

OAEP Is Secure under Key-Dependent Messages

Michael Backes^{1,2}, Markus Dürmuth¹, and Dominique Unruh¹

¹ Saarland University, Saarbrücken, Germany

{backes,duermuth,unruh}@cs.uni-sb.de

² Max-Planck-Institute for Software Systems, Saarbrücken, Germany

backes@mpi-sws.mpg.de

Abstract. Key-dependent message security, short KDM security, was introduced by Black, Rogaway and Shrimpton to address the case where key cycles occur among encryptions, e.g., a key is encrypted with itself. We extend this definition to include the cases of adaptive corruptions and arbitrary active attacks, called adKDM security incorporating several novel design choices and substantially differing from prior definitions for public-key security. We also show that the OAEP encryption scheme (using a partial-domain one-way function) satisfies the strong notion of adKDM security in the random oracle model. The OAEP construction thus constitutes a suitable candidate for implementing symbolic abstractions of encryption schemes in a computationally sound manner under active adversaries.

Keywords: Key-dependent message security, chosen ciphertext attacks, RSA-OAEP.

1 Introduction

Encryption schemes constitute the oldest and arguably the most important cryptographic primitive. Their security was rigorously studied very early, starting with Shannon's work for the information-theoretic case [31]. Computational definitions for public-key encryption were developed over time, in particular in [23,32,30,19]. For symmetric encryption, the first real definitions were, to the best of our knowledge, given in [19,28,8], using the same basic ideas as in public-key encryption. While these definitions seemed to take care of standard usage of encryption schemes, it was soon recognized that larger protocols might pose additional requirements on the encryption schemes, e.g., in multi-party computations with dynamic corruptions as in [7]. It was also recognized that in some cases, symmetric encryption initially seemed to be the appropriate method to use, but upon study other primitives such pseudorandom permutations [10,8] or authenticated encryption [12,9] proved to be better.

A specific additional requirement some larger protocols pose on encryption schemes is the ability to securely encrypt key-dependent messages. One speaks of key-dependent messages if a key K is used to encrypt a message m where m contains or depends on the key K (or the corresponding secret key in the case of public-key encryption). The first concrete use of this case seems to have been

in [15], where multiple private keys were used to encrypt one another in order to implement an all-or-nothing property in a credential system to discourage people from transferring individual credentials. Such key cycles also occur in implementations of disk encryption in, e.g., Windows Vista, that can store an encryption of its own secret keys to the disk in some situations. Key cycles also occur in some naively designed key exchange protocols of session keys given master keys shared among the two parties or with a key distribution center, where at the end of the protocol the newly exchanged key is “confirmed” by using it to encrypt or authenticate something that might include the master keys.

Another area that has brought additional requirements on cryptographic primitives, and in particular that of encryption with key cycles, is the use of formal methods or “symbolic cryptography”. Here the question is whether simple abstractions of cryptographic primitives exist that can be used by automated proof tools (model checkers or theorem provers) to prove or disprove a wide range of security protocols that use cryptography in a blackbox manner. The original abstractions used by this automation community are term algebras constructed from certain base types and cryptographic operators such as E and D for encryption and decryption. They are often called Dolev-Yao models after the first such abstraction [20]. As soon as one has a multi-user variant of such a model, the keys are terms, and from the term algebra side it is natural that keys can also be encrypted, i.e., most models simply assume that key cycles are allowed. Once cryptographic justification of such models was started in [2], it was recognized that key cycles had to be excluded from the original models to get cryptographic results. The same holds for later results [1,26,6,27,29,4,18,17].

Motivated primarily by symbolic cryptography, a definition of key-dependent message security (*KDM security*) was introduced in [13]. It generalizes the definition from [15] by allowing arbitrary functions of the keys (and not just individual keys) as plaintexts, and by considering symmetric encryption schemes. [13] also presents a definition and a construction (without proof) for the asymmetric case against *passive* attackers. In [5] it was shown that, in the case of symmetric encryption, an extension of the KDM definition that additionally allows for a limited revelation of secret keys of honest users, called DKDM security, is suitable for extending results about the justification of Dolev-Yao models to include protocols with key cycles. Full security in the presence of key-dependent messages has so far only been achieved in the random oracle model. In [24] and [25], the problem of implementing KDM secure symmetric encryption schemes without random oracles is investigated. There, solutions are given for relaxed variants of KDM security, e.g., security against a bounded number of queries or security with respect to a *single* key dependency function. No scheme is known, however, that fulfills any form of full-fledged KDM security (passive or active) without the use of random oracles. In [14], a scheme is presented that is secure if the key dependency functions are guaranteed to be affine. Extensions of KDM security for public-key encryption to active adversaries have not been proposed

yet, and establishing meaningful definitions for this case indeed raises non-trivial problems.

Our Contributions. We first propose a new definition of security under key-dependent messages, called *adKDM security*, that captures security against active attackers and adaptive corruptions in the case of public-key encryption. This definition incorporates several novel design choices and substantially differs from prior definitions for public-key security; in particular, it allows the adversary to iteratively construct nested encryptions without necessarily revealing inner encryptions, and it is required to keep track of the knowledge that the adversary maintains in an ideal setting.

We then investigate the OAEP encryption scheme and prove that it satisfies adKDM security in the random oracle model, assuming the partial-domain one-wayness of the underlying trapdoor-permutation. This in particular shows the OAEP construction to constitute a suitable candidate for soundly implementing symbolic abstractions of cryptography (so-called computational soundness). We leave it as an open problem for future work to prove that our definition of adKDM security is sufficient for a computational soundness result.

The need to incorporate key dependencies and the adaptive nature of adKDM security require substantial changes to the CCA2-security proof of OAEP. In particular, adKDM security does not allow for determining in advance which encryptions will be used as challenge encryptions. At the point of construction of these bitstrings, the adversary might not even know the challenge encryptions. Consequently, performing the reduction to the underlying assumption requires us to lazily construct them in order to decide as late as possible which encryption constitutes a challenge encryption.

2 Preliminaries

In this section, we present some definitions and conventions that will be used later on in the paper.

Notation. Let \oplus denote the XOR operation, and let \parallel denote concatenation. For a probabilistic algorithm B , let $y \leftarrow B(x)$ denote assigning the output of $B(x)$ to y . Let $\Pr[\pi : X]$ denote the probability that π holds after executing the instructions in X (which are of the form $y \leftarrow B(x)$). A function in n is negligible if it is in $n^{-\omega(1)}$. A function is non-negligible if it is not negligible. We formulate all our results for uniform adversaries, but they hold for nonuniform adversaries as well.

