

Random Oracles and Auxiliary Input*

Dominique Unruh**

Saarland University, Saarbrücken, Germany, unruh@cs.uni-sb.de

Abstract. We introduce a variant of the random oracle model where oracle-dependent auxiliary input is allowed. In this setting, the adversary gets an auxiliary input that can contain information about the random oracle. Using simple examples we show that this model should be preferred over the classical variant where the auxiliary input is independent of the random oracle.

In the presence of oracle-dependent auxiliary input, the most important proof technique in the random oracle model—lazy sampling—does not apply directly. We present a theorem and a variant of the lazy sampling technique that allows one to perform proofs in the new model almost as easily as in the old one.

As an application of our approach and to illustrate how existing proofs can be adapted, we prove that RSA-OAEP is IND-CCA2 secure in the random oracle model with oracle-dependent auxiliary input.

Keywords: Random oracles, auxiliary input, proof techniques, foundations.

Table of Contents

1	Introduction	2
1.1	Our results	5
1.2	Related work	6
1.3	Further applications	7
1.4	Organisation	7
1.5	Notation	7
2	Lazy sampling with auxiliary input	7
3	Example: one-wayness of the random oracle	11
4	Security amplification	12
5	OAEP encryption	14
6	Open questions	15
A	An auxiliary lemma	16
B	Proof of Lemma 6	17
C	Proof of Lemma 7	18
D	Proof of Lemma 8	18
E	Proof of Theorem 2	19
F	Example: given-preimage collision-resistance	20
G	Normal security notions and security amplification	21
H	Security of OAEP encryption	23

* This is the full version of a paper appearing at CRYPTO 2007

** Part of this work was done while the author was at the IAKS, University of Karlsruhe, Germany

1 Introduction

In [BR93] the following heuristic was advocated as a practical way to design cryptographic protocols:¹ To prove the security of a cryptographic scheme, one first introduces a *random oracle* \mathcal{O} , i.e., a randomly chosen function to which all parties including the adversary have access. One then proves the security of the scheme that uses the random oracle and subsequently replaces the random oracle by a suitably chosen function (or family of functions) H . The *random oracle heuristic* now states that if the scheme using \mathcal{O} is secure, the scheme using H is secure as well.

Unfortunately, a counter-example to this heuristic has been given in [CGH98]. It was shown that there exist public key encryption and signature schemes that are secure in the random oracle model but lose their security when instantiated with *any* function or family of functions. Nonetheless, the random oracle heuristic still is an important design guideline for implementing cryptographic schemes.

Furthermore, [Pas03] pointed out that zero-knowledge proofs in the random oracle model can lose their deniability when instantiated with a fixed function. In contrast to the result of [CGH98], this happens even for natural protocols. However, [Pas03] was able to give conditions under which this effect does not occur and gave a protocol that fulfilled these conditions.

Although the heuristic is known not to be sound in general, no practical scheme is known where it fails, and schemes that are proven to be secure using this heuristic tend to be simpler and more efficient than schemes that are shown to be secure in the standard model. As a consequence, schemes used in practise are often based on the random oracle heuristic, e.g., the RSA-OAEP encryption scheme, introduced in [BR95] and standardised in [PKC02], is one of the most widely used public-key encryption schemes, and its security is based on the random oracle heuristic.

In the light of the results of [CGH98] and [Pas03], and of the practical importance of the random oracle heuristic, it is important to try and learn what the exact limitations of the heuristic are, and, if possible, give criteria to distinguish those protocols in the random oracle model that become insecure when instantiated due to those limitations, and those protocols where we can at least hope—if not prove—that their instantiations are secure. The augmented definition of zero-knowledge by [Pas03] is an example of such a criterion.

In this paper, we uncover another such limitation of the random oracle world. We will see that there are natural schemes secure in the random oracle model that become insecure *with respect to auxiliary input* (or equivalently, with respect to nonuniform adversaries) when instantiated. As [Pas03] did for the deniability, we give augmented definitions for the random oracle model with auxiliary input that allow one to distinguish protocols that fail upon instantiation from those that do not (at least not due to the abovementioned limitation).

Although such a result does not imply the soundness of the random oracle model, it helps to better understand which protocol can reasonably be expected to be secure when instantiated with a fixed function.

We will now investigate the problem of auxiliary input in the random oracle model in more detail. An important concept in cryptology is the *auxiliary input*. The auxiliary input is a string that is given to the adversary at the beginning of the execution of some cryptographic protocol. This string is usually chosen nonuniformly and depends on all protocol inputs. In other words, the auxiliary input models the possibility that the adversary has some additional knowledge concerning the situation at the beginning of the protocol. This additional knowledge may, e.g., represent information acquired in prior protocol runs. It turns out that in many cases the presence of an auxiliary input is an essential concept for proving secure sequential composition. Therefore, most modern cryptographic schemes are designed to be secure even in the presence of an auxiliary input (given that the underlying complexity assumptions hold against nonuniform adversaries).

However, when we try to combine these two concepts, the random oracle model and the auxiliary input, undesirable effects may occur. We will demonstrate this by studying the definitions of two simple

¹ However, the basic idea seems to have already appeared earlier.

security notions: one-wayness and collision-resistance. First, consider the notion of a one-way function. A function f is one-way (with respect to auxiliary input) if for any polynomial-time adversary A and any auxiliary input z (more exactly, any sequence z_k of strings of polynomial-length), the following probability is negligible:

$$P(x \xrightarrow{\$} \{0,1\}^k, x' \leftarrow A(1^k, z, f(x)) : f(x') = f(x))$$

In other words, let x be a random element, give $f(x)$ to the adversary A , and then the probability shall be negligible in the security parameter k that the adversary outputs some preimage x' of $f(x)$. When we use the random oracle model, a random oracle \mathcal{O} is additionally introduced, and the function f is allowed to use the random oracle. E.g., if we set $f := \mathcal{O}$, one-wayness of f means that for any polynomial-time adversary A and any auxiliary input z , the following probability is negligible:

$$P(x \xrightarrow{\$} \{0,1\}^k, x' \leftarrow A^{\mathcal{O}}(1^k, z, \mathcal{O}(x)) : \mathcal{O}(x') = \mathcal{O}(x)). \quad (1)$$

Here \mathcal{O} is a randomly chosen function (with some given domain and range), and the adversary is given black-box access to \mathcal{O} . It is now easy to see that $f := \mathcal{O}$ is indeed secure in the above sense: The adversary can make at most a polynomial number of queries, and each query except $\mathcal{O}(x)$ returns a uniformly random image (exploiting this latter fact is later called the lazy sampling technique). From this fact one can conclude that the adversary must make an exponential number of queries to find a preimage of $\mathcal{O}(x)$, hence f is secure. The presence of the auxiliary input does not have noticeable impact on the proof. The random oracle heuristic now claims that $f := H$ is oneway for a sufficiently unstructured function H , even in the presence of auxiliary input. So far, nothing out of the ordinary has happened.

We now try to use the same approach for another security property: collision-resistance. Again, we set $f := \mathcal{O}$, and then collision-resistance of f means that for any polynomial-time adversary A and any auxiliary input z , the following probability is negligible:

$$P((x_1, x_2) \leftarrow A^{\mathcal{O}}(1^k, z) : x_1 \neq x_2 \text{ and } \mathcal{O}(x_1) = \mathcal{O}(x_2)). \quad (2)$$

This can again easily be proven using the lazy sampling technique: the answers to the adversary's queries are independent random values, and finding a collision requires two of these random values to be identical which happens only with negligible probability. Again, the auxiliary input does not help the adversary, since it does not contain any information on where a collision might be. We now use the random oracle heuristic, replace \mathcal{O} by some sufficiently unstructured function H , so that $f = H$, and then claim that f is collision-resistant in the presence of an auxiliary input. But this of course is impossible, since the auxiliary input may simply contain a collision of H , since H is a fixed function.²

Hence, the random oracle heuristic should not be applied to collision-resistance. On the other hand, we would like to prove the one-wayness of $f := \mathcal{O}$ in the random oracle model. We hence need a stronger variant of the random oracle heuristic that does not allow one to prove the collision-resistance of f , but still allows one to prove its one-wayness. An inspection of our proof above reveals the mistake we made: In the random oracle model, the auxiliary input was chosen before the random oracle, so it could not contain a collision. After instantiation, the function H was fixed, so the auxiliary input did depend on H and therefore could provide a collision. The random oracle heuristic should hence be recast as follows in the case of auxiliary input: When a scheme is secure in the random oracle model with *oracle-dependent* auxiliary input, it is still secure after replacing the random oracle by a sufficiently unstructured fixed function H , even in the presence of auxiliary input.

² If we replace \mathcal{O} by a family of functions, i.e., some parameter i is chosen at the beginning of the protocol, and then a function H_i is used, then the problem described here does not occur. Unfortunately, one is not always free to use such a family of functions. On one hand, the index has in some way to be chosen, and we do not want to leave that choice to the corrupted parties. On the other hand, practical applications usually instantiate the random oracle using a fixed function like SHA-1 or SHA-256 [SHS02].

It remains to clarify the formal meaning of oracle-dependent auxiliary input. Unfortunately, we cannot simply say: “for randomly chosen \mathcal{O} and every z ”. At least, the semantics underlying constructions like (1) and (2) get highly nontrivial in this case. Fortunately, there is another possibility. By an oracle function z we mean a function that returns a string $z^{\mathcal{O}}$ for each possible value of the random oracle \mathcal{O} . So formally, z is simply a function that maps functions to strings. Then a scheme is called secure in the random oracle model with oracle-dependent auxiliary input if for any polynomial-time adversary A and for any oracle function z into strings of polynomial length (that may depend on k , of course), the adversary cannot break the scheme *even when given $z^{\mathcal{O}}$ as auxiliary input*.

As with the traditional random oracle model, the exact form these definitions take depends on the security notion under consideration. For example, the one-wayness and the collision-resistance of $f := \mathcal{O}$ take the following form: for any polynomial-time adversary A and any oracle function z into strings of polynomial-length, we have

$$P(x \stackrel{\$}{\leftarrow} \{0,1\}^k, x' \leftarrow A^{\mathcal{O}}(1^k, z^{\mathcal{O}}, \mathcal{O}(x)) : \mathcal{O}(x') = \mathcal{O}(x)) \quad (3)$$

or

$$P((x_1, x_2) \leftarrow A^{\mathcal{O}}(1^k, z^{\mathcal{O}}) : x_1 \neq x_2 \text{ and } \mathcal{O}(x_1) = \mathcal{O}(x_2)), \quad (4)$$

respectively. However, we can give a simple guideline on how to transform a security definition in the random oracle model with an oracle-*independent* auxiliary input z into a security definition with oracle-*dependent* auxiliary input. First, one quantifies over oracle functions z instead of strings z . And then one replaces all occurrences of the string z by $z^{\mathcal{O}}$.

It is now easy to see that (4) is not negligible: let $z^{\mathcal{O}}$ encode a collision x_1, x_2 , and let the adversary output that collision. Since such a collision always exists (assuming a *length-reducing* f), this breaks the collision-resistance of f in the presence of oracle-dependent auxiliary input, as we would have expected.

On the other hand, we expect (3) to be negligible in the presence of oracle-dependent auxiliary input. However, it is not so easy to see whether there may not be some possibility to encode information about the random oracle in a string of polynomial length that allows one to find a preimage with non-negligible probability. Although one-wayness is one of the weakest conceivable security notions, proving its security with respect to oracle-dependent auxiliary input is quite difficult. (We encourage the reader to try and find an elementary proof for the one-wayness of f .³) The reason for this difficulty lies in the fact that it is not possible any more to apply the lazy sampling technique: given some information $z^{\mathcal{O}}$ on the random oracle \mathcal{O} , the images under the random oracle are not independently nor uniformly distributed any more. We therefore need new techniques if we want to be able to cope with oracle-dependent auxiliary input and to prove more complex cryptographic schemes secure in this model. Such techniques will be presented in this paper.

On nonuniform and uniform auxiliary input. In this work, we always consider nonuniform auxiliary inputs, that is, the auxiliary input is not required to be the result of an efficient computation. This is the most common modelling of auxiliary input in the cryptographic community. However, it is also possible to consider uniform auxiliary inputs: In this case, the auxiliary input is not an arbitrary sequence of strings, but is instead the output of a *uniform* probabilistic algorithm. The main motivation of the auxiliary input, namely to model information gained from prior executions of cryptographic protocols on the same data, and thus to allow for composability, is preserved by this uniform approach. (See [Gol93] for a detailed analysis.) The main disadvantage of the uniform approach is that definitions and proofs get more complicated due to the presence of another machine. This is why the nonuniform auxiliary input is more commonly used.

Applying the uniform approach to our setting, a uniform oracle-dependent auxiliary input would be the output of a polynomial-time oracle Turing machine Z with access to the random oracle. Since that Turing

³ We give a proof using the techniques from this paper in Lemma 10.

machine could only make a polynomial number of queries, using the lazy sampling technique would be easy: all positions of the random oracle that have not been queried by Z can be considered random.

However, if we use the random oracle heuristic to motivate the security of a protocol with respect to *uniform* auxiliary input, the result is incompatible with existing theorems and definitions in the *nonuniform* auxiliary input model. So to use the random oracle heuristic together with existing results, we either have to reprove all existing results for the uniform case, or we have to use *nonuniform* oracle-dependent auxiliary input. It is the latter approach we follow in this work.

