Composition Implies Adaptive Security in Minicrypt

Krzysztof Pietrzak*

Département d'informatique, Ecole Normale Supérieure, Paris, France
pietrzak@di.ens.fr

Abstract. To prove that a secure key-agreement protocol exists one must at least show $P \neq NP$. Moreover any proof that the sequential composition of two non-adaptively secure pseudorandom functions is secure against at least two adaptive queries must falsify the decisional Diffie-Hellman assumption, a standard assumption from public-key cryptography. Hence proving any of this two seemingly unrelated statements would require a significant breakthrough. We show that *at least one* of the two statements is true.

To our knowledge this gives the first *positive* cryptographic result (namely that composition implies some weak adaptive security) which holds in Minicrypt, but not in Cryptomania, i.e. under the assumption that one-way functions exist, but public-key cryptography does not.

1 Introduction

A pseudorandom function (PRF) is a function which cannot be distinguished from a uniformly random function by any efficient adversary. One can give different security definitions for PRFs depending on how the attacker can access the function: a *non-adaptive* adversary must choose all his queries to the function at once, whereas a (more powerful) *adaptive* adversary must only decide on the i'th query after receiving the $i-1$'th output. As a generalisation we define k-adaptive adversaries which can choose k blocks of queries to be made, where the k'th block must be chosen at once but only after receiving the outputs to the $k-1$'th block (in particular 1-adaptive means non-adaptive, and ∞-adaptive means adaptive). Consider the following two statements:

$\mathfrak{K}_k$**:** There exists a secure k-pass key-agreement protocol.
$\mathfrak{C}_k$**:** The sequential composition of two $(k-1)$-adaptively secure PRFs is k-adaptively secure.

The main result of this paper is that either composition of PRFs always increases the security in the sense that the cascade is k-adaptive secure whenever the components are $k-1$ secure OR that key agreement exists.

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: 10.1007/978-3-540-34547-3_36

* Most of this work was done while the author was a PhD student at ETH where he was supported by the Swiss National Science Foundation, project No. 200020-103847/1. Part of this work is supported by the Commission of the European Communities through the IST program under contract IST-2002-507932 ECRYPT.

S. Vaudenay (Ed.): EUROCRYPT 2006, LNCS 4004, pp. 328–338, 2006.
© Springer-Verlag Berlin Heidelberg 2006

Theorem 1. *For any $k \geq 2$:* $\mathfrak{C}_k \vee \mathfrak{K}_{2k-1}$

This theorem has a nice interpretation in terms of Impagliazzo's five possible worlds as described in the survey paper "A Personal View of Average-Case Complexity" [8]. Here "possible world" means that with our current knowledge we cannot rule out it as being reality. As each world does exists relative to an oracle, showing equivalence of two worlds would require non-relativizing techniques, and in the ten years that passed since this survey none has been resolved.[1] This five worlds are Algorithmica (where $P = NP$), Heuristica ($NP \neq P$ but NP is tractable on average), Pessiland (NP is hard on average but one-way functions do not exist), Minicrypt (one-way functions exist) and Cryptomania (Public-key cryptography exists, this is probably the real world). In this view, the theorem states that for any $k \geq 2$ the statement $\mathfrak{C}_k$ holds in Minicrypt but not in Cryptomania. As the naming suggests, Cryptomania is cryptographers paradise, but our result somewhat challenges this viewpoint, as cryptographers interested only in symmetric cryptography might well prefer to live in Minicrypt rather than in Cryptomania, as some results (in particular $\mathfrak{C}_k$) only can be found there.

But let us stress that there are known (black-box) constructions of adaptively secure PRFs from non-adaptively secure PRFs [4], but these constructions are inefficient as they need a linear (in the security parameter) number of calls to the underlying primitive on each invocation. Thus we do not show that adaptively secure PRF exists in Minicrypt (as this is known), but rather that here adaptive security can be achieved by probably most straight forward and efficient construction: cascading two functions.

We prove Theorem 1 by constructing a $2k-1$-pass key-agreement protocol from any pseudorandom functions which provides a counterexample for $\mathfrak{C}_k$, i.e. from any $(k-1)$-adaptively secure pseudorandom functions $\mathsf{F}(\cdot)$ and $\mathsf{G}(\cdot)$ where there exists an efficient k-adaptive D which can distinguish $\mathsf{G}(\mathsf{F}(\cdot))$ from a random function.

