

On the Security of OAEP

Alexandra Boldyreva¹ and Marc Fischlin²

¹College of Computing, Georgia Institute of Technology, USA

sasha@gatech.edu www.cc.gatech.edu/~aboldyre

²Darmstadt University of Technology, Germany

marc.fischlin@gmail.com www.fischlin.de

Abstract

Currently, the best and only evidence of the security of the OAEP encryption scheme is a proof in the contentious random oracle model. Here we give further arguments in support of the security of OAEP. We first show that partial instantiations, where one of the two random oracles used in OAEP is instantiated by a function family, can be provably secure (still in the random oracle model). For various security statements about OAEP we specify sufficient conditions for the instantiating function families that, in some cases, are realizable through standard cryptographic primitives and, in other cases, may currently not be known to be achievable but appear moderate and plausible. Furthermore, we give the first non-trivial security result about *fully* instantiated OAEP *in the standard model*, where both oracles are instantiated simultaneously. Namely, we show that instantiating both random oracles in OAEP by modest functions implies non-malleability under chosen plaintext attacks for random messages. We also discuss the implications, especially of the full instantiation result, to the usage of OAEP for secure hybrid encryption (as required in SSL/TLS, for example).

Keywords: OAEP, provable security, public-key encryption, random oracle model.

1 Introduction

OAEP is one of the most known and widely deployed asymmetric encryption schemes. It was designed by Bellare and Rogaway [5] as a scheme based on a trapdoor permutation such as RSA. OAEP is standardized in RSA's PKCS #1 v2.1 and is part of the ANSI X9.44, IEEE P1363, ISO 18033-2 and SET standards. The encryption algorithm of $\text{OAEP}^{G,H}[F]$ takes a public key f , which is an instance of a trapdoor permutation family F , and a message M , picks r at random and computes the ciphertext $C = f(s||t)$ for $s = G(r) \oplus M||0^{k_1}$ and $t = H(s) \oplus r$, where G and H are some hash functions. Despite its importance the only security results for OAEP are a proof of IND-CPA security assuming F is a one-way trapdoor permutation family [5] and a proof of IND-CCA2 security assuming F is partial one-way [15], both in the random oracle (RO) model, i.e., where G and H are idealized and modeled as random oracles [4]. However, such proofs merely provide heuristic evidence that breaking the scheme may be hard in reality (when the random oracles are instantiated with real functions).

A growing number of papers raised concerns regarding soundness of the controversial random oracle model [11, 18, 19, 16, 1, 13, 9, 20]. Moreover, most of the recent results question security of the practical schemes known to be secure in the RO model. For example, Dodis et al. [13] showed some evidence that the RSA Full Domain Hash signature scheme may not be secure in the standard model. Boldyreva and Fischlin [9]

showed that even presumably strong candidates like perfectly one-way hash functions (POWHFs) [10, 12] are insufficient to prove security of partial instantiations of OAEP (when only one of the two random oracles is instantiated with an instance of a POWHF).

The motivation of this work is to gather evidence of soundness of the OAEP design. Like the aforementioned works our goal is to go beyond the classical RO heuristic and study security of the scheme when one or all of its ROs are instantiated. Positive results in the direction of partial instantiations would give further evidence that breaking OAEP for good instantiations is hard, because breaking the scheme would then require to exploit interdependent weaknesses between the instantiations or the family F . Given the negative results of [9] it is unlikely to expect that the properties needed from the instantiating function families are weak or even easily realizable, even if one accepts weaker security stipulations than chosen-ciphertext security for partial or full instantiations. For example, although it seems plausible, it is currently not even known whether OAEP can be proven IND-CPA secure in the standard model assuming any reasonable properties of the instantiating functions.

Here we show that security proofs for instantiations of OAEP are indeed possible. For various security statements about OAEP we specify sufficient conditions on G and H that are certainly weaker than assuming that the functions behave as random oracles, yielding “positive” security statements regarding partially instantiated OAEP. Furthermore, we give the first non-trivial security results about fully instantiated OAEP in the standard model, where both oracles G and H are instantiated simultaneously. We next discuss these results in more detail.

THE OAEP FRAMEWORK. For better comprehension of our technical results we first reconsider the OAEP encryption scheme from a more abstract viewpoint. Let f be a random instance of a partial one-way trapdoor permutation family F , and the encryption algorithm computes a ciphertext as $C = f(s||t)$. Partial one-wayness [15] requires that it is hard to find the leading part of the pre-image $s||t$ under f and to output, say, s only. If we consider now for example a family $F_{t\text{-clear}}$ where each function is defined as $f \equiv g||\text{ID}$ such that $f(s||t) = g(s)||t$ for a trapdoor permutation g , then this family $F_{t\text{-clear}}$ is clearly partial one-way (and also a trapdoor permutation). Hence, this example describes a special case $\text{OAEP}^{G,H}[F_{t\text{-clear}}]$ for the partial one-way trapdoor permutation family $F_{t\text{-clear}}$ where each function outputs the t -part in clear. In particular, the security proof in the random oracle model for OAEP and general partial one-way families (including RSA as a special case) [15] carries over, but we outdo this by giving positive results of partial instantiation for such families $F_{t\text{-clear}}$.

Towards the standard-model security results for fully instantiated OAEP we take the above viewpoint one step further and look at $\text{OAEP}^{G,H}[F_{\text{lsb}||t\text{-clear}}]$ for families $F_{\text{lsb}||t\text{-clear}}$ where each function f outputs the k_1 least significant bits of $s = G(r) \oplus M||0^{k_1}$ (which equal those bits of $G(r)$) and t in clear. Since each function in $F_{\text{lsb}||t\text{-clear}}$ is also a member in $F_{t\text{-clear}}$ the partial instantiation results above remain true for $\text{OAEP}^{G,H}[F_{\text{lsb}||t\text{-clear}}]$.

We note that security of partial instantiations of $\text{OAEP}^{G,H}[F_{t\text{-clear}}]$ and of $\text{OAEP}^{G,H}[F_{\text{lsb}||t\text{-clear}}]$, although for qualified partial one-way trapdoor families, also have implications for the popular $\text{OAEP}^{G,H}[\text{RSA}]$ case. They show that any successful attacks on instantiations for RSA would have to take advantage of specific properties of the RSA function. Generic attacks which would also work for $F_{t\text{-clear}}$ or $F_{\text{lsb}||t\text{-clear}}$ are then ruled out.

PARTIAL INSTANTIATION RESULTS. Positive results about partial instantiations were first shown in [9] for the PSS-E encryption scheme. There it was also shown, however, that perfectly one-way hash functions cannot be securely used to instantiate either one of the ROs in OAEP. These negative results about partial instantiation through POWHFs hold for $\text{OAEP}^{G,H}[F_{t\text{-clear}}]$ as well. Yet we show that partial instantiations are possible by switching to other primitives.

To instantiate the G -oracle in $\text{OAEP}^{G,H}[F_{t\text{-clear}}]$ while preserving IND-CCA2 security (in the random oracle model), we introduce the notion of a near-collision resistant pseudorandom generator. For such a generator G it is infeasible to find different seeds $r \neq r'$ such that predetermined parts of the generator’s outputs $G(r)$, $G(r')$ match (they may differ on other parts). To be more precise for $\text{OAEP}^{G,H}[F_{t\text{-clear}}]$ the generator G is not allowed to coincide on the k_1 least significant bits, bequeathing this property to the values $s = G(r) \oplus M||0^{k_1}$ and $s' = G(r') \oplus M||0^{k_1}$ in the encryption process. We discuss that such pseudorandom generators can be derived from any one-way permutation.

Instantiating the H oracle in OAEP turns out to be more challenging. To this end we consider non-

malleable pseudorandom generators, where a given image of a seed r should not help significantly to produce an image of a related seed r' . Instantiating H through such a non-malleable pseudorandom generator the resulting scheme achieves NM-CPA security, where it is infeasible to convert a given ciphertext into one of a related message. Although this security notion for encryption schemes is not as strong as IND-CCA, it yet exceeds the classical IND-CPA security. That is, Bellare et al. [3] show that NM-CPA implies IND-CPA and is incomparable to IND-CCA1 security. Hence, NM-CPA security of schemes lies somewhere in between IND-CPA and IND-CCA2.¹

We also show that it is possible to extend the above result and to instantiate the H -oracle in $\text{OAEP}^{G,H}[F_{t\text{-clear}}]$ without even sacrificing IND-CCA2 security (again, for random oracle G). This however requires the very strong assumption for the pseudorandom generators which then must be non-malleable under chosen-image attacks. For such a generator non-malleability should even hold if the adversary can learn seeds of chosen images, and such generators resemble chosen-ciphertext secure encryption schemes already. Hence, we see this partial instantiation as a mere plausibility result that one can presumably instantiate oracle H and still have IND-CCA2 security. This is contrast to the results in [11] for example, showing that there are encryption schemes secure in the random oracle model but which cannot be securely realized for any primitive, not even for a secure encryption scheme itself.

As for the existence of non-malleable pseudorandom generators, we are not aware if they can be derived from standard cryptographic assumptions, and we leave this as an interesting open problem. We also remark that, while non-malleability under chosen-image attacks seems to be a rather synthetic property, plain non-malleability as required in the NM-CPA result appears to be a modest and plausible assumption for typical instantiation candidates like hash functions. For instance, it should not be easy to flip bits in given hash value, affecting bits in the pre-image in a reasonable way.

FULL INSTANTIATION RESULT. Our main result is a standard-model security proof for a fully instantiated OAEP. It is not very reasonable to expect a proof of IND-CCA2 security of OAEP in the standard model, even assuming very strong properties of instantiating functions (although we all would like to see such result). As we mentioned above, we are not aware if one can even show IND-CPA security of fully instantiated OAEP.

Nevertheless we show that OAEP in the standard model can be proven to satisfy a rather strong notion of security notion, namely $\$$ NM-CPA. It is slightly weaker than the standard non-malleability notion NM-CPA in that there is a restriction that an unknown random message is encrypted in the challenge ciphertext. A bit more formally this security notion $\$$ NM-CPA requires that given a public key and a ciphertext of a challenge message chosen uniformly at random from a large message space it is hard to compute a valid ciphertext of a message non-trivially related to the challenge message. Note that this is consistent with how asymmetric schemes are typically used to build hybrid encryption schemes, where the key of the symmetric scheme is derived from a random string encrypted with the public-key scheme. To appreciate the power of the $\$$ NM-CPA definition we note that it implies for example the notion of OW-CPA and, moreover, Bleichenbacher's attack [7] on PKCS #1 v1.5 is not possible for $\$$ NM-CPA secure schemes.² Thus our result provides better evidence that OAEP resists such attacks, and specifies what properties of the instantiating functions are sufficient for this.

For our full instantiation proof we consider $\text{OAEP}^{G,H}[F_{|s_b||t\text{-clear}}]$ where the t -part and the least significant bits of the s -part are output in clear. To achieve the $\$$ NM-CPA security notion under full instantiation of both oracles G and H in $\text{OAEP}^{G,H}[F_{|s_b||t\text{-clear}}]$ we need to augment the near-collision resistant generator G by a trapdoor property, allowing to invert images efficiently given the trapdoor information; such generators exist if trapdoor permutations exist. We again use a non-malleable pseudorandom generator H for instantiating H . Assuming that the generators above exist we show that $\text{OAEP}^{G,H}[F_{|s_b||t\text{-clear}}]$ is $\$$ NM-CPA.³

¹We mitigate the notion of NM-CPA such that the relation specifying related messages and the distribution over the messages must be fixed at the outset. This mildly affects the relationship to the IND notions, but we omit technical details in the introduction.

²Bleichenbacher's attack works by generating a sequence of ciphertexts from a given ciphertext and verifying validity of the derived ciphertexts by querying the decryption oracle. While requiring *adaptive* queries to recover the entire message, one can view the message in first derived ciphertext in such an attack as having a small (but not negligible) probability of being non-trivially related to the original (possibly random) message.

³Very recently, Brown [2] has shown that RSA-OAEP cannot be proven OW-CPA under certain security reductions. Our approach here does not fall under this kind of reductions and does not contradict his result. We provide more details in Section 3.2.

To give further evidence of the usefulness of the \S NM-CPA notion we finally show that we can derive a hybrid encryption scheme that is NM-CPA in the random oracle model from an asymmetric scheme secure in the sense of \S NM-CPA. For this, one encrypts a random string r with the asymmetric scheme and then runs r through an idealized key derivation process to obtain $K = G(r)$, modeled through a random oracle G . The actual message is then encrypted with a symmetric scheme for key K . The construction of such hybrid encryption schemes resembles the encryption method in SSL/TLS [17]. There, simply speaking, the client encrypts a random string under the server’s public key and then both parties derive the actual symmetric key K by hashing the random string iteratively. If one considers this hashing step as an idealized process then our results provide a security guarantee for this technique. Observe that this result is still cast in the random oracle model; yet it separates the security of the key derivation process from the security of the asymmetric encryption scheme and can be seen as a partial instantiation for the random oracles in the encryption algorithm.

PROSPECT. The random oracle model should provide confidence that the design of a cryptographic scheme is sound, even if a security proof in the standard model for this scheme is missing. The heuristic argument is that “good” instantiations of random oracles then give evidence that no “clever” attacks against a scheme work. But the well-known negative results about the random oracle principle have raised some doubts how much confidence this security heuristic really gives.

The approach we take here towards challenging the doubts is to trade security goals against partial or full instantiations of random oracles. Our “test case” OAEP shows that this is a viable way and gives more insights in “how clever” attacks against the instantiations would have to be. And while this still does not rule out the possibility of extraordinary attacks we see this as an important supplement to the random oracle heuristic and to the question how instantiating candidates should be selected, hopefully inciting other results along this direction.