Instantiating the random oracle with keyed families of functions. Above, we showed that the random oracle is (unsurprisingly) not collision-resistant in the presence of auxiliary input. It follows that we may not instantiate the random oracle with a fixed function if we need collision-resistance. On the other hand, replacing the random-oracle by a keyed family of functions may be secure, since the auxiliary input cannot encode a collision for each function. We do not claim that it is necessary to use oracle-dependent auxiliary input when instantiating with families of functions. Rather, oracle-dependent auxiliary input provides a tool for distinguishing the cases where the use of a single function⁴ is sufficient (e.g., in the case where we require only one-wayness) and where a keyed family of functions is necessary (e.g., in the case that we require collision-resistance). Since instantiating with a single function is much simpler (e.g., we do not have to worry about who chooses the key), and is the usual practice in real-world protocols, examining random oracle based protocols with respect to oracle-dependent auxiliary input may give additional insight into when instantiation with single functions is permitted and when we have to use keyed families. Another disadvantage of using a family of functions is that we have to ensure that the key is honestly generated, which may introduce additional difficulties if no trusted party is available for this task.

Designing special protocols. An alternative to the approach in this paper would be to systematically construct or transform a protocol so that it is secure with respect to oracle-dependent auxiliary input (instead of verifying a given protocol). However, here the same arguments as in the previous paragraph apply. First, we might not be interested in a new protocol, but might want to examine the security of an existing protocol (that possibly even has already been implemented). Further, efficiency considerations might prevent the use of more elaborate constructions.

1.1 Our results

We introduce and motivate the random oracle model with *oracle-dependent auxiliary input* (preceding section). In this model, the auxiliary input given to the adversary may depend on the random oracle.

In order to be able to prove security in the new model, we introduce a new variant of the lazy sampling technique that is applicable even in the presence of oracle-dependent auxiliary input. We show that one can replace the random-oracle \mathcal{O} by a new random oracle \mathcal{P} that is independent of the auxiliary input, except for a *presampling*. That is, a small fraction of the total random oracle \mathcal{P} is fixed (and dependent on the auxiliary input), while all other images are chosen independently and uniformly at random (and in particular are independent of the auxiliary input). In this new setting, lazy sampling is possible again: an oracle query that is not in the presampled set is given a random answer.

This also gives some insight into why some schemes are secure and some fail in the presence of oracle-dependent auxiliary input: Intuitively the protocols that fail are those for which you can have a “reason for a failure” (e.g., a collision) contained in a few entries of the random oracle.

As a technical tool, we also formulate a *security amplification* technique: for many security notions, security with respect to nonuniform polynomial-time adversaries implies security with respect to nonuniform adversaries whose running time is bounded by some suitable superpolynomial function f . This technique

⁴ Here, by a single function we mean that the function is not parametrised by a key that has to be known by all parties. However, the function may depend on the security parameter. Otherwise a property like collision-resistance trivially cannot be fulfilled by a single function, even against uniform adversaries. See also [Rog06] in this context.

is useful in the context of oracle-dependent auxiliary input, since some reduction proofs with presampling tend to introduce superpolynomial adversaries.

As an application of our techniques, we show that RSA-OAEP is IND-CCA2 secure in the random oracle model *with oracle-dependent auxiliary input*. Our proof closely follows the proof of [FOPS04] where the security of RSA-OAEP was shown in the classical random oracle model. This allows the reader to better compare the differences in the proof introduced by the oracle-dependent auxiliary input. However, we believe that the result does not only exemplify our techniques but is worthwhile in its own light: it gives the first evidence that RSA-OAEP as used in practical application (i.e., with the random oracle instantiated with a fixed function H), is secure *even in the presence of an auxiliary input*.

1.2 Related work

In [Wee06], the problem of composition of zero-knowledge proofs in the random-oracle model is investigated. It is shown that to guarantee sequential composition, oracle-dependent auxiliary input is necessary. Their definition of oracle-dependent auxiliary input is somewhat weaker than ours in that the machine z generating the auxiliary input is allowed only a polynomial number of queries to the random oracle (it is similar to uniform oracle-dependent auxiliary input in that respect). They give protocols that are secure with respect to that notion. It would be interesting to know whether the techniques developed here allow to show their protocols to be secure even with respect to our stronger notion of oracle-dependent auxiliary input.

In [GGKT05], it was shown that a random *permutation* is one-way with respect to oracle-dependent auxiliary input. They showed that the advantage of the adversary is in $2^{-\Omega(k)}$ which is essentially the same bound as we achieve for random *functions* in Section 3. However, their proof is specific to the property of one-wayness and does not generalise to our setting. According to [GGKT05], a similar result was shown for random *functions* in [Imp96]. However, their proofs apply only to the one-wayness of the random oracle, while our results imply that many more cryptographic properties of the random oracle are preserved in the presence of oracle-dependent auxiliary input.

In [Mau92, CM97, CCM02, DM02], unconditional security proofs in the bounded-storage model were investigated. In this model, one assumes that the adversary is computationally unlimited, but that it may only store a limited amount of data. One assumes that at the beginning of the protocol some large source of randomness (e.g., a random oracle) is available to all parties. The security of the protocol then roughly hinges on the following idea: The honest parties store some (small) random part of the source. Since the adversary does not know which part has been chosen, and since it may not store the whole source, with high probability the honest parties will find some part of the random source they both have information about, but that is unknown to the adversary. To prove the security in this model, it is crucial to show that the adversary cannot compress the source in a manner that contains enough information to break the protocol. This is very similar to the scenario investigated here, since the oracle-dependent auxiliary input can be seen as compressed information on the random oracle. Our results differ from those in the bounded-storage model in two ways: first, our results cover a more general case, since we consider the effect of auxiliary input on arbitrary protocols, while in the bounded-storage model a single protocol is analysed that is specially designed to extract information from the random source that cannot be extracted given only a part of the source. On the other hand, precisely due to the specialised nature of the protocols, the bounds achieved in the bounded-storage model are better than those presented here. In particular, there are protocols in the bounded storage model that are secure given a random source of polynomial size [DM02], while our results are—at least with the present bounds—only useful if the domain of the random oracle has superpolynomial size (cf. the exact bounds given by Theorem 2). It would be interesting to know whether our techniques can be used in the context of the bounded-storage model, and to what extent the techniques developed in the bounded-storage model can be applied to improve our bounds.

1.3 Further applications

Besides the application described above, namely to be able to use the random-oracle heuristic in the case of auxiliary input, our main result (the lazy sampling technique) may also be useful in other situations.

In [GGKT05] it was shown that a random permutation is one-way in the presence of oracle-dependent auxiliary input. This was the main ingredient for several lower bounds on black-box constructions using one-way permutations. Using our techniques, we might find lower bounds on black-box constructions based on other cryptographic primitives: namely, we would show that the random oracle (or a protocol using the random oracle) has a given security property X even in the presence of oracle-dependent auxiliary input. Then using techniques from [GGKT05], lower bounds on black-box constructions based on cryptographic primitives fulfilling X might be derived.

1.4 Organisation

In Section 1 we introduce and motivate the concept of oracle-dependent auxiliary input. In Section 2 we present the main result of this paper: a theorem that allows one to use the lazy sampling technique even in the presence of oracle-dependent auxiliary input. In Section 3 we give a simple example to show how to use the lazy sampling technique. In Section 4 we present the security amplification technique. This technique allows one to use superpolynomial adversaries in reduction proofs, which sometimes is needed when using the lazy sampling technique. In Section 5 we prove that RSA-OAEP is IND-CCA2 secure in the random oracle model with oracle-dependent auxiliary input.

1.5 Notation

For random variables A and B , we denote the Shannon-entropy of A by $H(A)$, and the conditional entropy of A given B by $H(A|B)$. The statistical distance between A and B is denoted $\Delta(A; B)$. The operator \log means the logarithm base 2. The variable k always denotes the security parameter. In asymptotic statements of theorems or definitions, some variables implicitly depend on the security parameter k . These variables are then listed at the end of the theorem/definition. We call a nonnegative function in k negligible, if it lies in $k^{-\omega(1)}$. We call a nonnegative function non-negligible if it is not negligible.

Let \mathcal{O} always denote the random oracle. Let *Domain* be the domain and *Range* the range of the random oracle, i.e., \mathcal{O} is a uniformly random function from *Domain* \rightarrow *Range*. In an asymptotic setting, \mathcal{O} , *Domain* and *Range* implicitly depend on the security parameter k . In this case we always assume $\#Domain$ and $\#Range$ to grow at least exponentially in k .

An *oracle function* g into X is a mapping from *Domain* \rightarrow *Range* into X . We write the image of some function \mathcal{O} under g as $g^{\mathcal{O}}$.

An *assignment* S is a list $S = (x_1 \rightarrow y_1, \dots, x_n \rightarrow y_n)$ with $x_i \in Domain$ and $y_i \in Range$ and with $x_i \neq x_j$ for $i \neq j$. The length of S is n . We call y_i the image of x_i under S . We write $x \in S$ if $x_i = x$ for some i . The image $\text{im } S$ is defined as $\text{im } S = \{y_1, \dots, y_n\}$.

2 Lazy sampling with auxiliary input

The main result of this paper is the following theorem which guarantees that we can replace a random oracle with oracle-dependent auxiliary input by a new random oracle that is independent of the auxiliary input with the exception of some fraction of its domain (which is *presampled*). In order to formulate the theorem, we first need to state what exactly we mean by an oracle with presampling:

Definition 1 (Random oracle with presampling). *Let $S = (x_1 \rightarrow y_1, \dots, x_n \rightarrow y_n)$ be an assignment. Then the random oracle \mathcal{P} with presampling S is defined as follows:*

When queried $x \in \text{Domain}$ with $x = x_i$ for some $i \in \{1, \dots, n\}$, the oracle returns y_i . If x has already been queried, the same answer is given again. Otherwise, a uniformly random element y is chosen from Range and returned.

We can now state the main theorem.

Theorem 2 (Lazy sampling with auxiliary input). *Let $f \geq 1$ and $q \geq 0$ be integers. Let z be an oracle function with finite range Z and $p := \log \#Z$.*

Then there is an oracle function S , such that $S^\mathcal{O}$ is an assignment of length at most f , so that the following holds: For any probabilistic oracle Turing machine A that makes at most q queries to the random oracle, it is

$$\Delta(A^\mathcal{O}(z^\mathcal{O}); A^\mathcal{P}(z^\mathcal{O})) \leq \sqrt{\frac{pq}{2f}}$$

where \mathcal{P} is the random oracle with presampling $S^\mathcal{O}$.

Before presenting the actual proof, we give a short sketch that is intended to serve as a guide through the rest of the proof. To ease comparison with the details given later, we provide forward references to the lemmas of the actual proof.

For any i , let J_i be the maximum amount of information that a sequence of i queries to the random oracle \mathcal{O} gives about the auxiliary input $z^\mathcal{O}$. Since $|z| = p$, $J_i \leq p$ for all i . Let F_i be the sequence of i queries that achieves this bound, that is, the mutual information between $z^\mathcal{O}$ and the oracle's answers to F_i is J_i .

Assume that the queries F_i have already been performed. Let G be a sequence of q queries. Then the answers to the queries F_i and G together contain at most J_{q+i} bits of information about z . Thus the answers to G contain at most $J_{q+i} - J_i$ bits of information about z beyond what is already known from the answers to F_i .

Consider the quantities $J_0, J_q, J_{2q}, \dots, J_{f+q}$ (assuming that q divides f). Since $J_0 \geq 0$ and $J_{f+q} \leq p$, there must be some $f' \leq f$ such that the $J_{f'+q} - J_{f'} \leq \frac{pq}{f}$. Thus, given the answers to $F := F_{f'}$, any sequence G of q queries reveals at most $\frac{pq}{f}$ bits about the auxiliary input z . In other words, the answers to G are almost independent of z (assuming that $\frac{pq}{f}$ is sufficiently small). Thus, if we fix the oracle \mathcal{P} to match the answers to F , but choose \mathcal{P} independently of z everywhere else, with q queries we cannot distinguish between \mathcal{P} and the original oracle \mathcal{O} . This gives Theorem 2 (except for the concrete bound $\sqrt{pq/2f}$).

In reality, however, the queries performed by A are adaptive, i.e., they depend on z and on the answer to prior queries. So we cannot talk about a fixed sequence G of queries made by A . To overcome this problem, we introduce the concept of an adaptive list (Definitions 3 and 4), which is a generalisation of a sequence of queries where the queries are allowed to be adaptive. When considering adaptive lists, it does not make immediate sense to speak about the mutual information between the answers to an adaptive list G and the auxiliary input $z^\mathcal{O}$. In Definition 5 we therefore introduce quantities $J(G)$ and $J(G|F)$ denoting the information that the answers to the adaptive list G contains about $z^\mathcal{O}$ (given the answers to the adaptive list F in the case of $J(G|F)$). For this quantity, we show that $J(F) \leq p$ (Lemma 7) and give a chain rule for the information contained in the concatenation of adaptive lists (Lemma 6). Then we can construct the sequence F as in the proof sketch above (Lemma 8). However, F is now an adaptive list. Finally, Theorem 2 is proven (page 10) by showing that the adversary A can be considered as an adaptive list G of length q , and therefore cannot distinguish the answers to queries outside F from uniform randomness. For convenience, in Corollary 9 we formulate an asymptotic version of Theorem 2.

We now give the details of the proof, broken down to several lemmas. First we have to define the concept of an adaptive list. To capture the possibility of adaptive queries, an adaptive list is formally just a deterministic oracle Turing machine. An adaptive list of length n makes n queries to the oracle and outputs an assignment containing the queries and the results of these queries. To be able to talk about the

concatenation of adaptive lists, we slightly extend this idea. An adaptive list takes an auxiliary input z , but also an assignment X . This assignment can be thought of as the queries made by an adaptive list executed earlier. So in a concatenation of two adaptive lists, the queries of the second adaptive list can depend on the results of the queries made by the first adaptive list. For definitional convenience, an adaptive list does not only output its queries, but also the queries received as input. An adaptive list expecting a queries as input and then making $b - a$ queries, we call an $a \rightarrow b$ adaptive list. We require that an adaptive list never repeats a query. Note that an adaptive list is indeed a generalisation of a non-adaptive sequence of queries: a sequence (x_1, \dots, x_n) corresponds to the $0 \rightarrow n$ adaptive list querying the positions x_1 to x_n and returning the results.

