

On the (Im)Possibility of Key Dependent Encryption*

Iftach Haitner**¹ and Thomas Holenstein***²

¹ Microsoft Research, New England Campus

² Department of Computer Science, Princeton University

Abstract. We study the possibility of constructing encryption schemes secure under messages that are chosen depending on the key k of the encryption scheme itself. We give the following separation results that hold both in the private and in the public key settings:

- Let \mathcal{H} be the family of $\text{poly}(n)$ -wise independent hash-functions. There exists no fully-black-box reduction from an encryption scheme secure against key-dependent messages to one-way permutations (and also to families of trapdoor permutations) if the adversary can obtain encryptions of $h(k)$ for $h \in \mathcal{H}$.
- There exists no reduction from an encryption scheme secure against key-dependent messages to, essentially, *any* cryptographic assumption, if the adversary can obtain an encryption of $g(k)$ for an *arbitrary* g , as long as the reduction’s proof of security treats both the adversary and the function g as black boxes.

Keywords: Key-dependent input, Black-box separations, One-way functions

1 Introduction

A cryptographic primitive is *key-dependent input secure*, or KDI-secure for short, if it remains secure also in case where the input depends on the secret key. In the case of encryption schemes, KDI-security means that the adversary can obtain, in addition to the usual queries, encryptions of $h(k)$, where k is the key of the scheme and h is chosen by the adversary from some (hopefully large) family of functions \mathcal{H} .

On a first look it might seem that by using the right design it is possible to prevent any “KDI-attacks” on the encryption scheme, and thus achieving such strong security would be only of pure theoretical interest. It turns out, however, that such attacks might “naturally” arise when considering complex systems.

* All omitted proofs can be found in [HH08].

** E-mail: iftach@microsoft.com. This work was performed while at Weizmann Institute of Science and at Microsoft Research, Silicon Valley Campus.

*** E-mail: tholenst@princeton.edu. This work was performed while at Microsoft Research, Silicon Valley Campus.

For instance, in the BitLocker disk encryption utility (used in Windows Vista), the disk encryption key can end up being stored on the disk and thus encrypted along with the disk contents. For more details on the importance of KDI-security, see [BHHO08] and references within.

In this work we study the possibility of obtaining such an encryption scheme both from one-way functions and from other hardness assumptions. In particular, we exclude different types of reductions from a KDI-secure encryption scheme to different hardness assumption. Intuitively, a *black-box reduction* of a primitive P to a primitive Q is a construction of P out of Q that ignores the internal structure of the implementation of Q and just uses it as a “subroutine” (i.e., as a black box). The reduction is *fully-black-box* (following [RTV04]) if the proof of security (showing that an adversary that breaks the implementation of P implies an adversary that breaks the implementation of Q) is also black-box (i.e., the internal structure of the adversary that breaks the implementation of P is ignored as well).

Our first result shows that there is no fully-black-box reduction from KDI-secure encryption schemes to one-way permutations, even if the KDI-security is only against the relatively small class of $\text{poly}(n)$ -wise independent hash-functions. When considering reduction from a KDI-secure encryption scheme, it is natural to ask whether the proof of security accesses the challenge function h in a black-box manner as well. Our second result, however, shows that under this restriction essentially no hardness assumption implies a KDI-secure encryption scheme.

1.1 Related Work

KDI security. The development of encryption secure against key-dependent inputs started by the works of Abadi and Rogaway [AR02]. They studied formal security proofs for cryptographic protocols (as described by [DY83]), and showed that these imply security by a reduction, as long as no *key-cycles* exist in the protocol, i.e., there is a partial order \preceq on the keys exists such that a message depending on k_1 would only be encrypted with k_2 if $k_1 \preceq k_2$. Since this is a restriction (even though it may be a very natural one), the community became aware that it would be desirable to create encryption schemes that provide security even in the existence of such key cycles. Consequently, Black, Rogaway, and Shrimpton [BRS02] define the (possibly stronger) notion of KDI-security for symmetric encryption schemes, and show how to obtain this notion in the random-oracle model. In such a scheme, an adversary can obtain encryptions of $h(k)$ under the key k , where h is given as a circuit to an encryption oracle. Such a scheme implies the security of the scheme under cycles as well. Independently of [BRS02], a notion of circular security has been defined by Camenisch and Lysyanskaya [CL01], considering asymmetric encryption schemes as well.

Recently, Halevi and Krawczyk [HK07] generalized the notion of KDI-security to other cryptographic primitives, such as pseudorandom functions. They also coined the name KDI for this sort of security (previously, it was named *key-dependent message security*). Their results in this setting are mainly for the

construction of pseudorandom functions. In addition, [HK07] shows that a *deterministic* encryption scheme cannot be KDI-secure. Independently and concurrently of [HK07], Hofheinz and Unruh [HU08] provided private-key encryption schemes secure under a limited class of KDI-attacks. The main limitation of their work is that the scheme only remains secure as long as $h(k)$ is significantly shorter than the key; also, after every application of the encryption scheme the key is updated. This makes the construction insufficient for the initial motivation of allowing key cycles in cryptographic protocols. Very recently, Boneh et al. [BHHO08] presented a public-key encryption scheme that is KDI-secure (assuming that the DDH assumption holds) against the family of affine transformations over the messages' domain. Their system remains secure also when key-cycles are allowed.

Black-box impossibility results. Impagliazzo and Rudich [IR89] showed that there is no black-box reduction of key-agreement to one-way permutations and additional work in this line followed (cf., [GKM⁺00,Rud88,Sim98]). Kim, Simon and Tetali [KST99] initiated a new line of impossibility results, by providing a lower bound on the *efficiency* any black-box reduction of universal one-way hash functions to one-way permutations, substantial additional work in this line followed (cf., [GGKT05,HHRS07,HK05,Wee07]). Dodis et al. [DOP05] (and also [Hof08]) give a black-box separation of a similar flavor to the one given in Theorem 2, in the sense that it excludes a large family of hardness assumptions.

1.2 Contributions of this Paper

In this paper we give two impossibility results for security proofs of constructions of KDI-secure private-key encryption schemes. However, since every public-key encryption scheme can be also viewed as a private-key scheme (i.e., both parties use the same private/public key), our impossibility results immediately extend to the public-key case. Our first result is a black-box separation of KDI encryption scheme from one-way permutations and from (even enhanced) family of trapdoor permutations.

Theorem 1. *Let (Enc, Dec) be an encryption scheme that is fully-black-box constructed from one-way permutations. Then there exists an efficient family \mathcal{H} of $\text{poly}(n)$ -wise independent hash functions such that the following holds: there exists no black-box reduction from breaking the KDI-security of (Enc, Dec) against \mathcal{H} to inverting one-way permutations. Furthermore, the above holds also with respect to (enhanced) families of trapdoor permutations.*

For our second result, we assume that the challenge function itself under which the scheme should be KDI-secure is treated by the proof of security as a black box. Moreover, the proof of security does not forward an access to the challenge function to a “third party”. We call such a reduction **strongly-black-box**.

Theorem 2 (informal). *There exists no reduction with strongly-black-box proof from the KDI-security of an encryption scheme to the security of “any cryptographic assumption”.*

We stress that the construction of the encryption scheme considered in Theorem 2 can be arbitrary. The formal statement of Theorem 2 is given in Section 4.

1.3 Interpretation of Our Results

So what should we think on the possibility of building KDI-secure encryption scheme given the above negative results? Let us start with Theorem 1 and let's first consider the fully-black-box restrictions. We remark that while quite a few black-box impossibility results of these types are known (see Section 1.1), there is not even a single known example where we have an impossibility result of the type given in Theorem 1, and yet a “non-black-box” reduction was found.³ We also remark that the reductions given in [BHHO08, HK07, HU08] are fully-black-box. The second issue is that we only rule out security against poly-wise independent hash function, where the value for “poly” is determined as a function of the encryption scheme. It seems, however, that in most settings one cannot limit the power of the queries used in the KDI attack (but merely assume that these functions should be efficiently computable). Typically, when designing an encryption scheme, the exact configuration of each of the systems in which the scheme is going to be used is unknown. These configurations, however, determine the challenge functions “used” in the KDI attacks.