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: a system is usually considered broken even if only a *non-uniform* advantage exists, whereas a protocol should be *uniform* and achieve its task with *overwhelming*[2] probability to be considered useful. The key-agreement protocol we construct uses D as a black-box, and only if D is *uniform* and has *noticeable* advantage in distinguishing $\mathsf{G}(\mathsf{F}(\cdot))$ from random, we will get a useful (as described above) key-agreement protocol. But if D in non-uniform, also the key-agreement protocol will be non-uniform. Furthermore if D has only *non-negligible* (but not noticeable) advantage, then our key-agreement protocol will only work (i.e. have overwhelming success probability) for infinitely many values of the security parameter (and not as usually for all).

[1] But several new worlds, in particular between Minicrypt and Cryptomania [3], have been added. Recently Harnik and Naor [5] proposed an interesting approach to show Minicrypt=Cryptomania. Wee investigates Pessiland in [17]. A classical result due to Rudich [15] oracle separates $\mathfrak{K}_k$ from $\mathfrak{K}_{k+1}$ for every k.

[2] $\tau(\cdot)$ is overwhelming if $1 - \tau(\cdot)$ is negligible.

1.1 What Is Known?

It is known that under the decisional Diffie-Hellman (DDH) assumption two-pass key-agreement (i.e. public-key encryption) exists [1, 2], and in [13] it is shown that under the same assumption $\neg\mathfrak{C}_2$ holds, i.e. that composition does not imply adaptive security.[3] Thus [13] shows a negative result for private-key systems under a standard assumption from public-key cryptography. By Theorem 1 this is not just an artificial property of the counterexample given in [13], but in fact any falsification of $\mathfrak{C}_2$ implies (and thus must either assume or unconditionally prove) the existence of the central public-key primitive key-agreement.

Interestingly the equivalent of $\mathfrak{C}_2$ in the information theoretic setting is true: the cascade of two functions, each having security ϵ against *non-adaptive* (computationally unbounded) distinguishers making at most q queries, has security 2ϵ against any *adaptive* distinguisher making q queries [11]. Therefore the reason why composition does imply adaptive security in the information-theoretic but probably not in the computational setting is closely related to the fact that public-key cryptography cannot exist in the information theoretic setting [16, 10] but is believed to exist in the real world [1]. We'll muse further on the implications of Theorem 1 in Section 4.

2 Basic Definitions

Throughout we denote by $n \in \mathbb{N}$ a security parameter. An algorithm is efficient if it can be implemented by a probabilistic Turing machine whose expected running time is polynomial in the input length (which for us will always mean polynomial in n). We use a SANS-SERIF font for efficient entities and a $\mathcal{CALLIGRAPHIC}$ font for idealised systems like uniform random functions.

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 $n \geq n_0$. And contrarily μ is non-negligible if for any $c > 0$ we have $\mu(n) \geq 1/n^c$ for infinitely many n.

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.

Unless stated otherwise, all characters that appear below are probabilistic efficient Turing machines.

Bit-Agreement. Bit-agreement is a protocol between two efficient parties, let's call them Amélie and Benoît . They get as a common input the security

[3] In [13] a $\mathsf{F}(\cdot)$ and $\mathsf{G}(\cdot)$ are constructed which are non-adaptively secure under the DDH assumption, but where *three* (and not two as required for $\neg\mathfrak{C}_2$) adaptive queries are enough to learn the whole key when querying $\mathsf{G}(\mathsf{F}(\cdot))$. But after two adaptive queries one already learns the key of G and thus can distinguish $\mathsf{G}(\mathsf{F}(\cdot))$ from random, and this is all we need to get $\neg\mathfrak{C}_2$. Previous to [13] is was already known that there is no black-box proof for $\mathfrak{C}_2$ as Myers [12] has constructed an oracle relative to which $\neg\mathfrak{C}_2$.

parameter n in unary (denoted 1^n) and can communicate over an authentic channel. Finally Amélie and Benoît output a bit b_{A} and b_{B} respectively. The protocol has correlation ϵ if for all n

$$\mathsf{Pr}[b_{\mathsf{A}} = b_{\mathsf{B}}] \geq \frac{1+\epsilon(n)}{2}$$

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

$$\mathsf{Pr}[\mathsf{E}(1^n, C) \to b_{\mathsf{A}}] \leq 1 - \frac{\delta(n)}{2}$$

Key-Agreement. If $\epsilon(\cdot)$ and $\delta(\cdot)$ are overwhelming then such a protocol achieves key-agreement. Any protocol which achieves bit-agreement with a noticeable correlation $\epsilon(\cdot)$ and overwhelming security $\delta(\cdot)$ can be turned into a key-agreement protocol by sequential composition, and using parallel repetition this can even be done without increasing the number of rounds [6, 7].