Definition 3 (Adaptive list). *Let $\#Domain \geq b \geq a \geq 0$. An $a \rightarrow b$ adaptive list M is defined as a deterministic oracle Turing machine that takes an assignment $X = (x_1 \rightarrow y_1, \dots, x_a \rightarrow y_a)$ and a string $z \in \Sigma^*$ as input and satisfies the following properties*

- $M = M(X, z)$ does not query the oracle at positions x_1, \dots, x_a .
- M never queries the oracle twice at the same position.
- M queries the oracle exactly $b - a$ times.
- Let x'_1, \dots, x'_{b-a} be the positions of the oracle calls made by M (in that order). Let $y'_i := \mathcal{O}(x'_i)$ be the corresponding oracle answers.
- Then M outputs the assignment $(x_1 \rightarrow y_1, \dots, x_a \rightarrow y_a, x'_1 \rightarrow y'_1, \dots, x'_{b-a} \rightarrow y'_{b-a})$.

We can now define simple operations on adaptive lists. The length of an adaptive list is the number of queries it makes, and the composition of two adaptive lists is the adaptive list that first queries the first list, and then executes the second, which gets the queries made by the first as input.

Definition 4 (Operations on adaptive lists). *Let M be an $a \rightarrow b$ adaptive list. Then the length $|M|$ is defined as $|M| := b - a$.*

Let N be an $a \rightarrow b$ adaptive list, and M some $b \rightarrow c$ adaptive list. Then the composition $M \circ N$ is defined as the oracle Turing machine that upon input of an assignment X and a string $z \in \Sigma^$ outputs $M^{\mathcal{O}}(N^{\mathcal{O}}(X, z), z)$.*

Obviously, $M \circ N$ is an $a \rightarrow c$ adaptive list, and $|M \circ N| = |M| + |N|$.

We can now define the quantity $J(M|N)$ for adaptive lists M, N . Intuitively, $J(M|N)$ denotes the information that the queries made by M (when executed after N) contain about the auxiliary input $z^{\mathcal{O}}$ beyond what is already known from the queries made by N . Since the results to the queries made by M should be uniformly random if they are independent of $z^{\mathcal{O}}$, we define $J(M|N)$ as the quantity by which the conditional entropy of M 's queries given N 's queries and $z^{\mathcal{O}}$ is lower than the hypothetical value of $|M| \cdot \log \#Range$.

Definition 5 (Information of an adaptive list). *Let N be some $0 \rightarrow b$ adaptive list, and M some $b \rightarrow c$ adaptive list. Let further \mathcal{O} be the random oracle and z a random variable (where z does not need to be independent of \mathcal{O}).*

Then the information $J(M|N)$ is defined by

$$J(M|N) := |M| \cdot \log \#Range - H(M \circ N^{\mathcal{O}}(z)|N^{\mathcal{O}}(z), z).$$

(Note that $J(M|N)$ implicitly depends on the joint distribution of \mathcal{O} and z .)

We write short $J(M)$ for $J(M|\emptyset)$ where \emptyset is the adaptive list making no queries.

We now give two simple properties of the information $J(M|N)$: a chain rule and an upper bound in terms of the auxiliary input's length.

Lemma 6 (Chain rule for the information). *Let N be some $0 \rightarrow b$ adaptive list, M_2 some $b \rightarrow c$ adaptive list, and M_1 some $c \rightarrow d$ adaptive list. Then*

$$J(M_1 \circ M_2|N) \geq J(M_1|M_2 \circ N) + J(M_2|N).$$

Lemma 7 (Bounds for the information). *Let z be a random variable with finite range Z and $p := \log \#Z$. Let F be some $0 \rightarrow b$ adaptive list. Then $J(F) \leq p$.*

The proofs of these lemmas as well as of the subsequent ones are given in Appendices B–D.

Let $J_i := \max_F J(F)$ where F ranges over all adaptive lists of length $|F| = i$. Choose F_i such that $J(F_i) = J_i$. Consider the quantities $J_0, J_q, J_{2q}, \dots, J_{f+q}$ (assuming that q divides f). Since $J_0 \geq 0$ and $J_{f+q} \leq p$ by Lemma 7, there must be some $f' \leq f$ such that $J_{f'+q} - J_{f'} \leq \frac{pq}{f} =: \varepsilon$. Defining $F := F_{f'}$ we get $J(G|F) \leq J(G \circ F_{f'}) - J(F_{f'}) \leq J_{q+f'} - J_{f'} \leq \varepsilon$ by Lemma 6.

By definition of $J(G|F)$, this implies that the results of the queries made by G are only ε away from the maximum possible entropy $|G| \cdot \log \#Range$. This implies using a result from [Kul67] that the statistical distance between those query-results and the uniform distribution is bounded by $\sqrt{\varepsilon/2}$, even when given the results of the queries made by F and the auxiliary input $z^\mathcal{O}$. This is formally captured by the following lemma which is the core of the proof of Theorem 2.

Lemma 8 (The adaptive list F). *Let $f, q \geq 1$ be integers. Let z be a random variable with finite range Z (z may depend on the random oracle \mathcal{O}) and $p := \log \#Z$. Let U_n denote the uniform distributions on n -tuples over $\#Range$ (independent of z and \mathcal{O}).*

Then there is an adaptive list F with $|F| \leq f$, such that for any $|F| \rightarrow \min\{|F| + q, \#Domain\}$ adaptive list G , it is

$$\Delta(\nabla G \circ F^\mathcal{O}(z), F^\mathcal{O}(z), z; U_{|G|}, F^\mathcal{O}(z), z) \leq \sqrt{\frac{pq}{2f}}.$$

Here $\nabla G \circ F$ denotes the oracle Turing machine that behaves as $G \circ F$ but only outputs the oracle answers that G got (instead of also outputting G 's input and G 's queries). More formally, if $G \circ F^\mathcal{O}(z) = (x_1 \rightarrow y_1, \dots, x_{|F|+|G|} \rightarrow y_{|F|+|G|})$, we have $\nabla G \circ F^\mathcal{O}(z) = (y_{|F|+1}, \dots, y_{|F|+|G|})$.

Using Lemma 8, proving the main Theorem 2 is easy. For some adversary A let $\mu := \Delta(A^\mathcal{O}(z^\mathcal{O}); A^\mathcal{P}(z^\mathcal{O}))$. By fixing the worst-case random-tape, we can make the adversary A deterministic. Then A 's output depends only on its input $z^\mathcal{O}$ and the answers to its oracle queries. So if we let A just output the queries it made, the statistical distance μ does not diminish. Further, if we give the presampled queries $S^\mathcal{O}$ as an additional input to A , we can assume A to make exactly q distinct queries, and not to query any $x \in S^\mathcal{O}$. But then A fulfils the definition of an adaptive list, so by Lemma 8 we have $\mu \leq \sqrt{\frac{pq}{2f}}$, which proves Theorem 2.

We give the full details of the proof in Appendix E.

An interesting question is whether the bound $\sqrt{pq/2f}$ on the statistical distance Δ achieved by Theorem 2 is tight. In particular, the bound falls only sublinearly with f , while we were unable to find a counterexample where Δ did not fall exponentially with f . So a tighter bound may be possible. However, this would need to use new proof techniques, since the approach in this paper uses an averaging argument that will at best give a bound that falls polynomially in f (cf. the computation of $J_{f'+q} - J_{f'}$ below Lemma 7 above.)

Finally, for convenience we state an asymptotic version of Theorem 2 that hides the exact bounds achieved there:

Corollary 9 (Lazy sampling with auxiliary input, asymptotic version). *For any superpolynomial function f and any polynomial q and oracle function z into strings of polynomial length, there is an oracle*

function S , such that $S^\mathcal{O}$ is an assignment of length at most f , so that for any probabilistic oracle Turing machine A making at most q queries, the following random variables are statistically indistinguishable:

$$A^\mathcal{O}(1^k, z^\mathcal{O}) \quad \text{and} \quad A^\mathcal{P}(1^k, z^\mathcal{O}).$$

Here \mathcal{P} is the random oracle with presampling $S^\mathcal{O}$.

(In this corollary, \mathcal{O} , z , and S depend implicitly on the security parameter k .)

Proof. Immediate from Theorem 2. □

3 Example: one-wayness of the random oracle

To give a first impression on how the lazy sampling technique is used in the random oracle model with oracle-dependent auxiliary input, we show a very simple result: If we let $f := \mathcal{O}$, then f is a one-way function.

In Appendix F, as a second example we show that $f := \mathcal{O}$ is given-preimage collision-resistant.

Lemma 10 (The random oracle is one-way). *Let $g := \mathcal{O}$ where \mathcal{O} denotes the random oracle. Then g is a one-way function in the random oracle model with oracle-dependent auxiliary input.*

More formally, for any probabilistic polynomial-time oracle Turing machine A and any oracle function z into strings of polynomial length, the following probability is negligible (in k):

$$\text{Adv}_A := P(x \xrightarrow{\$} \text{Domain}, x' \leftarrow A^\mathcal{O}(1^k, z^\mathcal{O}, \mathcal{O}(x)) : \mathcal{O}(x') = \mathcal{O}(x))$$

(In this lemma, \mathcal{O} , Domain , f , and z depend implicitly on the security parameter k .)

We present this proof in some detail, to illustrate how Theorem 2 or Corollary 9 can be used. Since these steps are almost identical in most situations, knowledge of this proof facilitates understanding of the proofs given later on.

Proof. Let $f := \min\{\sqrt{\#\text{Range}}, \sqrt{\#\text{Domain}}\}$. Let \tilde{A} be the oracle Turing machine that chooses a random x from Domain_k , then let $A(1^k, z^\mathcal{O}, \mathcal{O}(x))$ choose x' , and outputs 1 if and only if $\mathcal{O}(x') = \mathcal{O}(x)$. Then $\text{Adv}_A = P(\tilde{A}^\mathcal{O}(1^k, z^\mathcal{O}) = 1)$.

Since A is polynomial-time, \tilde{A} makes only a polynomial number of queries, so Corollary 9 applies to \tilde{A} , hence $\tilde{A}^\mathcal{O}(1^k, z^\mathcal{O})$ and $\tilde{A}^\mathcal{P}(1^k, z^\mathcal{O})$ are statistically indistinguishable (where \mathcal{P} is the random oracle with presampling $S^\mathcal{O}$, and S is as in Corollary 9). Then consider the following game:

$$\text{Game 1:} \quad x \xrightarrow{\$} \text{Domain}, x' \leftarrow A^\mathcal{P}(1^k, z^\mathcal{O}, \mathcal{P}(x)) : \mathcal{P}(x') = \mathcal{P}(x).$$

We call the probability that the last expression evaluates to true (i.e., that $\mathcal{P}(x') = \mathcal{P}(x)$) the advantage Adv_1 of the game. Since Adv_1 is the probability that $\tilde{A}^\mathcal{P}(1^k, z^\mathcal{O})$ outputs 1, $|\text{Adv}_A - \text{Adv}_1|$ is negligible.

(This step probably occurs at the beginning of virtually all proofs that use Theorem 2 or Corollary 9. We are now in the situation that with at most f exceptions, the oracle query $\mathcal{P}(x)$ returns a fresh random value, and can use standard techniques based on lazy sampling.)

We now modify A in the following way resulting in a machine A_2 : A_2 expects an assignment S as an additional argument. Whenever A would query the random oracle \mathcal{P} with a value x , A_2 first checks if $x \in S$. If so, A returns the image of x under S . Otherwise, A_2 queries its oracle. Then consider the following game:

$$\text{Game 2:} \quad x \xrightarrow{\$} \text{Domain}, y \leftarrow \mathcal{P}(x), x' \leftarrow A_2^\mathcal{P}(1^k, z^\mathcal{O}, y, S^\mathcal{O}) : y = \mathcal{P}(x')$$

Obviously, $\text{Adv}_1 = \text{Adv}_2$.

Since for some $x \notin S^\mathcal{O}$ (which happens with probability at least $1 - f/\#Domain$), the oracle \mathcal{P} returns a random $y \in \#Range$, the probability that $y \in \text{im } S^\mathcal{O}$ is at most $f/\#Domain + f/\#Range$. Furthermore, if $x' \in S^\mathcal{O}$ but $y \notin \text{im } S^\mathcal{O}$, the predicate $y = \mathcal{P}(x')$ will be false.

So $|\text{Adv}_2 - \text{Adv}_3| \leq P(y \in \text{im } S^\mathcal{O})$ is negligible for the following game 3:

$$\begin{aligned} \text{Game 3: } \quad & x \xleftarrow{\$} Domain, y \leftarrow \mathcal{P}(x), x' \leftarrow A_2^{\mathcal{P}}(1^k, z^\mathcal{O}, y, S^\mathcal{O}) : \\ & \text{if } (x' \in S^\mathcal{O}) \text{ then false else } y = \mathcal{P}(x'). \end{aligned}$$

Note that in game 3, A_2 never queries \mathcal{P} at a position in $S^\mathcal{O}$. Furthermore, the query $\mathcal{P}(x')$ is only executed if $x' \notin S$. So \mathcal{P} is only queried at a position in $S^\mathcal{O}$, if $x \in S^\mathcal{O}$, which has probability at most $f/\#Domain$. But when queried at positions outside $S^\mathcal{O}$, \mathcal{P} behaves like a normal random oracle (i.e., without presampling). We can therefore replace the oracle \mathcal{P} by a random oracle \mathcal{R} (independent of \mathcal{O}):

$$\begin{aligned} \text{Game 4: } \quad & x \xleftarrow{\$} Domain, y \leftarrow \mathcal{R}(x), x' \leftarrow A_2^{\mathcal{R}}(1^k, z^\mathcal{O}, y, S^\mathcal{O}) : \\ & \text{if } (x' \in S^\mathcal{O}) \text{ then false else } y = \mathcal{R}(x'). \end{aligned}$$

Then $|\text{Adv}_3 - \text{Adv}_4| \leq P(x \in S^\mathcal{O})$ is negligible.