In Theorem 2 we consider arbitrary constructions, but require only black-box access to the challenge functions. This additional restriction actually reflects three separate restrictions. The first is that the proof has only input/output access to the challenge function, the second is that the challenge function cannot be assumed to be efficient, and the third is that the reduction “knows” all the queries made to the challenge functions (we force the last restriction by disallowing the reduction to give a third party an handle to the challenge function). While the first two restrictions seem to be a real limitation on the generality of our second result, the third restriction is harmless in most settings. In particular, this is the case where the hardness assumption does not accept (even implicitly) handler to functions. This list includes all the “non-interactive hardness assumptions” such as one-way functions, factoring, DDH etc.⁴

1.4 Our Technique

In the proof of both our results, we are using the same oracle, Breaker, that helps us to break the KDI-security of every encryption scheme. Let (Enc, Dec)

³ The superiority of non-black-box techniques was demonstrated by Barak [Bar01] in the settings of zero-knowledge arguments for NP. In these settings, however, the black-box access is to the, possibly cheating, verifier and not to any underlying primitive.

⁴ The only exception we could think of for a reduction that benefits from passing the handler to the challenge function to a third party, is a reduction from one KDI-secure encryption scheme B to another KDI-secure encryption scheme A . In such a reduction, the security proof of scheme B typically forwards the challenge function to the security proof of scheme A .

be some fixed encryption scheme. On input (h, c) , where h is some *length doubling* function and c is a ciphertext, Breaker considers all possible keys, and returns the first key k for which $\text{Dec}(k, c) = h(k)$, or \perp if no such key exists. It is not hard to see that (Enc, Dec) is not KDI-secure, with respect to h , in the presence of Breaker. Therefore, our impossibility results follow if Breaker does not help to violate the underlying hardness assumption. For this, we need to assume that Breaker is only called with functions h that are chosen uniformly from the respective set of challenge functions. We ensure this by restricting the functions h for which Breaker performs the above computation to one that is randomly chosen (and then give the adversary access to it). Under this restriction, we manage to prove that Breaker does not help in breaking the assumption. Proving this is our main technical contribution (note that Breaker cannot be implemented efficiently) and we prove it differently in each of our separation results.

One-way permutations. Let π be a random permutation and let $(\text{Dec}^\pi, \text{Enc}^\pi)$ be a candidate of a KDI-secure encryption scheme given π as the one-way permutation. We find a polynomial $p(n)$ (which depends on Dec and Enc) and use a family of length-doubling $p(n)$ -wise independent hash-functions as the challenge functions. Imagine now that a call of A to Breaker helps A to invert π . Then, the behavior of Breaker must be very different for a large number of potential preimages of y , as otherwise the call gave roughly no information about the preimage of y . We show, however, that for all but a negligible fraction of the functions h , the behavior of Breaker will be the same for most possible preimages of y , no matter how the ciphertext is chosen.

Arbitrary assumptions. In this we use the family of *all* length-doubling functions as the challenge functions. The idea is that for a random h in the family, all calls to Breaker (done outside of the KDI game) are very likely to be answered with \perp . The reason is that for a fixed $k \in \{0, 1\}^t$, the probability that $\text{Dec}(k, c) = h(k)$ is roughly 2^{-2t} . This somewhat naive intuition is actually false, as it can fail in the following way: A picks a key k' itself, queries $h(k')$, encrypt this with k' itself, and gives the resulting ciphertext to Breaker. We prove, however, that this is essentially the only way in which the above intuition fails. Thus, instead of calling Breaker, A can as well check the keys on which h was previously queried, which can be done efficiently. We conclude that if there is a reduction with strongly-black-box proof of security from a KDI encryption scheme to a given hardness assumption, then the hardness assumption is false.

2 Preliminaries

2.1 Notation

We denote the concatenation of two strings x and y by $x \circ y$. If X is a random variable taking values in a finite set \mathcal{U} , then we write $x \leftarrow \mathcal{U}$ to indicate that x is selected according to the uniform distribution over \mathcal{U} . We often use probabilities where we choose an oracle π from some uncountable set of oracles at random. It

is possible to defined these using Lebesgue measure and an appropriate mapping of oracles to $[0, 1)$.

2.2 Many-wise Independence

We use standard facts on s -independence, as well as the following upper bound on the probability that many s -wise independent events occur, where each event has low probability.⁵ The proof is omitted in this version.

Lemma 1. *For $s, V \in \mathbb{N}$ let B_1, \dots, B_V be s -wise independent Bernoulli random variables with $\Pr[B_i = 1] \leq \frac{1}{V}$. If $\alpha > s$, then $\Pr[\sum_{i=1}^V B_j \geq \alpha] < \frac{\log(V)}{\alpha^{s-1}}$.*

2.3 Encryption Schemes and KDI Security

We define a (private-key) encryption scheme as a pair of an encryption and a decryption algorithm (Enc, Dec) . On security parameter n , the encryption algorithm Enc gets as input a key of length $t(n)$ and a message of length m , and outputs a ciphertext of length $\ell(n, m)$. The decryption algorithm Dec , gets a key and a ciphertext and outputs the message.⁶ Informally, an encryption scheme (Enc, Dec) is KDI-secure against a family of functions $\mathcal{H} \subseteq \{h : \{0, 1\}^t \rightarrow \{0, 1\}^*\}$, if no efficient adversary can distinguish between an oracle that correctly returns an encryption of $h(k)$, given input h , and one that returns an encryption of the all zero string of the same length. Note that if \mathcal{H} contains functions that map to constants, plain-text queries can be obtained as well.

Definition 1 (KDI-security). *Given an encryption scheme (Enc, Dec) and a key of length t , let $Q^{\text{Enc}, k}$ [resp. $\tilde{Q}^{\text{Enc}, k}$] be an algorithm that gets as input a function $h : \{0, 1\}^t \mapsto \{0, 1\}^{m(t)}$, and returns $\text{Enc}(k, h(k))$ [resp. $\text{Enc}(k, 0^{m(t)})$] (if the schemes is randomized, it returns a random encryption). We say that (Enc, Dec) is KDI-secure for a class of functions \mathcal{H} , if*

$$\left| \Pr_{k \leftarrow \{0, 1\}^{t(n)}} [A^{Q^{\text{Enc}, k}}(1^n) = 1] - \Pr_{k \leftarrow \{0, 1\}^{t(n)}} [A^{\tilde{Q}^{\text{Enc}, k}}(1^n) = 1] \right|$$

is negligible for every efficient algorithm A that only queries functions in \mathcal{H} .

2.4 Cryptographic Games

For reductions that treat the family \mathcal{H} of query-functions as black-box, we are able to prove a very strong impossibility result. In this case, we show that essentially no cryptographic assumption is sufficient to guarantee the KDI-security of

⁵ The usual bounds seem not strong enough in our setting, as they focus on the range where the probability of a single event is constant. Here, the probability of a single event decreases as the number of events increases, and we use a different bound.

⁶ In some definitions, a private-key encryption scheme also includes a key-generation algorithm “Gen”. We omit this since we are not concerned by polynomial factors, and in this case one can simply take the random-coins used by Gen as the private key.

the scheme. In order to do this, we first define the set of cryptographic assumptions we consider.⁷ For the sake of readability, however, we will not try to be as general as possible here. Yet, as far as we can see our definition still captures all natural hardness assumptions.

Definition 2 (cryptographic games). *A cryptographic game is a (possibly inefficient) random system Γ that on security parameter n interacts with an attacker A and may output a special symbol win. In case $\Gamma(1^n)$ outputs this symbol in an interaction with $A(1^n)$, we say that $A(1^n) \leftrightarrow \Gamma(1^n)$ wins. The game is secure if $\Pr[A(1^n) \leftrightarrow \Gamma(1^n) \text{ wins}]$ is negligible for all PPT A , where the probability is over the randomness of A and Γ .*

Examples: One might define the security of a one-way function f by the following game. On security parameter n , the system Γ selects a random $x \in \{0, 1\}^n$ and sends $y = f(x)$ to the adversary. Γ outputs win if A outputs $x' \in f^{-1}(y)$.

