key encryption, lossy trapdoor functions, correlation-secure trapdoor functions.

1 Introduction

In this paper, we describe new constructions of lossy trapdoor functions and correlation-secure trapdoor functions. These primitives are strengthened variants of the classical notion of trapdoor functions, and were introduced with the main

goal of enabling simple and black-box constructions of public-key encryption schemes that are secure against chosen-ciphertext attacks. At a high level, they are defined as follows:

Lossy trapdoor functions [25]: A collection of lossy trapdoor functions consists of two families of functions. Functions in one family are injective and can be efficiently inverted using a trapdoor. Functions in the other family are “lossy,” which means that the size of their image is significantly smaller than the size of their domain. The only computational requirement is that a description of a randomly chosen function from the family of injective functions is computationally indistinguishable from a description of a randomly chosen function from the family of lossy functions.

Correlation-secure trapdoor functions [26]: The classical notion of a one-way function asks for a function that is efficiently computable but is hard to invert given the image of a uniformly chosen input. Correlation security generalizes the one-wayness requirement by considering k -wise products of functions and any specified input distribution, not necessarily the uniform distribution. Given a collection of functions \mathcal{F} and a distribution \mathcal{C} over k -tuples of inputs, we say that \mathcal{F} is secure under \mathcal{C} -correlated inputs if the function $(f_1(x_1), \dots, f_k(x_k))$ is one-way, where f_1, \dots, f_k are independently chosen from \mathcal{F} and (x_1, \dots, x_k) are sampled from \mathcal{C} .

Lossy trapdoor functions were introduced by Peikert and Waters [25], who showed that they imply fundamental cryptographic primitives, such as trapdoor functions, collision-resistant hash functions, oblivious transfer, and CCA-secure public-key encryption. In addition, lossy trapdoor functions have already found various other applications, including deterministic public-key encryption [3], OAEP-based public-key encryption [17], “hedged” public-key encryption for protecting against bad randomness [1], security against selective opening attacks [2], and efficient non-interactive string commitments [22].

The notion of correlation security was introduced by Rosen and Segev [26], who showed that any collection of injective trapdoor functions that is one-way under a natural input distribution can be used to construct a CCA-secure public-key encryption scheme.¹ They showed that any collection of lossy trapdoor functions that are sufficiently lossy is in fact also correlation-secure. This result was recently refined by Mol and Yilek [19] who showed that even lossiness of any polynomial fraction of a single bit suffices.

These applications motivate us to investigate new constructions of lossy and correlation-secure functions. Such constructions would enable us to widen the basis upon which one can achieve the above cryptographic tasks in a simple and modular way.

¹ Any distribution where (x_1, \dots, x_k) are $(1 - \epsilon)k$ -wise independent, for a constant $\epsilon < 1$, can be used in their framework. In particular, this includes the case where x_1 is uniformly distributed and $x_1 = \dots = x_k$.

1.1 Our Contributions

We propose new and improved constructions of lossy and correlation-secure trapdoor functions based on well-established number-theoretic assumptions (some of which were previously not known to imply either of the primitives). By directly applying the results of [25,26,19], we obtain new CCA-secure public-key encryption schemes based on these assumptions. Concretely, we present the following constructions:

1. *Lossy trapdoor permutations based on the quadratic residuosity assumption.* Our construction relies on Rabin’s modular squaring function and is based solely on the quadratic residuosity assumption. More precisely, the function is defined as $f(x) = x^2 \cdot \delta_{r,s}(x) \bmod N$, where $N = PQ$ is an RSA modulus and $\delta_{r,s}(\cdot)$ is a function indexed by two public elements $r, s \in \mathbb{Z}_N$ serving two independent purposes. First, it extends the modular squaring function to a permutation over \mathbb{Z}_N . Second, $f(x)$ loses the information about the sign of x if and only if s is a quadratic residue. Therefore, under the quadratic residuosity assumption f has one bit of lossiness. We note that although a function with only one bit of lossiness (or, more generally, with only a non-negligible amount of lossiness) is not necessarily a (strong) one-way function, it nevertheless can be used as a building block for constructing even a CCA-secure public-key encryption scheme (see [19,26]).
2. *Lossy trapdoor functions based on the composite residuosity assumption.* Our construction is based on the Damgård-Jurik encryption scheme [8] with additional insights by Damgård and Nielsen [9,10]. The Damgård-Jurik scheme is based on computations in the group $\mathbb{Z}_{N^{s+1}}$, where $N = PQ$ is an RSA modulus and $s \geq 1$ is an integer (it contains Paillier’s encryption scheme [23] as a special case by setting $s = 1$). At a high level, each function is described by a pair (pk, c) , where pk is a public key for the encryption scheme, and c is either an encryption of 1 (injective mode) or an encryption of 0 (lossy mode). By using the homomorphic properties of the encryption scheme, given such a ciphertext c and an element x , it is possible to compute either an encryption of x in the injective mode, or an encryption of 0 in the lossy mode. We note that this construction was concurrently and independently proposed by Boldyreva et al. [3]. We also give an “all-but-one” version of the construction.
3. *Lossy trapdoor functions based on the d -Linear assumption.* Our construction both simplifies and generalizes the DDH-based construction of Peikert and Waters [25, Section 5]. (Recall that DDH is the 1-Linear assumption.) Let \mathbb{G} be a finite group of order p and choose an $n \times n$ matrix M over \mathbb{F}_p that has either rank d (lossy mode) or rank n (injective mode). We “encrypt” $M = (a_{ij})$ as the matrix $g^M = (g^{a_{ij}}) \in \mathbb{G}^{n \times n}$, where g is a generator of \mathbb{G} . If \vec{x} is a binary vector of length n , then given g^M we can efficiently evaluate the function $f_M(\vec{x}) = g^{M\vec{x}} \in \mathbb{G}^n$. If M has rank n , then given M we can efficiently invert f_M on the image of $\{0, 1\}^n$. On the other hand, if M has rank d and $p < 2^{n/d}$, then f is lossy. The d -Linear assumption implies that the lossy and injective modes cannot be efficiently distinguished. We also give an “all-but-one” version of the function f_M based on the DDH assumption.

4. *Correlation-secure trapdoor functions based on the hardness of syndrome decoding.* Our construction is based on Niederreiter’s coding-based encryption system [21] which itself is the dual of the McEliece encryption system [18]. Our trapdoor function is defined as $f(x) = Hx$, where H is a binary $(n - k) \times n$ matrix (of a certain distribution that allows for embedding a trapdoor) and x is bit string of small Hamming weight. We show that the function’s correlation security is directly implied by a result of Fischer and Stern [12] about the pseudorandomness of the function f . Interestingly, the related McEliece trapdoor function (which can be viewed as the dual of the Niederreiter function) is not correlation-secure.² It is however possible to extend the framework of correlation security in a natural way to obtain a direct construction of a CCA-secure encryption scheme from the McEliece trapdoor function. This was recently demonstrated by Dowsley et al [11] (who proposed the first coding-based encryption scheme that is CCA-secure in the standard model) and, for the related lattice case, independently by Peikert [24] and Goldwasser and Vaikuntanathan [14]. Our contribution is to show that the Niederreiter function admits a simple construction of correlation-secure trapdoor functions based on the same security assumptions as [11].³ The resulting CCA-secure encryption scheme is as efficient as the one from [11].