Definition 1 (Circuit). *A circuit is a Boolean circuit with $n_1 + \dots + n_t$ input bits ($t \geq 0$) and m output bits. The circuit may have arbitrary fan-in and fan-out, AND-, OR- and NOT-gates, and—in the case of an encryption scheme in the random-oracle model—gates for querying the random oracle(s). We assume that a circuit is always encoded by explicitly specifying all its gates and the numbers n_1, \dots, n_t, m . The evaluation $f(x_1, \dots, x_t)$ of a circuit f on bitstrings x_1, \dots, x_t*

is defined as follows: Let x'_i be the result of truncating or padding x_i with 0^* to the length n_i . Then $f(x_1, \dots, x_t)$ is the result of evaluating f with input $x'_1 || \dots || x'_t$.¹

Convention: Encryption is length-regular. For any encryption scheme, we impose the following assumption on the output of the encryption function Enc and the decryption function Dec: The length of the output of Enc depends only on the public key and the length of the message. The length of the output of Dec depends only on the *public* key and the length of the ciphertext. This can easily be achieved by suitable padding and encoding.

The OAEP scheme. The optimal asymmetric encryption padding (OAEP) scheme [11] constitutes a widely employed encryption scheme in the random oracle model based on a trapdoor 1-1 function.

Definition 2 (OAEP). Let k denote the security parameter and let k_0 and k_1 be functions such that $k_0, k_1, k - k_0 - k_1$ are superlogarithmic. Assume a 1-1 trapdoor function f with domain $\{0, 1\}^k = \{0, 1\}^{k-k_0} \times \{0, 1\}^{k_0}$. Let $G : \{0, 1\}^{k_0} \rightarrow \{0, 1\}^{k-k_0}$ and $H : \{0, 1\}^{k-k_0} \rightarrow \{0, 1\}^{k_0}$ denote random oracles. The public and secret key for the OAEP encryption scheme (Enc, Dec) consists of a public key and a trapdoor for f . An encryption $c = \text{Enc}(pk, m)$ with $|m| = k - k_0 - k_1$ is computed as $r \leftarrow \{0, 1\}^{k_0}$, $s := (m || 0^{k_1}) \oplus G(r)$, $t := r \oplus H(s)$, $c := f_{pk}(s || t)$.

A decryption $\text{Dec}(sk, c)$ is computed as $s || t := f_{sk}^{-1}(c)$, $r := t \oplus H(s)$, $m || z := s \oplus G(r)$ with $|s| = k - k_0$, $|t| = k_0$, $|m| = k - k_0 - k_1$ and $|z| = k_1$. If $z = 0^{k_1}$, the plaintext m is returned, otherwise the decryption fails with output \perp .

It has been shown in [21] that the OAEP scheme is IND-CCA2 secure in the random oracle model under the assumption that f fulfills the following Definition 3 of partial-domain one-wayness. They further showed that the RSA-trapdoor permutation, which is most commonly used for the OAEP scheme, is partial-domain one-way.

Definition 3 (Partial-Domain One-Wayness). A 1-1 function $f : S \times T \rightarrow \text{range } f$ with key generation KeyGen_f is partial-domain one-way if for any polynomial-time adversary A we have that

$$\Pr[s = s' : pk \leftarrow \text{KeyGen}_f, (s, t) \xleftarrow{\$} S \times T, s' \leftarrow A(pk, f_{pk}(s || t))]$$

is negligible in k , where $A, \text{KeyGen}_f, f, S, T$ depend on the the security parameter k . We sometimes call this probability the advantage of A .

3 The Definition of adKDM

We now present our definition of adKDM security. Since this definition incorporates several novel design choices and substantially differs from prior security

¹ Not granting a circuit access to the length of its arguments is not a restriction in our case, since this length will always be known in advance.

definitions for public key security, we do not immediately present the definition. Instead, we start with a direct adaption of an existing definition and show using an example why this adaption is not sufficient. We proceed with several plausible approaches for extending this adaption and explain why they fail. We finally present our definition of adKDM security and explain why it solves the problems observed with the tentative definitions discussed before.

Extending DKDM security. In [5] the security notion DKDM was proposed for the case of symmetric key-dependent encryptions. It is the strongest notion of KDM security considered so far; restating it one-to-one in the public-key setting would yield the following definition:²

Definition 4 (DKDM, public key setting – sketch). *The DKDM oracle maintains a sequence of key pairs pk_i, sk_i and a random challenge bit b . It answers to the following queries:*

- $pk(j)$: Return pk_j .
- $reveal(j)$ where j has not been used in an $enc(j, \cdot)$ query: Return sk_j .
- $enc(j, f)$ where f is a circuit and j has not been used in a $reveal(j)$ query: Compute $m_0 := f(sk_1, sk_2, \dots)$, $m_1 := 0^{|m_0|}$ and encrypt $c := \text{Enc}(pk_j, m_b)$. Return c .
- $dec(j, c)$ where c has not been returned by an $enc(j, \cdot)$ query with the same key index j : Return $\text{Dec}(sk_j, c)$.

A public key encryption scheme (Enc, Dec) is DKDM secure if no polynomial time adversary interacting with the DKDM-oracle guesses b with probability non-negligibly greater than $\frac{1}{2}$.

This definition is an almost immediate generalization of the IND-CCA definition to the multi-session setting (i.e., with several key pairs instead of only one). DKDM extends IND-CCA in two ways: First, the messages that are contained an $enc(\cdot, \cdot)$ encryption query may depend on all secret keys in the system. Second, one can reveal secret keys as long as the corresponding public keys have not been used for encrypting (otherwise one could decrypt a challenge ciphertext so that the definition cannot be met).

Although the notion of DKDM has been shown to be useful for soundness results for a specific class of protocols, it has obvious restrictions on the class of protocols considered. In particular, it is not allowed to reveal a key that has been used for encryption. The following simple protocol illustrates that this indeed constitutes a restriction: Alice holds two secret keys sk_1, sk_2 and a secret message m and sends the following messages to Bob:

$$c_1 := \text{Enc}(pk_1, \text{Enc}(pk_2, m \| sk_1 \| sk_2)), \quad c_2 := \text{Enc}(pk_2, \text{Enc}(pk_1, m \| sk_1 \| sk_2))$$

Then Bob chooses a value $i = 1, 2$ and Alice sends sk_i to Bob. We would intuitively expect the message m to stay secret since Bob learns at most one of the

² We have omitted one condition of their definition, namely that it should not be possible to generate a valid ciphertext without the knowledge of the secret key. This condition is not applicable to the public-key setting.

keys sk_1, sk_2 . However, a direct reduction against DKDM security fails. Namely, we have basically four possibilities to construct the messages c_1, c_2 by querying the DKDM oracle (note that enc denotes the query to the adKDM oracle while Enc is the encryption algorithm):

- (i) $c_1 := enc(1, g_1)$, $c_2 := enc(2, g_2)$ where g_1 and g_2 are circuits computing $Enc(pk_2, m \| sk_1 \| sk_2)$ and $Enc(pk_1, m \| sk_1 \| sk_2)$, respectively (given input (sk_1, sk_2)).
- (ii) $c_1 := Enc(pk_1, enc(2, g))$, $c_2 := Enc(pk_2, enc(1, g))$ where g computes $m \| sk_1 \| sk_2$.
- (iii) $c_1 := Enc(pk_1, enc(2, g))$, $c_2 := enc(2, g_2)$ where g and g_2 are as before.
- (iv) $c_1 := enc(1, g_1)$, $c_2 := Enc(pk_2, enc(1, g))$, where g and g_1 are as before.

