Weak Pseudorandom Functions in Minicrypt

Krzysztof Pietrzak[1] and Johan Sjödin[2] ⋆

[1] CWI Amsterdam
[2] ETH Zurich

Abstract. A family of functions is *weakly* pseudorandom if a random member of the family is indistinguishable from a uniform random function when queried on *random* inputs. We point out a subtle ambiguity in the definition of weak PRFs: there are natural weak PRFs whose security breaks down if the randomness used to sample the inputs is revealed. To capture this ambiguity we distinguish between *public-coin* and *secret-coin* weak PRFs.
We show that the existence of a secret-coin weak PRF which is *not* also a public-coin weak PRF implies the existence of two pass key-agreement (i.e. public-key encryption). So in Minicrypt, i.e. under the assumption that one-way functions exist but public-key cryptography does not, the notion of public- and secret-coin weak PRFs coincide.
Previous to this paper all positive cryptographic statements known to hold exclusively in Minicrypt concerned the adaptive security of constructions using non-adaptively secure components. Weak PRFs give rise to a new set of statements having this property. As another example we consider the problem of range extension for weak PRFs. We show that in Minicrypt one can beat the best possible range expansion factor (using a fixed number of distinct keys) for a very general class of constructions (in particular, this class contains all constructions that are known today).

1 Introduction

1.1 Weak Pseudorandom Functions

Informally, a pseudorandom function (PRF) is a function which cannot be distinguished from a uniform random function by any efficient distinguisher. PRFs have a wide range of applications in cryptography. Sometimes, however, the full power of a PRF is not needed and it is sufficient when the function cannot be distinguished when queried on random values. Such objects are referred to as *weak* PRFs.[3]

PRFs are black-box reducible to one-way functions. In particular, it was shown in [9] how to construct a PRF from any pseudorandom generator and in [11] a construction of a pseudorandom generator from any one-way function was introduced. Unfortunately those (black-box) reductions are not efficient enough to be practical[4]. In [23], Naor and Reingold gave a quite efficient construction of a PRF relying on a number-theoretic assumption. More precisely, assuming that the so called decisional Diffie-Hellman (DDH) assumption holds in some cyclic group $G = \langle g \rangle$ of order q, they proposed a PRF $(\mathbb{Z}_q)^{n+1} \times \{0,1\}^n \to G$ defined by

$$((k_0, \ldots, k_n), a) \mapsto g^{k_0 \cdot \Pi_{a_i=1} k_i},$$

⋆ This work was partially supported by the Zurich Information Security Center. It represents the views of the authors.

[3] Sometimes they are called PRFs secure under a *known-plaintext* attack (KPA).

[4] Even though progress has recently been made [10, 14] for achieving more efficient reductions of PRGs from one-way functions.

where $(k_0, \ldots, k_n)$ is the secret key and a_i denotes the i'th bit of a. As observed in [22], if one just wants to construct a weak PRF, then even a much simpler construction exists: let G be a cyclic group of order q, then

$$\mathsf{exp} : \mathbb{Z}_q \times G \to G \quad \text{defined as} \quad \mathsf{exp}(k, a) = a^k$$

is a weak PRF if the DDH assumption holds in G (the exponent k is the secret key). Note that compared to the Naor-Reingold construction, this construction has a much shorter key ($\lceil \log(q) \rceil$ compared to $(n+1)\lceil \log(q) \rceil$) and is more efficient (each evaluation requires one exponentiation, compared to one exponentiation and up to n multiplications).

1.2 Public-Coin vs. Secret-Coin weak PRFs

A standard choice for the group G used in exp would be a large subgroup of $\mathbb{Z}_p^*$ (where p is some large prime) of prime order q (which exists if and only if q divides the order $p-1$ of $\mathbb{Z}_p^*$). The DDH assumption is believed to hold in such groups if q is sufficiently large.[5] A natural way to sample an element from G is to choose an element $r \in \mathbb{Z}_q$ uniformly at random and set $a = g^r$ for some generator g of G.[6]

For $a_1, a_2, \ldots$ sampled this way, the tuples $(a_1, v_1), (a_2, v_2), \ldots$ computed by the weak PRF as $v_i = \mathsf{exp}(k, a_i) = a_i^k$ are indistinguishable from $(a_1, u_1), (a_2, u_2), \ldots$ where each u_i is a uniform random element in G. Now, assume that the distinguisher also gets to see the randomness used to sample the a_i's. Say r_1, r_2 are such that $a_1 = g^{r_1}$ and $a_2 = g^{r_2}$. Then one can easily distinguish $v_1 = a_1^k, v_2 = a_2^k$ from u_1, u_2 as $v_1^{r_2} = v_2^{r_1}$ $(= g^{r_1 r_2})$ but $u_1^{r_2} = u_2^{r_1}$ only holds with probability $1/q$.

Thus the security of the weak PRF exp completely breaks down if the randomness used to sample the random inputs is revealed. We will call such a weak PRF a *secret-coin* weak PRF, as opposed to a *public-coin* weak PRF which stays secure even if the random coins used to sample the inputs are revealed. Whether a weak PRF is a public-coin or just secret-coin weak PRF depends on the input-sampling algorithm, which hence must be part of the definition of the weak PRF.

1.3 Public-Coins=Secret-Coins in Minicrypt

As the function exp shows, the distinction between public- and secret-coin weak PRFs is meaningful and one can imagine many situations where the notion of a secret-coin weak PRF is not sufficient. In particular, this will always be the case when public randomness is used to sample the inputs.

It is not hard to see, that if a weak PRF is secret- but not public-coin, then the input-sampling algorithm must be a distributional one-way function [17].[7] We further show that any secret-coin weak PRF which is not also a public-coin weak PRF must be very artificial, in the sense that one can construct a public-key encryption scheme from it:

Theorem 1 *If there exists a secret-coin weak PRF which is not a public-coin weak PRF, then a (IND-CPA) secure public-key encryption scheme exists (see Section 1.4 for the disclaimer).*

[5] Say $p - 1 = 2q$ where $\log(p)$ is at least our security parameter.