1.2 Related Work

Most of the known constructions and applications of lossy and correlation-secure trapdoor functions are already mentioned above; here we include a few more. Besides their construction based on DDH, Peikert and Waters [25] also present a construction of lossy trapdoor functions based on the worst-case hardness of lattice problems. The construction does not enjoy the same amount of lossiness as their DDH-based one, but it still suffices for their construction of a CCA-secure public-key encryption scheme. The worst-case hardness of lattice problems is also used by Peikert [24] and by Goldwasser and Vaikuntanathan [14] to construct a CCA-secure encryption scheme using a natural generalization of correlation-secure trapdoor functions.

Kiltz et al. [17] show that the RSA trapdoor permutation is lossy under the \varPhi -Hiding assumption of Cachin et al. [6]. (Concretely, it has $\log_2(e)$ bits of lossiness, where e is the public RSA exponent.) Furthermore, they propose multi-prime hardness assumptions under which RSA has greater lossiness.

In concurrent and independent work, Mol and Yilek [19] propose a lossy trapdoor function based on the modular squaring function. Though this construction

² The McEliece trapdoor function is defined as $f'_H(x, e) := Hx \oplus e$, where H is a binary $k \times n$ matrix, x is a k -bit string and e is a error vector of small Hamming weight. Given H_1 , H_2 and two evaluations $y_1 = H_1x \oplus e$ and $y_2 = H_2x \oplus e$ one can reconstruct the unique x by solving $(H_1 \oplus H_2)x = y_1 \oplus y_2$ for x .

³ We remark that our construction of a correlation-secure trapdoor function from coding theory does not carry over to the lattice case since the “dual” of the one-way function used in [24,14] is not injective.

is related to ours, its security is based on the stronger assumption that a random two-prime RSA modulus is indistinguishable from a random three-prime RSA modulus. In another concurrent and independent work, Hemenway and Ostrovsky [15] generalize the framework of Peikert and Waters [25] to rely on any homomorphic hash proof system, a natural generalization of hash proof systems introduced by Cramer and Shoup [7]. Hemenway and Ostrovsky then show that homomorphic hash proof systems can be constructed based on either the quadratic residuosity assumption or the composite residuosity assumption. Their approach is significantly different than ours, and the resulting constructions seem incomparable when considering the trade-off between efficiency and lossiness.

1.3 Paper Organization

The remainder of this paper is organized as follows. In Section 2 we review the definitions of lossy and correlation-secure trapdoor functions. In Sections 3, 4, and 5 we present our constructions that are based on the quadratic residuosity assumption, the d -Linear assumption, and the hardness of syndrome decoding, respectively. Due to space constraints, the construction based on the composite residuosity assumption is only given in the full version [13].

2 Preliminaries

2.1 Lossy Trapdoor Functions

A *collection of lossy trapdoor functions* consists of two families of functions. Functions in one family are injective and can be efficiently inverted using a trapdoor. Functions in the other family are “lossy,” which means that the size of their image is significantly smaller than the size of their domain. The only computational requirement is that a description of a randomly chosen function from the family of injective functions is computationally indistinguishable from a description of a randomly chosen function from the family of lossy functions.

Definition 2.1 (Lossy trapdoor functions). A *collection of (n, ℓ) -lossy trapdoor functions* is a 4-tuple of probabilistic polynomial-time algorithms (G_0, G_1, F, F^{-1}) such that:

1. **Sampling a lossy function:** $G_0(1^n)$ outputs a function index $\sigma \in \{0, 1\}^*$.
2. **Sampling an injective function:** $G_1(1^n)$ outputs a pair $(\sigma, \tau) \in \{0, 1\}^* \times \{0, 1\}^*$. (Here σ is a function index and τ is a trapdoor.)
3. **Evaluation of lossy functions:** For every function index σ produced by G_0 , the algorithm $F(\sigma, \cdot)$ computes a function $f_\sigma : \{0, 1\}^n \mapsto \{0, 1\}^*$, whose image is of size at most $2^{n-\ell}$.
4. **Evaluation of injective functions:** For every pair (σ, τ) produced by G_1 , the algorithm $F(\sigma, \cdot)$ computes an injective function $f_\sigma : \{0, 1\}^n \mapsto \{0, 1\}^*$.
5. **Inversion of injective functions:** For every pair (σ, τ) produced by G_1 and every $x \in \{0, 1\}^n$, we have $F^{-1}(\tau, F(\sigma, x)) = x$.

- 6. Security:** The two ensembles $\{\sigma : \sigma \leftarrow \mathsf{G}_0(1^n)\}_{n \in \mathbb{N}}$ and $\{\sigma : (\sigma, \tau) \leftarrow \mathsf{G}_1(1^n)\}_{n \in \mathbb{N}}$ are computationally indistinguishable.

Note that ℓ can be a function of n . Note also that we do not specify the output of F^{-1} on inputs not in the image of f_σ .

A collection of all-but-one lossy trapdoor functions is a more general primitive. Such a collection is associated with a set B , whose members are referred to as *branches*. (If $B = \{0, 1\}$ then we obtain the previous notion of lossy trapdoor functions.) The sampling algorithm of the collection receives an additional parameter $b^* \in B$, and outputs a description of a function $f(\cdot, \cdot)$ together with a trapdoor τ and a set of lossy branches β . The function f has the property that for any branch $b \notin \beta$ the function $f(b, \cdot)$ is injective (and can be inverted using τ), and the function $f(b^*, \cdot)$ is lossy. Moreover, the description of f hides (in a computational sense) the set of lossy branches β .

Our definition is slightly more general than that of Peikert and Waters [25, Section 3.2], which allows only one lossy branch (i.e., $\beta = \{b^*\}$). We allow possibly many lossy branches (other than b^*), and require that given a description of a function and b^* it is computationally infeasible to find another lossy branch. The proof of security of the Peikert-Waters CCA-secure public-key encryption scheme [25, Section 4.3] can easily be adapted to our more general context. (We are currently not aware of other applications of all-but-one lossy trapdoor functions).