Then, depending on the value of i chosen by Bob, we have to issue $reveal(i)$. In cases (i) and (ii), no reveal query is allowed since queries of the forms $enc(1, \cdot)$ and $enc(2, \cdot)$ have been performed which excludes reveal queries $reveal(1)$ and $reveal(2)$ by Definition 4. Similarly, in case (iii) we are not allowed to query $reveal(2)$, and in case (iv) we are not allowed to query $reveal(1)$. Thus in order to perform the first step, we have to know in advance what the value of i will be and to construct c_1, c_2 as in case (iii) or (iv), respectively. Of course, in the present example it is possible to save the reduction proof by guessing i ; however, it is easy to thwart this possibility by performing many such games in parallel.³ A natural approach to extend the definition of DKDM to this case would be to allow to even reveal keys sk_j that are used in encryption queries $enc(j, \cdot)$. However, a query $enc(j, \cdot)$ returns an encryption c of the message m_b . So given the secret key sk_j , we could easily determine m_b from c and therefore the challenge bit b . Therefore, we will have to distinguish between two types of encryption queries: A normal encryption query $enc(j, f)$ will return the encryption of $m_0 := f(sk_1, \dots)$ irrespective of the value of b . A challenge encryption query $challenge(j, f)$ returns m_b where m_0 is as for $enc(j, f)$ and $m_1 := 0^{|m_0|}$. This leads to the following tentative definition:

Definition 5 (KDM security – tentative). *The oracle \mathcal{T} chooses a random bit b and accepts the following queries.*

- $pk(j)$ and $reveal(j)$: Return pk_j and sk_j , respectively. $dec(j, c)$: Return $Dec(sk_j, c)$.
- $enc(j, f(i_1, \dots, i_t))$ where f is a circuit: Compute $m_0 := f(sk_{i_1}, \dots, sk_{i_t})$ and return $Enc(pk_j, m_0)$.
- $challenge(j, f(i_1, \dots, i_t))$: Compute m_0 as before, $m_1 := 0^{|m_0|}$ and return $Enc(pk_j, m_b)$.

³ E.g., Alice sends $m_1^{(1)}, m_2^{(1)}, \dots, m_1^{(n)}, m_2^{(n)}$ with $m_1^{(\mu)} := Enc(pk_1^{(\mu)}, Enc(pk_2^{(\mu)}, m \| keys))$, $m_2^{(\mu)} := Enc(pk_2^{(\mu)}, Enc(pk_1^{(\mu)}, m \| keys))$ and $keys := sk_1^{(1)} \| sk_2^{(1)} \| \dots \| sk_1^{(n)} \| sk_2^{(n)}$. Then Bob chooses $i_1, \dots, i_n \in \{1, 2\}$ and Alice sends $sk_{i_1}^{(1)}, \dots, sk_{i_n}^{(n)}$. The fact that all keys are contained in each encryption also disables hybrid arguments. To the best of our knowledge, the security of this protocol cannot be reduced to DKDM security.

The oracle aborts in the following cases: $reveal(j)$ is queried but $challenge(j, \cdot)$ has been queried before. $challenge(j, \cdot)$ is queried but $reveal(j)$ has been queried. $dec(j, c)$ is queried but c was produced by $challenge(j, \cdot)$. A scheme is KDM secure if no polynomial-time adversary guesses b with probability noticeably larger than $\frac{1}{2}$.

This definition might look appealing, but it cannot be met: For example, one could encrypt a challenge plaintext under pk_1 via the query $challenge(1, m)$, then encrypt the key sk_1 under pk_2 via $c := enc(2, sk_1)$, and finally reveal sk_2 via $reveal(2)$.⁴ This sequence of queries is not forbidden by Definition 5. Now we can compute sk_1 from c using sk_2 and then decrypt the challenge encryption using sk_1 . This allows us to determine the bit b . Hence no encryption scheme can fulfill Definition 5. We hence have to relax the definition by excluding queries that would trivially allow to decrypt a challenge ciphertext. For this, we have to reject queries to the oracle that would allow the adversary to decrypt the challenge even in an ideal setting. For this, we keep track of the keys that the adversary can deduce from the queries made so far. We call this set *know* (the knowledge of the adversary) because it represents what the adversary knows *ideally*. The set *know* is inductively defined as follows: (a) If $reveal(j)$ has been queried, then $j \in know$. (b) If $j \in know$, and a $enc(j, f(i_1, \dots, i_t))$ has been queried, then $i_1, \dots, i_t \in know$. (c) If $enc(j, f(i_1, \dots, i_t))$ has been queried and returned the ciphertext c , and $dec(j, c)$ has subsequently been queried, then $i_1, \dots, i_t \in know$. Roughly, we say that the adversary knows all keys that either were revealed or are contained in ciphertexts it could decrypt using keys it knows. We can now relax Definition 5 by disallowing queries that would allow the adversary to know a secret key for a challenge encryption.

Definition 6 (KDM security – tentative). *KDM security is defined as in Definition 5 except that the oracle \mathcal{T} additionally aborts if a query would lead to the following situation: For some $j \in know$, a query $challenge(j, \cdot)$ has been performed (or is being performed).*

Introducing hidden encryptions. Definition 6, however, is still too weak to allow to adaptively choose which keys to reveal. In particular, the example protocol given above can still not be proven secure: When producing c_1, c_2 in a reduction proof, we have to decide which of the ciphertexts will be created by challenge encryptions ($challenge(\cdot, \cdot)$ queries) and which will be created by normal encryptions ($enc(\cdot, \cdot)$). Since we might have to invoke $reveal(1)$ later, we may not use $challenge(1, \cdot)$ queries, and since we might have to invoke $reveal(2)$, we may not use $challenge(2, \cdot)$ queries. But if no $challenge(\cdot, \cdot)$ query is issued, the oracle \mathcal{T} never uses the bit b and thus the adversary cannot guess b .⁵

Handling adaptive revelations of keys hence requires to further extend our approach. A closer inspection reveals why we failed to prove the security of the example protocol: We had two possible ways to construct the ciphertext

⁴ We use the shorthand m and sk_1 for the circuits outputting m and sk_1 , respectively.