(We have now succeeded in completely separating the oracle from the auxiliary input; \mathcal{R} is independent from $(z^\mathcal{O}, S^\mathcal{O})$. From here on, the proof is a standard proof of one-wayness of the random oracle. Note however, that $S^\mathcal{O}$ has a length that may be superpolynomial, so A_2 is not polynomially bounded any more. In our case, this does not pose a problem, since we only use the fact that A_2 uses a polynomial number of queries. In proof that additionally need computational assumptions, one might need additional tools which we present in Section 4.)

Consider the following game:

$$\begin{aligned} \text{Game 5: } \quad & x \xleftarrow{\$} Domain, y \xleftarrow{\$} Range, x' \leftarrow A_2^{\mathcal{R}}(1^k, z^\mathcal{O}, y, S^\mathcal{O}) : \\ & \text{if } (x' \in S^\mathcal{O}) \text{ then false else } y = \mathcal{R}(x'). \end{aligned}$$

Since A was polynomially bounded, there is a polynomial q bounding the number of oracle queries of A_2 . The probability that A_2 queries \mathcal{R} at position x is therefore at most $q/\#Domain$ (since x is randomly chosen and never used). Furthermore, game 4 and 5 only differ if A_2 queries \mathcal{R} at position x . So $|\text{Adv}_4 - \text{Adv}_5| \leq q/\#Domain$ is negligible.

Since A_2 makes at most q queries, the probability that one of these returns y is at most $q/\#Domain$. If x' returns a value x' it has not queried before, the probability that $y = \mathcal{R}(x')$ is at most $1/\#Domain$. So $\text{Adv}_5 \leq (q - 1)/\#Domain$ is negligible.

Collecting the bounds shown so far, we see that Adv_A is negligible. \square

In the preceding proof, we have only verified that the advantage of the adversary is negligible. By using Theorem 2 instead of Corollary 9 and computing the exact bounds, we even get $\text{Adv}_A \in 2^{-\Omega(k)}$ which is essentially the same bound as given in [Imp96] and [GGKT05].

4 Security amplification

When using a random oracle with presampling, reduction proofs sometimes run into situations where the adversary gets the presampling $S^\mathcal{O}$ as an input. Unfortunately, this presampling is usually of superpolynomial size, so the resulting adversary is not polynomial-time any more and a reduction to complexity assumptions relative to polynomial-time adversaries is bound to fail. (E.g., in the proof of Lemma 10 the

adversary A_2 was not polynomial-time in the security parameter any more. In that case however, this did not matter since we only used the polynomial bound on the number of queries made by A_2 .) An example of a situation where superpolynomial adversaries occur, and do pose a problem, is the proof that RSA-OAEP is secure with respect to oracle-dependent auxiliary input, cf. Section 5. One possibility is simply to assume a stronger security notion; in the case of RSA-OAEP one could use, e.g., the RSA-assumption against quasi-polynomial adversaries.

Fortunately, there is another way which allows to use standard assumptions (i.e., with respect to polynomial-time adversaries) in many cases. We show that for some kinds of security notions, security against polynomial-time adversaries implies security against adversaries with f -bounded runtime, where f is a suitably chosen *superpolynomial* function. Using this fact we can finish our reduction proof: Corollary 9 guarantees that for any superpolynomial function f' , we can replace the random oracle by a random oracle with presampling of length f' . We then choose f' to be the largest function such that all adversaries constructed in our proof are still f -bounded. Such an f' is still superpolynomial, so Corollary 9 applies. On the other hand, the resulting adversaries are efficient enough for the reduction to go through. This proof method is applied in Section 5 to show the security of RSA-OAEP.

Instead of giving a general proof of our *security amplification* technique, we give here a proof for the security notion of partial-domain one-wayness (Definition 11). The proof can easily be adapted to other security notions (in particular, our proof does not exploit how the advantage Adv is defined for this particular notion). In Appendix G we give a more general characterisation of the security notions for which security amplification is possible.⁵

Definition 11 (Partial-domain one-way). *A family of 1-1 functions $f_{pk} : B \times C \rightarrow D$ is partial-domain one-way, if for any nonuniform polynomial-time adversary A , the following advantage is negligible:*

$$\text{Adv}_{A,k} := P(pk \leftarrow K(1^k), (s, t) \xleftarrow{\$} B \times C, y \leftarrow f_{pk}(s, t), s' \leftarrow A(1^k, y) : s = s').$$

Here K denotes the index generation algorithm for the family f_{pk} of functions. Partial-domain one-way against f -bounded adversaries for some function f is defined analogously.

(In this definition, B , C , and D depend implicitly on the security parameter k .)

Lemma 12 (Security amplification for partial-domain one-wayness). *Let the family f_{pk} be partial-domain one-way (against polynomial-time nonuniform adversaries). Then there exists a superpolynomial function f such that f_{pk} is partial-domain one-way against f -bounded nonuniform adversaries.*

Proof. For $n \in \mathbb{N}$ let $\mu_n(k) := \max_{|A| \leq n} (\text{Adv}_{A,k})$ where A goes over all circuits of size at most n . Assume there was a polynomial p with integer coefficients (an integer polynomial for short) such that $\mu_{p(k)}(k)$ is not negligible in k . Then there is a nonuniform adversary A consisting of circuits A_k with $|A_k| \leq p(k)$ such that $\text{Adv}_{A,k} \geq \mu_{p(k)}(k)$ is non-negligible. Since A is polynomial-time, this contradicts the assumption that the f_{pk} are partial-domain one-way. Hence $\mu_p := \mu_{p(k)}(k)$ is negligible for all integer polynomials p .

We say that a function μ asymptotically dominates a function ν if for all sufficiently large k we have $\mu(k) \geq \nu(k)$. [Bel02] proves that for any countable set S of negligible functions, there is a negligible function μ^* that asymptotically dominates all $\mu \in S$.

Therefore, there is a negligible function μ^* , that asymptotically dominates μ_p for every integer polynomial p .

Let $f(k) := \max\{p \in \mathbb{N} : \mu_p(k) \leq \mu^*(k)\}$. Then $\mu_{f \times (k)}(k) \leq \mu^*(k)$ is negligible. So for any nonuniform f -bounded adversary A the advantage $\text{Adv}_{A,k}$ is negligible. Furthermore, we can show that

⁵ This includes one-wayness, partial-domain one-wayness, IND-CPA, IND-CCA2, black-box stand-alone security of function evaluations, UC (where the amplification concerns the running time of the environment), black-box zero-knowledge, arguments, black-box arguments of knowledge.

f is superpolynomial. Assume this is not the case. Then there an integer polynomial p such that $p > f$ infinitely often. But then $\mu_p > \mu^*$ holds infinitely often, in contradiction to the choice of μ^* (by definition of f). Thus f is superpolynomial. \square

5 OAEP encryption

In [FOPS04] it was shown that RSA-OAEP (introduced by [BR95]) is secure in the random oracle model under the RSA-assumption. However, their proof only covers the case that no auxiliary input is given (or at least that the auxiliary input is not oracle-dependent). In this section, we extend this result to encompass the case of oracle-dependent auxiliary input. On one hand, this gives a nontrivial example of the application of the lazy sampling technique in combination with the security amplification technique. On the other hand, this result is important in its own light, since it gives evidence that RSA-OAEP may be secure with respect to an auxiliary input, even when the random oracle has been instantiated with a fixed function.

To read this section, it is helpful to have at least basic knowledge of the OAEP construction and its proof from [FOPS04]. We recommend [FOPS04] as an introduction.

Theorem 13 (OAEP is secure with respect to oracle-dependent auxiliary input). *Let f_{pk} be a family of partial-domain one-way trapdoor 1-1 functions (with the property, that the elements of the domain of f_{pk} consist of two components each of superlogarithmic length).*

Then the OAEP encryption scheme based on f_{pk} is IND-CCA2 secure in the random oracle model with oracle-dependent auxiliary input.

This theorem implies that RSA-OAEP is IND-CCA2 secure under the RSA-assumption with respect to oracle-dependent auxiliary input, since in [FOPS04] it is shown that the RSA family of functions is partial-domain one-way.

At this point, we only describe on a high level, in what points our proof differs from the proof in [FOPS04]. In Appendix H, we reproduce the full proof of [FOPS04] and highlight our changes for comparison.

In [FOPS04], the proof has roughly the following outer form: First, the IND-CCA2 game is formulated for the special case of the OAEP encryption scheme. Then the game is rewritten in a series of small changes, to finally yield a plaintext extractor. If the first game had a non-negligible success probability (i.e., the OAEP encryption scheme was not IND-CCA2 secure), the plaintext extractor had, for some random ciphertext $f_{pk}(s, t)$, a non-negligible probability of outputting s . This breaks the assumption that f_{pk} is partial-domain one-way.

Our proof starts with the same game, except that the adversary now has access to an oracle-dependent auxiliary input $z^\mathcal{O}$. Then we can use Corollary 9 to replace the random oracle \mathcal{O} by a random oracle \mathcal{P} with presampling $S^\mathcal{O}$ of a yet to determine superpolynomial subexponential length f (similar to the first step in the proof of Lemma 10).⁶ In this new situation, for randomly chosen $x \in \text{Domain}$, with overwhelming probability, the oracle response $\mathcal{P}(x)$ is uniformly distributed. Using this fact, most of the rewriting steps in the sequence of games are the same as in [FOPS04], sometimes with slightly larger errors to account for the possibility of randomly choosing an $x \in S^\mathcal{O}$. Only in the construction of the plaintext extractor additional care has to be taken. Here the original argument uses that the answer to an oracle query can be assumed to be random if the adversary has not yet queried it. From this they conclude any ciphertext the decryption oracle would accept can also be decrypted by encrypting and comparing all oracles queries that have been made by the adversary so far. This does not hold any more since the auxiliary input $z^\mathcal{O}$ can supply additional information on the presampled queries $S^\mathcal{O}$. We thus have to change the plaintext extractor not only to encrypt all oracle queries but also all presampled queries $S^\mathcal{O}$. Therefore the plaintext extractor is not polynomial-time anymore, but instead a nonuniform machine with

⁶ The actual proof uses Theorem 2, but the asymptotic version is sufficient.

running time $p(f)$ for some polynomial p . We consequently do not directly obtain a contradiction to the partial-domain one-wayness, since therefore the plaintext extractor would have to be polynomial-time.

However, we can use the security amplification technique. By Lemma 12, there is a superpolynomial function f' such that f_{pk} is partial-domain one-way even against nonuniform f' -bounded adversaries. By choosing f small enough (but still superpolynomial), it is $p(f) \leq f'$, so the plaintext extractor is f' -bounded, and the fact that the plaintext extractor returns s for some $f_{pk}(s, t)$ with non-negligible probability is a contradiction.

6 Open questions

We have shown how to apply the lazy sampling technique to the case of oracle-dependent auxiliary input. Going further, the following open problems come to mind which we give in the form of several conjectures.

In Corollary 9, we require the length f of the presampling to be superpolynomial. This makes reduction proofs more difficult, in particular it necessitates the use of the security amplification technique. It would be preferable to be able to use a polynomial length f as stated by the following conjecture:

Conjecture 14 (Lazy-sampling with polynomial presampling) *For all polynomials p and q , there is a polynomial f , such that for any oracle function z into strings of length at most p , there is an oracle function S , where $S^\mathcal{O}$ is an assignment of length at most f , so that for any probabilistic oracle Turing machine A that makes at most q oracle queries the following random variables are statistically indistinguishable:*

$$A^\mathcal{O}(1^k, z^\mathcal{O}) \quad \text{and} \quad A^\mathcal{P}(1^k, z^\mathcal{O}).$$

Here \mathcal{P} is the random oracle with presampling $S^\mathcal{O}$.

(In this conjecture, \mathcal{O} , \mathcal{P} , z , and S depend implicitly on the security parameter k .)

Note that this is the strongest possible order of quantifier for which we can still expect the conjecture to be true. If f was independent of the length p of the auxiliary input, the conjecture could be disproven by setting $z^\mathcal{O} := \mathcal{O}(1), \dots, \mathcal{O}(f+1)$. If we try to choose f independent of the number of queries q , we can give the following counterexample: Let $z^\mathcal{O} := z_1, \dots, z_k$ with $z_i := \bigoplus_{\mu=1}^{2^i} \mathcal{O}(i)$. For $f < 2^l$ we can use 2^l queries to check the value of z_l . But if one of the queries $\mathcal{O}(1), \dots, \mathcal{O}(2^l)$ is not presampled, $\bigoplus_{\mu=1}^{2^i} \mathcal{P}(i)$ will be random and with high probability not equal to z_l .

A further interesting point is whether oracle-dependent auxiliary input does give actual computational power (beyond what is possible with a simple auxiliary input). It seems that information about some random data that the adversary could generate himself does not help more than just some auxiliary information. In that light, we state the following conjecture:

Conjecture 15 (Oracle-dependent auxiliary input does not give computational power) *For any polynomial-time adversary A , and for every oracle function z into strings of polynomial length there is a polynomial-time simulator S and an auxiliary input z' of polynomial length such that for any polynomial-time (or even unbounded) ITM B ,*

$$\langle A^\mathcal{O}(1^k, z^\mathcal{O}), B(1^k) \rangle \quad \text{and} \quad \langle S(1^k, z'), B(1^k) \rangle$$

are computationally (or even statistically) indistinguishable.

Here $\langle A, B \rangle$ denotes the output of B in an interaction of A and B .

(In this conjecture, \mathcal{O} , z , and z' depend implicitly on the security parameter k .)

A weaker, but still interesting variant would be to quantify over the ITM B first (i.e., to choose the simulator S in dependence of B). A stronger variant would require z and z' to be of same length.