2 Preliminaries

If S is a set then $x \stackrel{\$}{\leftarrow} S$ means that the value x is chosen uniformly at random from S . If \mathcal{A} is a deterministic (resp. randomized algorithm) with a single output then $x \leftarrow \mathcal{A}(y, z, \dots)$ (resp. $x \stackrel{\$}{\leftarrow} \mathcal{A}(y, z, \dots)$) means that the value x is assigned the output of \mathcal{A} for input (y, z, \dots) . An algorithm is called efficient if it runs in polynomial time in the input length (which, in our case, usually refers to polynomial time in the security parameter).

A function family $F = \bigcup_k F(1^k)$ consists of sets of functions $F(1^k) = \{f : \{0, 1\}^{m(k)} \rightarrow \{0, 1\}^{n(k)}\}$. It is called a family of trapdoor permutations if for each $f \in F(1^k)$ there exists f^{-1} such that $f(f^{-1}) \equiv \text{ID}$. We usually identify the functions f and f^{-1} simply with their descriptions, and write $(f, f^{-1}) \stackrel{\$}{\leftarrow} F(1^k)$ for the random choice of f (specifying also f^{-1}) from the family $F(1^k)$. Unless stated differently the minimal assumption about a function family in this paper is that it is one-way, and that it is efficiently computable.

2.1 The OAEP Framework

The OAEP encryption framework [5] is parameterized by integers k, k_0 and k_1 (where k_0, k_1 are linear functions of k) and makes use of a trapdoor permutation family F with domain and range $\{0, 1\}^k$ and two random oracles

$$G: \{0, 1\}^{k_0} \rightarrow \{0, 1\}^{k-k_0} \quad \text{and} \quad H: \{0, 1\}^{k-k_0} \rightarrow \{0, 1\}^{k_0}.$$

The message space is $\{0, 1\}^{k-k_0-k_1}$. The scheme $\text{OAEP}^{G,H}[F] = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ is defined as follows:

- The key generation algorithm $\mathcal{K}(1^k)$ picks a pair $(f, f^{-1}) \leftarrow F(1^k)$ at random. Let pk specify f and let sk specify f^{-1} .
- The encryption algorithm $\mathcal{E}(pk, M)$ picks $r \stackrel{\$}{\leftarrow} \{0, 1\}^{k_0}$, and computes $s \leftarrow G(r) \oplus (M \| 0^{k_1})$ and $t \leftarrow H(s) \oplus r$. It finally outputs $C \leftarrow f(s \| t)$.
- The decryption algorithm $\mathcal{D}(sk, C)$ computes $s \| t \leftarrow f^{-1}(C)$, $r \leftarrow t \oplus H(s)$ and $M \leftarrow s \oplus G(r)$. If the last k_1 bits of M are zeros, then it returns the first $k - k_0 - k_1$ bits of M , else it returns \perp .

The encryption scheme $\text{OAEP}^{G,H}[F]$ is IND-CCA2 secure in the RO model if the underlying trapdoor permutation family F is partial one-way [15].

As a side effect of the partial one-wayness result for OAEP [15] we can immediately conclude security of a particular OAEP variant, where we use partial one-way trapdoor permutation family $F_{t\text{-clear}}$ based on a trapdoor permutation function family F . Namely, each function $f_{t\text{-clear}} : \{0, 1\}^k \rightarrow \{0, 1\}^k$ in $F_{t\text{-clear}}$ is described by $f_{t\text{-clear}}(s||t) \equiv f(s)||\text{ID}(t) = f(s)||t$ for a one-way permutation $f : \{0, 1\}^{k-k_0} \rightarrow \{0, 1\}^{k-k_0}$, i.e., the t -part is output in clear. A random instance $(f_{t\text{-clear}}, f_{t\text{-clear}}^{-1}) \leftarrow F_{t\text{-clear}}(1^k)$ is sampled by picking $(f, f^{-1}) \leftarrow F(1^k)$ and setting $f_{t\text{-clear}}$ as above (the inverse $f_{t\text{-clear}}^{-1}$ is straightforwardly defined). Then $F_{t\text{-clear}}$ is clearly partial one-way and thus $\text{OAEP}^{G,H}[F_{t\text{-clear}}]$ IND-CCA2 secure in the random oracle model.

Analogously, we consider another important variant of OAEP where we also output the k_1 least significant bits $\text{lsb}_{k_1}(s)$ of s in clear and merely apply the trapdoor function f to the leading $k - k_0 - k_1$ bits of s . That is, a random function $f_{\text{lsb}||t\text{-clear}} : \{0, 1\}^k \rightarrow \{0, 1\}^k$ in $F_{\text{lsb}||t\text{-clear}}(1^k)$ is described by a random trapdoor permutation $f : \{0, 1\}^{k-k_0-k_1} \rightarrow \{0, 1\}^{k-k_0-k_1}$ and $f_{\text{lsb}||t\text{-clear}}(s||t) = f(s_{1\dots k-k_0-k_1})||\text{lsb}_{k_1}(s)||t$. Note that since $s = G(r) \oplus M||0^{k_1}$ this means that we output the least significant bits $\text{lsb}_{k_1}(G(r))$ of $G(r)$ and t in clear. For this reason we sometimes write $s||\gamma$ instead of s and denote by γ the k_1 bits $\text{lsb}_{k_1}(G(r))$ such that $f_{\text{lsb}||t\text{-clear}}(s||\gamma||t) = f(s)||\gamma||t$. $F_{\text{lsb}||t\text{-clear}}$ is clearly partial one-way and $\text{OAEP}^{G,H}[F_{\text{lsb}||t\text{-clear}}]$ is IND-CCA2 secure in the random oracle model.

In both cases we often identify $F_{t\text{-clear}}$ resp. $F_{\text{lsb}||t\text{-clear}}$ simply with the underlying family F and vice versa. In particular we often denote a random function from $F_{t\text{-clear}}$ or $F_{\text{lsb}||t\text{-clear}}$ simply by f . We call $F_{t\text{-clear}}$ resp. $F_{\text{lsb}||t\text{-clear}}$ *the induced family of F* .

RANDOM ORACLE INSTANTIATIONS. For an instantiation of the random oracle G in $\text{OAEP}^{G,H}[F]$ we consider a pair of efficient algorithms $\mathcal{G} = (\text{KGenG}, \text{G})$ where KGenG on input 1^k returns a random key K and the deterministic algorithm⁴ G maps this key K and input $r \in \{0, 1\}^{k_0}$ to an output string $\text{G}(K, r) = \text{G}_K(r)$ of $k - k_0$ bits. Then we write $\text{OAEP}^{\mathcal{G},H}[F]$ for the encryption scheme which works as defined above, but where the key pair (sk, pk) is now given by $sk = (f^{-1}, K)$ and $pk = (f, K)$ and where each evaluation of $G(r)$ is replaced by $\text{G}_K(r)$. We say that $\text{OAEP}^{\mathcal{G},H}[F]$ is a *partial G -instantiation of OAEP through \mathcal{G}* .

A *partial H -instantiation $\text{OAEP}^{G,\mathcal{H}}[F]$ of OAEP through \mathcal{H}* and partial instantiations of the aforementioned OAEP variations are defined accordingly. If we instantiate both oracles G, H simultaneously then we speak of a *full instantiation $\text{OAEP}^{\mathcal{G},\mathcal{H}}[F]$ of OAEP through \mathcal{G} and \mathcal{H}* .

2.2 Security of Encryption Schemes

In this section we review the relevant security notions for asymmetric encryption schemes $\text{AS} = (\mathcal{K}, \mathcal{E}, \mathcal{D})$. In addition to indistinguishability under chosen-plaintext and chosen-ciphertext attacks (IND-CPA, IND-CCA1, IND-CCA2) — see Appendix A for formal definitions — we occasionally also rely on the notions of non-malleability. This notion was introduced and formalized in [14, 3]. The most basic version is called NM-CPA and says that a ciphertext of a message M^* should not help to find a ciphertext of a related message M , where the distribution of message M^* is defined by an efficient distribution \mathcal{M} and related messages are specified by an efficient relation R , both chosen by the adversary.

Definition 2.1 (NM-CPA) *Let AS be an asymmetric encryption scheme. Then AS is called secure in the sense of NM-CPA if for every efficient algorithm \mathcal{A} the following random variables $\text{Exp}_{\text{AS},\mathcal{A}}^{\text{nm-cpa-1}}(k)$, $\text{Exp}_{\text{AS},\mathcal{A}}^{\text{nm-cpa-0}}(k)$ are computationally indistinguishable:*

⁴In general, the instantiating functions can be randomized. This requires some care with the decryption algorithms and possibly introduces new attacks. Since our results all hold with respect to deterministic algorithms this is beyond our scope here; see [9] for more details.

Experiment $\text{Exp}_{\text{AS},\mathcal{A}}^{\text{nm-cpa-1}}(k)$ $(pk, sk) \xleftarrow{\$} \mathcal{K}(1^k)$ $(\mathcal{M}, \text{state}) \xleftarrow{\$} \mathcal{A}(pk)$ $M^* \xleftarrow{\$} \mathcal{M}$ $C^* \xleftarrow{\$} \mathcal{E}_{pk}(M^*)$ $(R, C) \xleftarrow{\$} \mathcal{A}(\text{state}, C^*)$ $M \leftarrow \mathcal{D}_{sk}(C)$ Return 1 iff $(C \neq C^*) \wedge R(M^*, M)$	Experiment $\text{Exp}_{\text{AS},\mathcal{A}}^{\text{nm-cpa-0}}(k)$ $(pk, sk) \xleftarrow{\$} \mathcal{K}(1^k)$ $(\mathcal{M}, \text{state}) \xleftarrow{\$} \mathcal{A}(pk)$ $M^* \xleftarrow{\$} \mathcal{M}; M' \xleftarrow{\$} \mathcal{M}$ $C' \xleftarrow{\$} \mathcal{E}_{pk}(M')$ $(R, C) \xleftarrow{\$} \mathcal{A}(\text{state}, C')$ $M \leftarrow \mathcal{D}_{sk}(C)$ Return 1 iff $(C \neq C') \wedge R(M^*, M)$
---	---

It is assumed that the messages in the support of \mathcal{M} have equal length.

We note that the original definition of NM-CPA in [3] actually allows the adversary to output a vector of ciphertexts. Our results for OAEP merely hold with respect to binary relations and therefore we restrict the definition here to such relations. We remark that the aforementioned relationships of NM-CPA to the indistinguishability notions, e.g., that this notion is strictly stronger than the one of IND-CPA, hold for relations of arity two as well.

We define a weaker security notion is that of $\$$ NM-CPA where the adversary does not have the ability to choose a distribution over the messages, but where a random message is encrypted and the adversary tries to find a ciphertext of a related message.

Definition 2.2 ($\$$ NM-CPA) Let $\text{AS} = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an asymmetric encryption scheme and let \mathcal{M} for input 1^k describe the uniform distribution over all $\ell(k)$ bit strings for some polynomial ℓ . Then AS is called secure in the sense of $\$$ NM-CPA if for every efficient algorithm \mathcal{A} and for every efficient relation R the following random variables $\text{Exp}_{\text{AS},\mathcal{A},\mathcal{M},R}^{\$nm-cpa-1}(k)$, $\text{Exp}_{\text{AS},\mathcal{A},\mathcal{M},R}^{\$nm-cpa-0}(k)$ are computationally indistinguishable:

Experiment $\text{Exp}_{\text{AS},\mathcal{A},\mathcal{M},R}^{\$nm-cpa-1}(k)$ $(pk, sk) \xleftarrow{\$} \mathcal{K}(1^k)$ $M^* \xleftarrow{\$} \mathcal{M}(1^k)$ $C^* \xleftarrow{\$} \mathcal{E}_{pk}(M^*)$ $C \xleftarrow{\$} \mathcal{A}(pk, C^*, \langle R \rangle)$ $M \leftarrow \mathcal{D}_{sk}(C)$ Return 1 iff $(C \neq C^*) \wedge R(M^*, M)$	Experiment $\text{Exp}_{\text{AS},\mathcal{A},\mathcal{M},R}^{\$nm-cpa-0}(k)$ $(pk, sk) \xleftarrow{\$} \mathcal{K}(1^k)$ $M^* \xleftarrow{\$} \mathcal{M}(1^k); M' \xleftarrow{\$} \mathcal{M}(1^k)$ $C' \xleftarrow{\$} \mathcal{E}_{pk}(M')$ $C \xleftarrow{\$} \mathcal{A}(pk, C', \langle R \rangle)$ $M \leftarrow \mathcal{D}_{sk}(C)$ Return 1 iff $(C \neq C') \wedge R(M^*, M)$
--	---

While the notion of $\$$ NM-CPA is weaker than the one of NM-CPA —in addition to the restriction to uniformly distributed messages the relation is now fixed in advance— it yet suffices for example to show security in the sense of OW-CPA (where the adversary’s goal is to recover a random message in a given ciphertext) and it also covers Bleichenbacher’s attack on PKCS #1 v1.5. In Section 5 we also show that the notion of $\$$ NM-CPA is enough to derive NM-CPA security under an idealized key derivation function. Namely, one encrypts a random string r under the $\$$ NM-CPA public-key encryption scheme and then pipes r through a random oracle G to derive a key $K = G(r)$ for the symmetric scheme. In fact, one can view the SSL encryption method where the client sends an encrypted random key to the server and both parties derive a symmetric key through a complicated hash function operation as a special case of this method. Then this result about lifting $\$$ NM-CPA to NM-CPA security, together with the $\$$ NM-CPA security proof for the full instantiation of $\text{OAEP}_{\text{lsb}||\text{t-clear}}$, provides an interesting security heuristic (as long as the key derivation process behaves in an ideal way).