⁵ Again, this problem might be remedied by guessing in advance whether sk_1 or sk_2 will be needed, but see footnote 3 for an example where guessing does not work.

c_1 . Either (a) we could ask the oracle to produce $c'_1 := \text{Enc}(pk_2, m \| sk_1 \| sk_2)$ and encrypt it ourselves using pk_1 to produce c_1 . Or (b) we could request the ciphertext c_1 directly by sending to the oracle a circuit f that computes c'_1 from sk_1, sk_2 . In case (a), we are not allowed to reveal sk_2 since this would allow to decrypt c'_1 and thus reveal m . In case (b), if we were to reveal sk_1 this would allow to decrypt c_1 . As the plaintext c'_1 for c_1 has been produced using a circuit f from sk_1, sk_2 and m , the oracle has no way of knowing that c'_1 is actually an encryption of these values (this would require an analysis of the circuit to determine what it does) and thus has to consider the values sk_1, sk_2 and m to be leaked when c_1 is decrypted. Thus in case (b), we have to disallow the revelation of sk_1 . This analysis shows that we need a way to send the following instructions to the oracle: “First produce the ciphertext c'_1 as an encryption of $m \| sk_1 \| sk_2$ (where $m \| sk_1 \| sk_2$ is described by a suitable circuit). Do *not* return the value c'_1 (as otherwise we would be in case (a)). Then produce the ciphertext c_1 by encrypting c'_1 . Return c_1 .”

Given these instructions, the oracle has enough information to deduce that when revealing sk_1 , the message m is still protected by the encryption c'_1 using pk_2 (the details of this deduction process are discussed below). And if only sk_2 is revealed instead, c_1 cannot be a decryption and m is protected. Analogous reasoning applies to the construction of c_2 .

Hence we have to define an oracle \mathcal{T} that allows us to construct ciphertexts without revealing them. Instead, for each ciphertext we can adaptively decide whether to reveal it or whether we only use it inside other ciphertexts (that again may or may not be revealed). More concretely, whenever a query is issued to \mathcal{T} , instead of directly returning the result of that query, it is stored in some register $bits_h$ inside the oracle where h is a handle identifying the register. Only upon a special reveal query, the value $bits_h$ is returned to the adversary. A challenge encryption (i.e., one whose content depends on the challenge bit b) is then produced as follows: First produce a plaintext m (possibly using a circuit and depending on other hidden strings) and assign it to register $bits_{h_1}$. Then, depending on b , assign $bits_{h_1}$ or $0^{|bits_{h_1}|}$, respectively, to register $bits_{h_2}$ (using a special challenge query $h_2 \leftarrow C(h_1)$). Encrypt $bits_{h_2}$ using some key and assign the result to $bits_{h_3}$. Finally (optionally) reveal $bits_{h_3}$.⁶

These considerations lead to the following definition of the adKDM oracle (however, for the definition of adKDM security we will additionally define which sequences of queries are allowed):

Definition 7 (adKDM Oracle). *The adKDM oracle \mathcal{T} maintains two partial functions cmd and $bits$ (to increase readability we write $bits_h$ for $bits(h)$ and*

⁶ This is, of course, not the only possible way to model challenge encryptions. One could, e.g., use a special command for producing a challenge encryption. However, we believe that the approach of being able to make challenge values out of arbitrary messages allows for more direct reductions in proofs. E.g., in our example protocol we could directly model the fact that m is the value that should remain hidden by using oracle call $h' \leftarrow C(h)$ when $bits_h$ contains m and then using $bits_{h'}$ instead of $bits_h$ in subsequent encryptions.

cmd_h for $cmd(h)$), a set Φ , a sequence of secret/public key pairs sk_i, pk_i ($i \in \mathbb{N}$) (which are generated when first accessed), and a bit b (the challenge bit). The function cmd will store the structure of previous queries, the function $bits$ will store the corresponding bitstrings, and Φ will keep track of query results that are revealed to the adversary. We will refer to the elements in the domain of cmd and $bits$ as handles in the following. Upon the first activation, b is chosen uniformly from $\{0, 1\}$, $bits$ and cmd are initially undefined, and Φ is empty. The oracle responds to the following commands:

- Encryption: $h' \leftarrow E(j, h)$ where $cmd_{h'}$ has not been assigned, cmd_h has been assigned, and j is a key index: Set $bits_{h'} := \text{Enc}(pk_j, bits_h)$ and $cmd_{h'} := E(j, h)$.
- Decryption: $h' \leftarrow D(j, h)$ where $cmd_{h'}$ has not been assigned, cmd_h has been assigned, and j is a key index: Set $bits_{h'} := \text{Dec}(pk_j, bits_h)$, and $cmd_{h'} := D(j, h)$.
- Circuit evaluation: $h' \leftarrow F(f, h_1, \dots, h_t)$ where $cmd_{h'}$ has not been assigned, cmd_{h_i} has been assigned for all i , and f is a circuit with t arguments: Set $bits_{h'} := f(bits_{h_1}, \dots, bits_{h_t})$ and set $cmd_{h'} := F(f, h_1, \dots, h_t)$.
- Key request: $h' \leftarrow K(j)$ where $cmd_{h'}$ has not been assigned and j is a key index: Set $cmd_{h'} := K(j)$ and $bits_{h'} := sk_j$.
- Challenge: $h' \leftarrow C(h)$ where $cmd_{h'}$ has not been assigned and cmd_h has been assigned: Set $cmd_{h'} := C(h)$. If $b = 1$, set $bits_{h'} := bits_h$, otherwise set $bits_{h'} := 0^{|bits_h|}$.
- Reveal: $reveal(h)$ where cmd_h has been assigned: Add h to Φ and return $bits_h$.
- Public key request: $pk(j)$ where j is a key index: Return pk_j .

The above commands in particular allow to assign a constant c to a handle h' by issuing $h' \leftarrow F(f)$ where f is a nullary circuit that returns c . We abbreviate this as $h' \leftarrow F(c)$. Note that the length of every bitstring is always known to the adversary, because Enc , Dec , and all f are length-regular.

The knowledge of the adversary. If \mathcal{T} can be accessed in arbitrary ways, it is easy to determine b , e.g., querying $h_1 \leftarrow F(1)$, $h_2 \leftarrow C(h_1)$, $reveal(h_2)$ will return b . Thus we have to restrict the adversary to queries that will not trivially allow to deduce b . The necessary criteria are given below. In analogy to Definition 6 we do this by deriving a set *know* that characterizes what the adversary would ideally be able to know after the queries it performed. In contrast to Definition 6 the set *know* does not only contain keys, but the handles of all values produced by the oracle that the adversary would be able to know in an ideal setting. Intuitively, the knowledge *know* is defined by the following rules: All handles that the adversary requested (the set Φ) are considered known. If the decryption of a message is known, then that message is considered known.⁷ If a circuit evaluation is known, all its arguments are considered known. If a challenge is

⁷ It may seem surprising that by learning the result of a decryption we may learn something about the ciphertext. However, in fact we can get a single bit about the ciphertext, namely whether it is valid or not. Combining this with the application of circuits, we can in principle retrieve the full ciphertext.

known, the underlying message is considered known. If a key is known and an encryption of some message under that key is known, the message is considered known. And finally, if a decryption of some handle h_1 is known, and some handle h_2 evaluates to the same bitstring as h_1 , and that handle h_2 resulted from an encryption of some message m , then that message m is considered known.