Definition 2.2 (All-but-one lossy trapdoor functions). A collection of (n, ℓ) -all-but-one lossy trapdoor functions is a 4-tuple of probabilistic polynomial-time algorithms $(\mathsf{B}, \mathsf{G}, \mathsf{F}, \mathsf{F}^{-1})$ such that:

1. **Sampling a branch:** $\mathsf{B}(1^n)$ outputs a value $b \in \{0, 1\}^*$.
2. **Sampling a function:** For every value b produced by $\mathsf{B}(1^n)$, the algorithm $\mathsf{G}(1^n, b)$ outputs a triple $(\sigma, \tau, \beta) \in \{0, 1\}^* \times \{0, 1\}^* \times \{0, 1\}^*$ consisting of a function index σ , a trapdoor τ , and a set of lossy branches β with $b^* \in \beta$.
3. **Evaluation of lossy functions:** For every value b^* produced by $\mathsf{B}(1^n)$ and for every (σ, τ, β) produced by $\mathsf{G}(1^n, b^*)$, the algorithm $\mathsf{F}(\sigma, b^*, \cdot)$ computes a function $f_{\sigma, b^*} : \{0, 1\}^n \mapsto \{0, 1\}^*$, whose image is of size at most $2^{n-\ell}$.
4. **Evaluation of injective functions:** For any b^* and b produced by $\mathsf{B}(1^n)$ and for every (σ, τ, β) produced by $\mathsf{G}(1^n, b^*)$, if $b \notin \beta$, then the algorithm $\mathsf{F}(\sigma, b, \cdot)$ computes an injective function $f_{\sigma, b} : \{0, 1\}^n \rightarrow \{0, 1\}^*$.
5. **Inversion of injective functions:** For any b^* and b produced by $\mathsf{B}(1^n)$ and for every (σ, τ, β) produced by $\mathsf{G}(1^n, b^*)$, if $b \notin \beta$ then we have

$$\mathsf{F}^{-1}(\tau, b, \mathsf{F}(\sigma, b, x)) = x.$$

6. **Security:** For any two sequences $\{(b_n^*, b_n)\}_{n \in \mathbb{N}}$ such that b_n^* and b_n are distinct values in the image of $\mathsf{B}(1^n)$, the two ensembles $\{\sigma : (\sigma, \tau, \beta) \leftarrow \mathsf{G}(1^n, b_n^*)\}_{n \in \mathbb{N}}$ and $\{\sigma : (\sigma, \tau, \beta) \leftarrow \mathsf{G}(1^n, b_n)\}_{n \in \mathbb{N}}$ are computationally indistinguishable.

7. **Hiding lossy branches:** Any probabilistic polynomial-time algorithm \mathcal{A} that receives as input (σ, b^*) , where $b^* \leftarrow \mathsf{B}(1^n)$ and $(\sigma, \tau, \beta) \leftarrow \mathsf{G}(1^n, b^*)$, has only a negligible probability of outputting an element $b \in \beta \setminus \{b^*\}$ (where the probability is taken over the randomness of B , G , and \mathcal{A}).

2.2 Correlation-Secure Trapdoor Functions

A collection of efficiently computable functions is a pair of algorithms $\mathcal{F} = (\mathsf{G}, \mathsf{F})$, where G is a key-generation algorithm used for sampling a description of a function, and F is an evaluation algorithm used for evaluating a function on a given input. The following definition formalizes the notion of a *k-wise product*, which is a collection \mathcal{F}_k consisting of all k -tuples of functions from \mathcal{F} .

Definition 2.3 (k-wise product). Let $\mathcal{F} = (\mathsf{G}, \mathsf{F})$ be a collection of efficiently computable functions. For any integer k , we define the k -wise product $\mathcal{F}_k = (\mathsf{G}_k, \mathsf{F}_k)$ as follows:

- The key-generation algorithm G_k on input 1^n invokes k independent instances of $\mathsf{G}(1^n)$ and outputs $(\sigma_1, \dots, \sigma_k)$. That is, a function is sampled from \mathcal{F}_k by independently sampling k functions from \mathcal{F} .
- The evaluation algorithm F_k on input $(\sigma_1, \dots, \sigma_k, x_1, \dots, x_k)$ invokes F to evaluate each function σ_i on x_i . I.e., $\mathsf{F}_k(\sigma_1, \dots, \sigma_k, x_1, \dots, x_k) = (\mathsf{F}(\sigma_1, x_1), \dots, \mathsf{F}(\sigma_k, x_k))$.

A one-way function is a function that is efficiently computable but is hard to invert given the image of a uniformly chosen input. This notion extends naturally to one-wayness under any specified input distribution, not necessarily the uniform distribution. Specifically, we say that a function is one-way with respect to an input distribution \mathcal{I} if it is efficiently computable but hard to invert given the image of a random input sampled according to \mathcal{I} .

In the context of k -wise products, a straightforward argument shows that for any collection \mathcal{F} which is one-way with respect to some input distribution \mathcal{I} , the k -wise product \mathcal{F}_k is one-way with respect to the input distribution that samples k independent inputs from \mathcal{I} . The following definition formalizes the notion of one-wayness under correlated inputs, where the inputs for \mathcal{F}_k may be correlated.

Definition 2.4 (One-wayness under correlated inputs). Let $\mathcal{F} = (\mathsf{G}, \mathsf{F})$ be a collection of efficiently computable functions with domain $\{D_n\}_{n \in \mathbb{N}}$, and let \mathcal{C} be a distribution where $\mathcal{C}(1^n)$ is distributed over $D_n^k = D_n \times \dots \times D_n$ for some integer $k = k(n)$. We say that \mathcal{F} is one-way under \mathcal{C} -correlated inputs if \mathcal{F}_k is one-way with respect to the input distribution \mathcal{C} .

For the special case that distribution \mathcal{C} is the uniform k -repetition distribution (i.e., \mathcal{C} samples a uniformly random input $x \in D_n$ and outputs k copies of x), we say that \mathcal{F} is *one-way under k -correlated inputs*. Rosen and Segev [26, Theorem 3.3] show that a collection of (n, ℓ) -lossy trapdoor functions can be used to construct a collection \mathcal{F} that is one-way under k -correlated inputs for any $k < \frac{n-\omega(\log n)}{n-\ell}$.

3 A Construction Based on the Quadratic Residuosity Assumption

Our construction is based on the modular squaring function $x \mapsto x^2 \bmod N$, where $N = PQ$ for prime numbers $P \equiv Q \equiv 3 \pmod{4}$ (i.e., Blum integers). This is a 4-to-1 mapping on \mathbb{Z}_N^* , and the construction is obtained by embedding additional information in the output that reduces the number of preimages to either 2 (these are the lossy functions) or 1 (these are the injective functions) in a computationally indistinguishable manner. Although this results in one bit of lossiness when the functions are defined over \mathbb{Z}_N^* , all lossy trapdoor functions in a collection are required to share the same domain (i.e., the domain should depend only on the security parameter). We overcome this difficulty with a simple domain extension, which results in lossiness of $\log_2(4/3)$ bits.

For any odd positive integer N , we denote by $\text{JS}_N : \mathbb{Z} \rightarrow \{-1, 0, 1\}$ the Jacobi symbol mod N . We define functions $h, j : \mathbb{Z} \rightarrow \{0, 1\}$ as follows:

$$h(x) = \begin{cases} 1, & \text{if } x > N/2 \\ 0, & \text{if } x \leq N/2 \end{cases}$$

$$j(x) = \begin{cases} 1, & \text{if } \text{JS}_N(x) = -1 \\ 0, & \text{if } \text{JS}_N(x) = 0 \text{ or } 1 \end{cases}$$

We define h and j on \mathbb{Z}_N by representing elements of \mathbb{Z}_N as integers between 0 and $N - 1$.