If $\epsilon(\cdot)$ is only non-negligible (i.e. for any $c > 0 : \epsilon(n) \geq 1/n^c$ for all $n \in S_c \subset \mathbb{Z}$ where $|S_c|$ is infinite), then also the key-agreement protocol will only achieve correctness for security parameters $n \in S_c$ (one can choose any constant c here, the running time of the key-agreement protocol will then basically grows as n^{2c}).

Distinguisher. By a k-adaptive *distinguisher* we denote an efficient oracle algorithm which at the end of the computation outputs a decision bit. He may query the oracle an arbitrary number of times, but the queries must come in k blocks where he must settle for a whole block before reading any outputs on queries from that block.

This definition is not standard, but note that a 1-adaptive distinguisher is just a standard non-adaptive distinguisher and a ∞-adaptive distinguisher is a standard adaptive distinguisher.

As we only consider stateless systems (which always give the same answer on the same query) w.l.o.g we always *can and will assume that a distinguisher never makes the same query twice.* Moreover we *require the distinguishers themselves to be stateless.* This can be done w.l.o.g. if we always provide the previous outputs of the system queried as an input to the distinguisher when he must come up with the next query or the final decision bit (note that we need not to provide the previous inputs to the system as the distinguisher can compute this inputs himself).

Pseudorandom Function/Permutation. A pseudorandom function (PRF) is a pair of efficient algorithms F and $\mathsf{KeyGen}_{\mathsf{F}}$ where for any $n \in \mathbb{N}$ we have $\mathsf{KeyGen}_{\mathsf{F}} : 1^n \to \mathcal{K}_n$ and $\mathsf{F} : \mathcal{K}_n \times \{0,1\}^n \to \{0,1\}^n$. Let $\mathsf{F}_k(\cdot) \stackrel{\text{def}}{=} \mathsf{F}(k, \cdot)$. Let $\mathcal{R}_n : \{0,1\}^n \to \{0,1\}^n$ be a uniform random function, then F is ℓ-adaptive secure if for any efficient ℓ-adaptive distinguisher D

$$|\mathsf{Pr}[\mathsf{D}^{\mathsf{F}_k(\cdot)}(1^n) \to 1 | k \leftarrow \mathsf{KeyGen}_{\mathsf{F}}(1^n)] - \mathsf{Pr}[\mathsf{D}^{\mathcal{R}_n(\cdot)}(1^n) \to 1]| = \tau(n).$$

for some negligible τ. Pseudorandom permutations (PRP) are defined similarly, but here one additionally requires that for any k, $\mathsf{F}_k(\cdot)$ is a permutation.

Sequential Composition. For two functions F and G we denote by $\mathsf{G}\circ\mathsf{F}$ their sequential composition.

$$\mathsf{G}\circ\mathsf{F}(x) \stackrel{\text{def}}{=} \mathsf{G}(\mathsf{F}(x)).$$

For a set S we denote by $x \stackrel{\$}{\leftarrow} S$ that x is assigned a value from S uniformly at random.

3 The Reduction

In this section we prove the statement $\neg\mathfrak{C}_k \Rightarrow \mathfrak{K}_{2k-1}$ of Theorem 1. Actually, we only show that $\neg\mathfrak{C}_k$ implies a $(2k-1)$-pass *bitagreement* protocol with noticeable correlation and overwhelming security, but as said in the previous section, this is equivalent to $\mathfrak{K}_{2k-1}$.

For the clarity of exposition we prove only the special case $k = 2$ and we assume that $\neg\mathfrak{C}_2$ holds in a strong sense, namely that the cascade considered can be distinguished by an adversary which makes only two adaptive queries, this is a special case of a general 2-adaptive distinguisher which can make two blocks of arbitrary many queries (where he must settle for whole blocks at once). At the end of this section we will show how the reduction must be extended to cover the general case (and thus to prove Theorem 1).

Let $\mathsf{F}, \mathsf{KeyGen}_\mathsf{F}$ and $\mathsf{G}, \mathsf{KeyGen}_\mathsf{G}$ be two pseudorandom functions, each secure against non-adaptive distinguishers, but which can be distinguished with two adaptive queries. This means that there exists an efficient D and a non-negligible ϕ such that

$$\mathsf{Pr}[b_2 = 1] - \mathsf{Pr}[b_1 = 1] \geq \phi(n) \tag{1}$$

where b_1 and b_2 are bits whose distribution is defined by Games 1 and 2 below where D either queries the sequential composition (Game 1) or a random function (Game 2) with two adaptive queries.