[6] Alternatively one could sample random elements from $\mathbb{Z}_p^*$ until an element a is found where $a^q = a$, but this is less efficient (the expected number of tries is $q/(p-1)$), and only works if q^2 does not divide $p-1$ as otherwise there is more than just one subgroup of order q).

[7] A distributional one-way function is a function where no efficient algorithm can find a *random* pre-image. Distributional one-way functions are equivalent to one-way functions.

Thus it is not surprising that the DDH assumption required for the security of exp implies public-key encryption [4, 7]. So in Minicrypt (a name coined by Impagliazzo to denote the hypothetical world where one-way functions exist, but public-key cryptography does not [16]), the notion of public- and secret-coin weak PRFs coincide.

In order to prove Theorem 1, we show how to construct a public-key encryption scheme from any secret-coin weak PRF F and a distinguisher D which can distinguish F from a uniform random function when additionally to the random input/output pairs it is provided with the randomness used by the input sampling algorithm S (such a D exists as by assumption F is not a public-coin weak PRF).

1.4 Uniform vs. Non-Uniform and Negligible vs. Noticeable

There is a gap between what is generally considered a successful distinguisher (or any other kind of an adversary) and what one expects from a protocol like an encryption scheme: a system is usually considered broken even if only a *non-uniform* adversary exists, whereas a protocol should be *uniform* and achieve its task with *overwhelming*[8] probability to be considered useful. The encryption scheme we construct in order the prove Theorem 1 uses the distinguisher D as a black-box, and only if D is *uniform* and has *noticeable* advantage in distinguishing F from a uniform random function, we will get a useful (as described above) key-agreement protocol. But if D is non-uniform, also the key-agreement protocol will be non-uniform. Furthermore, if D has only *non-negligible* (but not noticeable) advantage, then our encryption scheme will only be secure for infinitely many values of the security parameter (and not as usual for all sufficiently large ones).

1.5 Range Extension for weak PRFs

The problem of range extension for weak PRFs is the following: given a weak PRF $f : \mathcal{K} \times \mathcal{X} \to \mathcal{X}$, construct a weak PRF $F : \mathcal{K}^t \times \mathcal{X} \to \mathcal{X}^s$ where F uses f as a black-box, t is the number of keys, and s is the so-called *expansion factor*. A trivial solution is to set

$$F((k_1, \ldots, k_t), x) = [f(k_1, x), \ldots, f(k_t, x)],$$

but this is not very satisfying as the expansion factor is only the number of keys (i.e. $s = t$). All efficient constructions for range extension of weak PRFs, that we are aware of [2, 21, 20], are of the form that the i'th output block y_i is computed as

$$y_i = f_{k_{i_q}} \circ f_{k_{i_{q-1}}} \circ \cdots \circ f_{k_{i_1}}(x)$$

for some $i_1, \ldots, i_q \in \{1, \ldots, t\}$. The construction from [20] achieves an expansion factor of $s = 2^t - 1$, and in [25] it is shown that this is tight: no construction (of the form as described above) with an expansion factor greater than $2^t - 1$ can be proven secure via a black-box reduction. We show that it is possible to beat this bound in Minicrypt. For this, consider the following construction which uses two keys and has an expansion factor of 4 (which is more than $2^t - 1 = 3$)

$$F((k_1, k_2), x) = [f(k_1, x), f(k_2, x), f(k_2, f(k_1, x)), f(k_1, f(k_2, x))]. \quad (1)$$

In [21], it has been claimed that this function is indeed a weak PRF, but in fact it is not as observed in [20]. The function $f = \mathsf{exp}$ is a simple counterexample, as in this case the last two values of the output in (1) are identical, namely

$$[\mathsf{exp}(k_2, \mathsf{exp}(k_1, x)), \mathsf{exp}(k_1, \mathsf{exp}(k_2, x))] = [x^{k_1 k_2}, x^{k_1 k_2}]. \quad (2)$$

[8] See Section 2 for a definition of negligible, noticeable and overwhelming.

In Section 4, we prove the following theorem:

Theorem 2 *If there exists a weak PRF f with superpolynomial domain size for which the construction given by equation (1) is not a weak PRF, then a secure public-key encryption scheme exists (disclaimer is below).*

The requirement that the domain size must be superpolynomial is necessary as otherwise the construction given in (1) is not a weak PRF even if f is a uniform random function. The reason is that the last two values of the output in (1), i.e. $[f(k_2, f(k_1, x)), f(k_1, f(k_2, x))]$, collide twice as often as for random elements. If now the domain size is not superpolynomial, collisions will occur with high probability after polynomially many input-output samples (allowing us to distinguish F from a uniform random function).

In order to prove Theorem 2, we show how to construct a two-pass key-agreement protocol from any weak PRF f and a distinguisher D which can (for random keys k_1, k_2) distinguish tuples computed as $[f(k_2, f(k_1, .)), f(k_1, f(k_2, .))]$ from random tuples. For this reduction, we have the same issue with uniform vs. non-uniform and negligible vs. noticeable, as discussed in Section 1.4.

1.6 Related Work

Adaptive Security in Minicrypt. The first positive cryptographic result proven to hold only in Minicrypt stated that the cascade of non-adaptively secure PRFs gives a construction with some weak form of adaptive security [24]. We have found more results since then, but all were about the adaptive security of some construction based on non-adaptively secure components.[9] The results of this paper show that weak PRFs give rise to a completely new class of statements that hold exclusively in Minicrypt.

Other Worlds. Wee [26] shows that some cryptography (i.e. "non-trivial" argument systems) is even possible in Pessiland, which is another of Impagliazzo's possible worlds [16] where not even one-way functions exist. Dent [3] explores the limits of cryptography in universes whose existence is conjectured in popular Science Fiction literature.