Fact 3.1 Let $N = PQ$ where $P \equiv Q \equiv 3 \pmod{4}$, and let $y \in \mathbb{Z}_N^*$ be a quadratic residue. Denote by $\{\pm x_0, \pm x_1\}$ the distinct solutions of the equation $x^2 = y \bmod N$. Then, $\text{JS}_P(-1) = \text{JS}_Q(-1) = -1$ and therefore

1. $\text{JS}_N(x_0) = \text{JS}_N(-x_0)$ and $\text{JS}_N(x_1) = \text{JS}_N(-x_1)$.
2. $\text{JS}_N(x_0) = -\text{JS}_N(x_1)$.

In particular, the four square roots of y take all four values of $(h(x), j(x))$.

The construction. We define a 4-tuple $\mathcal{F} = (\mathsf{G}_0, \mathsf{G}_1, \mathsf{F}, \mathsf{F}^{-1})$ (recall Definition 2.1) as follows:

1. **Sampling a lossy function:** On input 1^n the algorithm G_0 chooses an n -bit modulus $N = PQ$, where $P \equiv Q \equiv 3 \pmod{4}$ are $n/2$ -bit prime numbers. Then it chooses random $r \in \mathbb{Z}_N^*$ such that $\text{JS}_N(r) = -1$, and a random $s \in \mathbb{Z}_N^*$ such that $\text{JS}_N(s) = 1$ and s is a *quadratic residue*. The function index is $\sigma = (N, r, s)$.
2. **Sampling an injective function:** On input 1^n the algorithm G_1 chooses an n -bit modulus $N = PQ$, where $P \equiv Q \equiv 3 \pmod{4}$ are $n/2$ -bit prime numbers. Then it chooses random $r \in \mathbb{Z}_N^*$ such that $\text{JS}_N(r) = -1$, and a random $s \in \mathbb{Z}_N^*$ such that $\text{JS}_N(s) = 1$ and s is a *quadratic non-residue*. The function index is $\sigma = (N, r, s)$, and the trapdoor is $\tau = (P, Q)$.

- 3. Evaluation:** Given a function index $\sigma = (N, r, s)$ and $x \in \{0, 1\}^n$, the algorithm F interprets x as an integer in the set $\{1, \dots, 2^n\}$ and outputs

$$f_{N,r,s}(x) = \begin{cases} x^2 \cdot r^{j(x)} \cdot s^{h(x)} \bmod N, & \text{if } 1 \leq x < N \\ x, & \text{if } N \leq x \leq 2^n \end{cases}$$

- 4. Inversion:** Given a description of an injective function $\sigma = (N, r, s)$ together with its trapdoor $\tau = (P, Q)$ and $y = f_{N,r,s}(x)$, the algorithm F^{-1} retrieves x as follows. If $N \leq y \leq 2^n$, then the algorithm outputs y . Otherwise,

- (a) Find $j(x)$ by computing $\mathsf{JS}_N(f_{N,r,s}(x))$ (note that $\mathsf{JS}_N(f_{N,r,s}(x)) = \mathsf{JS}_N(x)$). Let $y' = yr^{-j(x)}$.
- (b) Find $h(x)$ by checking whether y' is a quadratic residue mod N (note that $h(x) = 1$ if and only if y' is not a quadratic residue). Let $y'' = y's^{-h(x)}$.
- (c) Find all square roots of y'' in \mathbb{Z}_N^* , and output the one that agrees with both $j(x)$ and $h(x)$. (We use Fact 3.1 if $y'' \in \mathbb{Z}_N^*$, and note that if $1 < \gcd(y'', N) < N$, then y'' has two square roots that are negatives of each other.)

We now prove that the above construction is indeed lossy based on the *quadratic residuosity assumption*. Let $\mathcal{J}_N = \{x \in \mathbb{Z}_N^* : \mathsf{JS}_N(x) = 1\}$, and let \mathcal{Q}_N be the subgroup of squares in \mathbb{Z}_N^* . Then the quadratic residuosity assumption states that the two distributions obtained by sampling uniformly at random from \mathcal{Q}_N and from $\mathcal{J}_N \setminus \mathcal{Q}_N$ are computationally indistinguishable.

Theorem 3.2. *Under the quadratic residuosity assumption, \mathcal{F} is a collection of $(n, \log_2(4/3))$ -lossy trapdoor functions.*

Proof. First, it follows from the correctness of the inversion algorithm that G_1 outputs permutations on the set $\{1, \dots, 2^n\}$. Next, we claim that G_0 outputs functions that are 2-to-1 on the set $\{1, \dots, N-1\}$. Suppose $y \in \mathcal{Q}_N$. Since s is a quadratic residue, Fact 3.1 implies that for each $(\eta, \iota) \in \{0, 1\}^2$ there is an $x_{\eta, \iota}$ satisfying

$$x_{\eta, \iota}^2 = ys^{-\eta}, \quad h(x_{\eta, \iota}) = \eta, \quad j(x_{\eta, \iota}) = \iota.$$

Then for each $\eta \in \{0, 1\}$ we have $f_{N,r,s}(x_{\eta,0}) = y$ and $f_{N,r,s}(x_{\eta,1}) = ry$. Thus each element in the set $\mathcal{Q}_N \cup r\mathcal{Q}_N$ has at least two preimages in \mathbb{Z}_N^* , and since this set has cardinality half that of \mathbb{Z}_N^* we deduce that $f_{N,r,s}$ is 2-to-1 on \mathbb{Z}_N^* . A similar argument shows that every square in the group $X_P = \{x \in \mathbb{Z}_N : \gcd(x, N) = P\}$ has two preimages in X_P , and the same for X_Q . Since $\{1, \dots, N-1\} = \mathbb{Z}_N^* \cup X_P \cup X_Q$, the function $f_{N,r,s}$ is 2-to-1 on this whole set.

Since N is an n -bit modulus (i.e., $2^{n-1} < N < 2^n$), the lossy functions are 2-to-1 on at least half of their domain, which implies that their image is of size at most $3/4 \cdot 2^n = 2^{n-\log_2(4/3)}$. In addition, descriptions of lossy functions and injective functions differ only in the element s , which is a random element of the subgroup of \mathbb{Z}_N^* with Jacobi symbol 1 that is a quadratic residue in the lossy case and a quadratic non-residue in the injective case. Therefore, the quadratic residuosity assumption implies that lossy functions are computationally indistinguishable from injective functions. \square

4 A Construction Based on the d -Linear Assumption

The d -*Linear assumption* [16,27] is a generalization of the decision Diffie-Hellman assumption that may hold even in groups with an efficiently computable d -linear map. The 1-Linear assumption is DDH, while the 2-Linear assumption is also known as the *Decision Linear* assumption [4]. The assumption is as follows:

Definition 4.1. Let $d \geq 1$ be an integer, and let \mathbb{G} be a finite cyclic group of order q . We say the d -Linear assumption holds in \mathbb{G} if the distributions

$$\begin{aligned} \{(g_1, \dots, g_d, g_1^{r_1}, \dots, g_d^{r_d}, h, h^{r_1+\dots+r_d}) : g_1, \dots, g_d, h \xleftarrow{\text{R}} \mathbb{G}, r_1, \dots, r_d \xleftarrow{\text{R}} \mathbb{Z}_q\}, \\ \{(g_1, \dots, g_d, g_1^{r_1}, \dots, g_d^{r_d}, h, h^s) : g_1, \dots, g_d, h \xleftarrow{\text{R}} \mathbb{G}, r_1, \dots, r_d, s \xleftarrow{\text{R}} \mathbb{Z}_q\} \end{aligned}$$

are computationally indistinguishable.