Game 1	Game 2	Game 3
$k_1 \leftarrow \mathsf{KeyGen}_\mathsf{F}(1^n)$		
$k_2 \leftarrow \mathsf{KeyGen}_\mathsf{G}(1^n)$		$k \leftarrow \mathsf{KeyGen}_\mathsf{G}(1^n)$
$x_1 \leftarrow \mathsf{D}(1^n)$	$x_1 \leftarrow \mathsf{D}(1^n)$	$z_1 \stackrel{\$}{\leftarrow} \{0,1\}^n$
$y_1 \leftarrow \mathsf{G}_{k_2}\circ\mathsf{F}_{k_1}(x_1)$	$y_1 \leftarrow \mathcal{R}_n(x_1)$	$y_1 \leftarrow \mathsf{G}_k(z_1)$
$x_2 \leftarrow \mathsf{D}(y_1)$	$x_2 \leftarrow \mathsf{D}(y_1)$	$z_2 \stackrel{\$}{\leftarrow} \{0,1\}^n$
$y_2 \leftarrow \mathsf{G}_{k_2}\circ\mathsf{F}_{k_1}(x_2)$	$y_2 \leftarrow \mathcal{R}_n(x_2)$	$y_2 \leftarrow \mathsf{G}_k(z_2)$
$b_1 \leftarrow \mathsf{D}(y_1, y_2)$	$b_2 \leftarrow \mathsf{D}(y_1, y_2)$	$b_3 \leftarrow \mathsf{D}(y_1, y_2)$

In Game 2 the y_1, y_2 are just uniform random values whereas in Game 3 the y_1, y_2 are computed by G on random inputs. From the non-adaptive security of G it also follows that for some negligible δ_{23}

$$|\mathsf{Pr}[b_2 = 1] - \mathsf{Pr}[b_3 = 1]| \leq \delta_{23}(n). \tag{2}$$

PROTOCOL BITAGREEMENT(n)

Amélie		Benoît
$b_A \xleftarrow{\$} \{0,1\}$		
$k_A \leftarrow \mathsf{KeyGen}_\mathsf{F}(1^n)$		$k_B \leftarrow \mathsf{KeyGen}_\mathsf{G}(1^n)$
$x_1 \leftarrow \mathsf{D}(1^n)$		
if $b_A = 0$ then $z_1 \leftarrow \mathsf{F}_{k_A}(x_1)$		
otherwise $z_1 \xleftarrow{\$} \{0,1\}^n$	$z_1 \rightarrow$	$y_1 \leftarrow \mathsf{G}_{k_B}(z_1)$
	$\leftarrow y_1$	
$x_2 \leftarrow \mathsf{D}(y_1)$		
if $b_A = 0$ then $z_2 \leftarrow \mathsf{F}_{k_A}(x_2)$		
otherwise $z_2 \xleftarrow{\$} \{0,1\}^n$	$z_2 \rightarrow$	$y_2 \leftarrow \mathsf{G}_{k_B}(z_2)$
		$b_B \leftarrow \mathsf{D}(y_1, y_2)$

Fig. 1. 3-pass BITAGREEMENT protocol from a 2-adaptive D

With such an F, G and D we can construct a bit-agreement protocol with non-negligible correlation and overwhelming security (and thus get key-agreement) as shown in Figure 1. If D is randomised we need Amélie and Benoît to use the same random coins for D in BITAGREEMENT. Here Amélie can simply choose the random coins initially and send them to Benoît .

Claim 1. BITAGREEMENT(n) *has correlation* $\phi - \delta_{23}$.

Proof. Note that if $b_A = 0$ ($b_A = 1$) then the distribution of b_B is the same as the distribution of b_1 (b_3) in game 1 (game 3), now as (1) and (2) imply

$$\mathsf{Pr}[b_3 = 1] - \mathsf{Pr}[b_1 = 1] \geq \phi(n) - \delta_{23}(n)$$

we get

$$\begin{aligned}\mathsf{Pr}[b_A = b_B] &= \mathsf{Pr}[b_A = 0]\mathsf{Pr}[b_B = 0|b_A = 0] + \mathsf{Pr}[b_A = 1]\mathsf{Pr}[b_B = 1|b_A = 1] \\ &= \frac{1 - \mathsf{Pr}[b_1 = 1]}{2} + \frac{\mathsf{Pr}[b_3 = 1]}{2} \\ &\geq \frac{1 + \phi(n) - \delta_{23}(n)}{2}\end{aligned}$$

□

Claim 2. BITAGREEMENT(n) *is* δ*-secure for an overwhelming* δ.