Win-Win. Our results can be viewed as "win-win" statements, where one shows that (at least) one of two "positive" cryptographic statements is true. For example we show that either every secret-coin weak PRF is a public-coin weak PRF or public-key crypto exists. Results of similar flavour have been given before, in particular Dziembowski shows that either "forward-secure storage" is possible or (a weak form of) oblivious transfer exists [6]. Dubrov and Ishai show that either every efficiently samplable distribution can be sampled using few random bits, or one-way functions imply collision-resistant hashing [5].

Public vs. Secret Coins. That differentiating between the public- and secret-coin variants of primitives is meaningful and important has been shown for at least two important primitives, namely collision resistant hash functions (CRHF) and trapdoor permutations (TDP).

In [15], Hsiao and Reyzin define public- and secret-coin families of CRHFs. The collision resistance of the latter requires the coins used to sample the function to be kept secret. They show that no black-box reduction from secret-coin to public-coin CRHFs exists.

[9] For example, in [19] it was shown that the four round Feistel-network with non-adaptively secure round functions is not a pseudorandom permutation in general. To be more precise, it was shown that there exists a non-adaptively secure function f (whose security is based on the *inverse* DDH assumption [1]) for which the four-round Feistel-network with f as round function can be distinguished from a random permutation with two adaptive queries. Here, one can show that any such counterexample implies a three round key-agreement protocol (this is unpublished).

The classical definition of a TDP states that it is hard to invert the permutation given a random element from the range of the permutation. In [8], Goldreich observes that for many applications this is not enough, and in fact *enhanced* TDPs are needed. Those have the property that it is hard to invert a random element even when given the random coins used to sample this element.

2 Basic Definitions

Throughout, let $n \in \mathbb{N}$ denote a security parameter. An entity (e.g. adversary) is efficient and uniform if it can be implemented by a probabilistic Turing machine whose running time is polynomial in the input length (which for us will always mean polynomial in n). It is efficient and non-uniform if it can be realized by a sequence of circuits (one for each n) of polynomial (in n) size.

For a set $\mathcal{X}$, let $x \xleftarrow{\$} \mathcal{X}$ denote that x is assigned a value from $\mathcal{X}$ uniformly at random. Let x^q denote the sequence $x_1, \ldots, x_q$. For a function $f : \mathcal{X} \to \mathcal{Y}$ and $x^q \in \mathcal{X}^q$, let $f(x^q)$ denote $f(x_1), \ldots, f(x_q)$.

Uniform Random Functions. $\mathbf{R}_{\mathcal{X},\mathcal{Y}}$ denotes a uniform random function $\mathcal{X} \to \mathcal{Y}$.

Negligible. A function $\mu : \mathbb{N} \to [0,1]$ is negligible if for any $c > 0$ there is an n_0 such that $\mu(n) \leq 1/n^c$ for all $n \geq n_0$. To the contrary, μ is non-negligible if for some $c > 0$ we have $\mu(n) \geq 1/n^c$ for infinitely many n. Throughout, $\mathsf{negl}(n)$ denotes a negligible function in n.

Overwhelming. A function $\tau(\cdot) : \mathbb{N} \to [0,1]$ is overwhelming if $1 - \tau(\cdot)$ is negligible.

Noticeable. A function $\phi : \mathbb{N} \to [0,1]$ is noticeable if for some $c > 0$ there is an n_0 such that $\phi(n) \geq 1/n^c$ for all $n \geq n_0$.

Note that non-negligible is not the same as noticeable. For example $\mu(n) \stackrel{\text{def}}{=} n \bmod 2$ is non-negligible but not noticeable.

Bit-Agreement. Bit-agreement is a protocol between two efficient parties, which we refer to as Alice and Bob . They get the security parameter n in unary (denoted 1^n) as a common input and can communicate over an authentic channel. Finally, Alice and Bob output a bit b_{A} and b_{B}, respectively. The protocol has correlation ϵ if for all n

$$\mathsf{P}[b_{\mathsf{A}} = b_{\mathsf{B}}] \geq (1 + \epsilon(n))/2.$$

Furthermore, the protocol is δ-secure if for any efficient adversary E, which can observe the whole communication C, and for all n

$$\mathsf{P}[\mathsf{E}(1^n, C) = b_{\mathsf{B}}] \leq 1 - \delta(n)/2.$$

Key-Agreement. If $\epsilon(\cdot)$ and $\delta(\cdot)$ are overwhelming then such a protocol achieves key-agreement. Any protocol which achieves bit-agreement with noticeable correlation $\epsilon(\cdot)$ and overwhelming security $\delta(\cdot)$ can be turned into a key-agreement protocol without increasing the number of rounds using parallel repetition and privacy amplification [12, 13].

If $\epsilon(\cdot)$ is only non-negligible (i.e. for any constant $c > 0$, $\epsilon(n) \geq 1/n^c$ for infinitely many n), then also the key-agreement protocol will only achieve correctness for infinitely many (and not for all sufficiently large) choices of the security parameter.

3 Public-Coin vs. Secret-Coin weak PRFs

Definition 1 (weak PRFs) *Consider a pair of efficient algorithms* $\mathsf{F}, \mathsf{KeyGen}$ *where for any* $n \in \mathbb{N}$ *we have*

$$\mathsf{KeyGen} : 1^n \to \mathcal{K}_n \qquad \mathsf{F} : \mathcal{K}_n \times \mathcal{X}_n \to \mathcal{Y}_n.$$

KeyGen *is the randomized key-generation algorithm which on input a security parameter* n *(and some uniform random bits) outputs a key from the keyspace* $\mathcal{K}_n$*. Let the random variables* X_i*,* Y_i*, and* Z_i *for* $1 \leq i \leq \ell$ *be defined by first sampling a key* $k \leftarrow \mathsf{KeyGen}(1^n)$ *and then setting (below we use the same random function* $\mathbf{R}_{\mathcal{X}_n,\mathcal{Y}_n}$ *for all* i*)*

$$X_i \xleftarrow{\$} \mathcal{X}_n \qquad Y_i \leftarrow \mathsf{F}(k, X_i) \qquad Z_i \leftarrow \mathbf{R}_{\mathcal{X}_n,\mathcal{Y}_n}(X_i).$$