2.3 Pseudorandom Generators

Typically, the minimal expected requirement when instantiating a random oracle is that the instantiating function describes a pseudorandom generator, consisting of the key generation algorithm KGen producing a public key K and the evaluation algorithm G mapping a random seed r with key K to the pseudorandom output. Usually the output of this generator should still look random when some side information $\text{hint}(r)$

about the seed r is given. This probabilistic function hint must be of course uninvertible, a weaker notion than one-wayness (cf. [10]).

We also incorporate into the definition the possibility that the key generation algorithm outputs some secret trapdoor information K^{-1} in addition to K . Given this information K^{-1} one can efficiently invert images. If this trapdoor property is not required we can assume that $K^{-1} = \perp$ and often omit K^{-1} in the key generator's output.

Definition 2.3 ((Trapdoor) Pseudorandom Generator) *Let KGen be an efficient key-generation algorithm that takes as input 1^k for $k \in \mathbb{N}$ and outputs a key K ; let G be an efficient deterministic evaluation algorithm that, on input K and a string $r \in \{0, 1\}^k$ returns a string of length $\ell(k)$. Then $\mathcal{G} = (\text{KGen}, \text{G})$ is called a pseudorandom generator (with respect to hint) if the following random variables are computationally indistinguishable:*

- *Let $K \leftarrow \text{KGen}(1^k)$, $r \xleftarrow{\$} \{0, 1\}^k$, $h \leftarrow \text{hint}(r)$, output $(K, \text{G}(K, r), h)$.*
- *Let $K \leftarrow \text{KGen}(1^k)$, $r \xleftarrow{\$} \{0, 1\}^k$, $h \leftarrow \text{hint}(r)$, $u \leftarrow \{0, 1\}^{\ell(n)}$, output (K, u, h) .*

Furthermore, if there is an efficient algorithm TdG such that for any $k \in \mathbb{N}$, any $(K, K^{-1}) \leftarrow \text{KGen}(1^k)$, any $r \in \{0, 1\}^k$ we have $\text{G}(K, \text{TdG}(K^{-1}, \text{G}(K, r))) = \text{G}(K, r)$ then $(\text{KGen}, \text{G}, \text{TdG})$ is called a trapdoor pseudorandom generator.

For our results about OAEP we often need further properties from the pseudorandom generator, including near-collision resistance and non-malleability. The former means that given a seed r it is hard to find a different seed r' such that $\text{G}(K, r)$ and $\text{G}(K, r')$ coincide on a predetermined set of bits (even if they are allowed to differ on the other bits). Non-malleability refers to generators where the generator's output for a seed should not help to produce an image of a related seed. We give precise definitions and details concerning existential questions on site.

3 Partial Instantiations for OAEP

In this section we prove security of partial instantiations of OAEP. Our results show that one can replace either one of the random oracle in OAEP by reasonable primitives and still maintain security (in the random oracle model).

3.1 Instantiating the G -Oracle for IND-CCA2 security

We first show how to construct a pseudorandom generator with a special form of collision-resistance. This property says that finding an input r' to a random input r , such that $\text{G}(K, r)$ and $\text{G}(K, r')$ coincide on the k least significant bits $\text{lsb}_k(\text{G}(K, r))$, $\text{lsb}_k(\text{G}(K, r'))$, is infeasible. According to comparable collision types for hash functions [6] we call this *near-collision resistance*.

Definition 3.1 (Near-collision Resistant Pseudorandom Generator) *A pseudorandom generator $\mathcal{G} = (\text{KGen}, \text{G})$ is called near-collision resistant (for the least significant k bits) if for any efficient algorithm \mathcal{C} the following holds: Let $K \leftarrow \text{KGen}(1^k)$, $r \leftarrow \{0, 1\}^k$, $r' \leftarrow \mathcal{C}(K, r)$. Then the probability that $r \neq r'$ but $\text{lsb}_k(\text{G}(K, r)) = \text{lsb}_k(\text{G}(K, r'))$ is negligible.*

Near-collision resistant generators can be built, for example, from one-way permutations via the well-known Yao-Blum-Micali construction [21, 8]. In that case, given a family G of one-way permutations the key generation algorithm $\text{KGen}_{\text{YBM}}(1^k)$ of this generator simply picks a random instance $g : \{0, 1\}^k \rightarrow \{0, 1\}^k$ of $G(1^k)$, and $\text{G}_{\text{YBM}}(g, r) = (\text{hb}(r), \text{hb}(g(r)), \dots, \text{hb}(g^{n-1}(r)), g^n(r))$ is defined through the hardcore bits hb of g . Since g is a permutation different inputs $r \neq r'$ yield different output parts $g^n(r) \neq g^n(r')$.

Given a near-collision resistant pseudorandom generator we show how to instantiate the G -oracle in $\text{OAEP}^{G, H}[F_{\text{t-clear}}]$ for the family $F_{\text{t-clear}}$ which is induced by a trapdoor permutation family F (i.e., where a member $f : \{0, 1\}^{k-k_0} \rightarrow \{0, 1\}^{k-k_0}$ of F is applied to the k -bit inputs such that the lower k_0 bits are output in clear).

Theorem 3.2 *Let $\mathcal{G} = (\text{KGenG}, \text{G})$ be a pseudorandom generator which is near-collision resistant (for the k_1 least significant bits). Let F be trapdoor permutation family and let $F_{\text{t-clear}}$ be the induced partial one-way trapdoor permutation family defined in Section 2.1. Then the partial \mathcal{G} -instantiation $\text{OAEP}^{\mathcal{G}, H}[F_{\text{t-clear}}]$ of OAEP through \mathcal{G} is IND-CCA2 in the random oracle model.*

The full proof appears in Appendix B. The idea is to gradually change the way the challenge ciphertext (encrypting one of two adversarially chosen messages, the hidden choice made at random) is computed in a sequence of games. We show that each of these steps does not change an adversary’s success probability of predicting the secret choice noticeably:

- Initially, in **Game**⁰ the challenge ciphertext $f(s^*)||t^*$ for message M^* is computed as in the scheme’s description by $s^* = \text{G}(K, r^*) \oplus M^*||0^{k_1}$ for the near-collision resistant generator G and $t^* = H(s^*) \oplus r^*$ for random oracle H .
- In **Game**¹ the ciphertext is now computed by setting $s^* = \text{G}(K, r^*) \oplus M^*||0^{k_1}$ as before, but letting $t^* = \omega \oplus r^*$ for a random ω which is independent of $H(s^*)$. Because H is a random oracle this will not affect the adversary’s success probability, except for the rare case that the adversary queries H about s^* .
- In **Game**², in a rather cosmetic change, we further substitute $t^* = \omega \oplus r^*$ simply for $t^* = \omega$, making the t -part independent of the generator’s pre-image r^* .
- in **Game**³ we use the pseudorandomness of generator G to replace $s^* = \text{G}(K, r^*) \oplus M^*||0^{k_1}$ by $s^* = u \oplus M^*||0^{k_1}$ for a random u .

Since ciphertexts in the last game are distributed independently of the actual message security of the original scheme follows, after a careful analysis that decryption queries do not help; this is the step where we exploit that H is still a random oracle and that \mathcal{G} is near-collision resistant. Namely, the near-collision resistance prevents an adversary from transforming the challenge ciphertext for values r^*, s^* into a valid one for the same s^* but a different r ; otherwise the least significant bits of $s^* = \text{G}(K, r^*) \oplus M^*||0^{k_1} = \text{G}(K, r) \oplus M^*||0^{k_1}$ would not coincide and the derived ciphertext would be invalid with high probability. Given this, the adversary must always use a “fresh” value s when submitting a ciphertext to the decryption oracle, and must have queried the random oracle H about s before (or else the ciphertext is most likely invalid). But then the adversary already “knows” $r = t \oplus H(s)$ —recall that for $F_{\text{t-clear}}$ the t -part is included in clear in ciphertexts—and therefore “knows” the (padded) message $M||z = s \oplus \text{G}(K, r)$ encapsulated in the ciphertext.

3.2 Instantiating the H -Oracle

To instantiate the H -oracle we introduce the notion of a non-malleable pseudorandom generator. For such a pseudorandom generator it should be infeasible to find for a given image $y^* = \text{H}_K(s^*)$ of a random s^* a different image $y = \text{H}_K(s)$ of a related value s , where the corresponding efficient relation $R(s^*, s)$ must be determined *before* seeing K and y^* .⁵ More precisely, we formalize non-malleability of a pseudorandom generator by the indistinguishability of two experiments. For any adversary \mathcal{B} it should not matter whether \mathcal{B} is given $f(s^*)$, $y^* = \text{H}_K(s^*)$ or $f(s^*)$, $y' = \text{H}_K(s')$ for an independent s' instead: the probability that \mathcal{B} outputs $f(s)$ and $y = \text{H}_K(s)$ such that s is related to s^* via relation R should be roughly the same in both cases.⁶

Definition 3.3 (Non-Malleable Pseudorandom Generator) *Assume $\mathcal{H} = (\text{KGenH}, \text{H})$ is a pseudorandom generator (which is pseudorandom with respect to $\text{hint}(x) = (f, f(x))$ for $(f, f^{-1}) \leftarrow F(1^k)$ from the trapdoor function family F). Then \mathcal{H} is called non-malleable with respect to hint if for any efficient algorithm \mathcal{B} and any efficient relation R the following random variables $\text{Exp}_{\mathcal{H}, \mathcal{B}, F, R}^{\text{nm-cma-1}}(k)$, $\text{Exp}_{\mathcal{H}, \mathcal{B}, F, R}^{\text{nm-cma-0}}(k)$ are computationally indistinguishable, where the experiments are defined as follows.*

⁵We are thankful to the people from the Ecrypt network for pointing out that a possibly stronger definition for adaptively chosen relations allows trivial relations over the images and cannot be satisfied.

⁶Adding the image under the trapdoor permutation uniquely determines the pre-image of the pseudorandom generator’s output and enables us to specify $R(s^*, s)$ via the pre-images. Since this also bundles the security of the trapdoor permutation and the generator, Brown’s recent impossibility result about security reductions for OAEP [2] does not apply.

Experiment $\text{Exp}_{G,B,F,R}^{nm\text{-cpa-1}}(k)$

$K \xleftarrow{\$} \text{KGenH}(1^k)$
 $(f, f^{-1}) \xleftarrow{\$} F$
 $s^* \xleftarrow{\$} \{0, 1\}^k$
 $y^* \xleftarrow{\$} \text{H}_K(s^*)$
 $(z, y) \xleftarrow{\$} \mathcal{B}(K, f, f(s^*), y^*)$
 $s \leftarrow f^{-1}(z)$
 Return 1 iff
 $R(s^*, s) \wedge \text{H}_K(s) = y \wedge s^* \neq s$

Experiment $\text{Exp}_{G,B,F,R}^{nm\text{-cpa-0}}(k)$

$K \xleftarrow{\$} \text{KGenH}(1^k)$
 $(f, f^{-1}) \xleftarrow{\$} F$
 $s^* \xleftarrow{\$} \{0, 1\}^k ; s' \xleftarrow{\$} \{0, 1\}^k$
 $y' \xleftarrow{\$} \text{H}_K(s')$
 $(z, y) \xleftarrow{\$} \mathcal{B}(K, f, f(s^*), y')$
 $s \leftarrow f^{-1}(z)$
 Return 1 iff
 $R(s^*, s) \wedge \text{H}_K(s) = y \wedge s^* \neq s$

Given a non-malleable pseudorandom generator we can prove NM-CPA security of the partial H -instantiation of OAEP, under the restriction that the adversarial chosen message distribution and relation are defined at the beginning of the attack via $(\mathcal{M}, R, \text{state}) \leftarrow \mathcal{A}(1^k)$ and thus depend only the security parameter. This relaxed notion still implies for example IND-CPA security (but for messages picked independently of the public key), is still incomparable to IND-CCA1 security, and also thwarts Bleichenbacher's attack. We call such schemes *NM-CPA for pre-defined message distributions and relations*.

Theorem 3.4 *Let F be a trapdoor permutation family and let $F_{t\text{-clear}}$ be the induced partial one-way trapdoor permutation family. Let $\mathcal{H} = (\text{KGenH}, \text{H})$ be a pseudorandom generator (with respect to $\text{hint}(x) = (f, f(x))$ for $(f, f^{-1}) \leftarrow F(1^k)$). Assume further that \mathcal{H} is non-malleable with respect to hint . Then the partial H -instantiation $\text{OAEP}^{G, \mathcal{H}}_{[F_{t\text{-clear}}]}$ through \mathcal{H} is NM-CPA for pre-defined message distributions and relations in the random oracle model.*

The proof idea is as follows. Assume that an attacker, given a ciphertext for some values r^*, s^* (which uniquely define the message in a ciphertext), tries to prepare a related ciphertext for some value $r \neq r^*$, without having queried random oracle G about r before. Then such a ciphertext is most likely invalid because with overwhelming probability the least significant bits of $s \oplus G(r)$ are not zero. Else, if $r = r^*$, then we must have $f(s) \neq f(s^*)$ and $s \neq s^*$, since the adversarial ciphertext must be different for a successful attack. But then the values $\text{H}(K, s^*)$ and $\text{H}(K, s)$ for different pre-images must be related via the ciphertext's relation, contradicting the non-malleability of the generator H . In any other case, if $r \neq r^*$ and r is among the queries to G , the random value $G(r^*)$ is independent of $G(r)$. So must be the messages $M^* || 0^{k_1} = s^* \oplus G(r^*)$ and $M || 0^{k_1} = s \oplus G(r)$, as required for non-malleability. Details can be found in Appendix C.