Proof. We must show that there is an overwhelming δ such that for all efficient D

$$\mathsf{Pr}[\mathsf{D}(z_1, y_1, z_2) \rightarrow b_A] \leq 1 - \frac{\delta(n)}{2}$$

We consider six more games which all define a distribution for the values (z_1, y_1, z_2). The distribution of (z_1, y_1, z_2) in game 4 and 9 is the same as in BITAGREEMENT(n) conditioned on $b_A = 0$ and $b_A = 1$ respectively.

Game 4	Game 5	Game 6
$k_1 \leftarrow \mathsf{KeyGen}_{\mathsf{F}}(1^n)$	$k_1 \leftarrow \mathsf{KeyGen}_{\mathsf{F}}(1^n)$	$k_1 \leftarrow \mathsf{KeyGen}_{\mathsf{F}}(1^n)$
$k_2 \leftarrow \mathsf{KeyGen}_{\mathsf{G}}(1^n)$		
$x_1 \leftarrow \mathsf{D}(1^n)$	$x_1 \leftarrow \mathsf{D}(1^n)$	$x_1 \leftarrow \mathsf{D}(1^n)$
$z_1 \leftarrow \mathsf{F}_{k_1}(x_1)$	$z_1 \leftarrow \mathsf{F}_{k_1}(x_1)$	$y_1 \xleftarrow{\$} \{0,1\}^n$
$y_1 \leftarrow \mathsf{G}_{k2}(z_1)$	$y_1 \xleftarrow{\$} \{0,1\}^n$	$x_2 \leftarrow \mathsf{D}(y_1)$
$x_2 \leftarrow \mathsf{D}(y_1)$	$x_2 \leftarrow \mathsf{D}(y_1)$	$z_1 \leftarrow \mathsf{F}_{k_1}(x_1)$
$z_2 \leftarrow \mathsf{F}_{k_1}(x_2)$	$z_2 \leftarrow \mathsf{F}_{k_1}(x_2)$	$z_2 \leftarrow \mathsf{F}_{k_1}(x_2)$

Game 7	Game 8	Game 9
		$k_2 \leftarrow \mathsf{KeyGen}_{\mathsf{G}}(1^n)$
$x_1 \leftarrow \mathsf{D}(1^n)$	$x_1 \leftarrow \mathsf{D}(1^n)$	$x_1 \leftarrow \mathsf{D}(1^n)$
$y_1 \xleftarrow{\$} \{0,1\}^n$	$z_1 \leftarrow \mathcal{R}_n(x_1)$	$z_1 \leftarrow \mathcal{R}_n(x_1)$
$x_2 \leftarrow \mathsf{D}(y_1)$	$y_1 \xleftarrow{\$} \{0,1\}^n$	$y_1 \leftarrow \mathsf{G}_{k_2}(z_1)$
$z_1 \leftarrow \mathcal{R}_n(x_1)$	$x_2 \leftarrow \mathsf{D}(y_1)$	$x_2 \leftarrow \mathsf{D}(y_1)$
$z_2 \leftarrow \mathcal{R}_n(x_2)$	$z_2 \leftarrow \mathcal{R}_n(x_2)$	$z_2 \leftarrow \mathcal{R}_n(x_2)$

With $\mathsf{Pr}_{Gi}[E]$ we denote the probability of the event E in game i, and δ_{ij} is defined by

$$|\mathsf{Pr}_{Gi}[\mathsf{D}(z_1, y_1, z_2) \to 1] - \mathsf{Pr}_{Gj}[\mathsf{D}(z_1, y_1, z_2) \to 1]| = \delta_{ij}(n)$$

Game 4 differs from Game 5 only by the computation of y_1 which is computed by G and random respectively. As G is non-adaptively secure (and a single query is always non-adaptive) δ_{45} is negligible. For the same reason δ_{89} is negligible. Game 6 differs from Game 7 only by the computation of z_1 and z_2 which in Game 6 are non-adaptively computed by F and in Game 7 by $\mathcal{R}$, so from F's non-adaptive security it follows that δ_{67} is also negligible. Finally δ_{56} and δ_{78} are 0 as Game 5 is equivalent to Game 6 (only the order of the commands is changed to emphasis that in Game 5 the F is in fact queried non-adaptively) and Game 7 is equivalent to Game 8.