The last rule merits some additional explanation: The adversary may, e.g., construct and reveal an encryption c (assigned to some handle h_2) of some m . Then it constructs a circuit f that evaluates to c (by hard-coding c into f) and assigns $h_1 \leftarrow F(f)$. Now h_1 and h_2 refer to the same bitstring. By revealing the decryption of h_1 , the adversary will then learn m . So after this sequence of queries, we have to ensure that m is considered known to the adversary. This is ensured by the last of the above rules. The following definition formally states the definition of the knowledge of the adversary.

Definition 8 (Knowledge). For partial functions $cmd, bits$ and a set Φ , we define the knowledge $know = know_{cmd, bits, \Phi}$ of the adversary to be inductively defined as follows:

- $\Phi \subseteq know$.
- If $h' \in know$ and $cmd_{h'} = D(j, h)$ then $h \in know$.
- If $h' \in know$ and $cmd_{h'} = F(f, h_1, \dots, h_t)$ then $h_1, \dots, h_t \in know$.
- If $h' \in know$ and $cmd_{h'} = C(h)$ then $h \in know$.
- If $h' \in know$ and $cmd_{h'} = D(j, h_1)$, $bits_{h_1} = bits_{h_2}$ and $cmd_{h_2} = E(j, h_3)$ then $h_3 \in know$.
- If $h'_1, h'_2 \in know$ and $cmd_{h'_1} = K(j)$ and $cmd_{h'_2} = E(j, h)$ then $h \in know$.

Note that $know$ can be efficiently computed given Φ , cmd , and $bits$ by adding handles to $know$ according to the rules in Definition 8 until $know$ does not grow any more. We are now ready to state the final definition of adKDM security. Intuitively, an encryption scheme is adKDM secure if the probability that the adversary guesses b correctly without performing a query that would even ideally allow it to retrieve a bitstring constructed using a $C(\cdot)$ query.

Definition 9 (Adaptive KDM Security (adKDM)). An encryption scheme (Enc, Dec) is adKDM secure if for any polynomial-time adversary A there is a negligible function μ such that the following holds:

$$\Pr[\text{Guess} \wedge \neg \text{Invalid}] \leq \frac{1}{2} + \mu(k)$$

where the events refer to an execution of A with input 1^k and oracle access to $\mathcal{T}_{(Enc, Dec)}$ and the events are defined as follows:

By **Guess** we denote the event that the adversary outputs b where b is the challenge bit.

By **Invalid** we denote the event that $h \in know_{cmd, bits, \Phi}$ with cmd_h being of the form $C(\cdot)$.

We will show that this definition can be met (at least in the random oracle model) in the next section. Clearly adKDM security implies DKDM security, since if we can only reveal keys that are not used for decrypting, the plaintexts of the challenge encryptions will never be in $know$.

Adaptive KDM security in the random oracle model. As the OAEP construction is formulated in the random oracle model, we need to know how Definition 9 needs to be adapted when used in the random oracle model. In this case, the adversary A is given access to the random oracle, and the circuits f passed to the adKDM oracle are allowed to contain invocations of the random oracle. Furthermore, the key generation, encryption, and decryption algorithms may contain invocations of the random oracle.

On simulation-based notions. We often motivated our design choices above by comparison with an ideal setting in which the adversary knows exactly the bitstrings associated with handles in *know*. This leads to the question whether it is possible to instead directly define security under key-dependent message attacks using a simulation-based definition, i.e., to define an ideal functionality that handles encryption and decryption queries in an ideal fashion. This approach has been successfully used to formulate IND-CCA security in the UC framework [16]. Their approach, however, strongly depends on the fact that the functionality only needs to output public keys and (fake) encryptions (secret keys are only implicitly present due to the ability to use the functionality to decrypt messages).⁸ It is currently unclear how this approach could be extended to a functionality that can output secret keys. (It is of course possible to define a functionality that outputs secret keys as long as no encryption queries have been performed for that key, but this lead to a definition that is too weak to handle, e.g., our example protocol and that would roughly correspond to Definition 4.) This difficulty persists if we do not use the strong UC model [16] but instead the weaker stand-alone model as in [22, Chapter 7]. Consequently, although a simulation-based definition of KDM security might be very useful, it is currently unknown how to come up with such a definition.

4 OAEP Is adKDM-Secure

We now prove the adKDM security of the OAEP scheme for a partial-domain one-way function. In particular, since the RSA permutation is partial-domain one-way under the RSA assumption [21], the adKDM security of RSA-OAEP follows.

Theorem 10 (OAEP is adKDM secure). *If f is a partial-domain one-way trapdoor 1-1 function, then the OAEP scheme (Enc, Dec) based on f is adKDM secure in the random oracle model.*

To show this theorem, we first define an alternative characterization of partial-domain one-wayness.

⁸ Technically, the reason is that a simulator has to be constructed that chooses the outputs of the functionality. As long as only public keys and ciphertexts are output, fake ciphertexts can be used since they cannot be decrypted. If the simulator had to generate secret keys, the fake ciphertexts could be decrypted and recognized.

Definition 11 (PD-Oracle). *The PD-oracle \mathcal{P}_f for a trapdoor 1-1 function $f : S \times T \rightarrow \text{range } f$ (that may depend on a security parameter) maintains sequences of public/secret key pairs sk_i, pk_i (generated on first use). It understands the following queries:*

- $pk(j)$ and $sk(j)$: Return pk_j or sk_j , respectively.
- $challenge(h, j)$: If h has already been used, ignore this query. Let $j_h := j$. Choose (s_h, t_h) uniformly from $S \times T$. Set $c_h := f_{pk_{j_h}}(s_h, t_h)$. Return c_h .
- $decrypt(h)$: Return (s_h, t_h) .
- $xdecrypt(c, j)$ where $(c, j) \neq (c_h, j_h)$ for all h . Check whether $f_{sk_j}^{-1}(c) = (s_h, t_h)$ for some h . If so, return (s_h, t_h) . Otherwise return \perp .
- $check(s)$: Return the first h with $s_h = s$. If no such h exists, return \perp .

By PDBreak we denote the event that a query $check(s)$ is performed such that

- The query returns $h \neq \perp$.
- No query $sk(j_h)$ and no query $decrypt(h)$ has been performed before the current query.

Lemma 12. *If f is partial-domain oneway, then for any polynomial-time adversary A querying \mathcal{P}_f we have that $\text{Pr}[\text{PDBreak}]$ is negligible in the security parameter.*

The proof is given in Appendix A. We additionally define a variant of the notion of knowledge as defined in Definition 8. We call this variant *lazy knowledge*.