To define the DDH hardness assumption one needs a bit more work.⁸ On security parameter n , the system Γ expects first a sequence of at least n ones, we denote the actual number received by α . The system Γ then sends A a description of an appropriately chosen group $\langle g \rangle$ of order $\Omega(2^n)$ and the generator g , as well as α randomly chosen triples $(g^{x_i}, g^{y_i}, g^{z_i})$, where $z_i = x_i y_i$ or a uniform random element, each with probability $\frac{1}{2}$. The attacker A wins, if the number of instances where he incorrectly predicts whether z_i was chosen independently of x_i and y_i , is at most $\frac{\alpha}{2} - \alpha^{2/3}$. Using [LJK07, Theorem 1], one can now show that winning the above is equivalent to the DDH assumption, we omit the details in this version.

2.5 Black-Box Reductions

A reduction from a primitive P to a primitive Q consists of showing that if there exists an implementation C of Q , then there exists an implementation M_C of P . This is equivalent to showing that for every adversary that breaks M_C , there exists an adversary that breaks C . Such a reduction is semi-black-box if it ignores the internal structure of Q 's implementation, and it is fully-black-box (using the terminology of [RTV04]) if it also has black-box proof of security. That is, the adversary for breaking Q ignores the internal structure of both Q 's implementation and of the (alleged) adversary breaking P . The following definition expands the above general discussion for the case of a fully-black-box reduction of a KDI-secure encryption scheme from a one-way permutation.

Definition 3 (fully-black-box reduction). *A fully-black-box reduction of a KDI-secure encryption scheme from a one-way permutation consists of polynomial-time oracle-aided algorithms $(\text{Enc}^{(\cdot)}, \text{Dec}^{(\cdot)})$ and a polynomial-time oracle-aided adversary $A_{\text{OWP}}^{(\cdot)}$, such that the following hold:*

⁷ We remark that definitions of similar spirit to the one below were previously used in [DOP05, Hof08].

⁸ The same argument can be applied for many other assumptions, but we refrain from formalizing this in order not to get bogged down in unrelated details.

- If f is a permutation, then $(\text{Enc}^f, \text{Dec}^f)$ is an encryption scheme.
- For any (possibly unbounded) A_{KDI} that breaks the KDI-security of the encryption scheme, $A_{\text{OWP}}^{f, A_{\text{KDI}}}$ inverts the permutation with non-negligible probability.

When considering reductions from a KDI-secure cryptosystem, it is natural to consider whether the proof of security accesses the challenge functions also as a black box. We say that a proof of KDI-security of a cryptosystem is **strongly-black-box**, if it treats the challenge function also as a black-box.

Definition 4 (strongly-black-box reduction). *A reduction from a KDI-secure encryption scheme to a cryptographic game Γ with strongly-black-box proof of security, consists of polynomial-time oracle-aided algorithms (Enc, Dec) and a polynomial-time oracle-aided adversary $A_\Gamma^{(\cdot)}$ such that the following holds:*

- (Enc, Dec) is an encryption scheme.
- For any adversary A_{KDI}^Q that breaks the KDI-security of (Enc, Dec) , the oracle-aided adversary $A_\Gamma^{(A_{\text{KDI}})}$ violates the security of Γ . Additionally, A_Γ treats the challenge functions provided by A_{KDI} as a black box.

The requirement that A_Γ treats the challenge function as black-box, means that A_Γ can only obtain evaluations of it at arbitrary chosen points and the reduction must work for every challenge function (not just efficiently computable ones). In addition, A_Γ does not provide Γ with a description of the function.⁹

2.6 Extending KDI-secure Encryption Schemes

We would like to make sure our impossibility results hold even for encryption schemes that encrypt messages of length one bit. For technical reasons, however, we will actually need to encrypt messages of length $2t$, where t is the key length. We therefore give a straightforward, but slightly tedious transformation that allows us to do that. (In fact, the following transformation does slightly more, in order to make the technical part in Sections 3 and 4 a bit easier.) We omit the proof of it in this version.

Proposition 1. *Let (Enc, Dec) be an encryption scheme for single bit messages. Assume (Enc, Dec) is KDI-secure for a given set $\mathcal{H} \subseteq \{\{0, 1\}^t \rightarrow \{0, 1\}\}$, then there exists an encryption scheme $(\text{Enc}_1, \text{Dec}_1)$ with the following properties:*

- (a) *The key length t_1 of $(\text{Enc}_1, \text{Dec}_1)$ equals the security parameter.*
- (b) *$(\text{Enc}_1, \text{Dec}_1)$ is defined for messages of arbitrary length.*
- (c) *$(\text{Enc}_1, \text{Dec}_1)$ is KDI-secure for $\mathcal{H}_1 := \{h : \{0, 1\}^t \rightarrow \{0, 1\}^* : \forall i, \forall \tau \in \{0, 1\}^{t_1-t} : h_{|i}(x, \tau) \in \mathcal{H}\}$, where $h_{|i}$ is the function that outputs the i 'th bit of the output of h .¹⁰*
- (d)

⁹ Alternatively, given A_Γ 's (partial) view, it is possible to (efficiently) list all the queries done to the challenge function during the execution.

¹⁰ Namely, \mathcal{H}_1 is the set of functions with the property that *every output bit* is described by a function in \mathcal{H} , after some appropriate padding of the input.

$(\text{Enc}_1, \text{Dec}_1)$ has perfect correctness. **(e)** $(\text{Enc}_1, \text{Dec}_1)$ has deterministic decryption. **(f)** If (Enc, Dec) has a strongly-black-box [resp. black-box] proof of security to a cryptographic game Γ , then $(\text{Enc}_1, \text{Dec}_1)$ has a strongly-black-box [resp. black-box] proof of security to Γ .

3 From One-way Permutations

In this section we prove Theorem 1, but we only give the proof for the case of one-way permutations. The proof for (enhanced) family of trapdoor permutations follows immediately using standard techniques (cf., [GT00,HHRS07]). Let $(\text{Enc}^{(\cdot)}, \text{Dec}^{(\cdot)})$ be an encryption scheme with oracle access to a one-way permutation. By Proposition 1, we can assume that the encryption scheme is always correct, has a deterministic decryption algorithm, defined on messages of any polynomial length and has a security parameter t equal to its key length. We let $\ell(t)$ be the length of an encryption of a message of length $2t$. In order to prove Theorem 1, we use the following inefficient algorithm $\text{Breaker}^{f,h}$.

Algorithm 3 $\text{Breaker}^{f,h}$.

Oracles: A function $f : \{0, 1\}^t \times \{0, 1\}^{\ell(t)} \mapsto \{0, 1\}^{2t}$ (defined for every $t \in \mathbb{N}$) and an infinite sequence of functions $h = \{h_t : \{0, 1\}^t \mapsto \{0, 1\}^{2t}\}_{t \in \mathbb{N}}$.

Input: A pair $(t, c) \in \mathbb{N} \times \{0, 1\}^*$.

Operation: Return the smallest $k \in \{0, 1\}^t$ such that $f(k, c) = h_t(k)$, or \perp if no such k exists.

Let $\Pi = \{\Pi_t\}_{t \in \mathbb{N}}$, where Π_t is the set of all possible permutations over $\{0, 1\}^t$, and let $\mathcal{H} = \{\mathcal{H}_t\}_{t \in \mathbb{N}}$, where h_t is a family of $(\ell(t) + t)$ -wise independent hash functions from $\{0, 1\}^t$ to $\{0, 1\}^{2t}$ with polynomial description size. We denote by $\pi = \{\pi_t\}_{t \in \mathbb{N}} \leftarrow \Pi$ [resp., $h = \{h_t\}_{t \in \mathbb{N}} \leftarrow \mathcal{H}$] the sequence of functions induced by selecting, for every $t \in \mathbb{N}$, π_t uniformly at random from Π_t [resp., h_t uniformly at random from \mathcal{H}_t]. In this section, we consider an instantiation of Breaker with $f = \text{Dec}^\pi$, where π is chosen at random from Π , and h chosen at random from \mathcal{H} . In Section 3.1, we show how to use $\text{Breaker}^{\text{Dec}^\pi, h}$ for violating the KDI-security of $(\text{Enc}^\pi, \text{Dec}^\pi)$, where in Section 3.2 we show that $\text{Breaker}^{\text{Dec}^\pi, h}$ does not help inverting a random π . We prove Theorem 1 in Section 3.3.