Replacing the H -oracle without violating IND-CCA2 security is more ambitious and we require a very strong assumption on the pseudorandom generator, called non-malleability under chosen-image attacks (where the adversary can also make inversion queries to the trapdoor pseudorandom generator). Since any pseudorandom generator with this property is already close to a chosen-ciphertext secure encryption scheme, we rather see this as an indication that a partial instantiation might be possible and that separation results as [11, 18, 19, 1, 16, 20, 9, 13] seem to be hard to find. The formal treatment of the following and the proof appear in Appendix D

Theorem 3.5 *Let F be trapdoor permutation family and let $F_{t\text{-clear}}$ be the induced partial one-way trapdoor permutation family defined in Section 2.1. Let $\mathcal{H} = (\text{KGenH}, \text{H}, \text{TdH})$ be a trapdoor pseudorandom generator which is non-malleable under chosen-image attacks (with respect to $\text{hint}(x) = (f, f(x))$ for $(f, f^{-1}) \leftarrow F_{t\text{-clear}}(1^k)$). Then the partial H -instantiation $\text{OAEP}^{G, \mathcal{H}}_{[F_{t\text{-clear}}]}$ through \mathcal{H} is IND-CCA2 in the random oracle model.*

4 Full Instantiation for OAEP

In this section we prove that there exists a full instantiation of $\text{OAEP}_{\text{lsb}||t\text{-clear}}$ which is secure in the sense of $\$$ NM-CPA in the standard model, implying for example that the scheme is OW-CPA. Recall that in $\text{OAEP}_{\text{lsb}||t\text{-clear}}$ we write $s || \gamma = G(s) \oplus M || 0^{k_1}$ instead of s to name the least significant bits explicitly.

To prove our result we need a near-collision resistant *trapdoor* pseudorandom generator, i.e., which combines near-collision resistance with the trapdoor property. Such generators can be easily built by using again the Blum-Micali-Yao generator, but this time by deploying a trapdoor permutation g instead of a one-way permutation, i.e., the generator's output for random r is given by $G_{\text{YBM}}(g, r) = (\text{hb}(r), \text{hb}(g(r)), \dots, \text{hb}(g^{n-1}(r)), g^n(r))$.

Letting K^{-1} contain the trapdoor information g^{-1} algorithm TdG can easily invert the k_1 least significant bits y of the output to recover a pre-image r .

To be precise we make use of two additional, specific properties of the Blum-Micali-Yao generator. First, we assume that recovering a pre-image is possible given the k_1 least significant bits only, i.e., without seeing the remaining part of the image. To simplify the proof we furthermore presume that the k_1 least significant bits of the generator's output are statistically close to uniform (over the choice of the seed).⁷ We simply refer to generators with the above properties as a *near-collision resistant trapdoor pseudorandom generator (for the least significant k bits)*.

Theorem 4.1 *Let F be trapdoor permutation family and let $F_{|sb||t-clear}$ be the induced partial one-way trapdoor permutation family. Let $\mathcal{G} = (\text{KGenG}, G)$ be a near-collision resistant trapdoor pseudorandom generator (for the k_1 least significant bits). Let $\mathcal{H} = (\text{KGenH}, H)$ be a generator which is pseudorandom and non-malleable with respect to $\text{hint}(s||\gamma) = (f, f(s)||\gamma)$ for $(f, f^{-1}) \leftarrow F(1^k)$. Then the full instantiation $\text{OAEP}^{\mathcal{G}, \mathcal{H}}[F_{|sb||t-clear}]$ through \mathcal{G} and \mathcal{H} is $\$NM\text{-CPA}$.*

The proof appears in Appendix E. The basic idea is similar to the one of NM-CPA security for the partial H -instantiation. The important difference is that the randomness of the encrypted message M in a ciphertext $f(s)||\gamma||t$ for $s||\gamma = G_K(r) \oplus M||0^{k_1}$ helps to overcome otherwise existing ‘‘circular’’ dependencies between \mathcal{G} and \mathcal{H} in the computations of ciphertexts (which, in the partial instantiation case, do not occur due to the fact that G is a random oracle).

5 Hybrid Encryption from $\$NM\text{-CPA}$ Schemes

We show that a public-key scheme which is secure in the sense of $\$NM\text{-CPA}$ (i.e., for pre-defined relations), together with an IND-CCA2 secure symmetric scheme suffices to build a NM-CPA secure hybrid scheme in the random oracle model (i.e., even for adaptively chosen message distributions and relations).

Construction 5.1 *Let $\text{AS} = (\mathcal{EK}_{asym}, \mathcal{E}_{asym}, \mathcal{D}_{asym})$ be an asymmetric encryption scheme and let $\text{SS} = (\mathcal{EK}_{sym}, \mathcal{E}_{sym}, \mathcal{D}_{sym})$ be a symmetric encryption scheme. Let G be a hash function mapping k -bit strings into the key space of the symmetric scheme. Then the hybrid encryption scheme $\text{AS}' = (\mathcal{EK}'_{asym}, \mathcal{E}'_{asym}, \mathcal{D}'_{asym})$ is defined as follows.*

- The key generation algorithm $\mathcal{EK}'_{asym}(1^k)$ outputs a key pair $(sk, pk) \stackrel{\$}{\leftarrow} \mathcal{EK}_{asym}(1^k)$.
- The encryption algorithm \mathcal{E}'_{asym} on input pk, M picks $r \stackrel{\$}{\leftarrow} \{0, 1\}^k$, computes $C_{asym} \stackrel{\$}{\leftarrow} \mathcal{E}_{asym}(pk, r)$, $C_{sym} \stackrel{\$}{\leftarrow} \mathcal{E}_{sym}(G(r), M)$ and returns (C_{asym}, C_{sym}) .
- The decryption algorithm \mathcal{D}'_{asym} on input (C_{asym}, C_{sym}) and sk computes $r \leftarrow \mathcal{D}_{asym}(sk, C_{asym})$, $M \leftarrow \mathcal{D}_{sym}(G(r), C_{sym})$ and returns M .

Theorem 5.2 *Let $\text{AS} = (\mathcal{EK}_{asym}, \mathcal{E}_{asym}, \mathcal{D}_{asym})$ be an asymmetric encryption scheme which is $\$NM\text{-CPA}$. Let $\text{SS} = (\mathcal{EK}_{sym}, \mathcal{E}_{sym}, \mathcal{D}_{sym})$ be an IND-CCA2 symmetric encryption scheme. Let G be a hash function and assume $\text{AS}' = (\mathcal{EK}'_{asym}, \mathcal{E}'_{asym}, \mathcal{D}'_{asym})$ is the hybrid encryption scheme defined according to Construction 5.1. Then AS' is NM-CPA secure in the random oracle model.*

The proof is in Appendix F and actually shows that the scheme is NM-CPA with respect to the stronger notion where the adversary outputs a sequence $\mathbf{C} = (C_1, \dots, C_m)$ of ciphertexts and the success is measured according to $R(M^*, \mathbf{M})$ for $\mathbf{M} = (M_1, \dots, M_m)$.

⁷It is easy to adapt the proof to the more general case of arbitrary distributions of the least significant bits, as long as they support extraction. But this would also require to change the definition of the non-malleable pseudorandom generator $G_{KG}(s||\gamma)$ to support arbitrary distributions on the γ -part.

Acknowledgments

We thank the anonymous reviewers for comments. Part of the work done while both authors were visiting Centre de Recerca Matemàtica (CRM) and Technical University of Catalonia (UPC), Barcelona, Spain, whose support is highly appreciated. The second author was also supported by the Emmy Noether Program Fi 940/2-1 of the German Research Foundation (DFG).

References

- [1] M. Bellare, A. Boldyreva and A. Palacio. An uninstantiable random-oracle-model scheme for a hybrid-encryption problem. In *Eurocrypt 2004*, Volume 3027 of *LNCS*, pp. 171–188. Springer-Verlag, 2004.
- [2] D. R. L. Brown. Unprovable Security of RSA-OAEP in the Standard Model. *Cryptology ePrint Archive, Report 2006/223*, 2006.
- [3] M. Bellare, A. Desai, D. Pointcheval and P. Rogaway. Relations Among Notions of Security for Public-Key Encryption Schemes. In *CRYPTO '98*, Volume 1462 of *LNCS*, pp. 26–45. Springer-Verlag, 1998.
- [4] M. Bellare and P. Rogaway. Random oracles are practical: a paradigm for designing efficient protocols. In *CCS '93*, pp. 62–73. ACM, 1993.
- [5] M. Bellare and P. Rogaway. Optimal asymmetric encryption – how to encrypt with RSA. In *Eurocrypt '94*, Volume 950 of *LNCS*, pp. 92–111. Springer-Verlag, 1995.
- [6] E. Biham and R. Chen. Near-Collisions of SHA-0. In *CRYPTO' 2004*, Volume 3152 of *LNCS*, pp. 290–305. Springer-Verlag, 2004.
- [7] D. Bleichenbacher. Chosen Ciphertext Attacks Against Protocols Based on the RSA Encryption Standard PKCS #1. In *CRYPTO '98*, Volume 1462 of *LNCS*, pp. 1–12. Springer-Verlag, 1998.
- [8] M. Blum and S. Micali. How to generate cryptographically strong sequences of pseudorandom bits. *Journal on Computing*, Volume 13, pp. 850–864, SIAM, 1984.
- [9] A. Boldyreva and M. Fischlin. Analysis of random-oracle instantiation scenarios for OAEP and other practical schemes. In *CRYPTO 2005*, Volume 3621 of *LNCS*, pp. 412–429. Springer-Verlag, 2005.
- [10] R. Canetti. Towards realizing random oracles: Hash functions that hide all partial information. In *CRYPTO '97*, Volume 1294 of *LNCS*. pp. 455–469. Springer-Verlag, 1997.
- [11] R. Canetti, O. Goldreich and S. Halevi. The random oracle methodology, revisited. In *STOC '98*, pp. 209–218. ACM, 1998.
- [12] R. Canetti, D. Micciancio and O. Reingold. Perfectly one-way probabilistic hash functions. In *STOC '98*, pp. 131–140. ACM, 1998.
- [13] Y. Dodis, R. Oliveira, and K. Pietrzak. On the generic insecurity of full-domain hash. In *CRYPTO 2005*, Volume 3621 of *LNCS*, pp. 449–466. Springer-Verlag, 2005.
- [14] D. Dolev, C. Dwork and M. Naor. Non-malleable cryptography. *Journal on Computing*, Vol. 30(2), pp. 391–437. SIAM, 2000.
- [15] E. Fujisaki, T. Okamoto, D. Pointcheval and J. Stern. RSA-OAEP is secure under the RSA assumption. In *CRYPTO 2001*, volume 2139 of *LNCS*, pp. 260–274. Springer-Verlag, 2001.
- [16] S. Goldwasser and Y. T. Kalai. On the (in)security of the Fiat-Shamir paradigm. In *FOCS 2003*. IEEE, 2003.
- [17] IETF-TLS Working Group. Transport Layer Security. <http://www.ietf.org/html.charters/tls-charter.html>, November 2005.

- [18] U. Maurer, R. Renner and C. Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In *TCC 2004*, volume 2951 of *LNCS*, pp. 21–39. Springer-Verlag, 2004.
- [19] J. Nielsen. Separating random oracle proofs from complexity theoretic proofs: The non-committing encryption case. In *CRYPTO 2002*, volume 2442 of *LNCS*, pp. 111–126. Springer-Verlag, 2002.
- [20] P. Paillier and D. Vergnaud. Discrete-Log-Based Signatures May Not Be Equivalent to Discrete Log. In *Asiacrypt 2005*, volume 3788 of *LNCS*, pp. 1–20. Springer-Verlag, 2005.
- [21] A. Yao. Theory and applications of trapdoor functions. In *FOCS '82*, pp. 80–91. IEEE, 1982.

A Encryption Schemes and their Security

An asymmetric encryption scheme $\text{AE} = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ is specified by three polynomial-time algorithms with the following functionalities. The randomized *key-generation* algorithm \mathcal{K} takes input 1^k , where k is the security parameter, and outputs a pair (pk, sk) consisting of a public key and a matching secret key, respectively. The randomized *encryption* algorithm \mathcal{E} takes input a public key pk and a message M , and outputs a ciphertext C . The deterministic *decryption* algorithm \mathcal{D} takes input a secret key sk and a ciphertext C , and outputs a message M or a special symbol \perp to indicate that the ciphertext is invalid. Associated to k is a *message space* $\text{MsgSp}(k)$ from which M is allowed to be drawn. For any $(pk, sk) \in [\mathcal{K}(1^k)]$, any $M \in \text{MsgSp}(k)$, it is required that $\mathcal{D}(sk, \mathcal{E}(pk, M)) = M$. The syntax of symmetric encryption schemes is very similar, except the same symmetric key K is used in place of public and secret keys ($pk = sk = K$) and the adversary is denied the key K .

Definition A.1 [Security of Asymmetric Encryption] *Let $\text{AE} = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an asymmetric encryption scheme. Consider experiments $\text{Exp}_{\text{AE}, \mathcal{A}, b}^{\text{enc-ind-cpa}}(k)$, $\text{Exp}_{\text{AE}, \mathcal{A}, b}^{\text{enc-ind-cca}}(k)$ associated to AE , a bit $b \in \{0, 1\}$ and an adversary \mathcal{A} . In both experiments \mathcal{A} is given input a public key pk and access to a left-right encryption oracle $\mathcal{O}_{\text{AE}}(pk, b, \cdot, \cdot)$, where pk and sk are matching keys generated via $(pk, sk) \xleftarrow{\$} \mathcal{K}(1^k)$. The oracle takes input two messages $M_0, M_1 \in \text{MsgSp}(k)$ of equal length and returns a ciphertext $C \xleftarrow{\$} \mathcal{E}(pk, M_b)$. In experiment $\text{Exp}_{\text{AE}, \mathcal{A}, b}^{\text{enc-ind-cca}}(k)$ the adversary is also given input a decryption oracle $\mathcal{D}(sk, \cdot)$. \mathcal{A} queries the oracle(s) on inputs of its choice⁸, with a restriction of not querying its decryption oracle on ciphertexts previously returned by the left-right encryption oracle. \mathcal{A} eventually stops and outputs a guess d which is also the output of the experiment. AE is said to be IND-CPA (resp. IND-CCA) secure if the the following*

$$\Pr[\text{Exp}_{\text{AE}, \mathcal{A}, 1}^{\text{enc-ind-atk}}(k) = 1] - \Pr[\text{Exp}_{\text{AE}, \mathcal{A}, 0}^{\text{enc-ind-atk}}(k) = 1] .$$

is negligible in k .