F *is a* weak pseudorandom function *secure if for every efficient distinguisher* D *and any polynomial* $\ell = \ell(n)$

$$|\mathsf{P}[\mathsf{D}(X^\ell, Y^\ell) = 1] - \mathsf{P}[\mathsf{D}(X^\ell, Z^\ell) = 1]| = \mathsf{negl}(n).$$

Definition 2 (public-coin and secret-coin weak PRFs) *Let* $\mathsf{F}, \mathsf{KeyGen}$ *be efficient algorithms as in the previous definition, and let* $\mathsf{Sample} : \{0,1\}^{s(n)} \to \mathcal{X}_n$ *be an efficient input sampling algorithm.*

Let the random variables R_i*,* X_i*,* Y_i*, and* Z_i *be defined for* $1 \leq i \leq \ell$ *by first sampling a key* $k \leftarrow \mathsf{KeyGen}(1^n)$ *and then setting*

$$R_i \xleftarrow{\$} \{0,1\}^{s(n)} \qquad X_i \leftarrow \mathsf{Sample}(R_i) \qquad Y_i \leftarrow \mathsf{F}(k, X_i) \qquad Z_i \leftarrow \mathbf{R}_{\mathcal{X}_n,\mathcal{Y}_n}(X_i).$$

The three algorithms $\mathsf{F}, \mathsf{KeyGen}, \mathsf{Sample}$ *are a* public-coin *weak PRF if for all efficient* D *and any polynomial* $\ell = \ell(n)$

$$|\mathsf{P}[\mathsf{D}(R^\ell, X^\ell, Y^\ell) = 1] - \mathsf{P}[\mathsf{D}(R^\ell, X^\ell, Z^\ell) = 1]| = \mathsf{negl}(n)$$

(i.e. the weak PRF, by Definition 1, stays secure even if the randomness used to sample the inputs is revealed). Furthermore, $\mathsf{F}, \mathsf{KeyGen}, \mathsf{Sample}$ *are referred to as a* secret-coin *weak PRF if for all efficient* D *and any polynomial* $\ell = \ell(n)$

$$|\mathsf{P}[\mathsf{D}(X^\ell, Y^\ell) = 1] - \mathsf{P}[\mathsf{D}(X^\ell, Z^\ell) = 1]| = \mathsf{negl}(n).$$

Clearly, every public-coin weak PRF is a secret-coin weak PRF. Also if $\mathsf{F}, \mathsf{KeyGen}$ is a weak PRF (by Definition 1) and the output of Sample is close to uniform, then $\mathsf{F}, \mathsf{KeyGen}, \mathsf{Sample}$ is a secret-coin weak PRF. Note that in the definitions above, secure means secure against efficient uniform adversaries. To get a (stronger) notion which implies security against non-uniform adversaries, one must just consider a sequence of poly-size circuits instead of the poly-time bounded Turing-machine D (cf. Section 1.4).

3.1 The Reduction

Let $(\mathsf{F}, \mathsf{KeyGen}, \mathsf{Sample})$ be a secret-coin weak PRF which is not a public-coin weak PRF. For $i = 1, 2, \ldots$ consider the random variables r_i, x_i, y_i, and u_i, defined by $k \leftarrow \mathsf{KeyGen}(1^n)$, $r_i \xleftarrow{\$} \{0,1\}^{s(n)}$, $x_i \leftarrow \mathsf{Sample}(r_i)$, $y_i \leftarrow \mathsf{F}(k, x_i)$, and $u_i \xleftarrow{\$} \mathcal{Y}_n$ (i.e. u_i is a random element from the range of F). Now, as F is not a public-coin weak PRF, there exist an efficient distinguisher D, a polynomial $q(.)$, and a non-negligible function $\phi(.)$ such that

$$\Pr[\mathsf{D}(r^q, x^q, y^q) = 1] - \Pr[\mathsf{D}(r^q, x^q, u^q) = 1] \geq \phi(n). \tag{3}$$

Further, as F is a secret-coin weak PRF we have for any efficient E that

$$|\Pr[\mathsf{E}(x^q, y^q) = 1] - \Pr[\mathsf{E}(x^q, u^q) = 1]| = \mathsf{negl}(n). \tag{4}$$

PROTOCOL BITAGREEMENT(n)

Alice		Bob
		$b_{\mathsf{B}} \xleftarrow{\$} \{0,1\}$
		$k \leftarrow \mathsf{KeyGen}(1^n)$
for $i = 1, \ldots, q = q(n)$ do		
$r_i \xleftarrow{\$} \{0,1\}^n$ $x_i \leftarrow \mathsf{Sample}(r_i)$ od;	$\xrightarrow{x^q}$	for $i = 1, \ldots, q$ do
		if $b_{\mathsf{B}} = 0$ then $z_i \leftarrow \mathsf{F}(k, x_i)$
		elseif $b_{\mathsf{B}} = 1$ then $z_i \leftarrow \mathbf{R}_{\mathcal{X}_n, \mathcal{Y}_n}(x_i)$ od;
	$\xleftarrow{z^q}$	
$b_{\mathsf{A}} \leftarrow \mathsf{D}(r^q, x^q, z^q)$		

Fig. 1. A bit-agreement protocol from a secret-coin weak PRF which is not a public-coin weak PRF.