Note that this conjecture is weaker than Conjecture 14. A random oracle with a presampling of polynomial length can be simulated using an auxiliary input of polynomial length (that simply contains the presampling).

Finally, it would be helpful to generalise the lazy sampling technique to more general types of random oracles, like random permutations (with or without access to the inverse), the generic group model, or ideal ciphers. If the following conjecture were true, the proof methods presented in this paper could be applied to those settings, too.

Conjecture 16 (Other types of oracles) *Let μ be a distribution on functions. Let \mathcal{O} be a random function chosen according to μ . Let S be an assignment. We say that f matches S if for any $x \in S$, $f(x)$ is the image of x under S . A random oracle \mathcal{P} with presampling S is a random function chosen according to μ under the condition that \mathcal{P} matches S .*

Then Corollary 9 holds in this extended setting.

(In this conjecture, μ , \mathcal{O} , \mathcal{P} , and S depend implicitly on the security parameter k .)

Note that Corollary 9 is a special case of this conjecture. If the Definition 3 of adaptive lists could be generalised to this case, so that Lemmas 6 and 7 still apply, the conjecture could be shown similarly to Corollary 9.

Most desirable would be a combination of Conjectures 14 and 16, of course.

Acknowledgements. We thank Michael Backes, Dennis Hofheinz, Yuval Ishai, Jörn Müller-Quade, Hoeteck Wee, and Jürg Wullschleger for valuable discussions. We further thank the anonymous referees for helpful comments.

A An auxiliary lemma

Lemma 17 (Conditional entropy). *Let A, B, C, U be discrete random variables.*

(i) $H(A, B|C) = H(A|B, C) + H(B|C)$.

(ii) $H(A|C) \leq H(A|B) + H(B|C)$.

(iii) *If A has range M , and U is the uniform distribution on M (and independent of B), and $H(A|B) \geq \log \#M - \varepsilon$, then $\Delta(A, B; U, B) \leq \sqrt{\varepsilon/2}$.*

Proof. Equality (i) is shown as

$$\begin{aligned} H(A, B|C) &= - \sum_{a,b,c} P(A = a, B = b, C = c) \log P(A = a, B = b|C = c) \\ &= - \sum_{a,b,c} P(A = a, B = b, C = c) \log P(A = a|B = b, C = c)P(B = b|C = c) \\ &= H(A|B, C) + H(B|C). \end{aligned}$$

Inequality (ii) follows from (i) by

$$H(A|C) \leq H(A, B|C) = H(A|B, C) + H(B|C) \leq H(A|C) + H(B|C).$$

We now show (iii). Let H_{KL} denote the Kullback-Leibler distance or relative entropy, defined as $H_{KL}(A; B) := \sum_x P(A = x) \log \frac{P(A=x)}{P(B=x)}$.

Then for any random variable X with range M , it is

$$\begin{aligned} H_{KL}(X; U) &= \sum_{x \in M} P(X = x) \log(P(X = x) \# M) \\ &= \sum_{x \in M} (P(X = x) \log P(X = x)) + \log \# M \\ &= \log \# M - H(X). \end{aligned}$$

Further, in [Kul67] it is shown that $\Delta(X; U) \leq \sqrt{H_{KL}(X; U)/2}$, so

$$\Delta(X; U) \leq \sqrt{(\log \# M - H(X))/2}. \quad (5)$$

Let A_b denote the random variable A conditioned on the event $B = b$. (I.e., $P(A_b = x) = P(A = x | B = b)$.) Then

$$\begin{aligned} H(A|B) &= - \sum_{a,b} P(A = a, B = b) \log P(A = a | B = b) \\ &= - \sum_b P(B = b) \sum_a P(A_b = a) \log P(A_b = a) = \sum_b P(B = b) H(A_b). \quad (6) \end{aligned}$$

Therefore,

$$\begin{aligned} \Delta(A, B; U, B) &= \sum_{x,b} |P(A = x, B = b) - P(U = x, B = b)| \\ &= \sum_{x,b} P(B = b) |P(A = x | B = b) - P(U = x | B = b)| \\ &\stackrel{(*)}{=} \sum_b P(B = b) \Delta(A_b; U) \\ &\stackrel{(5)}{\leq} \sum_b P(B = b) \sqrt{\frac{1}{2}(\log \# M - H(A_b))} \\ &\stackrel{(**)}{\leq} \sqrt{\frac{1}{2} \sum_b P(B = b) (\log \# M - H(A_b))} \\ &\stackrel{(6)}{=} \sqrt{\frac{1}{2}(\log \# M - H(A|B))} \leq \sqrt{\varepsilon/2}. \end{aligned}$$

Here $(*)$ uses the fact that U and B are independent, and $(**)$ is an application of Jensen's inequality. \square

B Proof of Lemma 6

Proof. Using the definition of $J(M|N)$, $|M|$ and Lemma 17 (ii), we get

$$\begin{aligned} J(M_1 \circ M_2 | N) &= |M_1| \log \# Range + |M_2| \log \# Range \\ &\quad - H(M_1 \circ M_2 \circ N^{\mathcal{O}}(z) | N^{\mathcal{O}}(z), z) \\ &\geq |M_1| \log \# Range + |M_2| \log \# Range \\ &\quad - H(M_1 \circ M_2 \circ N^{\mathcal{O}}(z) | M_2 \circ N^{\mathcal{O}}(z), z) \\ &\quad - H(M_2 \circ N^{\mathcal{O}}(z), z | N^{\mathcal{O}}(z), z) \\ &= J(M_1 | M_2 \circ N) + J(M_2 | N). \quad \square \end{aligned}$$

C Proof of Lemma 7

Proof. Let G be an arbitrary $b \rightarrow \#Domain$ adaptive list. (Such an adaptive list always exists, it simply has to query the oracle at all positions not already contained in its input.)

Since $G \circ F$ is an $0 \rightarrow \#Domain$ adaptive list, it outputs the complete random oracle \mathcal{O} . So $H(G \circ F^\mathcal{O}(z)|z) \geq H(\mathcal{O}|z)$.

Furthermore, by Lemma 17 (i), it is

$$H(G \circ F^\mathcal{O}(z)|z) = H(G \circ F^\mathcal{O}(z)|F^\mathcal{O}(z), z) + H(F^\mathcal{O}(z)|z).$$

Combining these inequalities, we get

$$H(F^\mathcal{O}(z)|z) \geq H(\mathcal{O}|z) - H(G \circ F^\mathcal{O}(z)|F^\mathcal{O}(z), z). \quad (7)$$

Since the random oracle is a uniformly distributed $Domain \rightarrow Range$ function, $H(\mathcal{O}) = \#Domain \log \#Range$. Since z has range Z , $H(z) \leq \log \#Z = p$. So

$$H(\mathcal{O}|z) = H(\mathcal{O}, z) - H(z) \geq H(\mathcal{O}) - H(z) \geq \#Domain \log \#Range - p. \quad (8)$$

Let $\nabla G \circ F$ denote the oracle Turing machine, that behaves as does $G \circ F$, but only outputs the oracle answers that G got (instead of also outputting G 's input and G 's queries). More formally, if $G \circ F^\mathcal{O}(z) = (x_1 \rightarrow y_1, \dots, x_{|F|+|G|} \rightarrow y_{|F|+|G|})$, it is $\nabla G \circ F^\mathcal{O}(z) = (y_{|F|+1}, \dots, y_{|F|+|G|})$.

Given the inputs of a deterministic oracle Turing machine and the answers to all its oracle queries (in the correct order, but without the positions at which the oracle has been queried), one can simulate the oracle Turing machine and calculate its output. So from $\nabla G \circ F^\mathcal{O}(z)$, $F^\mathcal{O}(z)$ and z one can calculate $G \circ F^\mathcal{O}(z)$, thus $H(\nabla G \circ F^\mathcal{O}(z)|F^\mathcal{O}(z), z) \geq H(G \circ F^\mathcal{O}(z)|F^\mathcal{O}(z), z)$.

Since $\nabla G \circ F^\mathcal{O}(z)$ outputs $\#Domain - |F|$ elements of $\#Range$, it is $H(\nabla G \circ F^\mathcal{O}(z)|F^\mathcal{O}(z), z) \leq (\#Domain - |F|) \log \#Range$, and thus also

$$H(G \circ F^\mathcal{O}(z)|F^\mathcal{O}(z), z) \leq (\#Domain - |F|) \log \#Range.$$

Combining this with (7) and (8), we get

$$H(F^\mathcal{O}(z)|z) \geq |F|(\log \#Range) - p.$$

So finally, $J(F) = |F| \log \#Range - H(F^\mathcal{O}(z)|z) \leq p$. □

D Proof of Lemma 8

Proof. For any $i \leq \#Domain$, let $J_i := \max_F J(F)$ where the maximum goes over all $0 \rightarrow i$ adaptive lists F . Note that the maximum exists, since there are only finitely many different $0 \rightarrow i$ adaptive lists F .⁷ For $i > \#Domain$, set $J_i := J_{\#Domain}$. By Lemma 7, $J_i \leq p$ for any $i \geq 0$.

We first show that there is an $f' \leq f$ such that $J_{f'+q} - J_{f'} \leq \frac{pq}{f}$. Assume this was not the case. Then (using $J_0 \geq 0$)

$$J_{\lfloor \frac{f}{q} \rfloor q + q} \geq \sum_{\mu=0}^{\lfloor \frac{f}{q} \rfloor} J_{q\mu+q} - J_{q\mu} > \left(\lfloor \frac{f}{q} \rfloor + 1 \right) \frac{pq}{f} \geq p,$$

⁷ Here, we consider two adaptive list as equal, if they implement the same function. Of course, there are infinitely many Turing machines implementing the same function. Then, since input and output domain are finite (these are assignments of a fixed length and a string z from the finite set Z), and the oracle is chosen from a finite set, there are only a finite number of such functions.

which is a contradiction to $J_i \leq p$.

Let F be a $0 \rightarrow f'$ adaptive list, such that $J(F) = J_{f'}$. Since $f' \leq f$ it is $|F| \leq f$.

For any $f' \rightarrow \min\{f' + q, \#Domain\}$ adaptive list G , it is $J(G \circ F) \geq J(G|F) + J(F)$ by Lemma 6. Further, $J(G \circ F) \leq J_{f'+q}$ (if $f' + q > \#Domain$, we use $J_{f'+q} = J_{\#Domain}$). So

$$J(G|F) \leq J(G \circ F) - J(F) \leq J_{f'+q} - J_{f'} \leq \frac{pq}{f}. \quad (9)$$

By Definition 5, it is

$$H(G \circ F^\mathcal{O}(z)|F^\mathcal{O}(z), z) = |G| \log \#Range - J(G|F).$$

Since the output of $\nabla G \circ F$ is contained in the output of $G \circ F$, we have

$$H(\nabla G \circ F^\mathcal{O}(z)|F^\mathcal{O}(z), z) \leq H(G \circ F^\mathcal{O}(z)|F^\mathcal{O}(z), z) = |G| \log \#Range - J(G|F).$$

Furthermore, the output of $\nabla G \circ F$ consists of $|G|$ elements of $\#Range$. So by Lemma 17 (iii), it follows

$$\Delta(\nabla G \circ F^\mathcal{O}(z), F(z), z; U_{|G|}, F^\mathcal{O}(z), z) \leq \sqrt{\frac{J(G|F)}{2}} \stackrel{(9)}{\leq} \sqrt{\frac{pq}{2f}}. \quad \square$$

E Proof of Theorem 2

Proof. If $q = 0$ the machine A never queries the oracle, so the inequality is trivially fulfilled. We can therefore assume $q \geq 1$.

Now let F be the adaptive list from Lemma 8 with $z := z^\mathcal{O}$. It is $f' := |F| \leq f$. Define $S^\mathcal{O}$ as $F^\mathcal{O}(z^\mathcal{O})$. Then $S^\mathcal{O}$ is an assignment of length $f' \leq f$.

To show the theorem, let A be any oracle Turing machine making at most q queries. Then set

$$\mu := \Delta(A^\mathcal{O}(z^\mathcal{O}); A^\mathcal{P}(z^\mathcal{O})).$$

We have to show that $\mu \leq \sqrt{\frac{pq}{2f}}$.

By fixing that random tape of A for which the above statistical distance is maximal, we can assume A to be deterministic. We now construct a new deterministic oracle Turing machine A_1 that takes an assignment S and a string z as input, and then runs A on input z , with the following difference: Whenever A queries the oracle at a position x_i , and $x_i \rightarrow y_i$ is contained in S , A receives the answer y_i . In all other cases A 's query is passed through to the oracle A_1 has access to. Finally, A_1 outputs A 's output.

Obviously, A_1 still makes at most q oracle queries. Furthermore, if A_1 has access to a random oracle \mathcal{R} independent of \mathcal{O} , A_1 effectively simulates \mathcal{P} for A by answering with the values prescribed by $S^\mathcal{O}$ if applicable, and with fresh randomness from \mathcal{R} otherwise. More formally, the random variables $A^\mathcal{P}(z^\mathcal{O})$ and $A_1^\mathcal{R}(S^\mathcal{O}, z^\mathcal{O})$ are identically distributed.

Now consider the case that A_1 has access to the random oracle \mathcal{O} . Since by construction of $S^\mathcal{O}$, for any $x \rightarrow y$ in the assignment $S^\mathcal{O}$, it is $\mathcal{O}(x) = y$, the simulated A always receives the answers \mathcal{O} would have given. So the random variables $A^\mathcal{O}(z^\mathcal{O})$ and $A_1^\mathcal{O}(S^\mathcal{O}, z^\mathcal{O})$ are identically distributed.