Using the triangle inequality we see that $\delta_{49} \leq \sum_{i=4}^{8} \delta_{i\,i+1}$ is negligible, and thus $\delta \stackrel{\text{def}}{=} 1 - \delta_{49}$ is overwhelming. We can now conclude the proof of the claim as

$$\begin{aligned}
&\mathsf{Pr}[\mathsf{D}(z_1, y_1, z_2) \to b_A] \\
&= \mathsf{Pr}[b_A = 0]\mathsf{Pr}[\mathsf{D}(z_1, y_1, z_2) \to 0 | b_A = 0] + \\
&\quad \mathsf{Pr}[b_A = 1]\mathsf{Pr}[\mathsf{D}(z_1, y_1, z_2) \to 1 | b_A = 1] \\
&= (1 - \mathsf{Pr}[\mathsf{D}(z_1, y_1, z_2) \to 1 | b_A = 0] + \mathsf{Pr}[\mathsf{D}(z_1, y_1, z_2) \to 1 | b_A = 1])/2 \\
&= (1 - \mathsf{Pr}_{G4}[\mathsf{D}(z_1, y_1, z_2) \to 1] + \mathsf{Pr}_{G9}[\mathsf{D}(z_1, y_1, z_2) \to 1])/2 \\
&\leq (1 + \delta_{49})/2 \\
&= 1 - \delta/2
\end{aligned}$$

This concludes the proof of $\neg\mathfrak{C}_k \Rightarrow \mathfrak{K}_{2k-1}$ for the case $k = 2$ with the additional assumption that the cascade can be broken by a distinguisher D which makes two adaptive queries (and not a general 2-adaptive distinguisher). □

We first explain how to adapt the reduction so that if works for any 2-adaptive distinguisher and not just for two adaptive queries. Then we show how to adapt it so that it works for any $k \geq 2$ which will then conclude the proof of Theorem 1.

PROTOCOL BITAGREEMENT(n)

	Amélie	Benoît
	$b_A \xleftarrow{\$} \{0,1\}$	
	$k_A \leftarrow \mathsf{KeyGen}_{\mathsf{F}}(1^n)$	$k_B \leftarrow \mathsf{KeyGen}_{\mathsf{G}}(1^n)$
for $i = 1$ to $k-1$ do		
	$X_i \leftarrow \mathsf{D}''(Y_1, \ldots, Y_{i-1})$	
	if $b_A = 0$then $Z_i \leftarrow \mathsf{F}_{k_A}(X_i)$	
	otherwise $Z_i \xleftarrow{\$} \{0,1\}^n$	$Z_i \rightarrow Y_i \leftarrow \mathsf{G}_{k_B}(Z_i)$
		$\leftarrow Y_i$
od;		
	$X_k \leftarrow \mathsf{D}''(Y_1, \ldots, Y_{k-1})$	
	if $b_A = 0$ then $Z_k \leftarrow \mathsf{F}_{k_A}(X_k)$	
	otherwise $Z_k \xleftarrow{\$} \{0,1\}^n$	$Z_k \rightarrow Y_k \leftarrow \mathsf{G}_{k_B}(Z_k)$
		$b_B \leftarrow \mathsf{D}''(Y_1, \ldots, Y_k)$

Fig. 2. $(2k-1)$-pass BITAGREEMENT protocol from a k-adaptive D''

REDUCTION FROM 2-ADAPTIVE D'. Let D' be any 2-adaptive distinguisher which can distinguish $\mathsf{F}_{k_1} \circ \mathsf{G}_{k_2}$ from random. From such a D' we can construct a 3-pass bitagreement protocol almost like from the D which made only two queries. If $q = q(n)$ denotes (an upper bound on) the size of the blocks requested by D', then just replace all occurrences of $x_1, x_2, y_1, y_2, z_1, z_2$ by appropriate q-tuples $X_1, X_2, Y_1, Y_2, Z_1, Z_2$ in the bitagreement protocol. For example replace $x_1 \leftarrow \mathsf{D}(1^n)$ with $X_1 = (x_1^1, x_1^2, \ldots, x_1^q)$ where $X_1 \leftarrow \mathsf{D}'(1^n)$, similarly replace $y_1 \leftarrow \mathsf{F}_{k_1} \circ \mathsf{G}_{k2}(x_1)$ by $Y_1 \leftarrow \mathsf{F}_{k_1} \circ \mathsf{G}_{k2}(X_1)$ and so on.

REDUCTION FROM k-ADAPTIVE D''. For any $k \geq 2$, let D'' be any k-adaptive distinguisher for $\mathsf{F}_{k_1} \circ \mathsf{G}_{k2}$ from random. To construct a bitagreement from such a distinguisher we can proceed similarly to the $k = 2$ case, only the number of rounds must be increased as now D'' must be fed with k and not just 2 input blocks.