3.1 Breaker Violates the KDI-Security of the Scheme

The following adversary uses $\text{Breaker}^{\text{Dec}^\pi, h}$ for breaking the KDI-security of $(\text{Enc}^\pi, \text{Dec}^\pi)$.

Algorithm 4 Algorithm $A_{\text{KDI}}^{\text{Breaker}^{\text{Dec}^\pi, h}, h}$.

Oracles: An infinite sequence of functions $h = \{h_t : \{0, 1\}^t \mapsto \{0, 1\}^{2t}\}_{t \in \mathbb{N}}$ and $\text{Breaker}^{\text{Dec}^\pi, h}$.

Input: Security parameter t .

Operation:

Step 1: Call $Q(h_t)$ [or $\tilde{Q}(h_t)$] to obtain an encryption c of $h_t(k)$ [or 0^{2t}].

Step 2: Call $\text{Breaker}^{\text{Dec}^\pi, h}(t, c)$ to obtain a candidate key k' or \perp .

Step 3: Output 1 iff Breaker did not return \perp .

Lemma 2. For every value of $\pi \in \Pi$, algorithm $A_{\text{KDI}}^{\text{Breaker}^{\text{Dec}^\pi, h}}$ breaks the KDI-security of the $(\text{Enc}^\pi, \text{Dec}^\pi)$ with probability one over a random choice of $h \in \mathcal{H}$.

Proof. Algorithm A_{KDI} only gives the wrong answer if the oracle is \tilde{Q} and Breaker does not return \perp . Assume now that the oracle is \tilde{Q} . Then, for any fixed k' and k we have $\Pr_{h_t}[h_t(k') = \text{Dec}(k', \text{Enc}(k, 0^{2t}))] = \frac{1}{2^{2t}}$, and using the union bound we have that for a fixed k the probability that Breaker does not return \perp is at most 2^{-t} .¹¹ Using an averaging argument, the probability that h_t is such that something else but \perp is returned with probability higher than $2^{-t/2}$ is at most $2^{-t/2}$.

Since the $h_t \in \mathcal{H}$'s are chosen independently from each other, the probability that there exists $t_0 \in \mathbb{N}$ for which A_{KDI} breaks the scheme for no $t > t_0$ is zero. We conclude that with probability one over the random choice of $h \in \mathcal{H}$, it holds that A_{KDI} breaks the KDI-security of (Enc, Dec) infinitely often.

3.2 Breaker Does not Invert Random Permutations

We prove the following upper bound on the probability that an algorithm with access to Breaker inverts a random permutation. In the following let $\mu_A(n)$ be an upper bound on number of π queries and the length of the maximal π query that A does on input $y \in \{0, 1\}^n$ (either directly or through the calls to $\text{Breaker}^{\text{Dec}^\pi, h}$), and let $\mu_{\text{Dec}}(t)$ the same bound with respect to the π queries that Dec does on input $(k, c) \in \{0, 1\}^t \times \{0, 1\}^{\ell(t)}$. We assume without loss of generality that both upper bounds are monotonically increasing, that $\mu_{\text{Dec}}(t) \geq t + \ell(t)$ and that $\mu_A(n) \geq n$. We also assume that $\mu_{\text{Dec}}(t) < 2^t$.

Lemma 3. Let A be an adversary that gets h as an auxiliary input¹² and has oracle access to π and $\text{Breaker}^{\text{Dec}^\pi, h}$. Then for every $y \in \{0, 1\}^n$ it holds that

$$\Pr_{\pi \leftarrow \Pi, h \leftarrow \mathcal{H}}[A_h^{(\pi, \text{Breaker})}(y) = \pi^{-1}(y)] < 3\mu_A(n)(2^{-\mu_{\text{Dec}}^{-1}(n)} + \mu_{\text{Dec}}(\mu_A(n))^2 2^{-n}),$$

where $\mu_{\text{Dec}}^{-1}(n) := \min \{t \in \mathbb{N} : \mu_{\text{Dec}}(t) \geq n\}$.

Applying the Borel-Cantelli lemma on the above we get the following corollary.

Corollary 1. Assume that A and Dec are polynomially bounded, then there exists a negligible function ε such that with probability one over the choice of π and h , $\Pr_{y \leftarrow \{0, 1\}^n}[A_h^{(\pi, \text{Breaker}^{\text{Dec}^\pi, h})}(y) = \pi^{-1}(y)] < \varepsilon(n)$ for large enough n .

¹¹ For this lemma, we are only using the “one-wise” independence of h .

¹² We handle the fact that h is an infinite object, by only providing A the (description of the) first $q(n)$ functions in the sequence, where $q(n)$ is an upper bound on the running-time of $A(y \in \{0, 1\}^n)$.

In Appendix A, we give a proof of a non-uniform version of Lemma 3 (the adversary can use an arbitrary additional non-uniform advice) using the technique introduced by Gennaro and Trevisan [GT00]. Here, we use a different technique that is similar to the one used by Simon [Sim98]. The main idea is to study what happens if π is modified slightly by mapping a second, randomly chosen element to y (the element that A tries to invert). We show that such a change will likely go unnoticed by $A(y)$, and it will not find the new preimage. After the change, however, both preimages of y are equally likely to be the original one, so $A(y)$ could not have found the original one either.¹³

For a given function $g : \{0, 1\}^n \mapsto \{0, 1\}^n$ and two strings $x^*, y \in \{0, 1\}^n$, we define the function $g_{|x^* \rightarrow y}$ as $g_{|x^* \rightarrow y}(x) := \begin{cases} y & \text{if } x = x^*, \\ g(x) & \text{otherwise.} \end{cases}$ We assume that

all calls to $\text{Dec}^{\pi_{|x^* \rightarrow y}}$ are well defined. In particular, if Dec queries $\pi_{|x^* \rightarrow y}$ both at position $\pi^{-1}(y)$ and at $x^* \neq \pi^{-1}(y)$ (and thus might act arbitrarily as it “notices” that $\pi_{|x^* \rightarrow y}$ is not a permutation), we assume it stops and outputs 0. We now wish to consider the elements $\{x^* \in \{0, 1\}^n\}$ for which $\text{Breaker}^{\text{Dec}^\pi, h}(t, c) \neq \text{Breaker}^{\pi_{|x^* \rightarrow y}, h}(t, c)$. The set $\text{Diff}^\pi(t, c, h, y)$ is a (possibly proper) superset of this set.

Definition 5 (Diff). *For an oracle function Dec , an infinite sequence of functions $h \in \mathcal{H}$, $t \in \mathbb{N}$, $c \in \{0, 1\}^*$ and $y \in \{0, 1\}^n$, we let $\text{Diff}^\pi(t, c, h, y) := \{x^* \in \{0, 1\}^n \mid \exists k \in \{0, 1\}^t : (\text{Dec}^\pi(k, c) \neq h_t(k) = \text{Dec}^{\pi_{|x^* \rightarrow y}}(k, c)) \vee (\text{Dec}^\pi(k, c) = h_t(k) \neq \text{Dec}^{\pi_{|x^* \rightarrow y}}(k, c))\}$.*

For $x^* \notin \text{Diff}^\pi(t, c, h, y)$, it holds that $\text{Breaker}^{\text{Dec}^\pi, h}(t, c) = \text{Breaker}^{\pi_{|x^* \rightarrow y}, h}(t, c)$. To see this, let $k_0 \neq \perp$ be the lexicographic smaller output of the two calls. Clearly, k_0 must be the output of both calls to Breaker . The next claim states that if h is uniformly chosen from \mathcal{H} , then $\text{Diff}^\pi(t, c, h, y)$ is very likely to be small for all possible c .