IND-CPA and IND-CCA security of symmetric encryption schemes is defined similarly, except that adversary is not given any key.

We adopt the convention that the *time complexity* of adversary \mathcal{A} is the execution time of the entire experiment, including the time taken for key generation, and computation of answers to oracle queries. The same convention will be used implicitly in other definitions of the paper.

B Proof of Theorem 3.2

In this section we present the formal proof of the IND-CCA2 security of the partial instantiation of G through a near-collision resistant pseudorandom generator.

Let \mathcal{A} be an arbitrary probabilistic polynomial-time algorithm. Let $\text{Game}_{\mathcal{A}, b}^0(k)$ denote the original attack of \mathcal{A} on the encryption scheme $\text{OAEP}^{\mathcal{G}, H}[F_{\text{t-clear}}]$ where message M_b for fixed bit b is encrypted in the challenge ciphertext. Let $\text{Game}_{\mathcal{A}, b}^1(k)$ denote the game where we replace $H(s^*)$ in the challenge ciphertext by a uniformly

⁸For simplicity in the analyses of asymmetric encryption schemes we assume that an adversary does *at most one* query to its left-right encryption oracle. It is well-known that this restriction does not change the asymptotic advantage of the adversary.

and independently distributed string ω^* . A rather syntactical change then allows us to replace $t^* = \omega^* \oplus r$ by $t^* = \omega^*$ in $\text{Game}_{\mathcal{A},b}^2$. In $\text{Game}_{\mathcal{A},b}^3(k)$ we furthermore replace the value $G(K, r)$ by a independent random string $u^* \leftarrow \{0, 1\}^{k-k_0}$ in the challenge ciphertext. All games are described formally in Figure B.

<p>Experiment $\text{Game}_{\mathcal{A},b}^0(k)$: $((f^{-1}, K), (f, K)) \xleftarrow{\\$} \mathcal{K}(1^k)$ $(M_0, M_1, \text{state}) \xleftarrow{\\$} \mathcal{A}^{H, \mathcal{D}(sk, \cdot)}(f, K)$ Compute ciphertext (C^*, t^*): Pick $r^* \xleftarrow{\\$} \{0, 1\}^{k_0}$ Compute $s^* \leftarrow G(K, r^*) \oplus M_b 0^{k_1}$ Compute $C^* \leftarrow f(s^*)$ Compute $t^* \leftarrow H(s^*) \oplus r^*$ $d \xleftarrow{\\$} \mathcal{A}^{H, \mathcal{D}(sk, \cdot) - \{(C^*, t^*)\}}((C^*, t^*), \text{state})$</p>	<p>Experiment $\text{Game}_{\mathcal{A},b}^1(k)$: $((f^{-1}, K), (f, K)) \xleftarrow{\\$} \mathcal{K}(1^k)$ $(M_0, M_1, \text{state}) \xleftarrow{\\$} \mathcal{A}^{H, \mathcal{D}(sk, \cdot)}(f, K)$ Compute ciphertext (C^*, t^*): Pick $r^* \xleftarrow{\\$} \{0, 1\}^{k_0}$ Compute $s^* \leftarrow G(K, r^*) \oplus M_b 0^{k_1}$ Compute $C^* \leftarrow f(s^*)$ Pick $\omega^* \xleftarrow{\\$} \{0, 1\}^{k_0}$ Compute $t^* \leftarrow \omega^* \oplus r^*$ $d \xleftarrow{\\$} \mathcal{A}^{H, \mathcal{D}(sk, \cdot) - \{(C^*, t^*)\}}((C^*, t^*), \text{state})$</p>
<p>Experiment $\text{Game}_{\mathcal{A},b}^2(k)$: $((f^{-1}, K), (f, K)) \xleftarrow{\\$} \mathcal{K}(1^k)$ $(M_0, M_1, \text{state}) \xleftarrow{\\$} \mathcal{A}^{H, \mathcal{D}(sk, \cdot)}(f, K)$ Compute ciphertext (C^*, t^*): Pick $r^* \xleftarrow{\\$} \{0, 1\}^{k_0}$ Compute $s^* \leftarrow G(K, r^*) \oplus M_b 0^{k_1}$ Compute $C^* \leftarrow f(s^*)$ Pick $\omega^* \xleftarrow{\\$} \{0, 1\}^{k_0}$ Compute $t^* \leftarrow \omega^*$ $d \xleftarrow{\\$} \mathcal{A}^{H, \mathcal{D}(sk, \cdot) - \{(C^*, t^*)\}}((C^*, t^*), \text{state})$</p>	<p>Experiment $\text{Game}_{\mathcal{A},b}^3(k)$: $((f^{-1}, K), (f, K)) \xleftarrow{\\$} \mathcal{K}(1^k)$ $(M_0, M_1, \text{state}) \xleftarrow{\\$} \mathcal{A}^{H, \mathcal{D}(sk, \cdot)}(f, K)$ Compute ciphertext (C^*, t^*): Pick $u^* \xleftarrow{\\$} \{0, 1\}^{k-k_0}$ Compute $s^* \leftarrow u^* \oplus M_b 0^{k_1}$ Compute $C^* \leftarrow f(s^*)$ Pick $\omega^* \xleftarrow{\\$} \{0, 1\}^{k_0}$ Compute $t^* \leftarrow \omega^*$ $d \xleftarrow{\\$} \mathcal{A}^{H, \mathcal{D}(sk, \cdot) - \{(C^*, t^*)\}}((C^*, t^*), \text{state})$</p>

Figure 1: Games in the Proof of Theorem 3.2: Shaded areas indicate the differences between the games. It is always assumed that the output (M_0, M_1, state) of \mathcal{A} in the first phase satisfies $|M_0| = |M_1|$.

Note that in $\text{Game}_{\mathcal{A},b}^3$ the distribution of the data is independent of bit b . Hence, the probabilities $\Pr[\text{Game}_{\mathcal{A},1}^3(k) = 1]$ and $\Pr[\text{Game}_{\mathcal{A},0}^3(k) = 1]$ for $b = 1$ and $b = 0$, respectively, are identical. Therefore,

$$\begin{aligned}
 & \Pr[\text{Game}_{\mathcal{A},1}^0(k) = 1] - \Pr[\text{Game}_{\mathcal{A},0}^0(k) = 1] \\
 &= \sum_{i=0}^2 \Pr[\text{Game}_{\mathcal{A},1}^i(k) = 1] - \Pr[\text{Game}_{\mathcal{A},1}^{i+1}(k) = 1] \\
 & \quad + \Pr[\text{Game}_{\mathcal{A},1}^2(k) = 1] - \Pr[\text{Game}_{\mathcal{A},0}^2(k) = 1] \\
 & \quad + \sum_{i=2}^0 \Pr[\text{Game}_{\mathcal{A},0}^{i+1}(k) = 1] - \Pr[\text{Game}_{\mathcal{A},0}^i(k) = 1]
 \end{aligned}$$

and it suffices to show that $\Pr[\text{Game}_{\mathcal{A},b}^i(k) = 1] - \Pr[\text{Game}_{\mathcal{A},b}^{i+1}(k) = 1]$ for $i = 0, 1, 2$ are negligible for any $b \in \{0, 1\}$. by flipping \mathcal{A} 's output bit we can always assume that the differences are positive. In the sequel we fix the bit b .

SIMULATING THE DECRYPTION ORACLE. We first describe how to simulate decryption queries in the games without knowing the secret key f^{-1} to f . This is accomplished through the random oracle mode and via one procedure D which works for all games. In addition to a ciphertext (C, t) this procedure gets the public data K, f and a list L_H , representing \mathcal{A} 's queries to random oracle H and the answers as input. The procedure

checks if there is exactly one pair (s, ω) in L_H such that $C = f(s)$; if so, then it computes $r \leftarrow t \oplus \omega$ and then $M || z \leftarrow s \oplus \mathsf{G}_{\text{ncr}}(K, r)$. It finally outputs M if $z = 0^{k_1}$. In any other case, if there is no unique entry in L_H or if $z \neq 0^{k_1}$ then it returns \perp .

We next prove that this decryption procedure may substitute the actual decryption oracle except with negligible simulation error probability in all games. More formally, this means that for every decryption request in the game we run D (on the list L_H of communication between \mathcal{A} and H up to this point) instead of \mathcal{D} . Let DecError_i denote the event that D returns a different answer than \mathcal{D} for the i -th decryption query in the corresponding game, given that the first $i - 1$ replies were identical. It then suffices to show that the probability of DecError_i is negligible for arbitrary i . Recall that we call a ciphertext valid iff \mathcal{D} returns a message $M \neq \perp$.

First note that collisions $(s, \omega), (s', \omega)$ for different $s \neq s'$ in the list L_H are unlikely and happen with negligible probability only at any point. This holds in all games, because the values ω are picked at random. So we can condition on the event that there are no such collision and analyze $\Pr[\text{DecError}_i]$ under this condition, i.e., it suffices to discuss the case of missing entries in L_H , as this is the only case when D 's behavior diverges; if there is a unique entry in L_H then D gives the same answer as the genuine decryption oracle.

BEHAVIOR IN GAME ZERO. Assume that \mathcal{A} submits some (C, t) to the decryption oracle with the i -th query in $\text{Game}_{\mathcal{A}, b}^0$ such that there is no matching entry in L_H . Let $s, r, M || z$ denote the unique values such that $f(s) = C$, $r = t \oplus H(s)$ and $M || z = \mathcal{G}_{\text{nm}}(r) \oplus s$. Let $\omega^*, r^*, M_b^* || 0^{k_1}$ denote the corresponding values for the challenge ciphertext (C^*, t^*) .

- If we are in the first phase of the game, before \mathcal{A} receives the challenge ciphertext, and there is no value for s in L_H , then $r = H(s) \oplus t$ is an unknown random value. The probability that the least significant bits of $\mathsf{G}(K, r) \oplus s$ equal 0^{k_1} is therefore negligibly close to 2^{-k_1} by the pseudorandomness of \mathcal{G} . Else it would be easy to construct a distinguisher.
- If we are in the second phase, after having received the challenge ciphertext, and $C \neq C^*$, then we have $s \neq s^*$ and $H(s)$ is again an unknown random value for which the same argument as for the first phase applies.
- If we are in the second phase and $C = C^*$ and thus $s = s^*$ but $t \neq t^*$, then we have $r \neq r^*$ as well. The equation $\mathsf{G}(K, r) \oplus M || z = s = s^* = \mathsf{G}_{\text{ncr}}(K, r^*) \oplus M_b^* || 0^{k_1}$ implies the equation $\text{lsb}_{k_1}(\mathsf{G}(K, r)) \oplus \text{lsb}_{k_1}(\mathsf{G}(K, r^*)) = z$ and in order for the actual decryption oracle to decrypt to a valid message, the least significant bits z must be zero. But the probability of finding such r^* is negligible by the near-collision resistance (otherwise it would be easy to construct a successful collision finder).

Hence, the probability of event DecError_i in $\text{Game}_{\mathcal{A}, b}^0$ is negligible and D simulates \mathcal{D} correctly with overwhelming probability.

BEHAVIOR IN GAMES ONE, TWO AND THREE. We address D 's behavior in experiment $\text{Game}_{\mathcal{A}, b}^1$. This case is even easier as the challenge ciphertext is now independent of H . That is, if the adversary submits a ciphertext (C, t) to the decryption oracle, without having queried H about $s = f^{-1}(C)$ before, then $r = t \oplus H(s)$ is an unknown random value. Hence, as in $\text{Game}_{\mathcal{A}, b}^0$, the probability that the least significant bits of $\mathsf{G}(K, r)$ equal those of s is negligible. The same argument also applies to $\text{Game}_{\mathcal{A}, b}^2$ and $\text{Game}_{\mathcal{A}, b}^3$. Note that we merely compare the behavior of the simulated decryption oracle and the original decryption procedure *in the corresponding game*, i.e., even if the challenge ciphertext is created with false values.

COMPARING GAMES ZERO AND ONE. We next show that \mathcal{A} 's output differs between $\text{Game}_{\mathcal{A}, b}^0$ and $\text{Game}_{\mathcal{A}, b}^1$ by a negligible probability only. The only difference between the games occurs if \mathcal{A} at some point queries H about the value s^* used in the challenge ciphertext (else the experiments are identical from \mathcal{A} ' viewpoint). Assume towards contradiction that this probability was non-negligible. Then we construct an algorithm \mathcal{B}_b (for fixed bit b), refuting the one-wayness of trapdoor permutation family F .

Algorithm \mathcal{B}_b is given (f, Z) as input, where f is a random function from F and $Z = f(Y)$ for a random Y . The goal of \mathcal{B}_b is to find Y . To achieve this goal \mathcal{B}_b picks $K \leftarrow \text{KGenG}(1^k)$ and starts to simulate \mathcal{A} on (K, f) . Using standard techniques, \mathcal{B}_b also simulates random oracle access to H , maintaining a list L_H for the communication between \mathcal{A} and the simulated oracle H . Every time \mathcal{A} queries H about a value s , then \mathcal{B}_b also checks if $f(s) = Z$. If so, \mathcal{B}_b stops immediately with output s .