It follows

$$\Delta(A_1^\mathcal{O}(S^\mathcal{O}, z^\mathcal{O}); A_1^\mathcal{R}(S^\mathcal{O}, z^\mathcal{O})) = \mu. \quad (10)$$

We now replace A_1 by a new deterministic oracle Turing machine A_2 with the following changes: A_2 never calls the oracle twice with the same argument. A_2 never queries the oracle at any position that occurs in the assignment S given as argument. And A_2 makes exactly q' queries to the oracle, where

$q' := \min\{q, \#Domain - f'\}$. (A_2 does not need to make more than $\#Domain - f'$ queries, because S has length f' , so f' positions of the random oracle are already known.) Then (10) also holds for A_2 .

Now we construct a new deterministic oracle Turing machine G that behaves as follows: On input S, z , the machine G runs A_2 on input S, z . Let x_1, \dots, x_n be the positions of the oracle queries the simulated A_2 makes (in the correct order), and let y_i be the corresponding answers. Finally, G outputs the assignment $(S, x_1 \rightarrow y_1, \dots, x_n \rightarrow y_n)$. Then G is an $f' \rightarrow f' + q'$ adaptive list.

Since A_2 is deterministic, given its input and the answers y_i to its oracle queries, we can calculate its output. Therefore

$$\Delta(G^\mathcal{O}(S^\mathcal{O}, z^\mathcal{O}), S^\mathcal{O}, z^\mathcal{O}; G^\mathcal{R}(S^\mathcal{O}, z^\mathcal{O}), S^\mathcal{O}, z^\mathcal{O}) \geq \Delta(A_2^\mathcal{O}(S^\mathcal{O}, z^\mathcal{O}); A_2^\mathcal{R}(S^\mathcal{O}, z^\mathcal{O})) = \mu.$$

Let now ∇G again be defined as in the proof of Lemma 7. As before, we can calculate the output of ∇G from its inputs and the results of its oracle queries (even without the positions of these queries). So it also holds

$$\Delta(\nabla G^\mathcal{O}(S^\mathcal{O}, z^\mathcal{O}), S^\mathcal{O}, z^\mathcal{O}; \nabla G^\mathcal{R}(S^\mathcal{O}, z^\mathcal{O}), S^\mathcal{O}, z^\mathcal{O}) \geq \mu.$$

Since ∇G only outputs the results of its oracle queries and never queries the oracle twice at the same position, and \mathcal{R} is independent of \mathcal{O} , the output of $\nabla G^\mathcal{R}$ consists of $|G|$ uniformly and independently chosen elements of $\#Range$ that are independent of $S^\mathcal{O}$ and $z^\mathcal{O}$. So the random variables $(\nabla G^\mathcal{R}(S^\mathcal{O}, z^\mathcal{O}), S^\mathcal{O}, z^\mathcal{O})$ and $(U_{|G|}, S^\mathcal{O}, z^\mathcal{O})$ are identically distributed. It follows

$$\Delta(\nabla G^\mathcal{O}(S^\mathcal{O}, z^\mathcal{O}), S^\mathcal{O}, z^\mathcal{O}; U_{|G|}, S^\mathcal{O}, z^\mathcal{O}) \geq \mu, \quad (11)$$

where $U_{|G|}$ is the uniform distribution on $|G|$ -tuples over $\#Range$. But since G is an $f' \rightarrow f' + q'$ adaptive list, and $f' + q' = \min\{|F| + q, \#Domain\}$, and since F has the properties stated in Lemma 8, we have

$$\Delta(\nabla G \circ F^\mathcal{O}(z^\mathcal{O}), F^\mathcal{O}(z^\mathcal{O}), z^\mathcal{O}; U_{|G|}, F^\mathcal{O}(z^\mathcal{O}), z^\mathcal{O}) \leq \sqrt{\frac{pq}{2f}}. \quad (12)$$

Since $F^\mathcal{O}(z^\mathcal{O}) = S^\mathcal{O}$, and $\nabla G \circ F^\mathcal{O}(z^\mathcal{O}) = \nabla G^\mathcal{O}(F^\mathcal{O}(z^\mathcal{O}), z^\mathcal{O}) = \nabla G^\mathcal{O}(S^\mathcal{O}, z^\mathcal{O})$, the left hand sides of (11) and (12) are equal, and therefore $\mu \leq \sqrt{\frac{pq}{2f}}$. \square

F Example: given-preimage collision-resistance

Lemma 18 (The random-oracle is given-preimage collision-resistant). *Let $g := \mathcal{O}$ where \mathcal{O} denotes the random oracle. Then g is given-preimage collision-resistant in the random oracle model with oracle-dependent auxiliary input.*

More formally, for any probabilistic polynomial-time oracle Turing machine A and any oracle function z into strings of polynomial length, the following probability is negligible (in k):

$$P(x \stackrel{\$}{\leftarrow} Domain, x' \leftarrow A^\mathcal{O}(1^k, z^\mathcal{O}, x) : x \neq x' \text{ and } \mathcal{O}(x) = \mathcal{O}(x'))$$

(In this lemma, \mathcal{O} , $Domain$, f , and z depend implicitly on the security parameter k .)

Proof. First, as we did in the proof of Lemma 10, we set f as $k^{\log k}$, and then can transform the game from the lemma into the following game (using Corollary 9):

$$\text{Game 1: } x \stackrel{\$}{\leftarrow} Domain, x' \leftarrow A^\mathcal{P}(1^k, z^\mathcal{O}, x) : x \neq x' \text{ and } \mathcal{P}(x) = \mathcal{P}(x')$$

Here \mathcal{P} is the random oracle with presampling $S^\mathcal{O}$ and the length of $S^\mathcal{O}$ is bounded by f . Then $|\text{Adv}_A - \text{Adv}_1|$ is negligible.

As in the proof of Lemma 10 we construct a new adversary A_2 that gets the assignment $S^\mathcal{O}$ and never queries \mathcal{P} at positions in $S^\mathcal{O}$. Then $\text{Adv}_1 = \text{Adv}_2$ is negligible for the following Game 2:

$$\text{Game 2: } x \stackrel{\$}{\leftarrow} \text{Domain}, x' \leftarrow A_2^{\mathcal{P}}(1^k, z^\mathcal{O}, x, S^\mathcal{O}) : x \neq x' \text{ and } \mathcal{P}(x) = \mathcal{P}(x')$$

We now construct the following game:

$$\begin{aligned} \text{Game 3: } x \stackrel{\$}{\leftarrow} \text{Domain}, x' \leftarrow A_2^{\mathcal{P}}(1^k, z^\mathcal{O}, x, S^\mathcal{O}) : \\ \text{if } (x' \in S^\mathcal{O}) \text{ then false else } x \neq x' \text{ and } \mathcal{P}(x) = \mathcal{P}(x') \end{aligned}$$

If $\mathcal{P}(x) \notin \text{im } S^\mathcal{O}$ and $x' \in S^\mathcal{O}$, it is never $\mathcal{P}(x) = \mathcal{P}(x')$, so $|\text{Adv}_2 - \text{Adv}_3|$ is bounded by $P(\mathcal{P}(x) \in \text{im } S^\mathcal{O})$ which is negligible (cf. the proof of Lemma 10).

In game 3, the oracle \mathcal{P} is only queried at a position in $S^\mathcal{O}$ if $x \in S^\mathcal{O}$. Since this happens only with negligible probability, we can replace \mathcal{P} by a random oracle \mathcal{R} which is independent of \mathcal{O} .

$$\begin{aligned} \text{Game 4: } x \stackrel{\$}{\leftarrow} \text{Domain}, x' \leftarrow A_2^{\mathcal{R}}(1^k, z^\mathcal{O}, x, S^\mathcal{O}) : \\ \text{if } (x' \in S^\mathcal{O}) \text{ then false else } x \neq x' \text{ and } \mathcal{R}(x) = \mathcal{R}(x') \end{aligned}$$

Then $|\text{Adv}_3 - \text{Adv}_4|$ is negligible.

Since A was polynomial-time, there is a polynomial bound q on the number of queries made by A_2 . Since \mathcal{R} returns random elements of Range , the probability that one of these queries (at a position other than x) returns $\mathcal{R}(x)$ is bounded by $q/\#\text{Range}$. Also the probability that $\mathcal{R}(x') = y$ is bounded by $1/\#\text{Domain}$ if $x' \neq x$ and x' was not queried before. So $\text{Adv}_4 \leq (q+1)/\#\text{Domain}$ which is negligible.

Collecting all bounds, we see that Adv_A is negligible. \square

G Normal security notions and security amplification

In Section 4 we have shown that for the notion of partial-domain one-wayness, a technique we called *security amplification* is applicable. The proof given there was quite independent of the details of the notion of partial-domain one-wayness, and it generalises to many other security definitions of similar form. We are now going to characterise a large class of security notions (called *normal security notions* henceforth) to which this proof technique applies.

Before presenting the definition of a normal security notion, which is quite abstract, we analyse an example security notion to identify its components. A proof system $\pi = (P, V)$ consisting of a prover P and a verifier V is black-box statistical zero-knowledge (for some language \mathcal{L}) if the following holds:

$$\exists \mathcal{S} \forall \mathcal{A}, x : \Delta(\langle P(1^k, x), \mathcal{A}_k(x) \rangle; \mathcal{S}^{\mathcal{A}_k}(1^k, x)) \text{ is negligible in } k.$$

Here \mathcal{S} is a uniform polynomial-time oracle Turing machine, \mathcal{A} a nonuniform polynomial-time adversary (i.e., a sequence of circuits \mathcal{A}_k), and x is a sequence of strings of polynomial length such that $x_k \in \mathcal{L}$ for all k .

We can now identify some structure in this definition: Innermost, we have the statement that some quantity (the statistical distance in this case) is negligible. We call this quantity the advantage of the adversary and write it $\text{Adv}_k^{\mathcal{S}}(\mathcal{A}_k, x_k) := \Delta(\langle P(1^k, x), \mathcal{A}_k(x) \rangle; \mathcal{S}^{\mathcal{A}_k}(1^k, x))$.

Further, we can group the quantifiers appearing in this definition into two types: First, we have the universal quantifiers for \mathcal{A} and x . If we want to consider variants of the notion with respect to another class of adversaries (e.g., f -bounded nonuniform adversaries), these are the quantifiers that have to be modified. The remaining quantifiers (in this case only $\exists \mathcal{S}$) we subsume as a function Φ on predicates. That is, for any predicate $P_{\mathcal{S}}$ depending on \mathcal{S} , let $\Phi(P_{\mathcal{S}}) := \text{true}$ if and only if $\exists \mathcal{S} : P_{\mathcal{S}}$. Then we can write the notion of

black-box statistical zero-knowledge as follows (where we give the simulator the name X better to match Definition 19 below):

$$\Phi(\forall(\mathcal{A}, x) \in \mathfrak{A} : \text{Adv}_k^X(\mathcal{A}_k, x_k) \text{ is negligible})$$

Here \mathfrak{A} is the set of all pairs (A, x) where A is a nonuniform polynomial-time adversary, and x a sequence of string $x_k \in \mathcal{L}$ of polynomial length. Note that we can parametrise the power of the adversaries by modifying the set \mathfrak{A} .

In this definition, there is no structural difference between \mathcal{A} and x any more. So we can simply write A instead of (\mathcal{A}, x) , and let $A_k := (\mathcal{A}_k, x_k)$. Then the definition finally has the form all normal security notions must have:

$$\Phi(\forall A \in \mathfrak{A} : \text{Adv}_k^X(A_k) \text{ is negligible})$$

The actual definition of a normal security notion (which introduces some additional constraints on the function Φ) is given by the following definition.

Definition 19 (Normal security notion). *In the scope of this section, a point-adversary is some object (the type depending on the security notion under consideration). An adversary A is a sequence A_k of point-adversaries ($k \in \mathbb{N}$ can be thought of as the security parameter). E.g., A could be a family of circuits. Or A could be an interactive Turing machine and then A_k is that Turing machine running with security parameter k . We assume that some nonnegative integer-valued function $|\cdot|$ (the size) is defined on the set of all point-adversaries. We call an adversary A nonuniformly f -bounded for some function f , if $|A_k| \leq f(k)$ for all sufficiently large $k \in \mathbb{N}$. We call an adversary nonuniformly polynomial, if he is p -bounded for some polynomial p .*

A function Φ that assigns a truth value (true/false) to any predicate parametrised on some value X we call a quantifier. A quantifier Φ is monotone when the following holds: If $\Phi(A_X) = \text{true}$ and $\forall X : A_X \Rightarrow B_X$, then $\Phi(B_X) = \text{true}$. We call Φ countably compact, if for any set \mathfrak{X} with $\Phi(X \in \mathfrak{X}) = \text{true}$ there exists a countable set $\mathfrak{X}' \subseteq \mathfrak{X}$ with $\Phi(X \in \mathfrak{X}') = \text{true}$.⁸

An advantage function Adv is a nonnegative real-valued family of functions on point-adversaries that is parametrised by $k \in \mathbb{N}$ (the security parameter) and X . (I.e., $\text{Adv}_k^X(A_k)$ is a nonnegative real number for an adversary A .)

A security notion is normal, if for any π (of a type that depends on the security notion, e.g. a protocol), there is a countably compact monotone quantifier Φ and an advantage function Adv such that for any set \mathfrak{A} of adversaries the statement “ π is secure with respect to adversaries in \mathfrak{A} ” can be written as

$$\Phi(\forall A \in \mathfrak{A} : \text{Adv}_k^X(A_k) \text{ is negligible in } k).$$

The following lemma gives some examples of normal security notions.

Lemma 20 (Normal security notions). *The following security notions are normal: One-wayness, partial-domain one-wayness, IND-CPA, IND-CCA2, black-box standalone-security of function evaluations, UC (with dummy-adversary and the environment considered as adversary), black-box zero-knowledge, arguments, black-box arguments of knowledge.*

Proof. Each of the notions can easily be recast into the form required by Definition 19. Then it is straightforward to check that the quantifier Φ is countably compact and monotone. \square

We can apply the security amplification technique to all normal security notions. In particular, the technique is applicable to all notions given in Lemma 20.

