Correlated Product Security From Any One-Way Function

Brett Hemenway[1][*], Steve Lu[2][**], and Rafail Ostrovsky[3][***]

[1] University of Michigan
[2] Stealth Software Technologies, Inc.
[3] UCLA

Abstract. It is well-known that the k-wise product of one-way functions remains one-way, but may no longer be when the k inputs are correlated. At TCC 2009, Rosen and Segev introduced a new notion known as Correlated Product secure functions. These functions have the property that a k-wise product of them remains one-way even under correlated inputs. Rosen and Segev gave a construction of injective trapdoor functions which were correlated product secure from the existence of Lossy Trapdoor Functions (introduced by Peikert and Waters in STOC 2008).

In this work we continue the study of correlated product security, and find many differences depending on whether the functions have trapdoors.

The first main result of this work shows that a family of correlated product secure functions (without trapdoors) can be constructed from any one-way function. Because correlated product secure functions are trivially one-way, this shows an equivalence between the existence of these two cryptographic primitives.

In the second main result of this work, we consider a natural decisional variant of correlated product security. Roughly, a family of functions is Decisional Correlated Product (DCP) secure if $f_1(x_1), \ldots, f_k(x_1)$ is indistinguishable from $f_1(x_1), \ldots, f_k(x_k)$ when $x_1, \ldots, x_k$ are chosen uniformly at random.

When considering DCP secure functions with trapdoors, we give a construction based on Lossy Trapdoor Functions, and show that any DCP secure function family with trapdoor satisfies the security requirements

for Deterministic Encryption as defined by Bellare, Boldyreva and O'Neill in CRYPTO 2007. In fact, we also show that definitionally, DCP secure functions with trapdoors are a strict subset of Deterministic Encryption functions by showing an example of a Deterministic Encryption function which according to the definition is not a DCP secure function.

Keywords: Correlated Product Security, Lossy Trapdoor Functions, Deterministic Encryption

1   Introduction

If f and g are one-way functions on some domain X, it follows immediately that $(x, y) \mapsto (f(x), g(y))$ is a one-way function. On the other hand, it is well-known that $x \mapsto (f(x), g(x))$ may not be. The RSA function provides a simple example of this observation. The RSA assumption posits that $x \mapsto x^e \mod n$ is a one-way function. Given $x^{e_1} \mod n$, and $x^{e_2} \mod n$, the extended euclidean algorithm provides an efficient means of computing $x^{\gcd(e_1, e_2)} \mod n$, so if $\gcd(e_1, e_2) = 1$, the map $x \mapsto (x^{e_1}, x^{e_2}) \mod n$, is trivially invertible, even though its constituents are believed to be one-way.

In [RS09], Rosen and Segev formalized the notion of Correlated Product (CP) Security. They called a family of one-way trapdoor functions CP secure if they remained one-way when evaluated on correlated (and in particular, repeated) inputs. Rosen and Segev were motivated by the construction of IND-CCA secure encryption based on Lossy Trapdoor Functions (LTDFs) given by Peikert and Waters in [PW08]. Rosen and Segev showed that CP security is exactly the property needed to prove security of the Peikert and Waters construction.

Correlated Product security is an appealing notion because it is easy to define and appears to be a significantly weaker property than the *statistical* lossiness requirement of Lossy Trapdoor Functions. Despite this appearance of relative simplicity there have been few examples of correlated product secure functions that are not Lossy Trapdoor Functions. The notable exceptions are the constructions given in [Pei09] and [FGK+10].

This work continues the study of Correlated Product Secure Functions. We introduce a natural decisional variant of correlated product security, and show how this notion of Decisional Correlated Product Security provides connections to many areas in cryptography.

1.1   Our Results

In this work, we introduce (in Section 3) the notion of Decisional Correlated Product (DCP) security, which strengthens the definition of Rosen and Segev. We argue that this is a natural stepping-stone between Lossy Trapdoor Functions and Correlated Product secure functions. Intuitively, these are families of functions such that for any k functions $f_1, \ldots, f_k$, the distributions $\{(f_1(x_1), \ldots, f_k(x_1))\}$

and $\{(f_1(x_1), \ldots, f_k(x_k))\}$ are indistinguishable when $x_1, \ldots, x_k$ are chosen uniformly at random. Like correlated product security, decisional correlated product security can be defined for distributions other than the repetition distribution. We have focused on the case of the repetition distribution because it is conceptually simple while still capturing the essence of the problem. The repetition distribution is also the distribution that is necessary for applications to IND-CCA encryption [PW08,RS09].

Our results can be divided into three categories.

1. **Connections to Correlated Product Security:**
   We begin by examining the connections between Correlated Product (CP) and Decisional Correlated Product (DCP) security.
   From the definition of DCP security, it is clear that a family of constant functions is DCP secure, so for non-trivial results, we either specify that the functions be (individually) one-way or that they be injective with large domain. It turns out that, under either one of these assumptions, these families can be shown to also be Correlated Product secure. This is proven in Section 4 as the following lemmas:

   Lemma 2. If $\mathcal{F} = (G, F)$ is a family of k-DCP secure functions with super-polynomial size domain that are injective, then $\mathcal{F}$ is k-correlated product secure.

   Lemma 3. If $\mathcal{F} = (G, F)$ is a family of k-DCP secure one-way functions, then $\mathcal{F}$ is k-correlated product secure.

2. **DCP Secure Functions Without Trapdoors:**
   Our first main result considers families of one-way functions that are DCP secure. We show that such families are automatically (plain) Correlated Product secure, and demonstrate a construction from any pseudorandom function family. Due to the celebrated fact that a PRF family can be constructed from any one-way function ([GGM86,ILL89,HILL99]), we obtain an equivalence between the existence of one-way functions, DCP secure one-way function families, and CP secure function families. This is proven in Section 5 as the following theorem:

   Theorem 1. The following statements are equivalent:
   (a) One-way functions exist.
   (b) k-DCP secure families of one-way functions exist.
   (c) k-CP secure families of one-way functions exist.

   Theorem 1 shows that without a trapdoor, correlated product security essentially, is no stronger than simple one-wayness. This is somewhat surprising given the results of Vahlis [Vah10] that show that Correlated Product secure functions with trapdoor cannot be constructed from enhanced one-way trapdoor permutations. It is also somewhat surprising since lossy functions (without trapdoor) have not proven to be significantly easier to construct than lossy trapdoor functions.