In order to prove Theorem 1, we must construct a two-pass public-key encryption scheme from the weak PRF. As discussed in Section 2, it is sufficient to construct a two-pass bit-agreement protocol with non-negligible (or noticeable, see the discussion in Section 1.4) correlation and overwhelming security. Such a protocol BITAGREEMENT is shown in Figure 1. The idea behind the protocol is quite simple: First, Alice samples some random strings, on which she invokes Sample to get random inputs to the secret-coin weak PRF F. Then she sends the inputs to Bob, who either return the outputs of F on these inputs or random values depending on his randomly chosen bit b_{B}. As Alice knows the randomness she used to sample the inputs, she can use the distinguisher D to get a guess b_{A} on b_{B} with non-negligible correlation, as shown in Claim 1 below. Furthermore, an adversary who does not know the randomness used to sample the inputs cannot distinguish the cases where Bob sends random values or values computed by F, as shown in Claim 2 below.

Claim 1 BITAGREEMENT(n) *has correlation $\phi(n)$, with ϕ as in (3).*

Proof.

$$\begin{aligned}\Pr[b_{\mathsf{A}} = b_{\mathsf{B}}] &= \Pr[b_{\mathsf{B}} = 1] \cdot \Pr[b_{\mathsf{A}} = 1 | b_{\mathsf{B}} = 1] + \Pr[b_{\mathsf{B}} = 0] \cdot \Pr[b_{\mathsf{A}} = 0 | b_{\mathsf{B}} = 0] \\ &= \frac{1}{2} + \frac{\Pr[b_{\mathsf{A}} = 1 | b_{\mathsf{B}} = 1] - \Pr[b_{\mathsf{A}} = 1 | b_{\mathsf{B}} = 0]}{2} \geq \frac{1}{2} + \frac{\phi(n)}{2}\end{aligned}$$

□

Claim 2 BITAGREEMENT(n) *is $1 - \mathsf{negl}(n)$ secure.*

Proof. For any efficient adversary E

$$\begin{aligned}\Pr[\mathsf{E}(x^q, z^k) = b_{\mathsf{B}}] &= \Pr[b_{\mathsf{B}} = 1] \cdot \Pr[\mathsf{E}(x^q, z^k) = 1 | b_{\mathsf{B}} = 1] + \Pr[b_{\mathsf{B}} = 0] \cdot \Pr[\mathsf{E}(x^q, z^k) = 0 | b_{\mathsf{B}} = 0]| \\ &= \frac{1}{2} + \frac{\Pr[\mathsf{E}(x^q, z^k) = 1 | b_{\mathsf{B}} = 1] - \Pr[\mathsf{E}(x^q, z^k) = 1 | b_{\mathsf{B}} = 0]}{2} = \frac{1}{2} + \mathsf{negl}(n),\end{aligned}$$

where the last step follows by (4). □

Proof (of Theorem 1). The theorem follows from Claim 1 and 2 and the fact that one can construct a key-agreement protocol from any bit-agreement protocol which has noticeable correlation and overwhelming security without increasing the number of rounds (via parallel repetition and privacy amplification [12, 13]).

4 Range Extension for Weak PRFs

PROTOCOL BITAGREEMENT2(n)

Alice		Bob
		$b_{\mathsf{B}} \xleftarrow{\$} \{0,1\}$
$k_{\mathsf{A}} \leftarrow \mathsf{KeyGen}(1^n)$		$k_{\mathsf{B}} \leftarrow \mathsf{KeyGen}(1^n)$
for $i = 1, \ldots, q = q(n)$ do		
$r_i \xleftarrow{\$} \mathcal{X}_n$ $s_i \leftarrow \mathsf{F}(k_{\mathsf{A}}, r_i)$ od;	$\xrightarrow{r^q, s^q}$	for $i = 1, \ldots, q$ do
		if $b_{\mathsf{B}} = 0$ then $t_i \leftarrow \mathsf{F}(k_{\mathsf{B}}, r_i)$
		$y_i \leftarrow \mathsf{F}(k_{\mathsf{B}}, s_i)$
		elseif $b_{\mathsf{B}} = 1$ then $t_i \xleftarrow{\$} \mathbf{R}_{\mathcal{X}_n, \mathcal{X}_n}(r_i)$
		$y_i \xleftarrow{\$} \mathbf{R}_{\mathcal{X}_n, \mathcal{X}_n}(s_i)$ od;
	$\xleftarrow{t^q, y^q}$	
for $i = 1, \ldots, q$ do $z_i \leftarrow \mathsf{F}(k_{\mathsf{A}}, t_i)$ od;		
$b_{\mathsf{A}} \leftarrow \mathsf{D}(r^q, s^q, t^q, y^q, z^q)$		

Fig. 2. A bit-agreement protocol from a weak PRF $\mathsf{F} : \mathcal{K}_n \times \mathcal{X}_n \to \mathcal{X}_n$ where the construction given by (1) is not a weak PRF.

Let $\mathsf{F} : \mathcal{K}_n \times \mathcal{X}_n \to \mathcal{X}_n, \mathsf{KeyGen} : 1^n \to \mathcal{K}_n$ denote a weak PRF as in Definition 1, with the additional property that the domain and range are identical, i.e. $\mathcal{X}_n = \mathcal{Y}_n$, and that the domain is of superpolynomial size, i.e. for all c there is an n_0 such that $|\mathcal{X}_n| \geq n^c$ for all $n > n_0$. In order to prove Theorem 2, we must show that if $\mathsf{G} : \mathcal{K}_n^2 \times \mathcal{X}_n \to \mathcal{X}_n^4$ defined as

$$\mathsf{G}((k, k'), x) = [\mathsf{F}(k, x), \mathsf{F}(k', x), \mathsf{F}(k', (\mathsf{F}(k, x)), \mathsf{F}(k, (\mathsf{F}(k', x))] \tag{5}$$

is *not* a weak PRF, then a two-pass key-agreement protocol exists. If G is not a weak PRF, there exists an efficient distinguisher D, a polynomial $q(.)$, and a non-negligible function $\phi(.)$, such that for $x^q \xleftarrow{\$} \mathcal{X}_n^q$, $u^q \leftarrow \mathbf{R}_{\mathcal{X}_n, \mathcal{X}_n^4}(x^q)$, $k \leftarrow \mathsf{KeyGen}(1^n)$, and $k' \leftarrow \mathsf{KeyGen}(1^n)$

$$\Pr[\mathsf{D}(x^q, \mathsf{G}((k, k'), x^q)) = 1] - \Pr[\mathsf{D}(x^q, u^q) = 1] \geq \phi(n). \tag{6}$$