Definition 13 (Lazy knowledge). *For partial functions $cmd, bits$ and a set Φ , we define the lazy knowledge $lknow = lknow_{cmd, bits, \Phi}$ of the adversary to be inductively defined as follows:*

- $\Phi \subseteq lknow$.
- If $h' \in lknow$ and $cmd_{h'} = D(j, h)$ then $h \in lknow$.
- If $h' \in lknow$ and $cmd_{h'} = F(f, h_1, \dots, h_t)$ then $h_1, \dots, h_t \in lknow$.
- If $h' \in lknow$ and $cmd_{h'} = C(h)$ then $h \in lknow$.
- If $h', h_1, h_2 \in lknow$, $cmd_{h'} = D(j, h_1)$, $bits_{h_1} = bits_{h_2}$ and $cmd_{h_2} = E(j, h_3)$ then $h_3 \in lknow$.
- If $h'_1, h'_2 \in lknow$ and $cmd_{h'_1} = K(j)$ and $cmd_{h'_2} = E(j, h)$ then $h \in lknow$.

The only change with respect to Definition 8 is that in the fifth rule we require that $h_1, h_2 \in lknow$. In Definition 13 all rules depend only on values $bits_h$ for which $h \in lknow$; thus one can efficiently compute $lknow$ without accessing $bits_h$ for values $h \notin lknow$ by adding handles to $lknow$ according to these rules until $lknow$ does not grow any further. We call this algorithm the *lazy knowledge algorithm*. Note that $lknow \subseteq know$.

Proof sketch (of Theorem 10). To prove Theorem 10 we give a sequence of games that transforms an attack against the adKDM security of the OAEP scheme into an attack against the PD-oracle. This proof sketch only contains the proof structure and highlights selected steps. The full proof is given in the full version [3].

GAME₁. The adversary A runs with access to the unmodified adKDM oracle \mathcal{T} . We assume that \mathcal{T} invokes an encryption oracle \mathcal{E} for encrypting and a

decryption oracle \mathcal{D} for decrypting. In particular, the encryption oracle \mathcal{E} performs the following actions in the i -th query:

$$r \stackrel{\$}{\leftarrow} \{0, 1\}^{k_0}, g := G(r), s := (m \| 0^{k_1}) \oplus g, h := H(s), t := r \oplus h, c := f_{pk}(s, t).$$

The decryption oracle \mathcal{D} acts as follows, assuming key index j and ciphertext c :

- $(s, t) := f_{pk_j}^{-1}(c), r := t \oplus H(s), (m, z) := s \oplus G(r)$ with $|m| = k - k_1 - k_0$ and $|z| = k_1$.
- If $z = 0^{k_1}$, return m , otherwise return \perp .

GAME₂. We change the encryption oracle to first choose the ciphertext c and then compute the values s, t, r, h, t, g from it, i.e., upon the i -th query the encryption oracle does the following:

$$(s, t) \stackrel{\$}{\leftarrow} \{0, 1\}^{k-k_0} \times \{0, 1\}^{k_0}, c \stackrel{\$}{\leftarrow} f_{pk}(s, t), r \stackrel{\$}{\leftarrow} \{0, 1\}^{k_0}, h := r \oplus t, g := (m \| 0^{k_1}) \oplus s$$

In particular, the values h and g are not retrieved from the oracles G and H any more. In order to keep the distribution of the values c, s, t, r, h, t, g consistent with the answers of the oracles G and H , the oracles G and H are additionally modified to return the values g and h chosen by the encryption oracle. We show that the probability of a successful attack is modified only by a negligible amount with respect to **GAME₁**.

GAME₃. We now change the definition of what constitutes a successful attack. In **GAME₁**–**GAME₂**, we considered it a successful attack if the adversary guessed the bit b chosen by the adKDM oracle \mathcal{T} without performing queries such that the knowledge in the sense of Definition 8 would contain a handle corresponding to a query of the form $C(\cdot)$; see Definition 9.

Now, in **GAME₃**, we consider it to be a successful attack if the adversary guessed b without performing queries such that the *lazy* knowledge in the sense of Definition 13 does not contain a handle corresponding to a query $C(\cdot)$. Since the lazy knowledge is a subset of the knowledge, this represents a weakening of the restrictions put on the adversary. Thus the probability of an attack in **GAME₃** is upper-bound by the probability of an attack in **GAME₂**.

GAME₄. This step is arguably the most important step in the proof. In **GAME₃**, bitstrings $bits_h$ associated to handles h are often computed but never used. For example, the adversary might perform a query $h \leftarrow E(\dots)$ and never use the handle h again. More importantly, however, even if the adversary performs a query $h' \leftarrow E(j, h)$ for that handle h , the value $bits_h$ does not need to be computed due to the following observation: The encryption oracle as introduced in **GAME₂** chooses the ciphertext c at random. The value g (which is the only value depending on the plaintext m) is only needed for suitably reprogramming the oracles G (namely such that $G(r) = g$). Thus we can delay the computation of g until G is queried at position r . Thus in case of a query $h' \leftarrow E(j, h)$, the value $m = bits_h$ is not needed for computing $bits_{h'}$. We use this fact to rewrite the whole game **GAME₃** such that it only

computes a value $bits_h$ when it is actually needed for computing some output sent to the adversary or for computing the lazy knowledge.

The bit b is only used in this game if a value $bits_h$ is computed that corresponds to a query $h \leftarrow C(\cdot)$. If this is not the case, the communication between the adversary and \mathcal{T} is independent of b . Hence, for proving that the probability of attack in the sense of GAME_3 is only negligibly larger than $\frac{1}{2}$ (which then shows Theorem 10), it is sufficient to show that only with negligible probability, a value $bits_h$ is computed such that h is not in the lazy knowledge. Namely, as long as no such value $bits_h$ is computed, the adversary cannot have a higher probability in guessing b than $\frac{1}{2}$ unless $h \in \text{know}$.

GAME_5 . Now we replace the decryption oracle by a plaintext extractor. More concretely, the decryption oracle performs the following steps when given a ciphertext c :

- (a) First, it checks whether $c = f_{pk}(s, t)$ for some pair (s, t) generated by the encryption oracle.⁹ Then values (s, t) are known such that $f_{pk}(s, t) = c$, and the oracle can decrypt c without accessing the secret key sk .
- (b) Otherwise, it checks whether for some s that has been computed by the encryption oracle, there exists a value t such that $f_{pk}(s, t) = c$. (Doing this efficiently requires the secret key; otherwise we had to iterate over all possible values t .) If so, reject the ciphertext.
- (c) Otherwise, for all values s, r that have been generated so far, compute $t := r \oplus H(s)$ and $(m, z) = s \oplus G(r)$. Then check whether $f_{pk}(s, t) = c$ and $z = 0^{k_1}$. If so, return m . Otherwise reject the ciphertext.