3. **DCP Secure Functions With Trapdoors:**

Our second main result considers DCP secure function families which also have trapdoor. We investigate the connection between this and other primitives. In Section 6, we show a construction of these one-way trapdoor DCP secure families from sufficiently lossy LTDFs. This is stated as the following theorem:

Theorem 2. Let $\epsilon(\lambda)$ be any function such that $1/2^{\epsilon(\lambda)}$ is negligible in λ. Let $\mathcal{F} = (G, F)$ be a family of LTDFs on domain $\{0, 1\}^\lambda$, with residual leakage[4] at most $\frac{\lambda + 2 - 2\log(1/\epsilon)}{k}$. Then functions of the form $F_s(h(x))$ form a family of k-DCP trapdoor functions, where h is an injective pairwise independent hash function.

Finally, in Section 7, we show that these families definitionally satisfy the security requirements of Deterministic Encryption, but the converse is not true in general. Using the notion of *PRIV1 security* for Deterministic Encryption, which we will recall later, we have:

Theorem 3. DCP secure function families with trapdoor are also PRIV1 secure deterministic encryption schemes.

1.2 Previous Work

In [PW08] Peikert and Waters introduced a new paradigm for constructing IND-CCA secure encryption based on the newly defined primitive Lossy Trapdoor Functions (LTDFs). Their construction of IND-CCA was natural and appealing, but LTDFs proved difficult to construct because of their strong statistical lossiness properties. Despite the power of LTDFs, in [PW08] they were able to give constructions from DDH and Lattice-based assumptions, and the authors of [BFO08] and [RS08,FGK+10] (independently) found identical efficient constructions of LTDFs from Paillier's Decisional Composite Residuousity Assumption.

In [RS09], Rosen and Segev examined which properties of LTDFs were necessary to construct IND-CCA secure encryption via the methods in [PW08]. With this goal, they defined Correlated Product secure functions, and gave a construction of IND-CCA secure encryption from Correlated Product secure functions with trapdoor paralleling the construction in [PW08]. One of the primary difficulties in constructing Lossy Trapdoor Functions is creating functions the necessary *statistical* lossiness property (i.e. that the image of the function is significantly smaller than the domain). Correlated Product secure functions do not have these statistical requirements, and thus should be easier to construct than LTDFs. This intuition was reinforced in [RS09] where they showed that LTDFs are Correlated Product secure, and showed a black-box separation in the opposite direction. Correlated Product secure functions remain difficult to realize, however, and the recent results of Vahlis [Vah10], show a black-box separation between (enhanced) one-way trapdoor permutations and Correlated Product Secure functions.

---

[4] Recall that the residual leakage is defined to be the average number of bits leaked about the input when the function is in lossy mode. In particular, the residual leakage is defined to be the log of the size of the image of the function in lossy mode.

In 2007, Bellare, Boldyreva, and O'Neill [BBO07] introduced a new notion known as Deterministic Encryption (DE). The deterministic property of the encryption affords the scheme many practical applications, such as searchable encryption, but at the same time requires new security definitions. Subsequent works [BFO08,BFOR08] demonstrate equivalences between various definitions of DE and show that the existence of a sufficiently lossy LTDFs imply the existence of deterministic encryption, which in turn implies the existence of IND-CCA secure cryptosystems.

The works [BFO08,BFOR08] show many different relationships between DE and other primitives. Indeed, they show that any LTDF is almost immediately a DE scheme, and show how a weaker notion of DE can be constructed from any one-way trapdoor permutation.

In [ABBC10] Acar et al. studied the notion of *cryptographic agility*, where families of cryptographic primitives are said to be agile if they remain secure when the same key is re-used across families. While this is also a notion regarding correlated security, it does not appear to be connected to DCP security. Cryptographic agility refers to the security of correlated *keys* across *different* families of primitives, while DCP security refers to the one-wayness of functions *from the same family* when evaluated on correlated *inputs*.

The notion of security under correlated inputs has been studied in other contexts as well. In [IKNP03], Ishai et al. defined the notion of *correlation robustness* and used correlation robust functions to efficiently extend the number of independent oblivious transfer pairs available in a secure multiparty protocol. Correlation robustness was then used to create cryptosystems secure under related key attacks [AHI11,GOR11]. The notion of correlation robustness is distinct from the notion of correlated product security that is studied in this work. Correlation robustness studies the security of *a single function* applied on correlated inputs, while correlated product security studies the notion of *different* functions applied to correlated inputs. This distinction makes the constructions and applications quite different between the two areas.

2 Preliminaries

If A is a PPT machine, then we use $a \xleftarrow{\$} A$ to denote running the machine A and obtaining an output, where a is distributed according to the internal randomness of A. For a PPT machine A, we use $\mathsf{coins}(A)$ to denote the distribution of the internal randomness of A. So the distributions $\{a \xleftarrow{\$} A\}$ and $\{r \xleftarrow{\$} \mathsf{coins}(A) : a = A(r)\}$ are identical. If R is a set, we use $r \xleftarrow{\$} R$ to denote sampling uniformly from R. If X and Y are families distributions indexed by a security paramete λ, we use $X \approx_s Y$ to mean the two distributions are statistically close i.e. the statistical distance between X and Y is negligible in λ. We use $X \approx_c Y$ to mean that the distributions are computationally close, i.e. no PPT distinguisher with oracle access to the distribution has a non-negligible distinguishing advantage. We will need an extension of the leftover hash lemma known as the Crooked Leftover Hash Lemma [BFO08].

Lemma 1 (Crooked Leftover Hash Lemma [BFO08]). *Let $\mathcal{H}$ be a pairwise independent hash family, such that for all $h \in \mathcal{H}$, $h : X \to X$. Let $f : X \to Y$, and let Z be any random variable independent of h and D_X a distribution over X such that the min entropy $\tilde{H}_\infty(D_X|Z) \geq \log |Y| + 2\log(1/\epsilon) - 2$. Then if we define $\Lambda_1 = \{h \xleftarrow{\$} \mathcal{H}; x \xleftarrow{\$} D_X : (h, f(h(x)), Z)\}$, and $\Lambda_2 = \{h \xleftarrow{\$} \mathcal{H}; y \xleftarrow{\$} Y : (h, f(h(U_X), Z))\}$, we have $\Delta(\Lambda_1, \Lambda_2) \leq \epsilon$.*

Notice that the Crooked Leftover Hash Lemma does *not* imply that $h(D_X)$ is close to U_X, and indeed this may not be the case.

2.1 Correlated Product Security

In this section, we review the definition of Correlated Product security, first defined in [RS09]. We begin by defining the k-wise product of a Function Family.