Furthermore, \mathcal{B}_b uses the simulated decryption oracle, procedure D, to answer decryption queries of \mathcal{A} without knowing f^{-1} . If \mathcal{A} outputs two challenge messages M_0, M_1 then \mathcal{B}_b picks ω^* at random and returns $(C^*, t^*) = (Z, \omega^*)$. If finally \mathcal{A} stops and has not triggered the event $f(s) = Z$, then \mathcal{B}_b stops with output \perp .

To analyze the success probability of \mathcal{B}_b note that \mathcal{B}_b prepares the challenge ciphertext part $C^* = Z = f(Y)$ with an independent random value Y (which \mathcal{B}_b tries to determine) instead of $f(G(K, r) \oplus M_b || 0^{k_1})$, and uses $t^* = \omega^*$ instead of $t^* = H(s^*) \oplus r$. We discuss that if the probability of \mathcal{A} asking H about $f^{-1}(C^*)$ drops significantly in this simulation, i.e., from noticeable in experiment $\text{Game}_{\mathcal{A}, b}^1$ to negligible here, then this refutes the pseudorandomness of \mathcal{G} via a distinguisher \mathcal{D}_b .

Consider the distinguisher \mathcal{D}_b , which on input (K, u) for random or pseudorandom u , picks f from F and starts the same simulation of \mathcal{A} as \mathcal{B}_b does (including the simulation of oracle H and decryption queries via D). Only to prepare the challenge ciphertext \mathcal{D}_b now computes $C^* = f(u \oplus M_b || 0^{k_1})$ and sets $t^* = \omega^*$ for random ω^* . If \mathcal{A} during the simulation at some point submits a query to H about $f^{-1}(C^*)$ then \mathcal{D}_b stops with output 1, else it returns 0.

If \mathcal{D}_b 's input u is pseudorandom then \mathcal{D}_b mimics \mathcal{A} 's attack in $\text{Game}_{\mathcal{A}, b}^1$ up to the H -query perfectly⁹ (since $t^* = \omega^*$ is also uniformly distributed, like $t^* = H(f^{-1}(C^*)) \oplus r$ for the unknown hash value). Hence, \mathcal{D}_b outputs 1 with noticeable probability. If, on the other hand, u is truly random, then $f^{-1}(C^*)$ has the same distribution as the random X and \mathcal{D}_b would output 1 (namely, if \mathcal{A} queries H) with negligible probability only. This would yield a contradiction to the pseudorandomness of \mathcal{G} .

Overall it follows that \mathcal{B}_b inverts f on Y with noticeable probability, in contradiction to the one-wayness of F . Therefore, our assumption about \mathcal{A} querying H about $f^{-1}(C^*)$ with noticeable probability in $\text{Game}_{\mathcal{A}, b}^1$ must have been wrong. But then \mathcal{A} 's output behavior cannot change noticeably in the two games.

COMPARING GAMES ONE AND TWO. Note that replacing $t^* = \omega^* \oplus r$ by $t^* = \omega^*$ cannot change \mathcal{A} 's output behavior at all.

COMPARING GAMES TWO AND THREE. If the adversary's output probability would change noticeably between the two games because of the substitution of $G(K, r)$ by u , then this would contradict the pseudorandomness of \mathcal{G} . Namely, construct a distinguisher \mathcal{D}_b as in the previous case, but this time \mathcal{D}_b on input K, u runs the simulation till the end (including generation of the challenge ciphertext from u), and finally copies \mathcal{A} 's output. If u was pseudorandom then \mathcal{D}_b 's output is almost identical (except for the negligible error in simulating decryption queries) to $\text{Game}_{\mathcal{A}, b}^1$, and if u was truly random then \mathcal{D}_b 's output is negligibly close to the one in $\text{Game}_{\mathcal{A}, b}^2$.

We conclude that the encryption scheme is IND-CCA2 in the RO model. \blacksquare

C Proof of Theorem 3.4

In this section we show that the partial instantiation of the H -oracle in OAEP through a non-malleable pseudorandom generator yields NM-CPA security (for pre-defined message distributions and relations).

Let \mathcal{A} be an arbitrary probabilistic polynomial-time algorithm attacking the encryption scheme in a NM-CPA scenario. As in the previous proof we consider a sequence of games where the fixed bit $b \in \{0, 1\}$ indicates whether the ciphertext encrypts message M^* ($b = 0$) or the independent message M' ($b = 1$).

- Let $\text{Game}_{\mathcal{A}, b}^0(k)$ denote the original attack of \mathcal{A} on the encryption scheme in experiment $\text{Exp}_{\text{AS}, \mathcal{A}}^{\text{nm-cpa-b}}$, where the adversary gets to see an encryption of message $M^* \leftarrow \mathcal{M}$ and $M' \leftarrow \mathcal{M}$, respectively, and where it tries to find a ciphertext of a related message to M^* (for both cases $b = 0$ and $b = 1$).
- In $\text{Game}_{\mathcal{A}, b}^1(k)$ we restrict \mathcal{A} 's access to random oracle G such that queries about the value r^* in the challenge ciphertext are answered with \perp instead (we write $G - \{r^*\}$ for this oracle).
- In $\text{Game}_{\mathcal{A}, b}^2(k)$ we change the output of the experiment if \mathcal{A} returns a ciphertext $C = f(s) || t$ such that the corresponding r value is different from r^* , yet \mathcal{A} has not asked oracle $G - \{r^*\}$ about r before; in that case we set the experiment's output to 0.

⁹Up to the negligible simulation error for decryption queries.

- In $\text{Game}_{\mathcal{A},b}^3(k)$ we further restrict the experiment's output, and define the output to be 0 if the adversary outputs a ciphertext with value $r = r^*$ (and without having asked $G - \{r^*\}$ about r , of course).
- In $\text{Game}_{\mathcal{A},b}^4(k)$ we are left with the case that the adversary outputs $r \neq r^*$ and has asked $G - \{r^*\}$ about the value r before. We now change the experiment by always encrypting message $M' \leftarrow \mathcal{M}$ (even for $b = 0$).

All games are described formally in Figure C. It is easy to see that the output of experiments $\text{Game}_{\mathcal{A},b}^0(k)$ for $b = 0, 1$ are identical distributions to the ones of $\text{Exp}_{\text{AS},\mathcal{A}}^{\text{nm-cpa-0}}$ and $\text{Exp}_{\text{AS},\mathcal{A}}^{\text{nm-cpa-1}}$, respectively, and that the output distribution of experiments $\text{Game}_{\mathcal{A},b}^4(k)$ for $b = 0, 1$ are identical.

To prove the theorem it remains to show that for both bits $b = 0, 1$ the transition from $\text{Game}_{\mathcal{A},b}^0(k)$ ($\equiv \text{Exp}_{\text{AS},\mathcal{A}}^{\text{nm-cpa-b}}$) to $\text{Game}_{\mathcal{A},b}^4(k)$ does not change the output behavior significantly. More precisely, it suffices to show that

$$\left| \Pr \left[\text{Game}_{\mathcal{A},b}^i(k) = 1 \right] - \Pr \left[\text{Game}_{\mathcal{A},b}^{i+1}(k) = 1 \right] \right|$$

is negligible for each $i = 0, 1, 2, 3$ and each $b = 0, 1$. By flipping the relation's output in \mathcal{A} 's attack we can always assume that each difference (without considering the absolute value) is positive. In the sequel fix a bit b .

COMPARING GAMES ZERO AND ONE. We show that the probability that \mathcal{A} asks G about r^* in experiment $\text{Game}_{\mathcal{A},b}^0(k)$ is negligible. Assume towards contradiction that this is not the case. Then it is easy to construct a successful distinguisher \mathcal{D}_b for the pseudorandomness of \mathcal{H} (with respect to the $\text{hint}(s^*)$ function which picks $(f, f^{-1}) \stackrel{\$}{\leftarrow} F$ and outputs $f, f(s^*)$).

Algorithm \mathcal{D}_b gets as input the generator's key K and f as well as a pair $(f(s^*), u^*)$ where $s^* \stackrel{\$}{\leftarrow} \{0, 1\}^{k-k_0-k_1}$ and u^* is either $\text{H}_K(s^*)$ or truly random. Algorithm \mathcal{D}_b next invokes a black-box simulation of \mathcal{A} on 1^k , simulating random oracle G by well-known techniques. When \mathcal{A} outputs a distribution \mathcal{M} , a relation R (and state) \mathcal{D}_b picks $r^* \leftarrow \{0, 1\}^{k_0}$ and computes $t^* = u^* \oplus r^*$. It continues the simulation of \mathcal{A} for $pk = (K, f)$ and (state, $f(s^*) || t^*$) until \mathcal{A} stops. If at some point during the simulation \mathcal{A} has submitted r^* to the (simulated) random oracle G , possibly before seeing the challenge ciphertext, then \mathcal{D}_b outputs 1, else \mathcal{D}_b returns 0.

It is clear that if \mathcal{A} 's probability of querying G about r^* would drop from noticeable to negligible in the case of a random u^* then \mathcal{D}_b would successfully distinguish pseudorandom and random inputs with respect to $\text{hint}(s^*) = (f, f(s^*))$. Hence \mathcal{A} must also have noticeable success probability in the case of random u^* . But for such a random u^* the value $t^* = u^* \oplus r^*$ hides r^* information-theoretically, and the probability that the i -th query to G equals r^* (given that the first $i - 1$ queries were all different from r^*) is at most $1/(2^{k_0} - i - 1)$. Therefore \mathcal{A} 's overall success probability remains negligible since the number of queries is polynomially bounded, contradicting our initial assumption.

COMPARING GAMES ONE AND TWO. The only difference between the two games lies in the case where \mathcal{A} outputs a ciphertext $f(s) || t$, implicitly specifying M, r , such that $r \neq r^*$ and \mathcal{A} has never asked $G - \{r^*\}$ about r . But then $G(r)$ is an unknown random value and the probability that the k_1 least significant bits of $G(r) \oplus M || 0^{k_1}$ equal those of s is at most 2^{-k_1} . We can thus neglect the contribution of this event to the output, without losing more than a negligible amount.

COMPARING GAMES TWO AND THREE. We claim that the probability of \mathcal{A} outputting a ciphertext of a related message for $r = r^*$ is negligible. This follows from the non-malleability of \mathcal{H} and the pseudorandomness of \mathcal{H} . Namely, in a first step we can replace the pair $(f(s^*), \text{H}_K(s^*))$ in the computation of the challenge ciphertext by a pair $(f(s^*), \text{H}_K(s'))$ for an independent s' . As we will describe below, by the non-malleability of \mathcal{H} the success probability of \mathcal{A} for $r = r^*$ will not change significantly. Then, due to the pseudorandomness of \mathcal{H} we can replace $\text{H}_K(s')$ by a pseudorandom value u' . But then r^* is information-theoretically hidden from \mathcal{A} and the probability that \mathcal{A} outputs a valid ciphertext for $r = r^*$ is negligible (and so must be the initial probability). Note that pseudorandomness of \mathcal{H} alone does not guarantee this but that some kind of non-malleability is necessary (cf. [9]).

We next formalize the above ideas. In $\text{Game}_{\mathcal{A}}^2(k)$ denote by SameRnd the event that \mathcal{A} outputs a ciphertext $f(s) || t$ such that $f(s) || t = \mathcal{E}_{pk}(M; r)$ is a valid ciphertext for M but different from $f(s^*) || t^*$, $r = r^*$ and

Experiment Game $^0_{\mathcal{A},b}(k)$:

$(pk, sk) \xleftarrow{\$} \mathcal{K}(1^k)$
 $(\mathcal{M}, R, \text{state}) \xleftarrow{\$} \mathcal{A}^G(1^k)$
 Pick $M^* \xleftarrow{\$} \mathcal{M}$
 if $b = 0$ then $M' \leftarrow M^*$ else $M' \xleftarrow{\$} \mathcal{M}$
 Pick $r^* \xleftarrow{\$} \{0, 1\}^k$
 Compute $C^* || t^* \leftarrow \mathcal{E}(pk, M'; r^*)$
 $C || t \xleftarrow{\$} \mathcal{A}^G(C^* || t^*, pk, \text{state})$
 $M \leftarrow \mathcal{D}(sk, C || t)$
 return $R(M^*, M)$

Experiment Game $^1_{\mathcal{A},b}(k)$:

$(pk, sk) \xleftarrow{\$} \mathcal{K}(1^k)$
 Pick $r^* \xleftarrow{\$} \{0, 1\}^k$
 $(\mathcal{M}, R, \text{state}) \xleftarrow{\$} \mathcal{A}^{G - \{r^*\}}(1^k)$
 Pick $M^* \xleftarrow{\$} \mathcal{M}$
 if $b = 0$ then $M' \leftarrow M^*$ else $M' \xleftarrow{\$} \mathcal{M}$
 Compute $C^* || t^* \leftarrow \mathcal{E}(pk, M'; r^*)$
 $C || t \xleftarrow{\$} \mathcal{A}^{G - \{r^*\}}(C^* || t^*, pk, \text{state})$
 $M \leftarrow \mathcal{D}(sk, C || t)$
 return $R(M^*, M)$

Experiment Game $^2_{\mathcal{A},b}(k)$:

$(pk, sk) \xleftarrow{\$} \mathcal{K}(1^k)$
 Pick $r^* \xleftarrow{\$} \{0, 1\}^k$
 $(\mathcal{M}, R, \text{state}) \xleftarrow{\$} \mathcal{A}^{G - \{r^*\}}(1^k)$
 Pick $M^* \xleftarrow{\$} \mathcal{M}$
 if $b = 0$ then $M' \leftarrow M^*$ else $M' \xleftarrow{\$} \mathcal{M}$
 Compute $C^* || t^* \leftarrow \mathcal{E}(pk, M'; r^*)$
 $C || t \xleftarrow{\$} \mathcal{A}^{G - \{r^*\}}(C^* || t^*, pk, \text{state})$
 $M \leftarrow \mathcal{D}(sk, C || t)$
 Let r be such that $C || t = \mathcal{E}(pk, M; r)$
 if $r \neq r^*$ and $r \notin \text{AskG}$ then
 return 0
 else
 return $R(M^*, M)$