Claim. Let A be an adversary with oracle access to π and $\text{Breaker}^{\text{Dec}^\pi, h}$, which gets h as an auxiliary input. Then, for every π and $y \in \{0, 1\}^n$:

$$\Pr_{h \leftarrow \mathcal{H}} \left[A_h^{(\pi, \text{Breaker}^{\text{Dec}^\pi, h})}(y) \text{ queries } \text{Breaker}^{\text{Dec}^\pi, h}(t, c) \text{ with } |\text{Diff}^\pi(t, c, h, y)| \geq \mu_{\text{Dec}}(\mu_A(n))^2 \right] < \mu_A(n) 2^{-\mu_{\text{Dec}}^{-1}(n)}$$

Proof. For $t \in \mathbb{N}$, $c \in \{0, 1\}^*$ and $k \in \{0, 1\}^t$, let $\mathcal{D}_{c, k}$ be the set of all possible images of $\text{Dec}^{\pi_{|x^* \rightarrow y}}(k, c)$, enumerating over all $x^* \in \{0, 1\}^n$ (i.e., the set $\mathcal{D}_{c, k} := \{\text{Dec}^{\pi_{|x^* \rightarrow y}}(k, c) : x^* \in \{0, 1\}^n\}$).¹⁴ We first note that $|\mathcal{D}_{c, k}| \leq \mu_{\text{Dec}}(t) + 1 \leq 2^t$ – in an execution of $\text{Dec}^\pi(k, c)$ at most $\mu_{\text{Dec}}(t)$ elements $x_1, \dots, x_{\mu_{\text{Dec}}(t)}$ are

¹³ The main difference between our approach and the one in [Sim98], is that we do not insist on keeping π a permutation. It turns out that this slackness makes our proof significantly simpler.

¹⁴ Note that the original image $\text{Dec}^\pi(k, c)$ is in $\mathcal{D}_{c, k}$.

queried, and only if $x^* \in \{x_1, \dots, x_{\mu_{\text{Dec}}(t)}\}$ the image of Dec can be changed. Let \mathcal{H}_t be the t 'th entry in \mathcal{H} . Applying Lemma 1 with $V = 2^t$, $s = t + \ell(t)$, $\alpha = \mu_{\text{Dec}}(t)$ and letting $B_k = 1$ iff $h_t(k) \in \mathcal{D}_{c,k}$, we have that $\Pr_{h_t \leftarrow \mathcal{H}_t} [|\{k \in \{0, 1\}^t : h_t(k) \in \mathcal{D}_{c,k}\}| \geq \mu_{\text{Dec}}(t)] < \frac{t}{\mu_{\text{Dec}}(t)^{t+\ell(t)-1}} \leq 2^{-t-\ell(t)}$.

We next show that $|\text{Diff}^\pi(t, c, h, y)| \leq \mu_{\text{Dec}}(t) \cdot |\{k \in \{0, 1\}^t : h_t(k) \in \mathcal{D}_{c,k}\}|$. We prove this by presenting an injective function ϕ from $\text{Diff}^\pi(t, c, h, y)$ to $\{1, \dots, \mu_{\text{Dec}}(t)\} \times \{k : h_t(k) \in \mathcal{D}_{c,k}\}$. If $x^* \in \text{Diff}^\pi(t, c, h, y)$, then there exists k_{x^*} such that $\text{Dec}^\pi(k_{x^*}, c) \neq h_t(k_{x^*}) = \text{Dec}^{\pi|_{x^* \rightarrow y}}(k_{x^*}, c)$ or $\text{Dec}^\pi(k_{x^*}, c) = h_t(k_{x^*}) \neq \text{Dec}^{\pi|_{x^* \rightarrow y}}(k_{x^*}, c)$ and therefore $k_{x^*} \in \{k : h_t(k) \in \mathcal{D}_{c,k}\}$. Furthermore, $\text{Dec}^\pi(k_{x^*}, c)$ must query π on x^* . Let i_{x^*} denote the index (i.e., position) of the query $\pi(x^*)$ among the π queries that $\text{Dec}^\pi(k_{x^*}, c)$ does, and let ϕ be the function that maps x^* to (i_{x^*}, k_{x^*}) . Since a pair (i, k) specifies a single x (the one queried at i 'th position in $\text{Dec}^\pi(k, c)$), it follows that ϕ is indeed injective.

We note that for $t > \mu_{\text{Dec}}^{-1}(n)$, it always holds that $x^* \notin \text{Diff}^\pi(t, c, h, y)$ (Dec cannot invoke π on such a long input). Combining the above, we have that with probability at least $1 - \mu_{\text{A}}(n)2^{-\mu_{\text{Dec}}^{-1}(n)}$ over the choice of h , for each of the at most $\mu_{\text{A}}(n)$ queries Breaker(t, c) that A(y) does, it holds that $|\text{Diff}(t, c, h, y)| \leq \mu_{\text{Dec}}(t)^2 \leq \mu_{\text{Dec}}(\mu_{\text{A}}(n))^2$.

For a given value of π and $x_0, x_1 \in \{0, 1\}^n$, let $\pi_{|x_0 \leftrightarrow x_1} := \pi_{|x_0 \rightarrow \pi(x_1), x_1 \rightarrow \pi(x_0)}$. We next show that with high probability, A(y) behaves exactly the same given the oracle π or $\pi_{|x^* \leftrightarrow y}$.

Definition 6 (trace). For a given oracle function Dec, the trace, $\text{tr}(\pi, h, y, r_{\text{A}})$, of an adversary A is the sequence of all queries A(y) makes to Breaker $^{\text{Dec}^\pi, h}$ and π (and their responses), when it uses r_{A} as its random-coins and gets h as an auxiliary input.

Claim. Let $y \in \{0, 1\}^n$ and let A be an adversary with oracle access to π and Breaker $^{\text{Dec}^\pi, h}$, and assume that A queries π on its output before returning it. Then, $\Pr_{\pi, h, r_{\text{A}}, x^* \leftarrow \{0, 1\}^n} [\text{tr}(\pi, h, y, r_{\text{A}}) \neq \text{tr}(\pi_{|x^* \leftrightarrow \pi^{-1}(y)}, h, y, r_{\text{A}})] < p(n)$, where $p(n) := 2\mu_{\text{A}}(n)(2^{-\mu_{\text{Dec}}^{-1}(n)} + \mu_{\text{Dec}}(\mu_{\text{A}}(n))^2 2^{-n})$.

Proof. Let \mathcal{X}_π be all the queries to π in $\text{tr}(\pi, h, y, r_{\text{A}})$, clearly $|\mathcal{X}_\pi| \leq \mu_{\text{A}}(n)$. Let further $\mathcal{X}_{\text{Diff}}$ be the union of the sets $\text{Diff}^\pi(t, c, h, y)$ for all calls (t, c) made by A to Breaker. If $x^* \notin \mathcal{X}_\pi \cup \mathcal{X}_{\text{Diff}}$, we have that $\text{tr}(\pi, h, y, r_{\text{A}}) = \text{tr}(\pi_{|x^* \rightarrow y}, h, y, r_{\text{A}})$ (cf., the remark after Definition 5). Claim 3.2 yields that with probability at least $1 - \mu_{\text{A}}(n)2^{-\mu_{\text{Dec}}^{-1}(n)}$ over the choice of h and r_{A} , it holds that $|\mathcal{X}_\pi \cup \mathcal{X}_{\text{Diff}}| \leq \mu_{\text{A}}(n)\mu_{\text{Dec}}(\mu_{\text{A}}(n))^2$. The union bound yields that $\Pr_{h, r_{\text{A}}, x^* \leftarrow \{0, 1\}^n} [\text{tr}(\pi, h, y, r_{\text{A}}) \neq \text{tr}(\pi_{|x^* \rightarrow y}, h, y, r_{\text{A}})] < \mu_{\text{A}}(n)(2^{-\mu_{\text{Dec}}^{-1}(n)} + \mu_{\text{Dec}}(\mu_{\text{A}}(n))^2 2^{-n})$.