Definition 1 (k-wise product). *Let $\mathcal{F} = (G, F)$ be a collection of efficiently computable functions. G is a (randomized) algorithm which takes as input a size parameter 1^λ and generates a key (or seed) s for F. Each function $F(s, \cdot)$ takes as input an element of some domain X and outputs some value in the range Y, both of which implicitly depend on the parameter λ. For notational purposes, we also write $F_s(\cdot) = F(s, \cdot)$.*
For $k \geq 2$, we define a family of k-wise products $\mathcal{F}^k = (G^k, F^k)$ as follows:

- ***Key Generation:***
  $G^k(1^\lambda)$ independently generates $s_i \xleftarrow{\$} G(1^\lambda)$, for $i = 1, \dots, k$.
- ***Evaluation:***
  To evaluate F^k on input $((s_1, \dots, s_k), (x_1, \dots, x_k))$, we define

$$F^k((s_1, \dots, s_k), (x_1, \dots, x_k)) = (F_{s_1}(x_1), \dots, F_{s_k}(x_k)).$$

Definition 2 (Correlated Product Security). *Let $\mathcal{F} = (G, F)$ be a collection of efficiently computable functions. Let $C_k = C_k(1^\lambda)$ be a distribution. We say that $\mathcal{F}$ is secure under C_k-correlated products if $\mathcal{F}^k$ is one-way with respect to the input distribution C_k.*

We remark that if the function family is very small, e.g. if it consists of only a single function, then correlated product security can be trivially satisfied, since $s_1 = \dots = s_k$ and hence $F_{s_1}(x) = \dots = F_{s_k}(x)$. This degenerate case only arises when considering CP security for functions without trapdoor. Throughout this work, we will focus on *decisional* correlated product security (Definition 3). We note that a family with fewer than k functions can *never* be k-DCP secure. Similarly, functions with trapdoor must also belong to a large (super-polynomial size) family. Since all of our results deal with DCP security or DCP security with trapdoor, we do not find it necessary to amend the definition of CP security explicitly require the size of the function family to be large.

For the remainder of the paper, we will focus on the case where C_k is the uniform k-repetition distribution, i.e. k copies of a uniformly chosen input. We refer

the reader to the Appendix for reminders of the definitions of the Discrete Log and DDH assumptions, Deterministic Encryption, Lossy Trapdoor Functions, and Pseudorandom Functions.

3   Decisional Correlated Product Security

In this work we introduce the notion of Decisional Correlated Product (DCP) security, which can be viewed as the decisional variant of Correlated Product security introduced in [RS09]. In [RS09], Rosen and Segev focused on the case where C_k was the uniform k-repetition distribution, i.e. C_k uniformly samples x and outputs k copies of x. We will also focus on the k-repetition distribution, although we will consider a decisional variant of the problem.

First, we remark that Correlated Product security seems to be a much stronger notion than simply one-wayness. For example, the map $f_e : x \mapsto x^e$ mod n, is one-way trapdoor permutation under the RSA assumption. However, given $f_{e_1}(x), f_{e_2}(x)$, if $\gcd(e_1, e_2) = 1$, we can immediately recover x, by using the extended Euclidean algorithm to calculate s, t such that $se_1 + te_2 = 1$, and noticing that $(x^{e_1})^s (x^{e_2})^t = x$. This example also shows that Decisional Correlated Product security does not follow immediately from Computational Correlated Product security, because if d_1, d_2, d_3 are relatively prime, and $e_i = ed_i$ for some fixed e, then $f_{e_1}, f_{e_2}, f_{e_3}$ will be Computationally Correlated Product secure under the RSA assumption, but will not be Decisional Correlated Product secure by a similar argument.

Definition 3 (Decisional Correlated Product Security). *Let $\mathcal{F} = (G, F)$ be a collection of efficiently computable functions. We say that $\mathcal{F}^k$ is k-wise Decisional Correlated Product secure if for all efficient PPT adversaries A,*

$$\left| \Pr\left[A^{\mathsf{indepdist}} = 1 \right] - \Pr\left[A^{\mathsf{repdist}} = 1 \right] \right| < \nu$$

for some negligible function ν, and where the games $\mathsf{indepdist}$ and $\mathsf{repdist}$ are defined as in Figure 1.

Independent	Repetition
$s_1 \xleftarrow{\$} G(1^\lambda), \ldots, s_k \xleftarrow{\$} G(1^\lambda)$	$s_1 \xleftarrow{\$} G(1^\lambda), \ldots, s_k \xleftarrow{\$} G(1^\lambda)$
$x_1 \xleftarrow{\$} X, \ldots, x_k \xleftarrow{\$} X$	$x \xleftarrow{\$} X$
$b \xleftarrow{\$} A(s_1, \ldots, s_k, F_{s_1}(x_1), \ldots, F_{s_k}(x_k))$	$b \xleftarrow{\$} A(s_1, \ldots, s_k, F_{s_1}(x), \ldots, F_{s_k}(x))$
Return b	Return b

Fig. 1. Decisional Correlated Product Security

To illustrate the power of this definition, we construct a very natural IND-CPA secure encryption from any family of 2-DCP secure injective trapdoor functions.

Let the public key be F_1, F_2, h where h is a pairwise independent hash function. Define encryption as $E(m, r) = (F_1(r), h(F_2(r)) \oplus m)$. To decrypt, we simply invert F_1 to recover r, from this we can recover $h(F_2(r))$ and recover the message. If F_i have domain $\{0, 1\}^\lambda$, and h maps from the range of F_i to $\{0, 1\}^{\lambda/2}$, then the leftover hash lemma tells us that $(F_1(r_1), h(F_2(r_2)) \oplus m)$ is statistically close to $(F_1(r_1), h(F_2(r_2)))$. So if y_0, y_1 are chosen from the repetition-distribution $(y_0, h(y_1) \oplus m)$ is a valid ciphertext, while if (y_0, y_1) are chosen from the independent distribution $(y_0, h(y_1) \oplus m)$ is independent of m, thus this scheme will be IND-CPA secure. We emphasize that this is not one of our main results, but simply an illustration of a natural construction that follows from this definition.