The construction of $(2k-1)$-pass bitagreement from a k-adaptive D'' is shown in Figure 2. It is straight forward (and we omit it) to adapt the Claims 1 and 2 and their proofs for this protocol.

4 Discussion

Does Theorem 1 $\mathfrak{C}_k \vee \mathfrak{K}_{2k-1}$ have any practical meaning? After all, DDH is believed to be true in the real world, so $\mathfrak{K}_2$ is true [1] and $\mathfrak{C}_2$ is wrong [13]. Even if someday $(2k-1)$-pass key-agreement turns out to be impossible, having $\mathfrak{C}_k$ instead is a cold comfort.

But one can see $\mathfrak{C}_k \vee \mathfrak{K}_{2k-1}$ as a positive result, even when assuming that DDH is true: Composition of k-adaptively secure pseudorandom functions implies $(k+1)$-adaptive security[4], *unless the pseudorandom functions themselves have some public-key functionality* in the sense that they can be turned into a key-agreement protocol by a black-box (BB for short) reduction. Of course that was more an intuitive argument than a result that can be actually applied. In the next section we prove a first positive composition result for PRFs whose security can be BB-reduced to the security of a one-way function.

4.1 Black-Box Breaks

Combining Theorem 1 with the Impagliazzo-Rudich result [9] that key-agreement cannot be BB-reduced to one-way functions we can prove a first positive result in the direction that composition sometimes does imply adaptive security (or rather, that the adaptive security cannot be broken in a generic way) even in the computational setting. Before we can state the theorem we first need some definitions.

$\mathsf{F}^{(\cdot)}$ is an oracle PRF whose k-adaptive security can be BB-reduced to the one-wayness of the oracle if the following is true: There exists an efficient $\mathsf{B}^{(\cdot)}$ such that for any (not necessarily efficient) k-adaptive adversary $\mathcal{A}^{(\cdot)}$ and any f (for simplicity we assume f is $\{0,1\}^* \to \{0,1\}^*$ and length preserving) for which

$$\left|\mathsf{Pr}[k \leftarrow \mathsf{KeyGen}^f_{\mathsf{F}}(n); \mathcal{A}^{\mathsf{F}^f_k} \to 1] - \mathsf{Pr}[\mathcal{A}^{\mathcal{R}_n} \to 1]\right|$$

is k-negligible (note that this means that $\mathcal{A}$ breaks the k-adaptive pseudorandomness of F^f), $\mathsf{B}^{\mathcal{A},f}$ breaks the one-wayness of f, this means that then also

$$\mathsf{Pr}[x \xleftarrow{\$} \{0,1\}^n; \mathsf{B}^{\mathcal{A},f}(f(x)) \in f^{-1}(x)]$$

is non-negligible. This definition of BB-reduction is standard and called a fully-BB reduction in the taxonomy from [14]. The definition of a BB-break given below is not standard.

We say that the k-adaptive security of $\mathsf{F}^{(\cdot)}$ can be *BB-broken* if there exists an efficient k-adaptive $\mathsf{C}^{(\cdot)}$ where

$$\left|\mathsf{Pr}[k \leftarrow \mathsf{KeyGen}^f_{\mathsf{F}}(n); \mathsf{C}^{\mathsf{F}^f_k, f} \to 1] - \mathsf{Pr}[\mathsf{C}^{\mathcal{R}_n, f} \to 1]\right|$$

is noticeable for all f; So C can distinguish F^f from $\mathcal{R}$ for every f, i.e. C breaks the the security of the construction $\mathsf{F}^{(\cdot)}$ and not some particular instantiation.

[4] And in particular composition of non-adaptively secure pseudorandom functions implies 2-adaptive security.

Note that if the k-adaptive security of $\mathsf{F}^{(\cdot)}$ can be BB-broken, then it obviously cannot be BB-reduced to the one-wayness of the oracle, but the converse is not true in general.