We now define three other functions $\widetilde{\mathsf{G}}, \mathsf{H}, \widetilde{\mathsf{H}}$ which will be used in the proof. The function $\widetilde{\mathsf{G}}$ is defined almost as G, but with $\mathsf{F}(k', .)$ replaced by a URF. The systems H and $\widetilde{\mathsf{H}}$ are defined like G

and $\widetilde{\mathsf{G}}$ respectively, but without the last term. For the rest of this section, we let $\mathbf{R}$ denote $\mathbf{R}_{\mathcal{X}_n,\mathcal{X}_n}$.

$$\widetilde{\mathsf{G}}((k,k'),x) = [\mathsf{F}(k,x), \mathbf{R}(x), \mathbf{R}(\mathsf{F}(k,x)), \mathsf{F}(k, \mathbf{R}(x))]$$
$$\mathsf{H}((k,k'),x) = [\mathsf{F}(k,x), \mathsf{F}(k',x), \mathsf{F}(k', \mathsf{F}(k,x))]$$
$$\widetilde{\mathsf{H}}((k,k'),x) = [\mathsf{F}(k,x), \mathbf{R}(x), \mathsf{F}(k', \mathbf{R}(x))]$$

Below we will show that if F is a weak PRF, then also $\widetilde{\mathsf{G}}$, H, and $\widetilde{\mathsf{H}}$ are weak PRFs. The idea behind the bit-agreement protocol given in Figure 2 is now quite simple: First, Alice and Bob each sample a random key for F. Then Bob flips a random coin b_B. If $b_\mathsf{B} = 0$, Alice and Bob together simulate an attack on G, and if $b_\mathsf{B} = 1$, they simulate an attack on $\widetilde{\mathsf{G}}$. As Alice can distinguish G from random (and thus from $\widetilde{\mathsf{G}}$), she can learn b_B with non-negligible advantage, as shown in Claim 5. However, an adversary Eve does not see the last term of G or $\widetilde{\mathsf{G}}$, as they are computed by Alice and not sent over to Bob. Hence, Eve only sees the outputs as they are given by H if $b_\mathsf{B} = 0$ and by $\widetilde{\mathsf{H}}$ if $b_\mathsf{B} = 1$. As H and $\widetilde{\mathsf{H}}$ are weak PRFs, Eve only has negligible advantage in distinguishing those two cases (and thus also in guessing b_B), as shown in Claim 6.

Claim 3 *If F is a weak PRF, then $\widetilde{\mathsf{G}}$ is a weak PRF.*

Proof. We have to show that for any polynomial $q(.)$ and $x_1, \ldots, x_q \xleftarrow{\$} \mathcal{X}_n$ ($q = q(n)$) the q four-tuples

$$[x_i, \mathsf{F}(k,x_i), \mathbf{R}(x_i), \mathbf{R}(\mathsf{F}(k,x_i)), \mathsf{F}(k, \mathbf{R}(x_i))] \tag{7}$$

are indistinguishable from random. Sample $x'_1, \ldots, x'_q \xleftarrow{\$} \mathcal{X}_n$, $x''_1, \ldots, x''_q \xleftarrow{\$} \mathcal{X}_n$, and consider the distribution

$$[x_i, \mathsf{F}(k,x_i), x'_i, x''_i, \mathsf{F}(k,x'_i)]. \tag{8}$$

As F is a weak PRF, the five tuples given by (8) are indistinguishable from random. We will now show that (8) is indistinguishable from (7). First note that $\mathbf{R}(x_1), \ldots, \mathbf{R}(x_q)$ has the same distribution as $x'_1, \ldots, x'_q$, unless $x_i = x_j$ for some $i \neq j$, as q is polynomial and $|\mathcal{X}|$ is superpolynomial, the probability of this event is negligible. So we can safely replace $\mathbf{R}(x_i)$ in (7) with x'_i in (8). Similarly, we can replace $\mathbf{R}(\mathsf{F}(k,x_i))$ with x''_i as this will make no difference unless $\mathsf{F}(k,x_i) = \mathsf{F}(k,x_j)$ or $\mathsf{F}(k,x_i) = x_j$ for some $i \neq j$, which only happens with negligible probability. □

Claim 4 *If F is a weak PRF, then H and $\widetilde{\mathsf{H}}$ are weak PRFs.*

Proof. That $\widetilde{\mathsf{H}}$ is weakly pseudorandom follows directly from the fact that $\widetilde{\mathsf{G}}$ is weakly pseudorandom (as shown by the previous claim). To show that H is weakly pseudorandom, we show that, for random $x_1, \ldots, x_q \in \mathcal{X}_n$, the tuples

$$[x_i, \mathsf{F}(k,x_i), \mathsf{F}(k',x_i), \mathsf{F}(k', \mathsf{F}(k,x_i))] \tag{9}$$

are indistinguishable from random. For this, it is sufficient by the triangle inequality, to show the following two facts. First, for random $x'_1, \ldots, x'_q \in \mathcal{X}_n$, the tuples

$$[x_i, x'_i, \mathsf{F}(k',x_i), \mathsf{F}(k',x'_i)] \tag{10}$$

are indistinguishable from (9), since F is a weak PRF and thus $\mathsf{F}(k,x_i)$ can be replaced by a random x'_i. That (10) is indistinguishable from random follows directly from the fact that F is a weak PRF. □

Claim 5 BitAgreement2(n) *has non-negligible correlation* $\phi(n) - \mathsf{negl}(n)$*, with* ϕ *as in (6).*