Remark. One of the appealing properties of the notion of k-DCP security is that it abstracts one of the most important properties of the DDH assumption. To see the parallel, recall a simple DDH-based PRG. The description of the function is the group $\mathcal{G}$, and two elements g, g^a, and $f(b) = (g^b, (g^a)^b)$. The first element of the output will be uniform if b is uniform, and the pair is indistinguishable from uniform by the DDH assumption. Now, it is easy to see that this construction will go through as before with an injective k-DCP family of functions. In particular, the description of the PRG will be $\mathcal{F}, s_1, \ldots, s_k$, and $f(x) = F_{s_1}(x), \ldots, F_{s_k}(x)$. If $F_{s_i}(\cdot)$ is a permutation, f will be a PRG with no modification. If the $F_{s_i}(\cdot)$ are merely injective, we will have to apply an extractor to "smooth" the output, but the proof of security remains exactly the same as in the DDH case. In fact, this observation can be generalized, the full version of this work contains a more detailed discussion of the parallel between DCP security and the DDH assumption.

The notion of *Decisional* Correlated Product security is clearly a stronger notion than the (Computational) Correlated Product security defined in [RS09] for *injective functions*. In the next section, we examine under what conditions DCP security implies CP security.

4 Relations to (Computational) Correlated Product Security

The notion of k-DCP security seems like a stronger requirement than Computational Correlated Product security, but we observe that if we do not put any requirements on the functions, then k-DCP security may be satisfied by trivial functions. For example the constant functions are trivially k-DCP for any $k \geq 2$. The following lemmas give sufficient conditions for when a k-DCP secure family is k-correlated product secure.

Lemma 2. *If $\mathcal{F} = (G, F)$ is a family of k-DCP secure functions with superpolynomial size domain and are injective, then $\mathcal{F}$ is k-correlated product secure.*

Proof. Let A be an efficient adversary that given $s_1, \ldots, s_k$, and $(F_{s_1}(x), \ldots, F_{s_k}(x))$, finds the inverse $(x_1', \ldots, x_k') = (x, x, \ldots, x)$ with non-negligible probability ϵ, we exhibit an efficient distinguisher D that uses A to break the k-DCP security of $\mathcal{F}$.

Algorithm 1 $D(s_1, \ldots, s_k, y_1, \ldots, y_k)$

$(x'_1, \ldots, x'_k) \xleftarrow{\$} A(s_1, \ldots, s_k, y_1, \ldots, y_k)$
if $x'_1 = x'_2 = \cdots = x'_k$ and $F_{s_i}(x'_i) = y_i$ for $i \in [k]$ **then**
   return 1
else
   return 0
end if

We must analyze the probability that D outputs 1 in the repdist and indepdist games.

$$\Pr[D^{\mathsf{repdist}} = 1] = \Pr[x'_1 = \cdots = x'_k \wedge F_{s_i}(x'_i) = y_i|$$
$$x \xleftarrow{\$} X, s_i \xleftarrow{\$} G(1^\lambda), y_i = F_{s_i}(x), \{x'_i\}_{i=1}^k \xleftarrow{\$} A(\{s_i\}_{i=1}^k, \{y_i\}_{i=1}^k)]$$
$$= \Pr[A \text{ successfully inverts}] = \epsilon.$$

$$\Pr[D^{\mathsf{indepdist}} = 1] = \Pr[x'_1 = \cdots = x'_k \wedge F_{s_i}(x'_i) = y_i|$$
$$x_i \xleftarrow{\$} X, s_i \xleftarrow{\$} G(1^\lambda), y_i = F_{s_i}(x_i), \{x'_i\}_{i=1}^k \xleftarrow{\$} A(\{s_i\}_{i=1}^k, \{y_i\}_{i=1}^k)]$$
$$= \Pr[x'_1 = \cdots = x'_k \wedge x'_i = x_i|$$
$$x_i \xleftarrow{\$} X, s_i \xleftarrow{\$} G(1^\lambda), y_i = F_{s_i}(x_i), \{x'_i\}_{i=1}^k \xleftarrow{\$} A(\{s_i\}_{i=1}^k, \{y_i\}_{i=1}^k)]$$
$$\leq \Pr[x_1 = x_2 | x_i \xleftarrow{\$} X] \leq \frac{1}{|X|}.$$

Thus the difference $|\Pr[D^{\mathsf{repdist}} = 1] - \Pr[D^{\mathsf{indepdist}} = 1]| \geq \epsilon - \frac{1}{|X|}$ is non-negligible, as $|X|$ is super-polynomial.

Next, we show that if a family $\mathcal{F} = (G, F)$ is a DCP secure, and each function is *individually* one-way, then the family is also Correlated Product secure.

Lemma 3. *If $\mathcal{F} = (G, F)$ is a family of k-DCP secure one-way functions, then $\mathcal{F}$ is k-correlated product secure.*

Proof. Suppose on the contrary that they were not. Let A be a PPT algorithm that breaks the correlated product security of (G, F), in particular given $\{s_1, \ldots, s_k, F_{s_1}(x_1), \ldots, F_{s_k}(x_1)\}$ A is able to find a pre-image $(x'_1, \ldots, x'_k)$ with some non-negligible probability ϵ, where the s_i are generated by G at random, and x_1 is chosen uniformly at random. We use A to build a PPT distinguisher D that can win in the k-DCP game.

We analyze the probability that D outputs 1. If indeed the inputs are correlated, i.e. $y_i = F_{s_i}(x_1)$, then A succeeds with probability ϵ and so D will output 1 with that probability.

On the other hand, if the inputs are random and independent, i.e. $y_i = F_{s_i}(x_i)$, then $(x_1, \ldots, x_k)$ is a uniformly chosen input from the product space. Because each $F_{s_i}(\cdot)$ is a one-way function, the product function $(F_{s_1}(\cdot), \ldots, F_{s_k}(\cdot))$

Algorithm 2 $D(s_1, \ldots, s_k, y_1, \ldots, y_k)$

$(x'_1, \ldots, x'_k) \xleftarrow{\$} A(s_1, \ldots, s_k, y_1, \ldots, y_k)$
if $F_{s_i}(x'_i) = y_i$ for $i \in [k]$ **then**
   return 1
else
   return 0
end if

---

is also one-way. Since the inputs are uncorrelated, the probability that A inverts it on a random value is negligible. Thus, in this case, D outputs 1 with only negligible probability.

This contradicts the k-DCP security of (G, F).

Many of the results in this work will focus on the case where the family $\mathcal{F}$ are in fact injective, or injective with trapdoor, and so the Correlated Product security will follow immediately from the DCP security of $\mathcal{F}$.

5   Equivalence of OWF and (Decisional) Correlated Product secure families of OWFs

In this section, we aim to prove the main theorem relating the existence of OWFs to that of (Decisional) Correlated Product secure OWF families.