⁸ Examples for countably compact monotone quantifiers are the existential quantifier ($\exists X \in M$), the universal quantifier ($\forall X \in M$) with a countable set M , the threshold quantifier (“there are at least n elements X such that . . .”), and combinations of these.

Theorem 21 (Security amplification). *For a normal security notion it holds, that if π is secure with respect to nonuniformly polynomial adversaries, then there is a superpolynomial function f , such that π is secure with respect to nonuniformly f -bounded adversaries.*

Proof. Let π be fixed. Then there are a monotone quantifier Φ and an advantage function Adv , such that for any set \mathfrak{A} of adversaries, π is secure with respect to \mathfrak{A} iff

$$\Phi(\forall A \in \mathfrak{A} : \text{Adv}_k^X(A_k) \text{ is negligible in } k).$$

Let \mathfrak{A}_P be the set of all nonuniformly polynomial adversaries. Assume that π is secure with respect to adversaries in \mathfrak{A}_P . Let \mathfrak{X} be the set of all X that satisfy $\forall A \in \mathfrak{A}_P : (\text{Adv}_k^X(A_k) \text{ is negligible in } k)$. Then $\Phi(X \in \mathfrak{X}) = \text{true}$. Since Φ is countably compact, there is a countable set $\mathfrak{X}' \subseteq \mathfrak{X}$ such that $\Phi(X \in \mathfrak{X}') = \text{true}$.

Fix some $X \in \mathfrak{X}'$.

For an $n \in \mathbb{N}$ let $\mu_n(k) := \sup_{|A_k| \leq p} (\text{Adv}_k^X(A_k))$ where A_k goes over all point-adversaries. Assume there was a polynomial p with integer coefficients (an integer polynomial for short) such that $\mu_{p(k)}(k)$ is not negligible in k . Then there is a sequence of point-adversaries A_k with $|A_k| \leq p(k)$ such that $\text{Adv}_k^X(A_k) \geq \mu_{p(k)} - 2^{-k}$. Since $\mu_{p(k)} - 2^{-k}$ is non-negligible, and the adversary A consisting of these A_k is in \mathfrak{A}_P , and $X \in \mathfrak{X}$, this is a contradiction to the definition of \mathfrak{X} . So μ_p is negligible for all integer polynomials p .

We say, a function μ asymptotically dominates a function ν , if for all sufficiently large k it is $\mu(k) \geq \nu(k)$. [Bel02] proves that for any countable set S of negligible functions, there is a negligible function μ^* that asymptotically dominates all $\mu \in S$.

Therefore, there is a negligible function μ^* , that asymptotically dominates μ_p for every integer polynomial p .

Let $f^X(k) := \max\{p \in \mathbb{N} : \mu_p(k) \leq \mu^*(k)\}$. Then $\mu_{f^X(k)}(k) \leq \mu^*(k)$ is negligible. So for any nonuniformly f^X -bounded adversary A the advantage $\text{Adv}_k^X(A_k) \leq \mu_{f^X}$ is negligible. Furthermore, we can show that f^X is superpolynomial. Assume this was not the case. Then there an integer polynomial p such that $p > f^X$ infinitely often. But then $\mu_p > \mu^*$ infinitely often, in contradiction to the choice of μ^* . So f^X is superpolynomial.

Again by the result of [Bel02], and since \mathfrak{X}' is countable, there is a negligible function ν that asymptotically dominates all $1/f^X$ with $X \in \mathfrak{X}'$. Set $f := 1/\nu$. Then for every $X \in \mathfrak{X}'$, f_X asymptotically dominates f . In particular, for any $X \in \mathfrak{X}'$, it holds that for any nonuniformly f -bounded adversary A the advantage $\text{Adv}_k^X(A_k)$ is negligible.

Let \mathfrak{A}_f be the set of all nonuniformly f -bounded adversaries. Then we have $X \in \mathfrak{X}' \Rightarrow \forall A \in \mathfrak{A}_f : (\text{Adv}_k^X(A_k) \text{ is negligible})$. Since furthermore $\Phi(X \in \mathfrak{X}') = \text{true}$ and Φ is monotone, it follows $\Phi(\forall A \in \mathfrak{A}_f : \text{Adv}_k^X(A_k) \text{ is negligible}) = \text{true}$. So π is secure with respect to nonuniformly f -bounded adversaries. \square

H Security of OAEP encryption

Definition 22 (IND-CCA2 security). *An encryption scheme \mathcal{E} is IND-CCA2 secure in the random oracle model with oracle-dependent auxiliary input if the following holds:*

For a given secret/public key pair (sk, pk) , let the decryption oracle $\mathcal{D}_{sk}^{\mathcal{O}}(c)$ return the decryption of c under \mathcal{E} , or \perp if the decryption fails.

For an adversary A and an oracle function z into strings of polynomial length, the advantage $\text{Adv}_A^{\text{IND}}$ is defined as

$$\begin{aligned} \text{Adv}_A^{\text{IND}} &:= 2P((pk, sk) \leftarrow K^{\mathcal{O}}(1^k), \\ &\quad (m_1, m_2) \leftarrow A^{\mathcal{O}, \mathcal{D}_{sk}^{\mathcal{O}}}(1^k, z^{\mathcal{O}}, pk), \\ &\quad b \stackrel{\$}{\leftarrow} \{0, 1\}, y^* \leftarrow E_{pk}^{\mathcal{O}}(m_b), \\ &\quad b' \leftarrow A^{\mathcal{O}, \mathcal{D}_{sk}^{\mathcal{O}}}(1^k, z^{\mathcal{O}}, pk, y^*) : \\ &\quad b = b') - 1. \end{aligned}$$

Here $K^{\mathcal{O}}(1^k)$ denotes the key generation algorithm of \mathcal{E} . Further, $E_{pk}^{\mathcal{O}}(m)$ is the encryption of m with respect to \mathcal{E} using public key pk . We assume that the adversary A is able to store an internal state between its two activations, and that the two messages m_1, m_2 are of fixed length.

Then the encryption scheme \mathcal{E} is secure if for any probabilistic polynomial-time adversary A and any oracle function z into strings of polynomial length, where the adversary A does not call $\mathcal{D}_{sk}^{\mathcal{O}}$ with the ciphertext y^* , the advantage $\text{Adv}_A^{\text{IND}}$ is negligible.

(In this definition, \mathcal{O} , z , and $\text{Adv}_A^{\text{IND}}$ depend implicitly on the security parameter k .)

Definition 23 (OAEP encryption). Let k_0 and k_1 be superlogarithmic functions of the security parameter k with $k_0 + k_1 < k$.

Let f_{pk} be a family of 1-1 trapdoor functions with domain $\{0, 1\}^{k-k_0} \times \{0, 1\}^{k_0}$. Let g_{sk} denote the inverse of f_{pk} if sk is a secret key corresponding to the public key pk .

The OAEP encryption scheme (for that family of functions f_{pk}) assumes two random oracles $G : \{0, 1\}^{k_0} \rightarrow \{0, 1\}^{k-k_0}$ and $H : \{0, 1\}^{k-k_0} \rightarrow \{0, 1\}^{k_0}$. (We can assume these to be part of a single random oracle \mathcal{O} by defining e.g. $G(x) := \mathcal{O}(1, x)$ and $H(x) := \mathcal{O}(1, x)$ where the output of \mathcal{O} is appropriately truncated.) To encrypt a message $m \in \{0, 1\}^{k-k_0-k_1}$, we set

$$\begin{aligned} r &\stackrel{\$}{\leftarrow} \{0, 1\}^{k_0} \\ s &\leftarrow (m \| 0^{k_1}) \oplus G(r) \\ t &\leftarrow r \oplus H(s) \\ c &\leftarrow f_{pk}(s, t) \end{aligned}$$

Then c is the encryption of m .

To decrypt a ciphertext c , we set

$$\begin{aligned} (s, t) &\leftarrow g_{sk}(c) \\ r &\leftarrow t \oplus H(s) \\ M &\leftarrow s \oplus G(r) \end{aligned}$$

If M has the form $m \| 0^{k_1}$, the decryption is m . Otherwise, the decryption fails (i.e., the ciphertext is rejected).

(In this definition, \mathcal{O} , \mathcal{P} , G , and H depend implicitly on the security parameter k .)

Proof (of Theorem 13). To allow the reader to see to what extent the proof of the OAEP security has to be changed to apply to the setting of oracle-dependent auxiliary input, we very closely follow the original proof from [FOPS04, Section 4].⁹ To highlight the differences, we use the following convention: The proof

⁹ Note that there are two proofs in [FOPS04], one in Section 4 and one in the appendix. The proof in the appendix is that from the conference version [FOPS01]. We choose the proof from Section 4 as our guideline, since it comes with more formal details and therefore seems better suited to show the applicability of our technique.

from [FOPS04] is given almost literally, and all changes necessary due to the oracle-dependent auxiliary input are typeset in *italic*.

To show the theorem, we proceed by giving a sequence of games and deriving bounds for the probabilities of given events. We start with the probability that the adversary breaks the OAEP encryption scheme and finally relate it to probability that the adversary breaks the partial-domain one-wayness property of the underlying trapdoor function.

For any adversary A , let $\varepsilon := \text{Adv}_A^{\text{IND}}$ denote the advantage of A in the IND-CCA2 game (cf. Definition 22). (The advantage ε depends implicitly on k .)

Let \mathfrak{f} be a some superpolynomial nonnegative integer-valued function with $\mathfrak{f} \geq 1$ such that $f \cdot 2^{-k_0}$ and $f \cdot 2^{k_1}$ are negligible. (We will fix a concrete choice of \mathfrak{f} below.)

GAME₀ A pair of keys (pk, sk) is generated by the key generation algorithm of the OAEP scheme (which is the same as that of the underlying trapdoor functions). Let $f := f_{pk}$ denote the trapdoor function corresponding to the public key pk , and let $g := g_{sk}$ denote its inverse. The adversary A is given the public key pk , the oracle dependent auxiliary input $z^\mathcal{O}$ and access to the decryption oracle $\mathcal{D}_{sk}^\mathcal{O}$ and the random oracle \mathcal{O} .¹⁰ Then adversary is expected to output a pair of messages (m_1, m_2) of length $k - k_0 - k_1$.

Now a challenge ciphertext y^* is produced as follows: First, a random bit $b \in \{0, 1\}$ is chosen. Then the ciphertext y^* is constructed as follows:

$$r^* \xleftarrow{\$} \{0, 1\}^{k_0}, s^* \leftarrow (m_b \| 0^{k_1}) \oplus G(r^*), t^* \leftarrow r^* \oplus H(s^*), x^* \leftarrow (s^*, t^*), y^* \leftarrow f(x^*).$$

The ciphertext y^* is then given to A . Additionally, A still has access to $z^\mathcal{O}$, pk , $\mathcal{D}_{sk}^\mathcal{O}$ and \mathcal{O} . However, A does not query the decryption of y^* from $\mathcal{D}_{sk}^\mathcal{O}$. The adversary A then outputs a bit b' . We denote by S_0 the event $b = b'$ and use a similar notation S_i for any **GAME_i** below. By definition of the OAEP scheme, it is $P(S_0) = \frac{1}{2} + \frac{\varepsilon}{2}$.

GAME_{1/2} Let q_G, q_H and q_D denote the number of queries made by A to the oracles G, H and $\mathcal{D}_{sk}^\mathcal{O}$, respectively. Let $S^\mathcal{O}$ be as in Theorem 2 an assignment of length at most \mathfrak{f} . Let \mathcal{P} be the random oracle with presampling $S^\mathcal{O}$. By virtue of Theorem 2 (and using the technique presented in the first step of the proof of Lemma 10), we can replace every occurrence of \mathcal{O} by an occurrence of \mathcal{P} (with exception of the auxiliary input $z^\mathcal{O}$, which stays $z^\mathcal{O}$ and does not become $z^\mathcal{P}$). By Theorem 2 it is $|P(S_0) - P(S_{1/2})| \leq \sqrt{\frac{pq}{2\mathfrak{f}}}$, where p is the length of the auxiliary input $z^\mathcal{O}$ and $q := q_G + q_H + 2q_D + 2$ is an upper bound on the number of oracle queries made in **GAME₀**.

To avoid too cluttered a notation, from now on, we denote by G and H the oracles that are part of \mathcal{P} (instead of being part of \mathcal{O} as in **GAME₀**).

GAME₁ We modify **GAME_{1/2}** by moving the generation of the seed r^* and its image $G(r^*)$ to the beginning of the game. More precisely, we randomly choose ahead of time some $r^+ \xleftarrow{\$} \{0, 1\}^{k_0}$ and $g^+ \xleftarrow{\$} \{0, 1\}^{k-k_0}$ and use r^+ instead of r^* and g^+ instead of $G(r^*)$. The game obeys the following two rules:

Rule 1. $r^* = r^+$ and $s^* = (m_b \| 0^{k_1}) \oplus g^+$. The other variables are generated as described above.

Rule 2. Whenever the random oracle G is queried at r^+ , the answer is g^+ .

Since we still choose r^+ randomly and independently, and since \mathcal{P} returns a uniformly chosen element when queried at position r^+ unless $r^+ \in S^\mathcal{O}$, the two games differ only if $r^+ \in S^\mathcal{O}$. This happens with probability at most $\mathfrak{f} \cdot 2^{-k_0}$. So we get

$$|P(S_1) - P(S_{1/2})| \leq \mathfrak{f} \cdot 2^{-k_0}.$$

¹⁰ Here we assume that G and H are part of the same random oracle \mathcal{O} . One could also define G and H as different random oracles and then use an auxiliary input $z^{G,H}$ and write the decryption oracle as $\mathcal{D}_{sk}^{G,H}$.