We can show that this plaintext extractor is a good simulation of the original decryption oracle (in particular, the adversary is able to produce an s triggering rejection in (b) only if the decryption would fail anyway). Thus the probability that a value $bits_h$ is computed such that h is not in the lazy knowledge does not increase by a non-negligible amount.

GAME_6 . In this final step, we modify GAME_5 not to generate the public/secret key pairs on its own, but to use the PD-oracle \mathcal{P} defined in Definition 11. In particular, we make the following changes:

- When the secret key sk_j is needed (for computing $bits_h$ for a $h \leftarrow K(j)$ query), query $sk(j)$ from \mathcal{P} .
- When producing a ciphertext $bits_{h'}$ (that are produced just to be random images of f_{pk}), use $\text{challenge}(h', j)$ where j is the corresponding key index.
- In the decryption oracle, for checking the condition (a) in GAME_5 , we distinguish two cases. If c was produced by the encryption oracle the decryption oracle sends a $\text{decrypt}(h)$ to \mathcal{P} where h is the query where c was produced. Otherwise it sends an $\text{xdecrypt}(c, j)$ query to \mathcal{P} where j is the index of the key used in the decryption query. In both cases, if the check in (a) would have succeeded, \mathcal{P} will send back a preimage (s, t) of c .
- The check (b) is performed by sending $\text{check}(s)$ to \mathcal{P} .

A case analysis reveals that if a value $bits_h$ is computed such that h is not in the lazy knowledge, then the event PDBreak (as in Definition 11) occurs. By

⁹ This does not imply that c has been generated by the encryption oracle since the encryption oracle might have used a different public key pk at that time.

Lemma 12 this can only happen with negligible probability. Thus no value $bits_h$ is computed such that h is not in the lazy knowledge, and therefore the advantage of the adversary is negligible (as discussed in GAME_4). \square

References

1. Abadi, M., Jürjens, J.: Formal eavesdropping and its computational interpretation. In: Proc. 4th International Symposium on Theoretical Aspects of Computer Software (TACS), pp. 82–94 (2001)
2. Abadi, M., Rogaway, P.: Reconciling two views of cryptography: The computational soundness of formal encryption. In: Watanabe, O., Hagiya, M., Ito, T., van Leeuwen, J., Mosses, P.D. (eds.) TCS 2000. LNCS, vol. 1872, pp. 3–22. Springer, Heidelberg (2000)
3. Backes, M., Dürmuth, M., Unruh, D.: OAEP is secure under key-dependent messages (2008), <http://www.infsec.cs.uni-sb.de/~unruh/publications/backes08oaep.html>
4. Backes, M., Pfitzmann, B.: Symmetric encryption in a simulatable Dolev-Yao style cryptographic library. In: Proc. 17th IEEE Computer Security Foundations Workshop (CSFW), pp. 204–218 (2004)
5. Backes, M., Pfitzmann, B., Scedrov, A.: Key-dependent message security under active attacks – BRSIM/UC-soundness of symbolic encryption with key cycles. In: Proc. of 20th IEEE Computer Security Foundation Symposium (CSF) (June 2007); Preprint on IACR ePrint 2005/421
6. Backes, M., Pfitzmann, B., Waidner, M.: A composable cryptographic library with nested operations (extended abstract). In: Proc. 10th ACM Conference on Computer and Communications Security, pp. 220–230 (January 2003); Full version in IACR Cryptology ePrint Archive 2003/015
7. Beaver, D., Haber, S.: Cryptographic protocols provably secure against dynamic adversaries. In: Rueppel, R.A. (ed.) EUROCRYPT 1992. LNCS, vol. 658, pp. 307–323. Springer, Heidelberg (1993)
8. Bellare, M., Desai, A., Jokipii, E., Rogaway, P.: A concrete security treatment of symmetric encryption. In: Proc. 38th IEEE Symposium on Foundations of Computer Science (FOCS), pp. 394–403 (1997)
9. Bellare, M., Namprempe, C.: Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 531–545. Springer, Heidelberg (2000)
10. Bellare, M., Rogaway, P.: Entity authentication and key distribution. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 232–249. Springer, Heidelberg (1994)
11. Bellare, M., Rogaway, P.: Optimal asymmetric encryption. In: De Santis, A. (ed.) EUROCRYPT 1994. LNCS, vol. 950, pp. 92–111. Springer, Heidelberg (1995)
12. Bellare, M., Rogaway, P.: Encode-then-encipher encryption: How to exploit nonces or redundancy in plaintexts for efficient constructions. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 317–330. Springer, Heidelberg (2000)
13. Black, J., Rogaway, P., Shrimpton, T.: Encryption-scheme security in the presence of key-dependent messages. In: Proc. 9th Annual Workshop on Selected Areas in Cryptography (SAC), pp. 62–75 (2002)
14. Boneh, D., Halevi, S., Hamburg, M., Ostrovsky, R.: Circular-secure encryption from decision diffie-hellman. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 108–125. Springer, Heidelberg (2008)

15. Camenisch, J., Lysyanskaya, A.: An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 93–118. Springer, Heidelberg (2001)
16. Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: Proc. 42nd IEEE Symposium on Foundations of Computer Science (FOCS), pp. 136–145 (2001); Extended version in Cryptology ePrint Archive, Report 2000/67
17. Canetti, R., Herzog, J.: Universally composable symbolic analysis of mutual authentication and key exchange protocols. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 380–403. Springer, Heidelberg (2006)
18. Cortier, V., Warinschi, B.: Computationally sound, automated proofs for security protocols. In: Proc. 14th European Symposium on Programming (ESOP), pp. 157–171 (2005)
19. Dolev, D., Dwork, C., Naor, M.: Nonmalleable cryptography. *SIAM Journal on Computing* 30(2), 391–437 (2000)
20. Dolev, D., Yao, A.C.: On the security of public key protocols. *IEEE Transactions on Information Theory* 29(2), 198–208 (1983)
21. Fujisaki, E., Okamoto, T., Pointcheval, D., Stern, J.: RSA-OAEP is secure under the RSA assumption. *Journal of Cryptology* 17(2), 81–104 (2004)
22. Goldreich, O.: *Foundations of Cryptography. Basic Applications*, vol. 2. Cambridge University Press, Cambridge (May 2004)
23. Goldwasser, S., Micali, S.: Probabilistic encryption. *J. Comput. Syst. Sci.* 28, 270–299 (1984)
24. Halevi, S., Krawczyk, H.: Security under key-dependent inputs. In: Proc. of the 14th ACM Conference on Computer and Communications Security (to appear, 2007); Preprint on IACR ePrint 2007/315
25. Hofheinz, D., Unruh, D.: Towards key-dependent message security in the standard model (August 2007); Preprint on IACR ePrint 2007/333
26. Laud, P.: Semantics and program analysis of computationally secure information flow. In: Proc. 10th European Symposium on Programming (ESOP), pp. 77–91 (2001)
27. Laud, P.: Symmetric encryption in automatic analyses for confidentiality against active adversaries. In: Proc. 25th IEEE Symposium on Security & Privacy, pp. 71–85 (2004)
28. Luby, M.: *Pseudorandomness and Cryptographic Applications*. Princeton Computer Society Notes, Princeton (1996)
29. Micciancio, D., Warinschi, B.: Soundness of formal encryption in the presence of active adversaries. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 133–151. Springer, Heidelberg (2004)
30. Rackoff, C., Simon, D.R.: Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 433–444. Springer, Heidelberg (1992)
31. Shannon, C.E.: Communication theory of secrecy systems. *Bell System Technical Journal* 28(4), 656–715 (1949)
32. Yao, A.C.: Theory and applications of trapdoor functions. In: Proc. 23rd IEEE Symposium on Foundations of Computer Science (FOCS), pp. 80–91 (1982)