Theorem 1. *The following statements are equivalent:*

1. *One-way functions exist.*
2. *k-DCP secure families of one-way functions exist.*
3. *k-CP secure families of one-way functions exist.*

To do this, we first show how to construct a DCP secure family of one-way functions from any pseudorandom function family. The idea is that a PRF family becomes DCP secure if we swap what we call the seed, and what we call the input. This idea has also been used in the past by Luby and Rackoff [LR89] to show the one-wayness of the UNIX-like password hash. If the PRF output is sufficiently long, then the resulting functions are also one-way, thus we have a family of DCP secure one-way functions. The exact lengths necessary are given in Lemma 5.

We then show that DCP secure one-way function families are also (ordinary) CP secure. This will follow directly from the fact that a product of one-way functions remain one-way under uniform independent inputs (Lemma 3). Finally, CP secure OWF families obviously are one-way, which completes the cycle of implications.

Let $(\mathsf{PRFGen}, \mathsf{PRF})$ be a PRF family, such that if $s \xleftarrow{\$} \mathsf{PRFGen}(1^\lambda)$, with $s \in \{0, 1\}^{w(\lambda)}$ then the domain of

$$\mathsf{PRF}(s, \cdot) : \{0, 1\}^{n(\lambda)} \to \{0, 1\}^{\ell(\lambda)}.$$

We can define a DCP family (G, F), by

- **Sampling:** $G(1^\lambda)$ outputs a uniform value $s \in \{0, 1\}^{n(\lambda)}$.
- **Evaluation:** For any $s \in \{0, 1\}^{n(\lambda)}$,

$$F_s(\cdot) : \{0, 1\}^{w(\lambda)} \to \{0, 1\}^{\ell(\lambda)}$$
$$x \mapsto \mathsf{PRF}(x, s).$$

Lemma 4. (G, F) *forms a k-Decisional Correlated Product secure function family for any $k = poly(\lambda)$.*

Proof. Define the distributions Λ_0, Λ_1 by sampling $s_1, \ldots, s_k \xleftarrow{\$} G(1^\lambda)$, and $x_1, \ldots, x_k \xleftarrow{\$} \{0, 1\}^{w(\lambda)}$

$$\Lambda_0 = \{s_1, \ldots, s_k, F_{s_1}(x_1), \ldots, F_{s_k}(x_1)\}$$
$$\Lambda_1 = \{s_1, \ldots, s_k, F_{s_1}(x_1), F_{s_2}(x_2), \ldots, F_{s_k}(x_k)\}$$

Thus we must show that any adversary who can distinguish Λ_0 from Λ_1 can distinguish the underlying Pseudorandom Function from a truly random function.

Now, by the definition of F, we have

$$\Lambda_0 = \{s_1, \ldots, s_k, F_{s_1}(x_1), \ldots, F_{s_k}(x_1)\} = \{s_1, \ldots, s_k, \mathsf{PRF}(x_1, s_1), \ldots, \mathsf{PRF}(x_1, s_k)\},$$
$$\Lambda_1 = \{s_1, \ldots, s_k, F_{s_1}(x_1), \ldots, F_{s_k}(x_k)\} = \{s_1, \ldots, s_k, \mathsf{PRF}(x_1, s_1), \ldots, \mathsf{PRF}(x_k, s_k)\}.$$

Now, it is clear that the security of the Pseudorandom Function gives

$$\Lambda_0 \approx_c \{s_1, \ldots, s_k, U_{\ell(\lambda)}, \ldots, U_{\ell(\lambda)}\} \approx_c \Lambda_1,$$

which gives the result.

Lemma 5. *If the size of the key space of F is a negligible fraction of the size of the output space, i.e. $1/2^{\ell(\lambda)-w(\lambda)}$ is negligible in λ, then (G, F) forms a family of one-way functions.*

Proof. Suppose to the contrary that for some key s, the function $F_s(\cdot)$ was not one-way. Let A be a PPT inverter that succeeds with non-negligible probability ϵ, i.e.
$$\Pr_x[F_s(z) = F_s(x) | z \leftarrow A(F_s(x))] = \epsilon$$

We use A to construct a PPT algorithm B that distinguishes between oracle access to PRF (with a randomly chosen seed x) and a truly random function $\mathcal{RO}$. The algorithm queries s on the oracle, and receives y, which is either $y = \mathsf{PRF}(x, s) = F_s(x)$ for some x, or a truly random value. The distinguisher B runs A on y, and receives some output x'. If it is the case that $F_s(x') = y$, then B outputs 1, otherwise B outputs 0.

We analyze the probabilities $\Pr[B^{\mathcal{RO}(\cdot)} = 1]$ and $\Pr_x[B^{\mathsf{PRF}(x,\cdot)} = 1]$. In the former case, the probability that a random value is in the range of $\mathsf{PRF}(s, \cdot)$ is $\frac{|Range|}{2^\ell} \leq \frac{2^w}{2^\ell}$ which we assumed to be negligible. On the other hand,

$$
\begin{aligned}
\Pr_x[B^{\mathsf{PRF}(x,\cdot)} = 1] &= \Pr_x[\mathsf{PRF}(z, s) = y | z \leftarrow A(y)] \\
&= \Pr_x[\mathsf{PRF}(z, s) = \mathsf{PRF}(x, s) | z \leftarrow A(\mathsf{PRF}(x, s))] \\
&= \Pr_x[F_s(z) = F_s(x) | z \leftarrow A(F_s(x))] = \epsilon
\end{aligned}
$$

This contradicts the pseudorandomness of PRF.

Corollary 1. *One-way functions imply k-DCP secure one-way function families.*

Proof. In Hastad, Impagliazzo, Levin and Luby [HILL99] it was shown that one-way functions imply PRGs, and in Goldreich, Goldwasser, Micali [GGM86] it was shown that PRGs imply the existence of PRF families with sufficiently long output, thus combining these results with our result, we have one-way functions imply k-DCP secure one-way functions.

Corollary 2. *One-way functions imply k-CP secure function families.*

Proof. This follows immediately from applying Lemma 3 to Corollary 1.