GAME₂ For this game we remove the Rule 2 in the GAME₁ and restore the calls to G . Therefore g^+ is used for the construction of x^* but does not appear anywhere else in the computation. Thus, since m_b appears in the calculation only in the form of $(m_b || 0^{k_1}) \oplus g^+$, the view of A does not depend on b . So $P(S_2) = \frac{1}{2}$.

GAME₁ and GAME₂ may only differ if $r^* = r^+$ is queried from G . Let AskG₂ denote the event that r^+ is queried from G in GAME₂ by the adversary A or by the decryption oracle (and similarly AskG _{i} for GAME _{i} below). Then

$$|P(S_2) - P(S_1)| \leq P(\text{AskG}_2).$$

GAME₃ Like we did in GAME₁, we now move the generation of s^* and $H(s^*)$ up front and make it independent of the other variables. More precisely, we now choose $s^+ \xleftarrow{\$} \{0,1\}^{k-k_0}$ and $h^+ \xleftarrow{\$} \{0,1\}^{k_0}$ randomly, and use s^+ instead of s^* , and h^+ instead of $H(s^*) = H(s^+)$. The game uses the following two rules:
Rule 1'. $g^+ = (m_b || 0^{k_1}) \oplus s^+$ and $t^* = r^+ \oplus h^+$.

Rule 2'. Whenever the random oracle H is queried at s^+ , the answer is h^+ .

Since $s^* = (m_b || 0^{k_1}) \oplus g^+$ for randomly chosen g^+ in GAME₁, replacing s^* by s^+ (and defining g^+ as prescribed by Rule 1') does not change the game. Furthermore, $H(s^+)$ is uniformly chosen by the random oracle *unless* $s^+ \in S^{\mathcal{O}}$. The probability that $s^+ \in S^{\mathcal{O}}$ is bounded by $\mathfrak{f} \cdot 2^{-k+k_0}$. So

$$|P(\text{AskG}_2) - P(\text{AskG}_3)| \leq \mathfrak{f} \cdot 2^{-k+k_0}.$$

GAME₄ We now remove Rule 2' from GAME₃ and thus restore the original meaning of $H(s^+)$. Then h^+ is used only in the computation of x^* but does not appear anywhere else. GAME₃ and GAME₄ differ only if s^+ is queried from H . Let AskH₄ denote the event that s^+ is queried from H . (And similarly AskH _{i} for GAME _{i} .) Then

$$|P(\text{AskG}_3) - P(\text{AskG}_4)| \leq P(\text{AskH}_4).$$

Furthermore, since $r^+ = t^* \oplus h^+$, and h^+ is never revealed, r^+ is independent of the adversary A 's view. So $P(\text{AskG}_4) \leq (q_G + q_D) \cdot 2^{-k_0}$ (where q_G is the number of G -queries made by A , and q_D the number of decryption oracle queries).

GAME₅ We again change the way the challenge ciphertext is generated. We now choose $t^+ \xleftarrow{\$} \{0,1\}^{k_0}$ and replace t^* by t^+ . Note that now the ciphertext $y^* = f(s^+, t^+)$ is a uniformly chosen image of f , independent of the other random variables A learns. Furthermore, since $t^* = h^+ \oplus r^+$ in GAME₄, and h^+ and r^+ were uniformly distributed and not used by the adversary, replacing t^* by t^+ does not change the game. Thus

$$P(\text{AskH}_5) = P(\text{AskH}_4).$$

GAME₆ We now change the decryption oracle \mathcal{D} . When \mathcal{D} is queried with a ciphertext $y = f(s, t)$, let $r := t \oplus H(s)$. If r has not been previously queried from G and r is not contained in the assignment $S^{\mathcal{O}}$ the modified decryption oracle rejects the ciphertext. Since $G(r)$ is uniformly distributed for $r \notin S^{\mathcal{O}}$, it only happens with probability 2^{-k_1} that $s \oplus G(r)$ has the form $m || 0^{k_1}$ (and otherwise the ciphertext is not accepted by the original oracle, either). The adversary makes at most q_D queries to \mathcal{D} , so

$$|P(\text{AskH}_6) - P(\text{AskH}_5)| \leq q_D \cdot 2^{-k_1}.$$

Note that the new decryption oracle does not query G any more, because the necessary query already has been made.

GAME₇ Now the decryption oracle additionally rejects cyphertexts y in the case that s (as defined in GAME₆) has not previously been queried from H , nor is $s \in S^{\mathcal{O}}$. This behaviour differs from the behaviour in GAME₆ only if r has been queried from G while $H(s)$ has not been asked and $s \notin S^{\mathcal{O}}$. But if $s \notin S^{\mathcal{O}}$, the value of $H(s)$ is uniformly distributed and so is $r = H(s) \oplus t$, so the probability that r has been

queried from G is at most $q_G \cdot 2^{-k_0}$ (the decryption oracle from GAME_8 does not query G any more). Since A makes at most q_D queries to the decryption oracle, we get

$$|P(\text{AskH}_7) - P(\text{AskH}_6)| \leq q_D q_G \cdot 2^{-k_0}.$$

GAME_8 Since the decryption oracle only decrypts some ciphertext y if s and r have already been queried from G or H , resp., or can be found in $S^\mathcal{O}$, we can replace the decryption oracle by the following algorithm: Let $\mathbf{G}\text{-List}$ denote the queries already made to G , and $\mathbf{H}\text{-List}$ those made to H . Then for each $r \in \mathbf{G}\text{-List} \cup S^\mathcal{O}$ and $s \in \mathbf{H}\text{-List} \cup S^\mathcal{O}$, check whether $f(s, t) = y$ where $t := r \oplus H(s)$. If so, the decryption can be completed without use of the secret key sk . Otherwise, the ciphertext is rejected. Since the decryption oracle from GAME_7 rejects all ciphertexts that do not pass the test $f(s, t) = y$ for some of the tested values for r, s , we have

$$P(\text{AskH}_8) = P(\text{AskH}_7).$$

There are at most q_H queries to H . $P(\text{AskH}_8)$ is the probability that one of these is made at the position s^+ . So when we randomly choose one of the queries to H and output it, with probability at least $P(\text{AskH}_8)/q_H$ we output an s satisfying $(s, t) = g(y^*)$. So we have constructed an adversary A^* against the partial-domain one-wayness of f_{pk} that has a success probability $\text{Adv}_{A^*}^{pdo} = P(\text{AskH}_8)/q_H$.

To conclude our proof, we have to see that a non-negligible advantage $\varepsilon = \text{Adv}_A^{\text{IND}}$ against the IND-CCA2 game implies a non-negligible advantage $\delta := \text{Adv}_{A^*}^{pdo}$ against the partial-domain one-wayness of f_{pk} , and that the adversary A^* is polynomial-time. The first is an easy calculation given the bounds we identified above, *but on the second we must fail: A^* has to consider all the superpolynomially many values in $S^\mathcal{O}$. However, we will use the security amplification technique presented in Section 4 to solve this problem.*

First, let us calculate ε from δ :

$$\begin{aligned} \frac{\varepsilon}{2} &= |P(S_0) - P(S_2)| \leq \sqrt{\frac{pq}{2f}} + f \cdot 2^{-k_0} + P(\text{AskG}_2) \\ P(\text{AskG}_2) &\leq f \cdot 2^{-k+k_0} + P(\text{AskG}_4) + P(\text{AskH}_4) \\ &\leq f \cdot 2^{-k+k_0} + (q_G + q_D) \cdot 2^{-k_0} + P(\text{AskH}_4) \\ P(\text{AskH}_4) &\leq q_D \cdot 2^{-k_1} + q_D q_G \cdot 2^{-k_0} + q_H \delta \end{aligned}$$

Here p is the length of the auxiliary input, and $q := q_G + q_H + 2q_D + 2$ the total number of queries to the random oracle. If δ is negligible, and q and p are polynomially bounded, and f is superpolynomial, and $f \cdot 2^{-k_0}, f \cdot 2^{-k_1}$ are negligible, then ε is negligible.

It is left to show that δ is indeed negligible under the assumption that f_{pk} is partial-domain one-way. *Unfortunately, the running time of A^* is only bounded by $\tau \leq p(f+k)$ for some polynomial p independent of f . The runtime τ is obviously not polynomially bounded for superpolynomial f . But Lemma 12 guarantees, that there is a superpolynomial function f' such that f_{pk} is partial-domain one-way even against f' -bounded nonuniform adversaries. We then can choose a superpolynomial f such that $\tau \leq f'$. We can further choose f small enough (but still superpolynomial) such that $f \leq 2^{-k_0/2}, 2^{-k_1/2}$. Then $f \cdot 2^{-k_0}, f \cdot 2^{-k_1}$ are negligible. For this choice of f , the adversary A^* has f' -bounded running time, so $\delta = \text{Adv}_{A^*}^{pdo}$ is negligible. Since δ is negligible, so is ε , which proves that the OAEP scheme is IND-CCA2 secure in the random oracle model with oracle-dependent auxiliary input. \square*

References

- [Bel02] Mihir Bellare. A note on negligible functions. *Journal of Cryptology*, 15(4):271–284, 2002. Online available at <http://eprint.iacr.org/1997/004>.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *1st ACM Conference on Computer and Communications Security, Proceedings of CCS 1993*, pages 62–73. ACM Press, 1993. Full version online available at <http://www.cs.ucsd.edu/users/mihir/papers/ro.ps>.
- [BR95] Mihir Bellare and Phillip Rogaway. Optimal asymmetric encryption—how to encrypt with RSA. In Alfredo de Santis, editor, *Advances in Cryptology, Proceedings of EUROCRYPT '94*, volume 950 of *Lecture Notes in Computer Science*, pages 92–111. Springer-Verlag, 1995. Extended version online available at <http://www.cs.ucsd.edu/users/mihir/papers/oa.ps>.
- [CCM02] Christian Cachin, Claude Crépeau, and Julien Marcil. Oblivious transfer with a memory-bounded receiver. In *34th Annual ACM Symposium on Theory of Computing, Proceedings of STOC 2002*, pages 493–502. ACM Press, 2002.
- [CGH98] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited. In *Thirtieth Annual ACM Symposium on Theory of Computing, Proceedings of STOC 1998*, pages 209–218. ACM Press, 1998. Preliminary version, extended version online available at <http://eprint.iacr.org/1998/011.ps>.
- [CM97] Christian Cachin and Ueli Maurer. Unconditional security against memory-bounded adversaries. In Burton S. Kaliski Jr., editor, *Advances in Cryptology, Proceedings of CRYPTO '97*, volume 1294 of *Lecture Notes in Computer Science*, pages 292–306. Springer-Verlag, 1997.
- [DM02] Stefan Dziembowski and Ueli M. Maurer. Tight security proofs for the bounded-storage model. In *34th Annual ACM Symposium on Theory of Computing, Proceedings of STOC 2002*, pages 341–350. ACM Press, 2002.
- [FOPS01] Eiichiro Fujisaki, Tatsuaki Okamoto, David Pointcheval, and Jacques Stern. RSA-OAEP is secure under the RSA assumption. In Joe Kilian, editor, *Advances in Cryptology, Proceedings of CRYPTO '01*, volume 2139 of *Lecture Notes in Computer Science*, pages 260–274. Springer-Verlag, 2001. Online available at <http://eprint.iacr.org/2000/061.ps>, full version is [FOPS04].
- [FOPS04] Eiichiro Fujisaki, Tatsuaki Okamoto, David Pointcheval, and Jacques Stern. RSA-OAEP is secure under the RSA assumption. *Journal of Cryptology*, 17(2):81–104, 2004. Online available at http://www.di.ens.fr/~pointche/Documents/Papers/2004_joc.pdf.
- [GGKT05] Rosario Gennaro, Yael Gertner, Jonathan Katz, and Luca Trevisan. Bounds on the efficiency of generic cryptographic constructions. *SIAM Journal on Computing*, 35(1):217–246, 2005. Online available at <http://dx.doi.org/10.1137/S0097539704443276>.
- [Gol93] Oded Goldreich. A uniform-complexity treatment of encryption and zero-knowledge. *Journal of Cryptology*, 6(1):21–53, 1993. Extended version online available at <http://www.wisdom.weizmann.ac.il/~oded/PS/uniform.ps>.
- [Imp96] Russel Impagliazzo. Very strong one-way functions and pseudo-random generators exist relative to a random oracle. Manuscript, 1996.
- [Kul67] Solomon Kullback. A lower bound for discrimination information in terms of variation (corresp.). *IEEE Transactions on Information Theory*, 13(1):126–127, 1967.
- [Mau92] Ueli Maurer. Conditionally-perfect secrecy and a provably-secure randomized cipher. *Journal of Cryptology*, 5(1):53–66, 1992. Online available at <ftp://ftp.inf.ethz.ch/pub/crypto/publications/Maurer92b.pdf>.
- [Pas03] Rafael Pass. On deniability in the common reference string and random oracle model. In Dan Boneh, editor, *Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003, Proceedings*, volume 2729 of *Lecture Notes in Computer Science*, pages 316–337. Springer, 2003.
- [PKC02] RSA Laboratories. *PKCS #1: RSA Cryptography Standard, Version 2.1*, 2002. Online available at <ftp://ftp.rsasecurity.com/pub/pkcs/pkcs-1/pkcs-1v2-1.pdf>.
- [Rog06] Phillip Rogaway. Formalizing human ignorance: Collision-resistant hashing without the keys. In *Vietcrypt 2006*, volume 4341 of *Lecture Notes in Computer Science*, pages 221–228. Springer-Verlag, 2006. Online available at <http://eprint.iacr.org/2006/281>.

- [SHS02] Federal Information Processing Standards Publications. *FIPS PUB 180-2: Secure Hash Standard (SHS)*, August 2002. Online available at <http://csrc.nist.gov/publications/fips/fips180-2/fips180-2withchangenotice.pdf>.
- [Wee06] Hoeteck Wee. Zero knowledge in the random oracle model, revisited. Manuscript, 2006.