Theorem 2. *If the k-adaptive security of the PRFs $\mathsf{F}^{(\cdot)}$ and $\mathsf{G}^{(\cdot)}$ can be BB-reduced to the one-wayness of the oracle, then the $(k+1)$-adaptive security of $\mathsf{G}^{(\cdot)} \circ \mathsf{F}^{(\cdot)}$ cannot be BB-broken.*

Proof. The proof is by contradiction: assume there is $(k+1)$-adaptive distinguisher $\mathsf{C}^{(\cdot)}$ which can distinguish $\mathsf{G}^f \circ \mathsf{F}^f$ from a random function with noticeable advantage for any f. With such a $\mathsf{C}^{(\cdot)}, \mathsf{F}^f, \mathsf{G}^f$ we can construct a key-agreement protocol.[5] The security of this protocol can be BB-reduced to the k-adaptive security of F^f and G^f whose security can again be BB-reduced to the one-wayness of f. So we have a BB-reduction from key-agreement to one-way functions which is not possible [9]. □

Note that the theorem does not claim that the $k+1$-adaptive security of $\mathsf{F}^f \circ \mathsf{G}^f$ can be BB-reduced to the one-wayness of f, but something weaker. Namely that there is no single efficient $\mathsf{C}^{(\cdot)}$ which breaks the $(k+1)$-adaptive security for all f.

4.2 Outlook

Are there other interesting statements that one can we prove to be true only under the assumption that public-key cryptography does not exist? It seems unlikely that our composition result is an isolated example.

As shown in Theorem 2 given such a statement one might well be able to prove a weaker version of it without making the (unlikely) assumption that public-key crypto does not exist. But what does "BB-broken" as used in Theorem 2 actually mean? Can one strengthen this theorem and replace "BB-broken" with "BB-reduced to the one-wayness of the oracle" or show that this is not possible.

Can we strengthen Theorem 1? For example can we show that key-agreement (via a BB-reduction) exists when the composition of two $(k-2)$-adaptive PRFs secure PRFs is k-adaptive secure[6]? We think this is not true,[7] but we believe that Theorem 2 holds with an infinite gap, i.e. where k-adaptive is replaced by non-adaptive and $(k+1)$-adaptive by adaptive. To show this one would have to show that there exists some statement $\mathfrak{L}$ such that $\mathfrak{L}$ is implied by the statement "the composition of two non-adaptive PRFs is not adaptively secure" and where $\mathfrak{L}$ cannot be BB-reduced to one-way functions.

[5] As shown in Section 3 for the special case $k = 1$ and where each of the $k+1$ blocks contained only one message.

[6] This is statement $\mathfrak{C}_k$ with an increased gap, i.e. $k-2$ instead of $(k-1)$.

[7] Because there seems to be an oracle relative to which no key-agreement exists and cascading $(k-2)$-adaptive PRFs does not give k-adaptive security. But we didn't check all details.

Acknowledgments

I'd like to thank Ueli Maurer for insightful discussions on this topic and Thomas Holenstein for several clarifying conversations on key- and bit-agreement.

References

1. Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, IT-22(6):644–654, 1976.
2. 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.
3. Yael Gertner, Sampath Kannan, Tal Malkin, Omer Reingold, and Mahesh Viswanathan. The Relationship between Public Key Encryption and Oblivious Transfer. In *FOCS*, pages 325–335, 2000.
4. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.
5. Danny Harnik and Moni Naor. On the Compressibility of NP instances and Cryptographic Applications, 2005. Manuscript.
6. Thomas Holenstein, 2005. Personal Communication.
7. Thomas Holenstein. *Immunization of key-agreement schemes, PhD.thesis.* PhD thesis, ETH Zürich, 2006. to appear.
8. Russell Impagliazzo. A personal view of average-case complexity. In *Structure in Complexity Theory Conference*, pages 134–147, 1995.
9. 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.
10. Ueli M. Maurer. Secret key agreement by public discussion from common information. *IEEE Transactions on Information Theory 39(3)*, pages 733-742, 1993
11. Ueli Maurer, Krzysztof Pietrzak, and Renato Renner. Indistinguishability Amplification, 2006. Manuscript.
12. Steven Myers. Black-box composition does not imply adaptive security. In *Advances in Cryptology — EUROCRYPT 04*, volume 3027 of *Lecture Notes in Computer Science*, pages 189–206, 2004.
13. Krzysztof Pietrzak. Composition does not imply adaptive security. In *Advances in Cryptology — CRYPTO '05*, volume 3621 of *Lecture Notes in Computer Science*, pages 55–65, 2005.
14. Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Notions of reducibility between cryptographic primitives. In *TCC*, pages 1–20, 2004.
15. Steven Rudich. The use of interaction in public cryptosystems (extended abstract). In *CRYPTO*, pages 242–251, 1991.
16. Claude E. Shannon. A mathematical theory of communication. *Bell Systems Technical Journal*, 27:373–423 and 27:623–656, 1948.
17. Hoeteck Wee. Finding pessiland. In *TCC*, pages 429–442, 2006.