Since every Correlated Product secure function family is trivially a one-way function family, we have
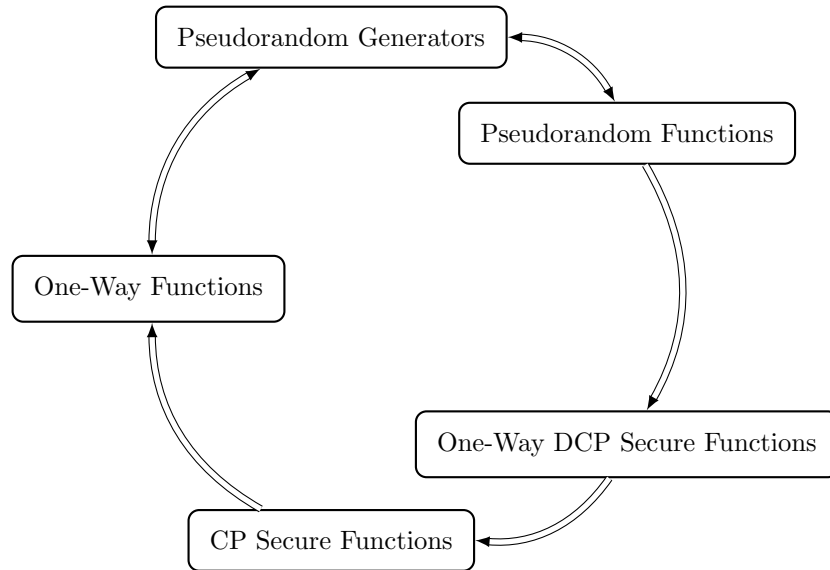
Since pseudo-random synthesizers [NR95,Rei98] are equivalent to one-way functions, we also achieve an equivalence between DCP secure functions and synthesizers. In the full version of this work, we give a direct proof that every family of pseudo-random synthesizers is immediately DCP secure.

In [BHK11], Braverman, Hassidim and Kalai introduced the notion of leakage-resilient random-input PRFs. A leakage-resilient random-input PRF is a pseudo-random function which remains pseudo-random when queried on *random* inputs (i.e. it is a weak PRF) even when partial information about the seed is leaked. Applying our construction to a leakage-resilient random-input PRF, we obtain a family of functions which is decisionally correlated product secure for any distribution $(X_1, \ldots, X_n)$ where that satisfies $\tilde{H}_\infty(X_i | X_1, \ldots, X_{i-1}) > \lambda$. Notice that the repetition distribution does not have this property, so by applying our construction to leakage-resilient random-input PRFs, we achieve DCP security for a completely different class of distributions.

6 DCP with trapdoor from Lossy Trapdoor Functions

In the preceding sections, we examined DCP secure functions without trapdoors, and showed that one-way DCP secure functions *without trapdoor* could be constructed from any one-way function. Now, we show constructions of DCP with trapdoor. In particular, in this section, we show that lossy trapdoor functions with sufficient lossiness imply DCP secure injective trapdoor functions.

Theorem 2. *Let $\mathcal{H}$ be a family of invertible[5] pairwise independent hash functions with $h : \{0,1\}^\lambda \to \{0,1\}^\lambda$. Let $\epsilon(\lambda)$ be any function such that $1/2^{\epsilon(\lambda)}$ is negligible in λ. Let $\mathcal{F} = (G, F)$ be a family of LTDFs on domain $\{0,1\}^\lambda$, where the lossy mode has residual leakage $r \leq \frac{\lambda + 2 - 2\log(1/\epsilon)}{k}$, for some integer k. Define $\hat{\mathcal{F}} = (\hat{G}, \hat{F})$ by*

- *$\hat{G}(1^\lambda)$, samples $s \xleftarrow{\$} G(1^\lambda)$, and $h \xleftarrow{\$} \mathcal{H}$, and outputs the function index h, s.*
- *Given a function index (h, s) and an input x, $\hat{F}_{h,s}(x) = F_s(h(x))$.*

Then $\hat{\mathcal{F}}$ is a k-DCP secure injective trapdoor function.

Proof. To prove the claim, we must show that distributions

$$\{h_1, s_1, \ldots, h_k, s_k, \hat{F}_{h_1,s_1}(x_1), \ldots, \hat{F}_{h_k,s_k}(x_1)\}$$
$$\text{and}$$
$$\{h_1, s_1, \ldots, h_k, s_k, \hat{F}_{h_1,s_1}(x_1), \hat{F}_{h_k,s_k}(x_k)\}$$

are computationally indistinguishable, where $h_1, s_1, \ldots, h_k, s_k \xleftarrow{\$} \hat{G}(1^\lambda)$, and $x_1, \ldots, x_k$ are sampled uniformly at random from the domain $\{0,1\}^\lambda$.

---

[5] We remark that this is not a strong restriction, and the natural construction $h(x) = ax + b$ over a finite field yields a collection of invertible pairwise independent hash functions.

If the function $F_s(\cdot)$ is in lossy mode, it has image size at most $2^{\frac{\lambda+2-2\log(1/\epsilon)}{k}}$, so if x is chosen uniformly from $\{0,1\}^\lambda$, then

$$\tilde{H}_\infty(x|\hat{F}_{h_1,s_1}(x),\ldots,\hat{F}_{h_{k-1},s_{k-1}}(x)) \geq \lambda - (k-1)\frac{\lambda+2-2\log(1/\epsilon)}{k}$$

$$= \lambda - (\lambda+2-2\log(1/\epsilon)) + \frac{\lambda+2-2\log(1/\epsilon)}{k}$$

$$= \frac{\lambda+2-2\log(1/\epsilon)}{k} + 2\log(1/\epsilon) - 2.$$

By the Crooked Leftover Hash Lemma, we have

$$\Delta\Big(\big\{h_1,s_1,\ldots,h_k,s_k,F_{s_1}(h_1(x_1)),\ldots,F_{s_k}(h_k(x))\big\},$$
$$\big\{h_1,s_1,\ldots,h_k,s_k,F_{s_1}(h_1(U_\lambda)),F_{s_2}(h_2(x)),\ldots,F_{s_k}(h_k(x))\big\}\Big) < \epsilon.$$

Repeating this argument a total of k times, we have

$$\Delta\Big(\big\{h_1,s_1,\ldots,h_k,s_k,F_{s_1}(h_1(x)),\ldots,F_{s_k}(h_k(x))\big\},$$
$$\big\{h_1,s_1,\ldots,h_k,s_k,F_{s_1}(h_1(U_\lambda)),\ldots,F_{s_k}(h_k(U_\lambda))\big\}\Big) < k\epsilon.$$

Since ϵ was assumed to be negligible, so is $k\epsilon$. Thus when the s_i are chosen to be lossy keys, the two distributions $\{h_1,s_1,\ldots,h_k,s_k,\hat{F}_{h_1,s_1}(x_1),\ldots,\hat{F}_{h_k,s_k}(x_1)\}$ and $\{h_1,s_1,\ldots,h_k,s_k,\hat{F}_{h_1,s_1}(x_1),\hat{F}_{h_k,s_k}(x_k)\}$ are statistically indistinguishable. The computational indistinguishability of lossy and injective keys implies that when the s_i are injective keys, the two distributions are computationally indistinguishable. Thus $(\hat{G},\hat{F})$ forms a family of k-DCP secure trapdoor functions.