For any $d \geq 1$, the d -linear assumption implies the $(d+1)$ -linear assumption [16, Lemma 3].

Peikert and Waters [25, Section 5] give lossy and all-but-one lossy trapdoor functions based on the DDH assumption. In the Peikert-Waters construction, the function index is an ElGamal encryption of an $n \times n$ matrix M which is either the zero matrix (lossy mode) or the identity matrix (injective mode) using a finite cyclic group \mathbb{G} of order p . The DDH assumption in \mathbb{G} implies that these two encryptions cannot be distinguished. The construction can be generalized to d -linear assumptions using generalized ElGamal encryption, but such schemes are less efficient since ElGamal based on the d -Linear assumption produces $d+1$ group elements per ciphertext (see e.g. [27]).

Our construction is based on the following basic observation from linear algebra: if M is an $n \times n$ matrix over a finite field \mathbb{F}_p and \vec{x} is a length- n column vector, then the map $f_M : \vec{x} \mapsto M\vec{x}$ has image of size $p^{\text{Rk}(M)}$. If we restrict the domain to only *binary* vectors (i.e., those with entries in $\{0, 1\}$), then the function f_M is injective when $\text{Rk}(M) = n$, and its inverse can be computed by $f_M^{-1} : \vec{y} \mapsto M^{-1}\vec{y}$. If on the other hand we have $\text{Rk}(M) < n/\log_2(p)$, then f_M is not injective even when the domain is restricted to binary vectors, since the image is contained in a subgroup of size less than 2^n .

By performing the above linear algebra “in the exponent” of a group of order p , we can create lossy trapdoor functions based on DDH and the related d -Linear assumptions. In particular, for any n the size of the function index is the same for all d .

We will use the following notation: we let \mathbb{F}_p denote a field of p elements and $\text{Rk}_d(\mathbb{F}_p^{n \times n})$ the set of $n \times n$ matrices over \mathbb{F}_p of rank d . If we have a group \mathbb{G} of order p , an element $g \in \mathbb{G}$, and a vector $\vec{x} = (x_1, \dots, x_n) \in \mathbb{F}_p^n$, then we define $g^{\vec{x}}$ to be the column vector $(g^{x_1}, \dots, g^{x_n}) \in \mathbb{G}^n$. If $M = (a_{ij})$ is an $n \times n$ matrix over \mathbb{F}_p , we denote by g^M the $n \times n$ matrix over \mathbb{G} given by $(g^{a_{ij}})$. Given a matrix $M = (a_{ij}) \in \mathbb{F}_p^{n \times n}$ and a column vector $\mathbf{g} = (g_1, \dots, g_n) \in \mathbb{G}^n$, we define \mathbf{g}^M by

$$\mathbf{g}^M = \left(\prod_{j=1}^n g_j^{a_{1j}}, \dots, \prod_{j=1}^n g_j^{a_{nj}} \right).$$

Similarly, given a matrix $\mathbf{S} = (g_{ij}) \in \mathbb{G}^{n \times n}$ and a column vector $\vec{x} = (x_1, \dots, x_n) \in \mathbb{F}_p^n$, we define $\mathbf{S}^{\vec{x}}$ by

$$\mathbf{S}^{\vec{x}} = \left(\prod_{j=1}^n g_{1j}^{x_j}, \dots, \prod_{j=1}^n g_{nj}^{x_j} \right).$$

With these definitions, we have $(g^M)^{\vec{x}} = (g^{\vec{x}})^M = g^{(M\vec{x})}$.

The construction. For any positive integer d and any real number $\epsilon \in (0, 1)$, we define a 4-tuple $\mathcal{F} = (\mathsf{G}_0, \mathsf{G}_1, \mathsf{F}, \mathsf{F}^{-1})$ (recall Definition 2.1) as follows:

1. **Sampling a lossy function:** On input 1^n , the algorithm G_0 chooses at random a $\lceil \epsilon n/d \rceil$ -bit prime p , a group \mathbb{G} of order p , and a generator g of \mathbb{G} . Then it chooses a matrix $M \xleftarrow{\text{R}} \text{Rk}_d(\mathbb{F}_p^{n \times n})$ and computes $\mathbf{S} = g^M \in \mathbb{G}^{n \times n}$. The function index is $\sigma = \mathbf{S}$.
2. **Sampling an injective function:** On input 1^n , the algorithm G_1 chooses at random a $\lceil \epsilon n/d \rceil$ -bit prime p , a group \mathbb{G} of order p , and a generator g of \mathbb{G} . Then it chooses a matrix $M \xleftarrow{\text{R}} \text{Rk}_n(\mathbb{F}_p^{n \times n})$ and computes $\mathbf{S} = g^M \in \mathbb{G}^{n \times n}$. The function index is $\sigma = \mathbf{S}$, and the trapdoor is $\tau = (g, M)$.
3. **Evaluation:** Given a function index \mathbf{S} and $x \in \{0, 1\}^n$, we interpret x as a binary column vector $\vec{x} = (x_1, \dots, x_n) \in \mathbb{F}_2^n$. The algorithm F computes the function $f_{\mathbf{S}}(x) = \mathbf{S}^{\vec{x}}$.
4. **Inversion:** Given a function index \mathbf{S} , a trapdoor $\tau = (g, M)$, and a vector $\mathbf{g} \in \mathbb{G}^n$, we define $\mathsf{F}^{-1}(\tau, \mathbf{g})$ as follows:
 - (a) Compute $\mathbf{h} = (h_1, \dots, h_n) \leftarrow \mathbf{g}^{M^{-1}}$.
 - (b) Let $x_i = \log_g(h_i)$ for $i = 1, \dots, n$.
 - (c) Output $\vec{x} = (x_1, \dots, x_n)$.

Theorem 4.1. Suppose $\epsilon n > d$. If the d -Linear assumption holds for \mathbb{G} , then the above family is a collection of $(n, (1 - \epsilon)n)$ -lossy trapdoor functions.

Proof. We first note that in the lossy case, when M is of rank d , the image of $f_{\mathbf{S}}$ is contained in a subgroup of \mathbb{G}^n of size $p^d < 2^{\epsilon n}$. The condition $\epsilon n > d$ guarantees $p \geq 3$, so when M is of rank n the function $f_{\mathbf{S}}$ is in fact injective. It is straightforward to verify that the inversion algorithm performs correctly for injective functions. Finally, by [20, Lemma A.1], the d -Linear assumption implies that the matrix \mathbf{S} when M is of rank n is computationally indistinguishable from the matrix \mathbf{S} when M is of rank d . \square

Note that the system's security scales with the bit size of p , i.e., as $\epsilon n/d$. In addition, note that the discrete logarithms in the inversion step can be performed efficiently when \vec{x} is a binary vector. (Here we take advantage of the fact that the output of F^{-1} is unspecified on inputs not in the image of F .)