Experiment Game $^3_{\mathcal{A},b}(k)$:

$(pk, sk) \xleftarrow{\$} \mathcal{K}(1^k)$
 Pick $r^* \xleftarrow{\$} \{0, 1\}^k$
 $(\mathcal{M}, R, \text{state}) \xleftarrow{\$} \mathcal{A}^{G - \{r^*\}}(1^k)$
 Pick $M^* \xleftarrow{\$} \mathcal{M}$
 if $b = 0$ then $M' \leftarrow M^*$ else $M' \xleftarrow{\$} \mathcal{M}$
 Compute $C^* || t^* \leftarrow \mathcal{E}(pk, M'; r^*)$
 $C || t \xleftarrow{\$} \mathcal{A}^{G - \{r^*\}}(C^* || t^*, pk, \text{state})$
 $M \leftarrow \mathcal{D}(sk, C || t)$
 Let r be such that $C || t = \mathcal{E}(pk, M; r)$
 if $(r \neq r^*$ and $r \notin \text{AskG})$ or $r = r^*$ then
 return 0
 else
 return $R(M^*, M)$

Experiment Game $^4_{\mathcal{A},b}(k)$:

$(pk, sk) \xleftarrow{\$} \mathcal{K}(1^k)$
 Pick $r^* \xleftarrow{\$} \{0, 1\}^k$
 $(\mathcal{M}, R, \text{state}) \xleftarrow{\$} \mathcal{A}^{G - \{r^*\}}(1^k)$
 Compute ciphertext (C^*, t^*) :
 Pick $M^* \xleftarrow{\$} \mathcal{M}$
 Pick $M' \xleftarrow{\$} \mathcal{M}$
 Compute $C^* || t^* \leftarrow \mathcal{E}(pk, M'; r^*)$
 $C || t \xleftarrow{\$} \mathcal{A}^{G - \{r^*\}}(C^* || t^*, pk, \text{state})$
 $M \leftarrow \mathcal{D}(sk, C || t)$
 Let r be such that $C || t = \mathcal{E}(pk, M; r)$
 if $(r \neq r^*$ and $r \notin \text{AskG})$ or $r = r^*$ then
 return 0
 else
 return $R(M^*, M)$

Figure 2: Games in the Proof of Theorem 3.4: Shaded areas indicate the differences between the games. It is always assumed that the support of \mathcal{M} consists of equal-length messages and that $C^* || t^* \neq C || t$. Let AskG denote the queries \mathcal{A} made to its oracle.

$R_{\mathcal{A}}(M^*, M) = 1$ for the pre-defined relation $R_{\mathcal{A}}$. Now consider the probability of event **SameRnd** when we slightly change the experiment by tweaking the computation of the challenge ciphertext $f(s^*)||t^*$ for $t^* = \mathsf{H}_K(s^*) \oplus r^*$ as follows. We instead pick an independent s' and compute t^* as $t^* = \mathsf{H}_K(s') \oplus r^*$ (but keep the value $f(s^*)$). Denote the corresponding experiment by $\text{Game}_{\mathcal{A},b}^2(k)$ and the event by **SameRnd'**.

We claim that the probability of **SameRnd** in the altered game changes only negligibly, due to the non-malleability of \mathcal{H} . Specifically, construct the following algorithm \mathcal{B}_b from \mathcal{A} , executing either $\text{Game}_{\mathcal{A},b}^2(k)$ or $\text{Game}_{\mathcal{A},b}^2(k)$, and the following relation $R_{\mathcal{B}}$.

Since we consider the case of a pre-defined distributions \mathcal{M} and relations $R_{\mathcal{A}}$ in \mathcal{A} 's attack we can specify a pre-defined relation $R_{\mathcal{B}}$ based on \mathcal{M} and $R_{\mathcal{A}}$ as follows: This relation works by simulating \mathcal{A} 's program in the first phase, up to the step where \mathcal{A} outputs \mathcal{M} , R and **state**. The description of the relation includes all random oracle queries made by \mathcal{A} and the answers (including the pre-selected value r^*), the descriptions of \mathcal{M} and R as well as **state** and the fixed bit b , and two random messages $M^*, M' \leftarrow \mathcal{M}(1^k)$. The relation $R_{\mathcal{B}}$ for input s^*, s then lets $\gamma = s^* \oplus M^*||0^{k_1}$ if $b = 0$ and $\gamma = s^* \oplus M'||0^{k_1}$ if $b = 1$; then let $M||z = s \oplus \gamma$ and output 1 iff $z = 0^{k_1}$ and $R_{\mathcal{A}}(M^*, M) = 1$.

This algorithm \mathcal{B}_b gets as input a tuple $(f, f(s^*), y)$ where $f \xleftarrow{\$} F$, $s^* \xleftarrow{\$} \{0, 1\}^k$ and y is either $\mathsf{H}_K(s^*)$ or $\mathsf{H}_K(s')$ for an independent $s' \xleftarrow{\$} \{0, 1\}^{k-k_0}$. It also receives the description of the relation $R_{\mathcal{B}}$, including the random oracle queries to $G - \{r^*\}$ and the values r^*, M^*, M' . Algorithm \mathcal{B}_b starts a black-box simulation of \mathcal{A} for input (f, K) and **state** as well as $f(s^*)||t^*$ for $t^* = y \oplus r^*$; algorithm \mathcal{B}_b also When the encryption attacker \mathcal{A} finally outputs $f(s)||t$ we let \mathcal{B}_b output $f(s), t \oplus r^*$.

For the analysis consider first the case that \mathcal{B}_b 's input y equals $\mathsf{H}_K(s^*)$. It is easy to see that \mathcal{B}_b perfectly simulates \mathcal{A} 's view in $\text{Game}_{\mathcal{A},b}^2(k)$. Hence, if **SameRnd** occurs in the simulation—which happens with the same probability as in $\text{Game}_{\mathcal{A},b}^2(k)$ —then we have for \mathcal{B}_b 's output $(R_{\mathcal{B}}, f(s), t \oplus r^*)$ that $r = r^*$ and thus $G(r^*) = G(r)$ for the unspecified value $G(r^*)$. Therefore, and since the ciphertext $f(s)||t$ of M under r is valid,

$$\begin{aligned} t \oplus r^* &= \mathsf{H}(s) \oplus r \oplus r^* = \mathsf{H}(s), & R_{\mathcal{A}}(M^*, M) &= 1 & \text{and} \\ M||0^{k_1} &= s \oplus G(r) = s \oplus G(r^*) = s \oplus (s^* \oplus M^*||0^{k_1}) = s \oplus \gamma & & & \text{if } b = 0 \\ M||0^{k_1} &= s \oplus G(r) = s \oplus G(r^*) = s \oplus (s^* \oplus M'||0^{k_1}) = s \oplus \gamma & & & \text{if } b = 1 \end{aligned}$$

Furthermore, for $r = r^*$ the values s, s^* and thus the ciphertexts must be different. Hence \mathcal{B}_b 's attack returns 1 with the same probability as **SameRnd** in $\text{Game}_{\mathcal{A},b}^2(k)$ occurs.

Now consider the case that y equals $\mathsf{H}_K(s')$ for an independent value s' . As above it follows that \mathcal{B}_b 's attack returns 1 with the same probability as **SameRnd'** in $\text{Game}_{\mathcal{A},b}^2(k)$ happens. By the non-malleability assumption the probability for **SameRnd'** in $\text{Game}_{\mathcal{A},b}^2(k)$ must be negligibly close to **SameRnd** in $\text{Game}_{\mathcal{A}}^2(k)$.

In the next step we transform $\text{Game}_{\mathcal{A},b}^2(k)$ into an experiment $\text{Game}_{\mathcal{A},b}^2(k)$ where, instead of using $\mathsf{H}_K(s')$ for the computation of the challenge ciphertext, we pick a truly random string u' , i.e., now we set $t^* = u' \oplus r^*$. Denote by **SameRnd''** the event that \mathcal{A} outputs $(R_{\mathcal{A}}, f(s)||t)$ in this game such that $f(s)||t = \mathcal{E}_{pk}(M; r)$ is a valid ciphertext for M but different from $f(s^*)||t^*$, and that $R_{\mathcal{A}}(M^*, M) = 1$ and $r = r^*$.

We claim that the difference of the probabilities for **SameRnd** and **SameRnd''** is negligible. Consider the following distinguisher \mathcal{D}_b against the pseudorandomness. It gets as input a value (K, y) where y is either $\mathsf{H}_K(s')$ for a random s' or equals a random value u' . Then \mathcal{D}_b picks $(f, f^{-1}) \leftarrow F$ and $r^* \leftarrow \{0, 1\}^{k_0}$ and starts a black-box simulation of \mathcal{A} for input (f, K) . In particular, \mathcal{D}_b simulates random oracle access to $G - \{r^*\}$ as \mathcal{B}_b . If \mathcal{A} at some points generates \mathcal{M}, R and **state** then \mathcal{D}_b picks $M^* \xleftarrow{\$} \mathcal{M}(1^k)$ and $s^* \leftarrow \{0, 1\}^{k-k_0}$ and continues the simulation for **(state, $f(s^*)||t^*$)** where $t^* = y \oplus r^*$ for the given y . When \mathcal{A} finally returns $f(s)||t$ then \mathcal{D}_b tries to decrypt $f(s)||t$ to M, r and returns 1 if and only if decryption succeeds, $f(s)||t \neq f(s^*)||t^*$, $r = r^*$ and $R_{\mathcal{A}}(M^*, M) = 1$.

Note that in contrast to \mathcal{B}_b the distinguisher \mathcal{D}_b here knows the inverse function f^{-1} to f ; this is necessary to evaluate the relation $R_{\mathcal{A}}$ on the messages at the end. But then it is easy to verify that \mathcal{D}_b returns 1 in the experiment above for $y = \mathsf{H}_K(s')$ with the same probability that **SameRnd'** in $\text{Game}_{\mathcal{A},b}^2(k)$ happens, and for random $y = u'$ with the same probability that **SameRnd''** in $\text{Game}_{\mathcal{A},b}^2(k)$ occurs. By the pseudorandomness both probabilities must be negligibly close then.

The final step is to observe that the probability of SameRnd'' in $\text{Game}_{\mathcal{A},b}^2(k)$ is negligible. This is because the distribution of the challenge ciphertext is now completely independent of r^* and we can thus think of r^* as drawn only after \mathcal{A} has output the valid ciphertext $f(s)||t$. But the probability of $r = r^*$ is at most 2^{-k_0} in this case. It follows that the probability of event SameRnd in $\text{Game}_{\mathcal{A},b}^2(k)$ is negligible. This proves that the output distributions in $\text{Game}_{\mathcal{A},b}^2(k)$ and $\text{Game}_{\mathcal{A},b}^3(k)$ are negligibly close.

COMPARING GAMES THREE AND FOUR. For $b = 1$ it is clear that the two games are identical. Let $b = 0$. Since \mathcal{A} is only granted access to oracle $G - \{r^*\}$ the distributions of s^* in both games are identical. On the other hand, the adversary's message M is already determined upon giving the final output (because $r \neq r^*$ has been submitted to G before). Hence, this message M is independent of the message encrypted in the challenge ciphertext, no matter whether this is M^* or M' .

This proves that the encryption scheme is secure in the NM-CPA sense. \blacksquare

D Instantiating the H -Oracle for IND-CCA2 Security

In this section we show that the partial instantiation of H achieves IND-CCA2 security, as long as the instantiating pseudorandom generator is non-malleable under chosen-image attacks (i.e., where the adversary is allowed to make inversion queries to the trapdoor pseudorandom generator):

Definition D.1 Assume $\mathcal{H} = (\text{KGenH}, \text{H}, \text{TdH})$ is a trapdoor pseudorandom generator (which is pseudorandom with respect to $\text{hint}(x) = (f, f(x))$ for $(f, f^{-1}) \leftarrow F(1^k)$ from the trapdoor function family F). Then \mathcal{H} is called non-malleable under chosen-image attacks with respect to hint if for any efficient algorithm \mathcal{B} and any efficient relation R the following random variables $\text{Exp}_{\mathcal{H},\mathcal{B},F,R}^{\text{nm-cia-1}}(k)$, $\text{Exp}_{\mathcal{H},\mathcal{B},F,R}^{\text{nm-cia-0}}(k)$ are computationally indistinguishable, where the experiments are defined as follows.

<p>Experiment $\text{Exp}_{\mathcal{G},\mathcal{B},F,R}^{\text{nm-cia-1}}(k)$</p> <p>$(K, K^{-1}) \xleftarrow{\\$} \text{KGenH}(1^k)$</p> <p>$(f, f^{-1}) \xleftarrow{\\$} F$</p> <p>$s^* \xleftarrow{\\$} \{0, 1\}^k$</p> <p>$y^* \xleftarrow{\\$} \text{H}_K(s^*)$</p> <p>$(z, y) \xleftarrow{\\$} \mathcal{B}^{\text{TdH}(K^{-1}, \cdot) - \{y^*\}}(K, f, f(s^*), y^*)$</p> <p>$s \leftarrow f^{-1}(z)$</p> <p>Return 1 iff</p> <p style="text-align: center;">$R(s^*, s) \wedge \text{H}_K(s) = y \wedge s^* \neq s$</p>	<p>Experiment $\text{Exp}_{\mathcal{G},\mathcal{B},F,R}^{\text{nm-cia-0}}(k)$</p> <p>$K \xleftarrow{\\$} \text{KGenH}(1^k)$</p> <p>$(f, f^{-1}) \xleftarrow{\\$} F$</p> <p>$s^* \xleftarrow{\\$} \{0, 1\}^k$; $s' \xleftarrow{\\$} \{0, 1\}^k$</p> <p>$y' \xleftarrow{\\$} \text{H}_K(s')$</p> <p>$(z, y) \xleftarrow{\\$} \mathcal{B}^{\text{TdH}(K^{-1}, \cdot) - \{y'\}}(K, f, f(s^*), y')$</p> <p>$s \leftarrow f^{-1}(z)$</p> <p>Return 1 iff</p> <p style="text-align: center;">$R(s^*, s) \wedge \text{H}_K(s) = y \wedge s^* \neq s$</p>
---	--

Although it seems to be moot to plug in a “chosen-ciphertext secure” function to obtain a chosen-ciphertext encryption scheme, our goal is to show that one can instantiate the random oracle in principle, providing a possibly non-optimized feasibility result. As the counterexamples show such results are far from trivial, e.g., the CGH encryption scheme [11] in the random oracle model cannot be instantiated with *any* function, not even with a secure encryption scheme.