7 Decisional Correlated Product Security is Deterministic Encryption

In this section, we examine the consequences of DCP secure functions, again *with trapdoor*. We show that any 2-DCP secure functions with trapdoor are – almost without modification – a PRIV1 secure uniform deterministic encryption. The notion of PRIV1 security is the original definition of security for deterministic encryption put forward in [BBO07]. PRIV1 security is the natural relaxation of the notion of semantically secure encryption to the deterministic setting. Recall that a cryptosystem is semantically secure if for any function $f(\cdot)$, an adversary's probability of calculating $f(m)$ remains essentially unchanged if the adversary is given access to an encryption $E(m)$. PRIV1 security requires that for any function $f(\cdot)$ *which is independent of the public key*, an adversary's ability to calculate $f(m)$ remains essentially unchanged whether he has access to the public key, or the public key *and* an encryption $E(m)$. See [BBO07] for the formal definition of PRIV1 security.

We follow the terminology of [BFOR08], where a uniform deterministic encryption is one which is only guaranteed to be secure against message adversaries that choose messages from the uniform distribution, instead of simply any high min-entropy distribution.

Let $\mathcal{F} = (G, F)$ be a family of 2-Decisional Correlated Product secure Functions.

We can define a (Uniform) Deterministic Encryption by

KeyGen:	Encryption:	Decryption:
$(s,t) \xleftarrow{\$} G(1^\lambda)$	$E(pk, m) = F_{pk}(m)$	$D(sk, c) = F_t^{-1}(c)$
$pk = s,\ sk = t$		

Fig. 2. Decisional Correlated Product Secure functions with trapdoor are PRIV1 secure

Theorem 3. *The scheme outlined in Figure 2 is BB-CSS secure.*

Proof. First, we recall the notion of BB-CSS (Balanced Boolean Comparison-based Semantic Security) as defined in [BFOR08]. This is similar to the Comparison Semantic Security PRIV1, outlined by the games privreal and privideal, except that the side information t is required to be a balanced boolean function, i.e. $\Pr[t = 0] \approx \Pr[t = 1] \approx \frac{1}{2}$.

For simplicity, we assume that $\Pr[t = 0] = \Pr[t = 1] = \frac{1}{2}$, but it is easy to see that if the distributions are only negligibly close to $\frac{1}{2}$ then the argument goes through as well.

Notice that in this setting *any* adversary has a $\frac{1}{2}$ chance of winning in the privideal game since his view is independent of the actual side information, thus it is enough to consider the adversary's probability of winning in the privreal game.

Now, suppose there exists an adversary $A = (A_m, A_g)$, such that $(m, t) \xleftarrow{\$} A_m(1^\lambda)$, where m is uniform on X the domain of f_s, and t is uniform on $\{0, 1\}$. The guessing adversary A_g on input pk, c outputs a guess t'. If $c = E(Pk, m)$, then $\Pr[t = t'] = \frac{1}{2} + \epsilon$.

We show how to use A to create a distinguisher D that can distinguish the 2-repetition distribution from the 2-independent distribution. The algorithm D takes as input the description of two functions s_0, s_1, and two outputs y_0, y_1, which come from either the repetition distribution (in which case $y_i = F_{s_i}(x)$) or the independent distribution (in which case $y_i = F_{s_i}(x_i)$, for two independently sampled x_i). The distinguisher D is described by Algorithm 3.

Now, we must analyze the probability that D succeeds. If y_0, y_1 were generated from the repetition distribution, then since A_g succeeds with probability $\frac{1}{2} + \epsilon$, the probability that D guesses "repetition" is $(\frac{1}{2} + \epsilon)^2 + (\frac{1}{2} - \epsilon)^2 = \frac{1}{2} + 2\epsilon^2$. If y_0, y_1 were generated from the independent distribution, because the side information is a balanced boolean function, the probability that the t_0, t_1 that would have been generated by A_m are equal is $\frac{1}{2}$. Intuitively, this should mean

Algorithm 3 $D(s_0, s_1, y_0, y_1)$

$t'_0 \overset{\$}{\leftarrow} A_g(s_0, y_0)$
$t'_1 \overset{\$}{\leftarrow} A_g(s_1, y_1)$
if $t'_0 = t'_1$ **then**
   return Repetition
else
   return Independent
end if

---

the probability that D correctly guesses "independent" is just $\frac{1}{2}$. This is in fact the case, because

$\Pr[D \text{ correctly guesses independent}]$

$$= \frac{1}{2}\Pr[D \text{ guesses independent}|t_0 = t_1] + \frac{1}{2}\Pr[D \text{ guesses independent}|t_0 \neq t_1]$$

$$= \frac{1}{2}\left(2\left(\frac{1}{2}+\epsilon\right)\left(\frac{1}{2}-\epsilon\right)\right) + \frac{1}{2}\left(\left(\frac{1}{2}+\epsilon\right)^2 + \left(\frac{1}{2}-\epsilon\right)^2\right) = \frac{1}{2}.$$

Thus the probability that D is correct is $\frac{1}{2} + \epsilon^2$.

Corollary 3. *The scheme outlined above is PRIV1 secure.*

Proof. In [BFOR08], they show that BB-CSS security (Comparison based Semantic Security against Balanced Boolean side information) implies B-CSS security (Comparison based Semantic Security against any Boolean side information), which in turns implies A-CSS which is security against Arbitrary side information. A-CSS security is the terminology in [BFOR08] for PRIV1 security. The only thing to do is to notice that both proofs in [BFOR08] go through unchanged when the adversaries are restricted to be uniform adversaries.

Remark. We note that if the function family $\mathcal{F} = (G, F)$ were assumed to be Decisional Correlated Product (DCP) secure when the inputs were chosen not uniformly, but simply from some high min-entropy distribution, the same proof would go through to show PRIV1 security against any (not necessarily uniform) adversary A_m.