Finally, if $\text{tr}(\pi, h, y, r_{\text{A}})$ and $\text{tr}(\pi_{|x^* \leftrightarrow \pi^{-1}(y)}, h, y, r_{\text{A}})$ are different, then one of $\text{tr}(\pi, h, y, r_{\text{A}}) \neq \text{tr}(\pi_{|x^* \rightarrow y}, h, y, r_{\text{A}})$ or $\text{tr}(\pi_{|x^* \rightarrow y}, h, y, r_{\text{A}}) \neq \text{tr}(\pi_{|x^* \leftrightarrow \pi^{-1}(y)}, h, y, r_{\text{A}})$ must hold. Since $\pi_{|x \leftrightarrow \pi^{-1}(y)}$ is also a permutation, and x^* is a uniformly chosen element given $\pi_{|x^* \leftrightarrow \pi^{-1}(y)}$ and y , the inequality obtained before states that both these events have probability at most $p(n)/2$.

Proof. (of Lemma 3) Assuming without loss of generality that A queries its output, the probability that $A^{(\pi, \text{Breaker}^{\text{Dec}^\pi, h})}(y)_h = \pi^{-1}(y)$ is at most the probability that the traces $\text{tr}(\pi, h, y, r_A)$ and $\text{tr}(\pi|_{x^* \leftrightarrow \pi^{-1}(y)}, h, y, r_A)$ are different plus 2^{-n} (to handle the case $x^* = \pi^{-1}(y)$). By Claim 3.2, the latter probability is at most $2^{-n} + p(n)$.

3.3 Putting it Together

Proof. (of Theorem 1) Assume that there exists a black-box proof of security from breaking the KDI-security of $(\text{Enc}^\pi, \text{Dec}^\pi)$ using a poly-wise independent hash function to breaking the hardness of π , and let $M^{(\cdot)}$ be the algorithm for inverting π as guaranteed by this proof of security. Lemma 2 yields that $A_{\text{KDI}}^{\text{Breaker}^{\text{Dec}^\pi, h}, h}$ breaks the KDI-security of $(\text{Enc}^\pi, \text{Dec}^\pi)$ with probability one over the choice of h . Thus, $M_{\text{KDI}}^{\text{Breaker}^{\text{Dec}^\pi, h}, h}$ needs to break the one-way property of π with probability one over the choice of h as well. However, since $A^{\pi, B}$ can be efficiently emulated by an algorithm \tilde{A} with oracle access to $\text{Breaker}^{\text{Dec}^\pi, h}$ and π , and given h as an auxiliary input, Corollary 1 yields that with probability one over the choice of π and h , algorithm $A^{\pi, B}$ does not break the one-wayness of π , and a contradiction is derived.

4 From Arbitrary Assumptions

In this section, we rule out the existence of reductions with strongly-black-box proof of security from the KDI-security of an encryption scheme, to a very large class of hardness assumptions. That is, we prove the following theorem.

Theorem 5 (formal restatement of Theorem 2). *There exists no reduction with strongly-black-box proof of security from the KDI-security of an encryption scheme to any cryptographic game.*

Let (Enc, Dec) be an encryption scheme. As in Section 3, we use Proposition 1 and assume without loss of generality that the encryption scheme is always correct, has deterministic decryption algorithm, is defined on messages of any length, and has security parameter t equal to the key length. We let $\ell(t)$ be the length of an encryption of a message of length $2t$.

Consider an instantiation of Breaker (Algorithm 3) with $f = \text{Dec}$ and $h \in \mathcal{H} = \{\mathcal{H}_t\}_{t \in \mathbb{N}}$, where \mathcal{H}_t is the set of all possible function from $\{0, 1\}^t$ to $\{0, 1\}^{2t}$. As in Section 3, we have that there exists an efficient algorithm, with oracle access to $\text{Breaker}^{\text{Dec}, h}$ and h , that breaks the KDI-security of (Enc, Dec) with probability one over the choice of $h \in \mathcal{H}$. The following Lemma states that in many settings, having oracle access to $\text{Breaker}^{\text{Dec}, h}$ does not yield any significant additional power.

Lemma 4. *Let $A^{\text{Breaker}^{\text{Dec}, h}, h}$ be an algorithm with oracle access to $\text{Breaker}^{\text{Dec}, h}$ and h , and let $t_A(n)$, for security parameter n , be a polynomial-time computable upper bound on the running-time of $A^{\text{Breaker}^{\text{Dec}, h}, h}$.*

Then for every polynomial computable function $\delta : \mathbb{N} \mapsto [0, 1]$, there exists an algorithm \tilde{A}_δ^h , which has oracle access only to h , runs in time $\text{poly}(1/\delta(n), t_A(n), n)$ and uses random-coins of the same length as $A^{\text{Breaker}^{\text{Dec},h,h}}$ such that the following holds. If $A^{\text{Breaker}^{\text{Dec},h,h}}$ and \tilde{A}_δ^h are using the same random-coins, then $A^{\text{Breaker}^{\text{Dec},h,h}}(1^n) = \tilde{A}_\delta^h(1^n)$ with probability $1 - \delta(n)$ over a random choice of h .

Proof. Algorithm \tilde{A} emulates A , while remembering all query and answer pairs to h . When A queries $\text{Breaker}^{\text{Dec},h}(t, c)$, algorithm \tilde{A} distinguishes two cases:

Case 1: $t < \log(t_A(n)) + \log(1/\delta(n))$. \tilde{A} fully emulates $\text{Breaker}^{\text{Dec},h}$. Namely, \tilde{A} evaluates $h_t(k)$ for all $k \in \{0, 1\}^t$ and returns the first one for which $\text{Dec}(k, c) = h_t(k)$. It returns \perp if no such k exists.

Case 2: $t \geq \log(t_A(n)) + \log(1/\delta(n))$. \tilde{A} checks all the previous queries to h of length t in lexicographic order. If for one of those queries it holds that $\text{Dec}(k, c) = h_t(k)$, it returns k , otherwise it returns \perp .

The bound on the running-time of \tilde{A} is clear, in the following we show \tilde{A} emulates A well. We first note that in Case 1, \tilde{A} always returns the same answer that $\text{Breaker}^{\text{Dec},h}$ would. To handle Case 2, let $k \in \{0, 1\}^t$ and assume that the query $h_t(k)$ was not perviously asked by A . Since h is length doubling, the probability over the choice of h that $\text{Dec}(k, c) = h_t(k)$ is 2^{-2t} . Using a union bound we have that the probability, over the choice of h , that \tilde{A} returns a value different from what $\text{Breaker}^{\text{Dec},h}$ would (i.e., \tilde{A} returns \perp where $\text{Breaker}^{\text{Dec},h}$ finds a consistent key) is at most $2^{-\log(t_A(n)) - \log(1/\delta(n))} = \delta(n)/t_A(n)$. Since there are at most $t_A(n)$ calls to $\text{Breaker}^{\text{Dec},h}$, the probability that in any of those \tilde{A} returns a wrong value is at most $\delta(n)$, which proves the lemma.

Proof. (of Theorem 5) The proof follows the lines of the one of Theorem 1, but we need to work a little harder for proving that having access to h does not give the adversary additional power.¹⁵

Assume that there exists a strongly-black-box proof of security from (Enc, Dec) to a cryptographic game Γ and let $M^{(\cdot)}$ be the algorithm for breaking Γ as guaranteed by this proof of security. It easily follows from the proof of Lemma 2 that also in the setting of this section there is an efficient algorithm $A_{\text{KDI}}^{\text{Breaker}^{\text{Dec},h,h}}$, with oracle access to $\text{Breaker}^{\text{Dec},h}$ and h breaking the KDI-security of (Enc, Dec) with probability one over the choice of h . Thus, $M_{\text{KDI}}^{A_{\text{KDI}}^{\text{Breaker}^{\text{Dec},h,h}}}$ breaks Γ with probability 1 over the choice of h . Namely,

$$\Pr_h \left[\Pr_{r_A, r_\Gamma} [M_{\text{KDI}}^{A_{\text{KDI}}^{\text{Breaker}^{\text{Dec},h,h}}} \leftrightarrow \Gamma(1^n) \text{ wins}] > \frac{1}{p_h(n)} \text{ for infinitely many } n \right] = 1, \quad (1)$$

where r_A and r_Γ denote the random-coins of A and Γ , respectively, and p_h is some polynomial that may depend on h . In the following we first remove the dependence of the polynomial p_h from h . For this let $\varepsilon(n) := \Pr_{h, r_A, r_\Gamma} [M_{\text{KDI}}^{A_{\text{KDI}}^{\text{Breaker}^{\text{Dec},h,h}}}$

¹⁵ One gets this property “for free”, when the underlying hardness assumption is inverting a random permutation.