Proof. For $x^q \xleftarrow{\$} \mathcal{X}_n^q$ and $u^q \leftarrow \mathbf{R}_{\mathcal{X}_n,\mathcal{X}_n^q}(x^q)$

$$\begin{aligned}
\Pr[b_{\mathsf{A}} = b_{\mathsf{B}}] &= \Pr[b_{\mathsf{B}} = 1] \cdot \Pr[b_{\mathsf{A}} = 1 | b_{\mathsf{B}} = 1] + \Pr[b_{\mathsf{B}} = 0] \cdot \Pr[b_{\mathsf{A}} = 0 | b_{\mathsf{B}} = 0] \\
&= \frac{1}{2} + \frac{\Pr[b_{\mathsf{A}} = 1 | b_{\mathsf{B}} = 1] - \Pr[b_{\mathsf{A}} = 1 | b_{\mathsf{B}} = 0]}{2} \\
&= \frac{1}{2} + \frac{\Pr[\mathsf{D}(x^q, \mathsf{G}((k_{\mathsf{A}}, k_{\mathsf{B}}), x^q)) = 1] - \Pr[\mathsf{D}(x^q, \widetilde{\mathsf{G}}((k_{\mathsf{A}}, k_{\mathsf{B}}), x^q)) = 1]}{2} \\
&= \frac{1}{2} + \frac{\Pr[\mathsf{D}(x^q, \mathsf{G}((k_{\mathsf{A}}, k_{\mathsf{B}}), x^q)) = 1] - \Pr[\mathsf{D}(x^q, u^q) = 1] \pm \mathsf{negl}(n)}{2} \\
&\geq \frac{1}{2} + \frac{\phi(n) \pm \mathsf{negl}(n)}{2},
\end{aligned}$$

where the second last step follows as $\widetilde{\mathsf{G}}$ is a weak PRF (as shown in Claim 3). □

Claim 6 BitAgreement2(n) *is* $1 - \mathsf{negl}(n)$ *secure.*

Proof. Consider any efficient adversary E who can observe the communication $C = \{r^q, s^q, t^q, y^q\}$ between Alice and Bob . If $b_{\mathsf{B}} = 0$, then C has the same distribution as generated by H, and if $b_{\mathsf{B}} = 1$, then C has the same distribution as generated by $\widetilde{\mathsf{H}}$. The security now follows as an efficient E cannot distinguish H from $\widetilde{\mathsf{H}}$, since these are both weak PRFs (as shown in Claim 4). More formally, for $x^q \xleftarrow{\$} \mathcal{X}_n^q$

$$\begin{aligned}
\Pr[\mathsf{E}(C) = b_{\mathsf{B}}] &= \Pr[b_{\mathsf{B}} = 1] \cdot \Pr[\mathsf{E}(C) = 1 | b_{\mathsf{B}} = 1] + \Pr[b_{\mathsf{B}} = 0] \cdot \Pr[\mathsf{E}(C) = 0 | b_{\mathsf{B}} = 0]| \\
&= \frac{1}{2} + \frac{\Pr[\mathsf{E}(C) = 1 | b_{\mathsf{B}} = 1] - \Pr[\mathsf{E}(C) = 1 | b_{\mathsf{B}} = 0]}{2} \\
&= \frac{1}{2} + \frac{\Pr[\mathsf{E}(x^q, \widetilde{\mathsf{H}}(k_{\mathsf{A}}, x^q)) = 1] - \Pr[\mathsf{E}(x^q, \mathsf{H}(k_{\mathsf{A}}, x^q)) = 1]}{2} = \frac{1}{2} \pm \mathsf{negl}(n).
\end{aligned}$$

□

Proof (of Theorem 2). The theorem follows from Claim 5 and 6, and the fact that one can construct a key-agreement protocol from any bit-agreement protocol which has noticeable correlation and overwhelming security without increasing the number of rounds (via parallel repetition and privacy amplification [12, 13]).

5 Can we Efficiently Deconstruct "Useful" Properties?

Black Box Falsification. What does it mean that some statement "holds in Minicrypt"? Trivially, it means that in order to falsify the statement, we must assume the existence of something at least as strong as key-agreement. As observed in [24], the fact that no black-box reduction one-way functions to key-agreement exists [18], implies that no "black-box falsification" for such statements can exist (in [24] this was called a "black-box break"). E.g. by Theorem 2, there is no black-box reduction from one-way functions to a weak PRF and a distinguisher, such that the distinguisher breaks the security of the construction given by (1) when instantiated with this weak PRF.[10]

[10] Let us stress that black-box reductions for one-way functions to PRFs do exist [11, 9].

Deconstructing the Homomorphic Property. We showed that the statements of the theorems from this paper are non-trivial, by showing that they do no longer hold outside Minicrypt, or more precisely under the standard DDH assumption (which is false in Minicrypt as it implies key-agreement). These counterexamples use the homomorphic property of the group, i.e. that $(x^a)^b = (x^b)^a$. This is eminent in (2) which shows that the construction given by (1) does not give secure range extension outside of Minicrypt. Usually, a weak PRF which is homomorphic is a very useful thing to have, but clearly not if we want to use this PRF in the construction given in (1). In order to safely use a weak PRF in (1), all we have to make sure is, that the protocol from Figure 2 is NOT a secure bit-agreement protocol. Intuitively, that should not be too hard.

Open Problem 1 *Is there an efficient construction ϕ, such that for any weak PRF F, $\phi(\mathsf{F})$ is a weak PRF but the protocol from Figure 2 is NOT a secure bit-agreement protocol when instantiated with $\phi(\mathsf{F})$ (and thus (1) is a secure range extension for $\phi(\mathsf{F})$).*

We can solve the above problem by first constructing a PRG from the weak PRF and then using the GGM reduction [9] to get a regular PRF (note that (1) is trivially secure for PRFs), but this is not really efficient: each evaluation of the PRF would make a linear (in the input length) number of invocations to the weak PRF. The above problem is somewhat antipodal to questions usually asked in cryptography, where one tries to construct something useful, like asking "can we construct key-agreement from PRFs via black-box reductions" (the answer is no [18]). Whereas here we are looking for a way to make some particular construction insecure. Being more ambitious, we can ask if we can efficiently destroy any property of a weak PRF which could be used to get a key-agreement protocol.