Remark. On the other hand, there is an example (outlined below) of a PRIV1 secure uniform DE scheme that is not n-DCP secure (treating the public key as the seed, key generation as G, and encryption as F), where n is the size of the message. This does not preclude the construction of a DCP secure family from such a DE scheme, but instead shows that these two notions are not *definitionally* equivalent. To see that a PRIV1 secure DE need not be n-DCP secure, take any IND-CPA secure (randomized) encryption scheme, and transform it into a "leaky" scheme that leaks the first bit of randomness used in encryption by simply taking an extra dummy bit of randomness and revealing it in the ciphertext.

The construction of uniform DE from one-way trapdoor permutations given in [BFOR08] makes use of an IND-CPA secure (randomized) encryption scheme.

Without fully reproducing the [BFOR08] construction, we only need to point out that the first bit of randomness is the hard-core predicate defined by the dot product of the message and a vector from the public key. If the "leaky" encryption of the same message under n different public keys is revealed, the message can be reconstructed using linear algebra. This immediately breaks (Decisional) Correlated Product security.

8 Conclusion and Open Problems

In this work we suggested a new primitive, the decisional variant of Correlated Product (DCP) secure functions. We argue that this primitive has many appealing properties. To this end, we show a parallel between Correlated Product security and DCP and the Discrete Log Problem and its decisional variant DDH. We also show how to construct simple primitives from DCP such as PRGs and IND-CPA secure encryption.

Our main results examine two main cases: DCP functions with trapdoor and without trapdoor. We show that DCP secure functions (and CP secure functions) without trapdoor are equivalent to one-way functions. This is a somewhat surprising result since notions of correlated product security appear to be much stronger than simple one-wayness. When examining DCP secure functions with trapdoor, we show that they are implied by Lossy Trapdoor Functions, and that DCP secure functions are immediately a Deterministic Encryption scheme.

An interesting line of future research would be to develop further constructions of DCP secure functions with trapdoor. A second line of research would be a closer examination of the connections between DCP security and deterministic encryption. For example, we know that DCP secure functions are deterministic encryption, but it would be interesting to see how the security is affected by auxiliary information, e.g. along the lines of [BS11].

References

ABBC10. Tolga Acar, Mira Belenkiy, Mihir Bellare, and David Cash. Cryptographic agility and its relation to circular encryption. In *EUROCRYPT '10*, volume 6110 of *LNCS*, pages 403–422, 2010.

AHI11. Benny Applebaum, Danny Harnik, and Yuval Ishai. Semantic security under related-key attacks and applications. In *ITCS '11*, 2011.

BBO07. Mihir Bellare, Alexandra Boldyreva, and Adam O'Neill. Deterministic and efficiently searchable encryption. In *CRYPTO '07*, volume 4622 of *Lecture Notes in Computer Science*, pages 535–552. Springer Berlin / Heidelberg, 2007.

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 '08*, volume 5157 of *Lecture Notes in Computer Science*, pages 335–359. Springer, 2008.

BFOR08. Mihir Bellare, Marc Fischlin, Adam O'Neill, and Thomas Ristenpart. De-
         terministic encryption: Definitional equivalences and constructions without
         random oracles. In David Wagner, editor, *CRYPTO '08*, volume 5157 of
         Lecture Notes in Computer Science, pages 360–378. Springer, 2008.

BHK11.   Mark Braverman, Avinatan Hassidim, and Yael Tauman Kalai. Leaky
         pseudo-entropy functions. In *ITCS '11*, 2011.

BS11.   Zvika Brakerski and Gil Segev. Better security for deterministic public-key
         encryption: The auxiliary-input setting. In *CRYPTO '11*, 2011.

FGK+10. David Mandell Freeman, Oded Goldreich, Eike Kiltz, Alon Rosen, and Gil
         Segev. More constructions of lossy and correlation-secure trapdoor func-
         tions. In *Public Key Cryptography 2010 (PKC 2010)*, Lecture Notes in
         Computer Science, pages 279–295, 2010.

GGM86.   Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct
         random functions. *Journal of the ACM*, 33(4):792–807, 1986.

GOR11.   Vipul Goyal, Adam O'Neill, and Vanishree Rao. Correlated-input secure
         hash functions. In *TCC 2011*, volume 6597, page 182, 2011.

HILL99. Johan Hastad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A
         pseudorandom generator from any one-way function. *SIAM J. Comput.*,
         28(4):1364–1396, 1999.

IKNP03. Yuval Ishai, Joe Kilian, Kobbi Nissim, and Erez Petrank. Extending obliv-
         ious transfers efficiently. In *CRYPTO '03*, 2003.

ILL89.   Russell Impagliazzo, Leonid A. Levin, and Michael Luby. Pseudo-random
         generation from one-way functions (extended abstract). In *STOC 89*, pages
         12–24, 1989.

LR89.   Michael Luby and Charles Rackoff. A study of password security. *Journal
         of Cryptology*, 1(3):151–158, 1989.

NR95.   Moni Naor and Omer Reingold. Synthesizers and their application to the
         parallel construction of psuedo-random functions. In *FOCS '95*, pages 170–
         181, 1995.

Pei09.   Chris Peikert. Public-key cryptosystems from the worst-case shortest vector
         problem: extended abstract. In *STOC '09: Proceedings of the 41st annual
         ACM symposium on Theory of computing*, pages 333–342, New York, NY,
         USA, 2009. ACM.

PW08.   Chris Peikert and Brent Waters. Lossy trapdoor functions and their appli-
         cations. In *STOC '08: Proceedings of the 40th annual ACM symposium on
         Theory of computing*, pages 187–196, New York, NY, USA, 2008. ACM.

Rei98.   Omer Reingold. *Pseudo-Random Synthesizers Functions and Permutations*.
         PhD thesis, The Weizmann Institute of Science, 1998.

RS08.   Alon Rosen and Gil Segev. Efficient lossy trapdoor functions based on the
         composite residuosity assumption. http://eprint.iacr.org/2008/134, 2008.

RS09.   Alon Rosen and Gil Segev. Chosen-ciphertext security via correlated prod-
         ucts. In *TCC '09: Proceedings of the 6th Theory of Cryptography Confer-
         ence on Theory of Cryptography*, pages 419–436, Berlin, Heidelberg, 2009.
         Springer-Verlag.

Vah10.   Yevgeniy Vahlis. Two is a crowd? a black-box separation of one-wayness
         and security under correlated inputs. In *TCC '10*, volume 5978 of *Lecture
         Notes in Computer Science*, pages 165–182. Springer Berlin / Heidelberg,
         2010.