$\leftrightarrow \Gamma(1^n)$ wins] = $\mathbb{E}_h \left[\Pr_{r_A, r_\Gamma} [M_{\text{KDI}}^{\text{Breaker}^{\text{Dec}, h, h}} \leftrightarrow \Gamma(1^n)$ wins] \right]. We show that ε is non negligible. Using Markov's inequality we get for every $n \in \mathbb{N}$ that $\Pr_h \left[\Pr_{r_A, r_\Gamma} [M_{\text{KDI}}^{\text{Breaker}^{\text{Dec}, h, h}} \leftrightarrow \Gamma(1^n)$ wins] < $n^2 \varepsilon(n)$ \right] > $1 - \frac{1}{n^2}$, and therefore¹⁶ $\Pr_h \left[\Pr_{r_A, r_\Gamma} [M_{\text{KDI}}^{\text{Breaker}^{\text{Dec}, h, h}} \leftrightarrow \Gamma(1^n)$ wins] < $n^2 \varepsilon(n)$ for all $n > 2$ \right] \geq 1 - \sum_{n=2}^{\infty} \frac{1}{n^2} > \frac{1}{3}. Combining this with Equation (1) we get that $\Pr_h \left[\frac{1}{p(n)} < \Pr_{r_A, r_\Gamma} [M_{\text{KDI}}^{\text{Breaker}^{\text{Dec}, h, h}} \leftrightarrow \Gamma(1^n)$ wins] < $n^2 \varepsilon(n)$ for infinitely many n \right] > \frac{1}{3}, which implies that there is a polynomial $p(n)$ such that $\varepsilon(n) > \frac{1}{p(n)}$ infinitely often.

In order to finish the proof, we will now to apply Lemma 4 on $M_{\text{KDI}}^{\text{Breaker}^{\text{Dec}, h, h}}$. Recall that Lemma 4 was proved only in the stand alone settings, where in particular no interaction with a random system is considered. Since the proof of security of (Enc, Dec) is strongly-black-box, we have that Γ does not access, through interaction with $M_{\text{KDI}}^{\text{Breaker}^{\text{Dec}, h, h}}$, the function h . Therefore, Γ 's answers are determined by the output behavior of $M_{\text{KDI}}^{\text{Breaker}^{\text{Dec}, h, h}}$ and the proof of Lemma 4 goes through also in this setting. Hence, Lemma 4 yields, letting $\delta(n) = \frac{1}{2p(n)}$, the existence of an efficient algorithm \widetilde{M}^h with oracle access *only* to h , such that $\Pr_{h, r_A, r_\Gamma} [\widetilde{M}^h \leftrightarrow \Gamma(1^n)$ wins] > $\frac{1}{2p(n)}$ for infinitely many n 's.

Our final step is to emulate \widetilde{M}^h , where rather than accessing h we randomly chooses the answer of each time one is requested (and cache it). The latter emulation breaks the cryptographic assumption with probability at least $\frac{1}{2p(n)}$ for infinitely many n 's and since it is also efficient, it implies that Γ is not secure.

5 Applying Our Technique to Other Primitives

It seems tempting to try and use the above Breaker also to show the impossibility of constructing other KDI-secure primitives. Consider for instance pseudorandom functions or permutations that are supposed to be secure even if the adversary can obtain its value on a function of its secret key. Halevi and Krawczyk [HK07] show that a deterministic construction cannot exist, but give a construction in case the permutation has an additional public parameter (i.e., *salt*) chosen after the challenge function is fixed. Their construction, however, compresses (e.g., maps n bits to $n/2$).

It is indeed possible to generalize our techniques to this case, as long for as the pseudorandom functions are injective for *every key*. In this case, Breaker finds a key k such that $f^{k, r}(h(k)) = c$, where f is the pseudorandom function and r is the random salt. The reason this method fails if the construction compresses (as the one given by Halevi and Krawczyk [HK07]), is that Breaker as defined above does not seem to give useful information about the key anymore, since it is unlikely that the correct key is the lexicographically smallest.

¹⁶ The σ -additivity of the measure implies that the event in the next probability is measurable.

It seems that we also cannot utilize our Breaker for the general case of length increasing (non-injective) pseudorandom functions (or equivalently, for the case that we are allowed to make several KDI queries). Consider the question whether a given pseudorandom function is constant on a negligible fraction of the keys (e.g., on a single key k it holds that $f^{k,r}(\cdot) := 0^\ell$). Deciding whether a given function has this property or not might be infeasible. Yet, using for instance the Breaker of Section 4, we can easily find the right answer: ask the Breaker on $(h, 0^\ell)$, where h is a random hash function, and answer “Yes” is the Breaker finds some consistent key. Thus, in this setting our Breaker gives us an extra power that we cannot emulate.

Acknowledgments

We are very grateful to Oded Goldreich, Jonathan Hoch, Gil Segev, Omer Reingold and Udi Wieder for useful discussions. We thank the anonymous referees for many useful comments.

References

- [AR02] M. Abadi and P. Rogaway. Reconciling two views of cryptography (the computational soundness of formal encryption). *JoC*, 15(2):103–127, 2002.
- [Bar01] B. Barak. How to go beyond the black-box simulation barrier. In *42nd FOCS*, pages 106–115. IEEE Computer Society, 2001.
- [BHHO08] D. Boneh, S. Halevi, M. Hamburg, and R. Ostrovsky. Circular-secure encryption from decision Diffie-Hellman. In *CRYPTO 2008*.
- [BRS02] J. Black, P. Rogaway, and T. Shrimpton. Encryption-scheme security in the presence of key-dependent messages. *Selected Areas in Cryptography 2002*.
- [CL01] J. Camenisch and A. Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. *EUROCRYPT 2001*.
- [CW79] J. L. Carter and M. N. Wegman. Universal classes of hash functions. *JCSS*, 18(2):143–154, 1979.
- [DOP05] Y. Dodis, R. Oliveira, and K. Pietrzak. On the generic insecurity of the full domain hash. In *CRYPTO 2005*.
- [DY83] D. Dolev and A. C. Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2):198–208, 1983.
- [GGKT05] R. Gennaro, Y. Gertner, J. Katz, and L. Trevisan. Bounds on the efficiency of generic cryptographic constructions. *S. J. on Comp*, 35(1):217–246, 2005.
- [GKM⁺00] Y. Gertner, S. Kannan, T. Malkin, O. Reingold, and M. Viswanathan. The relationship between public key encryption and oblivious transfer. *FOCS 2000*.
- [GT00] R. Gennaro and L. Trevisan. Lower bounds on the efficiency of generic cryptographic constructions. In *FOCS 2000*.
- [HHRS07] I. Haitner, J. J. Hoch, O. Reingold, and G. Segev. Finding collisions in interactive protocols – A tight lower bound on the round complexity of statistically-hiding commitments. *FOCS 2007*.
- [HH08] I. Haitner and T. Holenstein. On the (Im)Possibility of Key Dependent Encryption (full version). <http://eprint.iacr.org/2008/164>