We now describe the extension of the system to all-but-one lossy trapdoor functions, in the case where the parameter d in the above construction is equal to 1. Let I_n denote the $n \times n$ identity matrix. For any real number $\epsilon \in (0, 1)$, we define a 4-tuple $\mathcal{F} = (\mathsf{G}_0, \mathsf{G}_1, \mathsf{F}, \mathsf{F}^{-1})$ (recall Definition 2.2) as follows:

1. **Sampling a branch:** On input 1^n , the algorithm B outputs a uniformly distributed $b \in \{1, \dots, 2^{\lfloor \epsilon n \rfloor}\}$.
2. **Sampling a function:** On input 1^n and a lossy branch b^* , the algorithm G chooses at random a $\lceil \epsilon n \rceil$ -bit prime p , a group \mathbb{G} of order p , and a generator g of \mathbb{G} . Then it chooses a matrix $A \xleftarrow{\text{R}} \text{Rk}_1(\mathbb{F}_p^{n \times n})$. Let $M = A - b^* I_n \in \mathbb{F}_p^{n \times n}$ and $\mathbf{S} = g^M \in \mathbb{G}^{n \times n}$. The function index is $\sigma = \mathbf{S}$, the trapdoor is $\tau = (g, M)$, and the set of lossy branches is $\beta = \{b^*, b^* - \text{Tr}(A)\}$.
3. **Evaluation:** Given a function index \mathbf{S} , a branch b , and an input $x \in \{0, 1\}^n$, we interpret x as a binary column vector $\vec{x} = (x_1, \dots, x_n)$. The algorithm F computes the function $f_{\mathbf{S}, b}(\vec{x}) = \mathbf{S}^{\vec{x}} * g^{b\vec{x}}$, where $*$ indicates the component-wise product of elements of \mathbb{G}^n .
4. **Inversion:** Given a function index \mathbf{S} , a trapdoor $\tau = (g, M)$, a branch b , and a vector $\mathbf{g} \in \mathbb{G}^n$, we define $\mathsf{F}^{-1}(\tau, b, \mathbf{g})$ as follows:
 - (a) If $M + bI_n$ is not invertible, output \perp .
 - (b) Compute $\mathbf{h} = (h_1, \dots, h_n) \leftarrow \mathbf{g}^{(M+bI_n)^{-1}}$.
 - (c) Let $x_i = \log_g(h_i)$ for $i = 1, \dots, n$.
 - (d) Output $\vec{x} = (x_1, \dots, x_n)$.

Theorem 4.2. Suppose $\epsilon n > 1$. If the DDH assumption holds for \mathbb{G} , then the above family is a collection of $(n, (1 - \epsilon)n)$ -all-but-one lossy trapdoor functions.

Proof. We first observe that if A is the rank 1 matrix computed by $\mathsf{G}(1^n, b^*)$, then

$$f_{\mathbf{S}, b}(\vec{x}) = g^{(A - (b^* - b)I_n)\vec{x}}. \quad (4.1)$$

We now verify each property of Definition 2.2. Properties (1) and (2) are immediate. To verify property (3), note that (4.1) implies that $f_{\mathbf{S}, b^*}(\vec{x}) = g^{A\vec{x}}$. Since A has rank 1, the image of $f_{\mathbf{S}, b^*}$ is contained in a subgroup of \mathbb{G}^n of size $p < 2^{\epsilon n}$.

To check property (4), we observe that the condition $\epsilon n > 1$ guarantees $p \geq 3$, so when $A - (b^* - b)I_n$ is invertible the function $f_{\mathbf{S}, b}$ is injective. The condition $A - (b^* - b)I_n$ being not invertible is equivalent to $(b^* - b)$ being an eigenvalue of A . Since A has rank 1, its eigenvalues are 0 and $\text{Tr}(A)$. Thus $(b^* - b)$ is an eigenvalue of A if and only if $b \in \beta$, and $f_{\mathbf{S}, b}$ is injective for all $b \notin \beta$. It is straightforward to verify that the inversion algorithm performs correctly whenever $b \notin \beta$, so property (5) holds.

Properties (6) and (7) follow from the DDH assumption for \mathbb{G} . We show property (6) by constructing a sequence of games:

Game₀: This is the real security game. The adversary is given b_0, b_1 , and $g^{A - b_\omega I_n}$ for $\omega \xleftarrow{\text{R}} \{0, 1\}$ and $A \xleftarrow{\text{R}} \text{Rk}_1(\mathbb{F}_p^{n \times n})$, and outputs a bit ω' . The adversary wins if $\omega' = \omega$.

Game₁: The same as Game₀, except the challenge is $g^{A' - b_\omega I_n}$ for some full rank matrix $A' \xleftarrow{\text{R}} \text{Rk}_n(\mathbb{F}_p^{n \times n})$.

Game₂: The same as Game₁, except the challenge is $g^{U - b_\omega I_n}$ for some uniform matrix $U \xleftarrow{\text{R}} \mathbb{F}_p^{n \times n}$.

Game₃: The same as Game₂, except the challenge is g^U .

Since the Game_3 challenge is independent of ω , the advantage of any adversary playing Game_3 is zero. We now show that if the DDH assumption holds for \mathbb{G} , then for $i = 0, 1, 2$, no polynomial-time adversary \mathcal{A} can distinguish Game_i from Game_{i+1} with non-negligible advantage.

- $i = 0$: Any algorithm that distinguishes Game_0 from Game_1 can be used to distinguish the distributions $\{g^A : A \xleftarrow{\text{R}} \text{Rk}_1(\mathbb{F}_p^{n \times n})\}$ and $\{g^{A'} : A' \xleftarrow{\text{R}} \text{Rk}_n(\mathbb{F}_p^{n \times n})\}$. By [5, Lemma 1], any algorithm that distinguishes these distributions can solve the DDH problem in \mathbb{G} .
- $i = 1$: Since the proportion of full-rank matrices to all matrices in $\mathbb{F}_p^{n \times n}$ is $(p - 1)/p$, even an unbounded adversary can distinguish Game_1 from Game_2 with probability at most $1/p$.
- $i = 2$: Since the matrix U is uniform in $\mathbb{F}_p^{n \times n}$, the matrix $U - b_\omega I_n$ is also uniform in $\mathbb{F}_p^{n \times n}$, so Game_2 and Game_3 are identical.

We conclude that for any b_0, b_1 , no polynomial-time adversary can win Game_0 with non-negligible advantage.