Open Problem 2 *Is there an efficient construction ϕ, such that for any weak PRF F, $\phi(\mathsf{F})$ is a weak PRF and any construction which is secure in* Minicrypt *when instantiated with a weak PRF, is secure (in the real world) when instantiated with $\phi(\mathsf{F})$?*

The general problem can be summarized as follows. We know several constructions for extending the range of weak PRFs, getting public-coin weak PRFs from secret-coin weak PRFs[11], and achieving adaptive security from non-adaptive PRFs [24, 19], that one can show to be secure in Minicrypt, but which are (under standard assumptions) not secure in the real world. As the constructions are secure in Minicrypt, each weak/non-adaptive PRF for which the construction actually is insecure can be used to construct a key-agreement protocol (via some particular black-box construction), and must hence have a lot of structure. The question is whether there is an efficient way to modify the weak/non-adaptive PRF, such that it still keeps its original security guarantee (of being a weak/non-adaptive PRF), but cannot be used anymore for the key-agreement protocols. Then this modified PRF can safely be used in the efficient constructions that have been proven to be secure in Minicrypt.

References

1. Feng Bao, Robert H. Deng, and Huafei Zhu. Variations of Diffie-Hellman problem. In *ICICS '03*, volume 2836 of *LNCS*, pages 301–312. Springer, 2003.
2. Ivan Damgård and Jesper B. Nielsen. Expanding pseudorandom functions; or: From known-plaintext security to chosen-plaintext security. In *Advances in Cryptology — CRYPTO '02*, volume 2442 of *LNCS*, pages 449–464. Springer, 2002.

[11] This construction simply uses the secret-coin weak PRF as the public-coin weak PRF.

3. A. Dent. Cryptography in a hitchhiker's universe. *Journal of Craptology*, 4, 2007.
4. Whitfield Diffie and Martin Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, IT-22(6):644–654, 1976.
5. Bella Dubrov and Yuval Ishai. On the randomness complexity of efficient sampling. In *Proc, 38th ACM Symposium on the Theory of Computing (STOC)*, pages 711–720, 2006.
6. Stefan Dziembowski. On forward-secure storage. In *Advances in Cryptology — CRYPTO '06*, LNCS, pages 251–270. Springer, 2006.
7. Taher El-Gamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, 31(4):469–472, 1985.
8. Oded Goldreich. *Foundations of Cryptography – Volume II – Basic Applications*. Cambridge University Press, 2004.
9. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.
10. Iftach Haitner, Danny Harnik, and Omer Reingold. Efficient pseudorandom generators from exponentially hard one-way functions. In *Automata, Languages and Programming — ICALP '06 (Part II)*, volume 4052 of *LNCS*, pages 228–239. Springer, 2006.
11. Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
12. Thomas Holenstein, 2005. Personal Communication.
13. Thomas Holenstein. *Immunization of key-agreement schemes*. PhD thesis, ETH Zürich, 2006. ISBN 3-86628-088-2.
14. Thomas Holenstein. Pseudorandom generators from one-way functions: A simple construction for any hardness. In *TCC'06*, volume 3876 of *LNCS*, pages 443–461. Springer, 2006.
15. Chun-Yuan Hsiao and Leonid Reyzin. Finding collisions on a public road, or do secure hash functions need secret coins? In *Advances in Cryptology — CRYPTO '04*, volume 3152 of *LNCS*, pages 92–105. Springer, 2004.
16. Russell Impagliazzo. A personal view of average-case complexity. In *Structure in Complexity Theory Conference*, pages 134–147, 1995.
17. Russell Impagliazzo and Michael Luby. One-way functions are essential for complexity based cryptography (extended abstract). In *IEEE Symposium on the Foundations of Computer Science (FOCS) '89*, pages 230–235, 1989.
18. Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *Proc, 21th ACM Symposium on the Theory of Computing (STOC)*, pages 44–61, 1989.
19. Ueli Maurer, Yvonne Anne Oswald, Krzysztof Pietrzak, and Johan Sjödin. Luby-Rackoff ciphers from weak round functions? In *Advances in Cryptology — EUROCRYPT '06*, volume 4004 of *LNCS*, pages 391–408. Springer, 2006.
20. Ueli M. Maurer and Johan Sjödin. A fast and key-efficient reduction of chosen-ciphertext to known-plaintext security. In *Advances in Cryptology — EUROCRYPT '07*, volume 4515 of *LNCS*, pages 498–516. Springer, 2007.
21. Kazuhiko Minematsu and Yukiyasu Tsunoo. Expanding weak PRF with small key size. In *Information Security and Cryptology — ICISC '05*, volume 3935 of *LNCS*, pages 284–298. Springer, 2005.
22. Moni Naor, Benny Pinkas, and Omer Reingold. Distributed pseudo-random functions and KDCs. In *Advances in Cryptology — EUROCRYPT '99*, volume 1592 of *LNCS*, pages 327–346. Springer, 1999.
23. Moni Naor and Omer Reingold. Number-theoretic constructions of efficient pseudo-random functions. *J. of the ACM*, 51(2):231–262, 2004.
24. Krzysztof Pietrzak. Composition implies adaptive security in minicrypt. In *Advances in Cryptology — EUROCRYPT '06*, volume 4004 of *LNCS*, pages 328–338. Springer, 2006.
25. Krzysztof Pietrzak and Johan Sjödin. Domain extension for weak PRFs; the good, the bad, and the ugly. In *Advances in Cryptology — EUROCRYPT '07*, volume 4515 of *LNCS*, pages 517–533. Springer, 2007.
26. Hoeteck Wee. Finding pessiland. In *Theory of Cryptograpy — TCC '06*, volume 3876 of *LNCS*, pages 429–442. Springer, 2006.