- [HK05] O. Horvitz and J. Katz. Bounds on the efficiency of "black-box" commitment schemes. In *ICALP '05*, pages 128–139, 2005.
- [HK07] S. Halevi and H. Krawczyk. Security under key-dependent inputs. In *14th ACM CCS*, 2007.
- [Hof08] D. Hofheinz. Possibility and impossibility results for selective decommitments. Technical Report 2008/168, eprint.iacr.org, April 2008.
- [HU08] D. Hofheinz and D. Unruh. Towards key-dependent message security in the standard model. In *EUROCRYPT 2008*.
- [IJK07] R. Impagliazzo, R. Jaiswal, and V. Kabanets. Chernoff-type direct product theorems. In *CRYPTO 2007*, pages 500–516, 2007.
- [IR89] R. Impagliazzo and S. Rudich. Limits on the provable consequences of one-way permutations. In *STOC 1989*.
- [KST99] J. H. Kim, D. R. Simon, and P. Tetali. Limits on the efficiency of one-way permutation-based hash functions. In *FOCS 1999*.
- [RTV04] O. Reingold, L. Trevisan, and S. P. Vadhan. Notions of reducibility between cryptographic primitives. In *TCC 2004*.
- [Rud88] S. Rudich. *Limits on the Provable Consequences of One-Way Functions*. PhD thesis, U.C. Berkeley, 1988.
- [Sim98] D. Simon. Finding collisions on a one-way street: Can secure hash functions be based on general assumptions? In *EUROCRYPT '98*, 1998.
- [Wee07] H. Wee. One-way permutations, interactive hashing and statistically hiding commitments. In *TCC 2007*.

A Gennaro-Trevisan Style Proof of Lemma 3

In this section we prove an alternative non-uniform version of Lemma 3.

Lemma 5 (non-uniform version of Lemma 3). *Let A be a non-uniform adversary that gets h as an auxiliary input and has oracle access to π and $\text{Breaker}^{\text{Dec}^\pi, h}$. Assume that A and Dec satisfy the bounds $\mu_A(n) = \mu_{\text{Dec}}(n) = n^C$ for some $C \in \mathbb{N}$. For $\varepsilon(n) = 2^{-n^{1/(2C)}}$ we have that $\Pr_{h \leftarrow \mathcal{H}, \pi \leftarrow \Pi, y \leftarrow \{0,1\}^n} [A^{(\pi, \text{Breaker}^{\text{Dec}^\pi, h})}(y, h) = \pi^{-1}(y)] < 3\varepsilon$.*

Our main tool is the following lemma.

Lemma 6. *Let A be a non-uniform adversary that gets h as an auxiliary input, and has oracle access to π and Breaker , and assume $\Pr_y [A^{(\pi, \text{Breaker}^{\text{Dec}^\pi, h})}(y, h) = \pi^{-1}(y) \wedge \overline{\text{Bad}(y)}] > \varepsilon(n)$, where $\text{Bad}(y)$ is the event that $A(y, h)$ makes a query $\text{Breaker}^{\text{Dec}^\pi, h}(t, c)$ for which $|\text{Diff}^\pi(t, c, h, y)| \geq \mu_{\text{Dec}}(\mu_A(n))^2$. Then, π can be described using $\log((2^n - s(n))!) + 2s(n) \log\left(\frac{\varepsilon 2^n}{s(n)}\right) + \mu_A(n)^2 \mu_{\text{Dec}}(\mu_A(n))$ bits, where $s(n) = \varepsilon(n) 2^n / (2\mu_A(n)(\mu_{\text{Dec}}(\mu_A(n)))^2)$.*

We omit the proof of Lemma 5 in this version, it is obtained from Lemma 6.

Proof. (of Lemma 6) We assume w.l.o.g. $\varepsilon(n) > 2\mu_A(n)2^{-n}$, since otherwise the statement is trivial. Our description of π consists of the following parts: the description of a set $S \subseteq \{0,1\}^n$, the description of the image of S under π (which

roughly corresponds to the y 's on which A succeeds in inverting), the description of the permutation that π implies if restricted on $\{0, 1\}^n \setminus S$ (i.e., the elements not in S) and finally the description of $\{h_m \in \mathcal{H} \mid m \leq \mu_A(n)\}$.

The description of S and the image of S both require $\log\left(\binom{2^n}{|S|}\right) \leq |S| \log\left(e \frac{2^n}{|S|}\right)$ bits. The description of the permutation requires at most $\log((2^n - |S|)!)$ bits. To store the functions h_m takes $\mu_A(n)^2 \mu_{\text{Dec}}(\mu_A(n))$ bits, for some appropriate family \mathcal{H} . Thus, in total $\log((2^n - |S|)!) + 2|S| \log\left(e \frac{2^n}{|S|}\right) + \mu_A(n)^2 \mu_{\text{Dec}}(\mu_A(n))$ bits are sufficient. In the following we succeed in making S as big as $|S| = \varepsilon(n)2^n / (2\mu_A(n)(\mu_{\text{Dec}}(\mu_A(n)))^2)$, which implies our description size of π .

Defining the set S . We use the following, inefficient, algorithm to create S : we start by letting $I = \{y \in \{0, 1\}^n : A(y) = \pi^{-1}(y) \wedge \text{Bad}(y)\}$ and iteratively do the following. First, remove the lexicographic smallest element y from I and add $\pi^{-1}(y)$ to S . Next, emulate $A^{(\pi, \text{Breaker}^{\text{Dec}^\pi, h})}(y, h)$ and remove all queries A makes whose answers are in I from I (without putting them into S). In addition, for each query to $\text{Breaker}^{\text{Dec}^\pi, h}(t, c)$ done by $A^{(\pi, \text{Breaker}^{\text{Dec}^\pi, h})}(y, h)$, remove the images $\pi(x)$ from I for all $x^* \in \text{Diff}^\pi(t, c, h, y)$ (note that y itself is not removed from I , as we already removed it). Once the emulation is over, repeat with the next element. Since for every emulation we remove at most $\mu_A(n) + \mu_A(n) \cdot \mu_{\text{Dec}}(\mu_A(n))^2$ elements from I before moving another element to S , we have that $|S| \geq \varepsilon(n)2^n / 2\mu_A(n)(\mu_{\text{Dec}}(\mu_A(n)))^2$.

The reconstruction of π . We now show that we can reconstruct π from the given information. For this, we first reconstruct the oracle outside of S from the given information. Then pick the lexicographic smallest element $y \in S$ whose preimage is not yet known, and emulate $A^{\pi, \text{Breaker}^{\text{Dec}^\pi, h}}(y, h)$. We first consider the queries $\pi(x)$ done by $A(y, h)$. The definition of S yields that we either know the answer for this query, or we are guaranteed that $\pi(x) = y$ (and we can stop the emulation). So it is left to consider the queries $\text{Breaker}^{\text{Dec}^\pi, h}(t, c)$. We note that if $k \in \{0, 1\}^t$ is the value we should return as the answer of $\text{Breaker}^{\text{Dec}^\pi, h}(t, c)$, then the answers to all π -queries made by $\text{Breaker}^{\text{Dec}^\pi, h}(t, c)$ when it calls $\text{Dec}^\pi(k, c)$ (see Algorithm 3) should be known, where the only exception is a query on $\pi^{-1}(y)$ if such occurs. Therefore, we try all candidates x^* (i.e., the elements whose image we don't know at this point) for $\pi^{-1}(y)$, and emulate $\text{Dec}^{\pi|x^* \rightarrow y}(k, c)$. The latter emulation succeeds unless a query is made whose answer we don't know. In this case, we know by the above observation that the current pair (k, x^*) is not the one we are looking for, and we can safely move to the next candidate for x^* . Finally, note that if a successful emulation of $\text{Dec}^{\pi|x^* \rightarrow y}(k, c) = h(k)$ done by $\text{Breaker}^{\text{Dec}^\pi, h}(t, c)$ satisfies $\text{Dec}^{\pi|x^* \rightarrow y}(k, c) = h(k)$, then k must be the correct answer to $\text{Breaker}^{\text{Dec}^\pi, h}(t, c)$. The reason for that x^* would be in $\text{Diff}^\pi(t, c, h, y)$, and therefore cannot be in S . All in all, we can emulate $A(y, h)$'s run correctly and obtain the correct $\pi^{-1}(y)$ as the output of A .