Finally, to demonstrate property (7) we show that any adversary \mathcal{A} that produces an element of β given \mathbf{S} and b^* can be used to compute discrete logarithms in \mathbb{G} , contradicting the DDH assumption. Choose a matrix $A \xleftarrow{\text{R}} \text{Rk}_1(\mathbb{F}_p^{n \times n})$, and let $A'(X)$ be the $n \times n$ matrix over $\mathbb{F}_p[X]$ that is the matrix A with the first row multiplied by X . For any value $X = t \neq 0$, the matrix $A'(t)$ is uniformly distributed in $\text{Rk}_1(\mathbb{F}_p^{n \times n})$.

Let (g, g^t) be a discrete logarithm challenge for \mathbb{G} . For any b^* we compute the matrix $\mathbf{S} = g^{A'(t)-b^*I_n}$ and give (\mathbf{S}, b^*) to the adversary \mathcal{A} . If the adversary outputs $b \in \beta$ with $b \neq b^*$, then we can compute $\text{Tr}(A'(t))$ since this is the only nonzero eigenvalue of $A'(t)$. If a_{ii} is the i th diagonal entry of A , this gives us an equation

$$a_{11}t + a_{22} + \cdots + a_{nn} = \lambda. \quad (4.2)$$

Since $a_{11} = 0$ with probability $1/p$, we can solve for t with all but negligible probability. \square

If we choose any integer $d \geq 1$ and repeat the above construction with p a $\lceil \epsilon n/d \rceil$ -bit prime and A a rank d matrix, then we expect to obtain an all-but-one lossy trapdoor function under the d -Linear assumption. Indeed, the proofs of properties (1)–(6) carry through in a straightforward way. However, the above proof of property (7) does not seem to generalize. In particular, the generalization of (4.2) is the equation $\det(A'(t) - \lambda I_n) = 0$, which can be written as $ut + v = 0$ for some (known) $u, v \in \mathbb{F}_p$. When $d = 1$ the element $u = a_{11}$ is independent of λ , so we can conclude that it is nonzero with high probability; however when $d \geq 2$ this is not the case. We thus leave as an open problem the completion of the proof for $d \geq 2$.

5 Correlated Input Security from Syndrome Decoding

Our construction is based on Niederreiter's coding-based encryption system [21] which itself is the dual of the McEliece encryption system [18].

Let $0 < \rho = \rho(n) < 1$ and $0 < \delta = \delta(n) < 1/2$ be two functions in the security parameter n . We set the domain $D_{n,\delta}$ to be the set of all n -bit strings with Hamming weight δn . Note that D_n is efficiently samplable (see e.g. [12]). The Niederreiter trapdoor function $\mathcal{F} = (\mathsf{G}, \mathsf{F}, \mathsf{F}^{-1})$ is defined as follows.

- **Key generation:** On input 1^n the algorithm G chooses at random a non-singular binary $\rho n \times \rho n$ matrix S , a $(n, n - \rho n, \delta n)$ -linear binary Goppa code capable of correcting up to δn errors (given by its $\rho n \times n$ binary parity check matrix G), and a $n \times n$ permutation matrix P . It sets $H := SGP$, which is a binary $\rho n \times n$ matrix. The description of the function is $\sigma = H$, the trapdoor is $\tau = (S, G, P)$.
- **Evaluation:** Given a description H of a function and $x \in \{0, 1\}^n$ with Hamming weight δn , the algorithm F computes the function $f_H(x) = Hx \in \{0, 1\}^{\rho n}$.
- **Inversion:** Given the trapdoor (S, G, P) and $y = Hx$, the algorithm F^{-1} computes $S^{-1}y = GPx$, applies a syndrome decoding algorithm for G to recover $\hat{y} = Px$, and computes $x = P^{-1}\hat{y}$.

The Niederreiter trapdoor function can be proved one-way under the indistinguishability and syndrome decoding assumptions which are indexed by the parameters $0 < \rho < 1$ and $0 < \delta < 1/2$.

Indistinguishability assumption. The binary $\rho n \times n$ matrix H output by $\mathsf{G}(1^n)$ is computationally indistinguishable from a uniform matrix of the same dimensions.

Syndrome decoding assumption. The collection of functions which is defined as $f_U(x) := Ux$ for a uniform $\rho n \times n$ binary matrix U is one-way on domain $D_{n,\delta}$.

Choosing the weight δ to be close to the Gilbert-Warshamov bound is commonly believed to give hard instances for the syndrome decoding problem. The Gilbert-Warshamov bound for a $(n, k, \delta n)$ linear code with $\delta < 1/2$ is given by the equation $k/n \leq 1 - H_2(\delta)$, where $H_2(\delta) := -\delta \log_2 \delta - (1 - \delta) \log_2(1 - \delta)$. It is therefore assumed that the syndrome decoding assumption holds for all $0 < \delta < 1/2$ satisfying $H_2(\delta) < \rho$ [12]. Note that one-wayness also implies that the cardinality of $D_{n,\delta}$ is super-polynomial in n .

The following theorem was proved in [12].

Theorem 5.1 ([12]). *If the syndrome decoding assumption holds for $\tilde{\rho}$ and δ then the ensembles $\{(M, Mx) : M \xleftarrow{\text{R}} \{0, 1\}^{\tilde{\rho}n \times n}; x \xleftarrow{\text{R}} D_{n,\delta}\}_{n \in \mathbb{N}}$ and $\{(M, y) : M \xleftarrow{\text{R}} \{0, 1\}^{\tilde{\rho}n \times n}; y \xleftarrow{\text{R}} \{0, 1\}^{\tilde{\rho}n}\}_{n \in \mathbb{N}}$ are computationally indistinguishable.*

This theorem implies that the Niederreiter trapdoor function is one-way under k -correlated inputs.

Theorem 5.2. *Suppose ρ, δ , and k are chosen such that $\tilde{\rho} := \rho k < 1$, and the indistinguishability and the syndrome decoding assumptions hold for parameters $\tilde{\rho}$ and δ . Then the Niederreiter trapdoor function is one-way under k -correlated inputs.*

Proof. Fix a probabilistic polynomial-time adversary \mathcal{A} that plays the security game for one-wayness under k -correlated inputs. Define

$$\varepsilon = \Pr[\mathcal{A}(H_1, \dots, H_k, H_1(x), \dots, H_k(x)) = x],$$

where $H_i \xleftarrow{\text{R}} \mathsf{G}(1^n)$ and $x \xleftarrow{\text{R}} D_{n,\delta}$. We now exchange all the matrices H_i for uniform matrices U_i of the same dimension. By the indistinguishability assumption and a hybrid argument, we have that

$$\begin{aligned} & \left| \Pr[\mathcal{A}(H_1, \dots, H_k, H_1(x), \dots, H_k(x)) = x] \right. \\ & \quad \left. - \Pr[\mathcal{A}(U_1, \dots, U_k, U_1(x), \dots, U_k(x)) = x] \right| \in \text{negl}(n). \end{aligned}$$

For $\tilde{\rho} := \rho k$, define the $\tilde{\rho}n \times n$ matrix U by concatenating the columns of the matrices U_i . Then the distributions $(U_1, \dots, U_k, U_1(x), \dots, U_k(x))$ and (U, Ux) are identical. Since $H_2(\delta) \leq \rho/k = \tilde{\rho}$ we can apply Theorem 5.1 to obtain

$$|\Pr[\mathcal{A}(U, Ux) = x] - \Pr[\mathcal{A}(M, u_{\tilde{\rho}n}) = x]| \in \text{negl}(n),$$

where $u_{\tilde{\rho}n}$ is a uniform bit-string in $\{0,1\}^{\tilde{\rho}n}$. Observing that $\Pr[\mathcal{A}(U, u_{\tilde{\rho}n}) = x] = 1/|D_{n,\delta}| \in \text{negl}(n)$ (since the Niederreiter function is assumed to be one-way) implies that ε is negligible. \square

We remark that the above proof implies that the Niederreiter trapdoor function has linearly many hard-core bits, which greatly improves efficiency of the CCA-secure encryption scheme obtained by using the construction from [26].