A Proof of Lemma 12

Proof. Given an adversary A against the PD-Oracle \mathcal{P} we construct an adversary B against partial-domain one-wayness of the underlying function f as follows.

The machine B that implements the PD-oracle with slight changes: Let q be an upper bound on the number of queries performed by A . Then B gets as input a key pair pk^*, sk^* , values $(s^*, t^*) \in S \times T$ and a value c^* . Let j^* be the i_1 -th key index that is used in A 's queries, and let h^* the i_2 -th handle that is used in a query of the form $challenge(h, j^*)$. Then B answers to A 's queries as follows (for simplicity, if we write f_{pk}^{-1} we mean an application of the secret key sk):

- $pk(j)$: If $j = j^*$, return pk^* , otherwise return pk_j .
- $sk(j)$: If $j = j^*$, return sk^* , otherwise return sk_j .
- $challenge(h, j)$: If h has already been used, ignore this query.
 - If $h = h^*$ (and thus also $j = j^*$) then set $c_h := c^*$ and return c_h .
 - If $h \neq h^*$ then choose (s_h, t_h) uniformly from $S \times T$. Set $c_h := f_{pk_{j_h}}(s_h, t_h)$. Return c_h .
- $decrypt(h)$: If $h = h^*$, return (s^*, t^*) . Otherwise return (s_h, t_h) .
- $xdecrypt(c, j)$ where $(c, j) \neq (c_h, j_h)$ for all h . This is equivalent to the following:
 - If $j \neq j^*$ then check whether $f_{pk_j}^{-1}(c) = (s_h, t_h)$ for some $h \neq h^*$ or $f_{pk_{j^*}}(f_{pk_j}^{-1}(c)) = c_{h^*}$. If so, return $f_{pk_j}^{-1}(c)$. Otherwise, return \perp .
 - If $j = j^*$ then test if $f_{pk_j}(s_h, t_h) = c$ for any $h \neq h^*$. If such an h exists, output (s_h, t_h) . Otherwise, return \perp .
- $check(s)$: If $s = s_h$ for some h , return the first h with $s_h = s$. If $sk(j^*)$ or $decrypt(h^*)$ has been queried, check whether $s = s^*$. If so, return h^* .

We claim that this machine B behaves identically to the PD-oracle \mathcal{P} until the event $PDBreak$ occurs and that A 's view is independent of i_1, i_2 until the event $PDBreak$ occurs (assuming that the inputs sk^*, pk^* are an honestly generated key pair, (s^*, t^*) is uniformly distributed on $S \times T$ and $c^* = f_{pk^*}(s^*, t^*)$). For the queries pk , sk , $challenge$, and $decrypt$ this is straightforward. In the case of $xdecrypt$ we distinguish two cases: For $j \neq j^*$, the check performed is equivalent to checking whether $f_{pk_j}^{-1}(c) = (s_h, t_h)$ for some $h \neq h^*$ or $f_{pk_{j^*}}^{-1}(c) = (s^*, t^*)$ and then returning h or h^* , respectively. Thus in this case the answer to the query $xdecrypt$ is the same as that the PD-oracle \mathcal{P} would give. For $j = j^*$, in comparison to \mathcal{P} , the check whether $f_{pk_j}(s^*, t^*) = c$ is missing. However, if this check held true, we would have that $(c, j) = (c^*, j^*)$ which is excluded. To see that the query $check(s)$ gives the same answers in B and \mathcal{P} until $PDBreak$ occurs, note that the only case where $check(s)$ would give another answer in \mathcal{P} is when $s = s^*$ but neither $sk(j^*)$ nor $decrypt(h^*)$ have been queried. However, in this case h^* would be returned in \mathcal{P} , thus $PDBreak$ occurs.¹⁰ So altogether, we have that B behaves identically to \mathcal{P} and A 's view is independent of i_1, i_2 until the event $PDBreak$ occurs. By $PDBreak_{i_1', i_2'}$, denote the event that $check(s)$ is queried with $s = s_h$ where h is the i_2' -th handle used by A , and no query $sk(j_h)$ or $decrypt(h)$ has been performed where j_h is the i_1' -th key index used by A . Obviously, if

¹⁰ In slight abuse of notation, we denote by $PDBreak$ not the event that $h \neq \perp$ is returned without a query of $sk(j_h)$ or $decrypt(h)$, but that some $check(s)$ is queried such that $s = s_h$ and no query $sk(j_h)$ or $decrypt(h)$ has been performed. Since for \mathcal{P} these are equivalent, it is enough to show the lemma w.r.t. this slightly changed definition.

PDBreak occurs, then $\text{PDBreak}_{i'_1, i'_2}$ occurs for some $i'_1, i'_2 \in \{1, \dots, q\}$. Since the view of A is independent of i_1, i_2 , we have that $\Pr[\text{PDBreak}_{i_1, i_2}] \geq \frac{1}{q^2} \Pr[\text{PDBreak}]$. So it is enough to show that $\Pr[\text{PDBreak}_{i_1, i_2}] =: \varepsilon$ is negligible. Observe that in the description of B , in case of the event $\text{PDBreak}_{i_1, i_2}$ the inputs sk^*, s^*, h^* are never accessed. So if we run B with the inputs sk^*, s^*, h^* set to \perp , $\text{PDBreak}_{i_1, i_2}$ still occurs with probability at least ε . Further, $\text{PDBreak}_{i_1, i_2}$ implies that $\text{check}(s)$ is called an s satisfying $f^{-1}(c^*) = \perp$. So if let B output one of the values s used in $\text{check}(s)$ queries (randomly chosen), we break the partial-domain one-wayness of f with probability at least ε/q . Thus by contradiction, ε must be negligible. Thus $\Pr[\text{PDBreak}]$ is negligible in an execution of B and thus also in one of \mathcal{P} . \square