Acknowledgements. We thank Ivan Damgård and Chris Peikert for useful discussions. David Mandell Freeman's research was conducted at CWI and Universiteit Leiden, Netherlands, and supported by a National Science Foundation International Research Fellowship, with additional support from the Office of Multidisciplinary Activities in the NSF Directorate for Mathematical and Physical Sciences. Oded Goldreich is partially supported by the Israel Science Foundation (grant No. 1041/08). Eike Kiltz is supported by the research program Sentinels. Sentinels is being financed by Technology Foundation STW, the Netherlands Organization for Scientific Research (NWO), and the Dutch Ministry of Economic Affairs. Alon Rosen is partially supported by the Israel Science Foundation (grant No. 334/08). Gil Segev is supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities.

References

1. Bellare, M., Brakerski, Z., Naor, M., Ristenpart, T., Segev, G., Shacham, H., Yilek, S.: Hedged public-key encryption: How to protect against bad randomness. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 232–249. Springer, Heidelberg (2009)

2. Bellare, M., Hofheinz, D., Yilek, S.: Possibility and impossibility results for encryption and commitment secure under selective opening. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 1–35. Springer, Heidelberg (2009)
3. Boldyreva, A., Fehr, S., O’Neill, A.: On notions of security for deterministic encryption, and efficient constructions without random oracles. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 335–359. Springer, Heidelberg (2008)
4. Boneh, D., Boyen, X., Shacham, H.: Short group signatures. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 41–55. Springer, Heidelberg (2004)
5. Boneh, D., Halevi, S., Hamburg, M., Ostrovsky, R.: Circular-secure encryption from decision Diffie-Hellman. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 108–125. Springer, Heidelberg (2008)
6. Cachin, C., Micali, S., Stadler, M.: Computationally private information retrieval with polylogarithmic communication. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 402–414. Springer, Heidelberg (1999)
7. Cramer, R., Shoup, V.: Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 45–64. Springer, Heidelberg (2002)
8. Damgård, I., Jurik, M.: A generalisation, a simplification and some applications of Paillier’s probabilistic public-key system. In: Kim, K.-c. (ed.) PKC 2001. LNCS, vol. 1992, pp. 119–136. Springer, Heidelberg (2001); Full version (with additional co-author Nielsen, J. B.), available at www.daimi.au.dk/~ivan/GenPaillier_finaljour.ps
9. Damgård, I., Nielsen, J.B.: Perfect hiding and perfect binding universally composable commitment schemes with constant expansion factor. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 581–596. Springer, Heidelberg (2002)
10. Damgård, I., Nielsen, J.B.: Universally composable efficient multiparty computation from threshold homomorphic encryption. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 247–264. Springer, Heidelberg (2003)
11. Dowsley, R., Müller-Quade, J., Nascimento, A.C.A.: A CCA2 secure public key encryption scheme based on the McEliece assumptions in the standard model. In: Fischlin, M. (ed.) RSA Conference 2009. LNCS, vol. 5473, pp. 240–251. Springer, Heidelberg (2009)
12. Fischer, J.-B., Stern, J.: An efficient pseudo-random generator provably as secure as syndrome decoding. In: Maurer, U.M. (ed.) EUROCRYPT 1996. LNCS, vol. 1070, pp. 245–255. Springer, Heidelberg (1996)
13. Freeman, D.M., Goldreich, O., Kiltz, E., Rosen, A., Segev, G.: More constructions of lossy and correlation-secure trapdoor functions. Cryptology ePrint Archive, Report 2009/590 (2009), <http://eprint.iacr.org/2009/590>
14. Goldwasser, S., Vaikuntanathan, V.: New constructions of correlation-secure trapdoor functions and CCA-secure encryption schemes. Manuscript (2008)
15. Hemenway, B., Ostrovsky, R.: Lossy trapdoor functions from smooth homomorphic hash proof systems. Electronic Colloquium on Computational Complexity, Report TR09-127 (2009)
16. Hofheinz, D., Kiltz, E.: Secure hybrid encryption from weakened key encapsulation. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 553–571. Springer, Heidelberg (2007)
17. Kiltz, E., O’Neill, A., Smith, A.: Lossiness of RSA and the instantiability of OAEP. Manuscript (2009)
18. McEliece, R.J.: A public-key cryptosystem based on algebraic coding theory. DSN Prog. Rep., Jet Prop. Lab., 114–116 (January 1978)

19. Mol, P., Yilek, S.: Chosen-ciphertext security from slightly lossy trapdoor functions. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 296–311. Springer, Heidelberg (2010), <http://eprint.iacr.org/2009/524>
20. Naor, M., Segev, G.: Public-key cryptosystems resilient to key leakage. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 18–35. Springer, Heidelberg (2009), Full version <http://eprint.iacr.org/2009/105>.
21. Niederreiter, H.: Knapsack-type cryptosystems and algebraic coding theory. Problems of Control and Information Theory [Problemy Upravlenija i Teorii Informacij] 15, 159–166 (1986)
22. Nishimaki, R., Fujisaki, E., Tanaka, K.: Efficient non-interactive universally composable string-commitment schemes. In: Pieprzyk, J., Zhang, F. (eds.) ProvSec 2009. LNCS, vol. 5848, pp. 3–18. Springer, Heidelberg (2009)
23. Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 223–238. Springer, Heidelberg (1999)
24. Peikert, C.: Public-key cryptosystems from the worst-case shortest vector problem. In: 41st ACM Symposium on Theory of Computing, pp. 333–342 (2009)
25. Peikert, C., Waters, B.: Lossy trapdoor functions and their applications. In: 40th ACM Symposium on Theory of Computing, pp. 187–196 (2008), Full version <http://eprint.iacr.org/2007/279>
26. Rosen, A., Segev, G.: Chosen-ciphertext security via correlated products. In: Reingold, O. (ed.) Theory of Cryptography. LNCS, vol. 5444, pp. 419–436. Springer, Heidelberg (2009)
27. Shacham, H.: A Cramer-Shoup encryption scheme from the Linear assumption and from progressively weaker Linear variants. Cryptology ePrint Archive, Report 2007/074 (2007), <http://eprint.iacr.org/2007/074>