We now prove Theorem 3.5 stating that the partial H -instantiation through a pseudorandom generator which is non-malleable under chosen-image attacks is IND-CCA2:

Proof of Theorem 3.5: The proof follows similar to the one of the IND-CCA2 secure G -instantiation.

That is, for an arbitrary probabilistic polynomial-time algorithm \mathcal{A} we let $\text{Game}_{\mathcal{A},b}^0(k)$ denote the original attack of \mathcal{A} on the encryption scheme $\text{OAEP}^{G,\mathcal{H}}[F_{\text{t-clear}}]$ where message M_b for fixed bit b is encrypted in the challenge ciphertext. Let $\text{Game}_{\mathcal{A},b}^1(k)$ denote the game where we replace $G(r^*)$ in the challenge ciphertext by a uniformly and independently distributed string ω^* , making the challenge ciphertext independent of b (as formally carried out in $\text{Game}_{\mathcal{A},b}^2(k)$). All games are described formally in Figure D. We remark that we use the chosen-image pseudorandomness of \mathcal{H} for showing that the adversarial behavior in the two games cannot differ significantly.

<p>Experiment Game$^0_{\mathcal{A},b}(k)$: $((f^{-1}, K), (f, K)) \xleftarrow{\\$} \mathcal{K}(1^k)$ $(M_0, M_1, \text{state}) \xleftarrow{\\$} \mathcal{A}^{G, \mathcal{D}(sk, \cdot)}(f, K)$ Compute ciphertext (C^*, t^*): Pick $r^* \xleftarrow{\\$} \{0, 1\}^{k_0}$ Compute $s^* \leftarrow G(r^*) \oplus M_b 0^{k_1}$ Compute $C^* \leftarrow f(s^*)$ Compute $t^* \leftarrow \text{H}(K, s^*) \oplus r^*$ $d \xleftarrow{\\$} \mathcal{A}^{G, \mathcal{D}(sk, \cdot) - \{(C^*, t^*)\}}((C^*, t^*), \text{state})$</p>	<p>Experiment Game$^1_{\mathcal{A},b}(k)$: $((f^{-1}, K), (f, K)) \xleftarrow{\\$} \mathcal{K}(1^k)$ $(M_0, M_1, \text{state}) \xleftarrow{\\$} \mathcal{A}^{G, \mathcal{D}(sk, \cdot)}(f, K)$ Compute ciphertext (C^*, t^*): Pick $r^* \xleftarrow{\\$} \{0, 1\}^{k_0}$ Pick $\omega^* \xleftarrow{\\$} \{0, 1\}^{k-k_0}$ Compute $s^* \leftarrow \omega^* \oplus M_b 0^{k_1}$ Compute $C^* \leftarrow f(s^*)$ Compute $t^* \leftarrow \text{H}(K, s^*) \oplus r^*$ $d \xleftarrow{\\$} \mathcal{A}^{G, \mathcal{D}(sk, \cdot) - \{(C^*, t^*)\}}((C^*, t^*), \text{state})$</p>
<p>Experiment Game$^1_{\mathcal{A},b}(k)$: $((f^{-1}, K), (f, K)) \xleftarrow{\\$} \mathcal{K}(1^k)$ $(M_0, M_1, \text{state}) \xleftarrow{\\$} \mathcal{A}^{G, \mathcal{D}(sk, \cdot)}(f, K)$ Compute ciphertext (C^*, t^*): Pick $r^* \xleftarrow{\\$} \{0, 1\}^{k_0}$ Pick $\omega^* \xleftarrow{\\$} \{0, 1\}^{k-k_0}$ Compute $s^* \leftarrow \omega^*$ Compute $C^* \leftarrow f(s^*)$ Compute $t^* \leftarrow \text{H}(K, s^*) \oplus r^*$ $d \xleftarrow{\\$} \mathcal{A}^{G, \mathcal{D}(sk, \cdot) - \{(C^*, t^*)\}}((C^*, t^*), \text{state})$</p>	

Figure 3: Games in the Proof of Theorem 3.5: Shaded areas indicate the differences between the games. It is always assumed that the output (M_0, M_1, state) of \mathcal{A} in the first phase satisfies $|M_0| = |M_1|$.

It again suffices to show that $\Pr [\text{Game}_{\mathcal{A},b}^i(k) = 1] - \Pr [\text{Game}_{\mathcal{A},b}^{i+1}(k) = 1]$ is negligible for any $i = 0, 1$ and $b \in \{0, 1\}$, as the probabilities $\Pr [\text{Game}_{\mathcal{A},1}^2(k) = 1]$ and $\Pr [\text{Game}_{\mathcal{A},0}^2(k) = 1]$ for $b = 1$ and $b = 0$, respectively, are identical.

SIMULATING THE DECRYPTION ORACLE. We again first describe how to simulate decryption queries in the games without knowing the secret key f^{-1} to f . This is accomplished through the random oracle model and via one procedure D which works for all games. In addition to a ciphertext (C, t) this procedure gets the public data K, f and a list L_G , representing \mathcal{A} 's queries to random oracle G and the answers as input. The procedure also gets access to the function $\text{TdH}(K^{-1}, \cdot)$, initialized with the matching secret key K^{-1} to K . To answer a decryption request (C, t) procedure D browses through each entry (r, ω) in the list L_G and does the following (where (C^*, t^*) denotes the challenge ciphertext):

- Compute $h = t \oplus r$ and submit h to $\text{TdH}(K^{-1}, \cdot)$.
- Check if the answer s matches C , i.e., if $f(s) = C$. If so, compute $M || z = s \oplus \omega$.
- If $z = 0^{k_1}$ then return M .
- In any other case, if there is no matching entry (r, ω) , or if there is one but $z \neq 0^{k_1}$, then D returns \perp .

Note that D never puts a G -query during this procedure. Observe also that, if there is a corresponding entry in L_G , then it is unique.

We next prove that this decryption procedure may substitute the actual decryption oracle except with negligible simulation error probability in all games. More formally, this means that for every decryption request in the game we run D (on the list L_G of communication between \mathcal{A} and G up to this point) instead of \mathcal{D} . Let DecError_i denote the event that D returns a different answer than \mathcal{D} for the i -th decryption query in the corresponding game, given that the first $i-1$ replies were identical. It then suffices to show that the probability of DecError_i is negligible for arbitrary i . Recall that we call a ciphertext valid iff \mathcal{D} returns a message $M \neq \perp$.

BEHAVIOR IN GAME ZERO. Assume that \mathcal{A} submits some (C, t) to the decryption oracle with the i -th query in $\text{Game}_{\mathcal{A},b}^0$ such that there is no matching entry in L_G . Let $s, r, M||z$ denote the unique values such that $f(s) = C$, $r = t \oplus H(K, s)$ and $M||z = G(r) \oplus s$. Let $\omega^*, r^*, M_b^*||0^{k_1}$ denote the corresponding values for the challenge ciphertext (C^*, t^*) .

- If we are in the first phase of the game, before \mathcal{A} receives the challenge ciphertext, and there is no value for r in L_G , then $G(r)$ is an unknown random value. The probability that the least significant bits of $G(r) \oplus s$ equal 0^{k_1} is therefore 2^{-k_1} .
- If we are in the second phase, after having received the challenge ciphertext, and $C = C^*$ and thus $s = s^*$ but $t \neq t^*$, then $r \neq r^*$ as well. Hence, $G(r)$ is an unknown independent random value and the probability that the lower bits of $G(r) \oplus s$ are 0^{k_1} is 2^{-k_1} .
- If we are in the second phase and $C \neq C^*$ and thus $s \neq s^*$, then the generator's algorithm TdH returns this value s . In this case the answer of D is identical to the one of \mathcal{D} .

Hence, the probability of event DecError_i in $\text{Game}_{\mathcal{A},b}^0$ is negligible and D simulates \mathcal{D} correctly with overwhelming probability.

BEHAVIOR IN GAMES ONE AND TWO. To analyze the behavior of D in experiment $\text{Game}_{\mathcal{A},b}^1$ we note that, from the adversary's point of view, the games are identical if \mathcal{A} never queries G about r^* . In particular, conditioning on the fact that \mathcal{A} never submits r^* to G , the substitution procedure D works almost perfectly in $\text{Game}_{\mathcal{A},b}^1$, too (exploiting the fact that D never queries G).

Assume that the probability that \mathcal{A} queries G about r^* in $\text{Game}_{\mathcal{A},b}^0$ was noticeable. Then we claim that this remains noticeable if we replace $G(r^*)$ in the challenge ciphertext by a random ω^* and, likewise, substitute $H(s^*)$ for a random value u^* . Suppose towards contradiction that this was not the case. Then we derive a contradiction to the non-malleability and pseudorandomness of the generator as follows.

In the first step we show that replacing $H(K, s^*)$ by $H(K, s')$ for an independent s' cannot change \mathcal{A} 's success probability significantly due to the non-malleability of \mathcal{H} . For this consider the following adversary \mathcal{B}_b (with fixed bit b) against the non-malleability. Let $R(s^*, s)$ be the pre-defined relation which outputs 1 if and only if the least significant bit of s equals 1. After outputting a description of this relation, \mathcal{B}_b gets $K, \text{hint}(s^*) = (f, f(s^*)), y^*$ for random s^* as input, where either $y^* = H(K, s^*)$ or $y^* = H(K, s')$ for an independent s' . Adversary \mathcal{B}_b is also allowed to query oracle $\text{TdH}(K^{-1}, \cdot)$ for values different from y^* .

\mathcal{B}_b selects some $r^* \leftarrow \{0, 1\}^{k_0}$ at the beginning. It starts an emulation of \mathcal{A} by simulating the random oracle G as usual and using procedure D to answer decryption queries. If at any point during this simulation \mathcal{A} queries G about r^* then \mathcal{B}_b immediately stops with output s such that $\text{lsb}_1(s) = 1$. By this \mathcal{B}_b simulates \mathcal{A} 's attack in the first phase up to the step where \mathcal{A} outputs M_0, M_1 . Then \mathcal{B}_b prepares a ciphertext (C^*, t^*) as $C = f(s^*)$ and $t^* = y^* \oplus r^*$. It returns (C^*, t^*) to \mathcal{A} and continues the simulation (as before, answering decryption queries as before and stopping with output s such that $\text{lsb}_1(s) = 1$ if \mathcal{A} asks G about r^*).

If \mathcal{B}_b 's input y^* equals $G(K, s^*)$, then the probability of \mathcal{A} asking G about r^* is identical to the one in $\text{Game}_{\mathcal{A},b}^0$. Up to the point where \mathcal{A} queries G about r^* the value $G(r^*) \oplus M_b||0^{k_1}$ is random, as is the given s^* . Hence, the adversary's view in this simulation is identical to the one in the game. Overall, \mathcal{B}_b outputs 1 in this case with noticeable probability by assumption. If \mathcal{B}_b 's input y^* , on the other hand, is for an independent s' , then the probability of returning 1 drops to negligible by assumption. But then \mathcal{B}_b would successfully refute the non-malleability of the generator.

In the next step we replace the value $H(K, s')$ by a truly random value u^* and we again show that this cannot affect \mathcal{A} 's probability of querying G about r^* noticeably. Suppose again that this was not the case and that the probability would drop to negligible. Then we construct a successful distinguisher \mathcal{D}_b against the pseudorandomness of \mathcal{H} .

Algorithm \mathcal{D}_b gets K, y^* as input, where either $y^* = H(K, s')$ or y^* is truly random. \mathcal{D}_b chooses $(f, f^{-1}) \xleftarrow{\$} F(1^k)$ and $r^* \leftarrow \{0, 1\}^{k_0}$ at the beginning. It starts an emulation of \mathcal{A} by simulating the random oracle G as usual and using procedure D to answer decryption queries, with the important difference that it does not use procedure $\text{TdH}(K^{-1}, \cdot)$ but instead uses f^{-1} to derive the pre-image and then compares it to the given value. Analogously to \mathcal{B}_b distinguisher \mathcal{D}_b stops with output 1 if \mathcal{A} queries G about r^* during the simulation. Thus \mathcal{D}_b simulates \mathcal{A} 's attack in the first phase up to the step where \mathcal{A} picks messages M_0, M_1 . Then \mathcal{D}_b computes a ciphertext (C^*, t^*) as $C = f(s^*)$ for random s^* and $t^* = y^* \oplus r^*$. It returns (C^*, t^*) to \mathcal{A} and continues the simulation (as before, answering decryption queries as before and stopping with output 1 if \mathcal{A} asks G about r^*).

If \mathcal{D} 's input y^* is pseudorandom, then the probability of \mathcal{A} asking G about r^* is identical to the one in $\text{Game}_{\mathcal{A}, b}^0$ when we use $H(K, s')$. Namely, up to the point where \mathcal{A} queries G about r^* the value $G(r^*) \oplus M_b || 0^{k_1}$ is random, as is s^* . Hence, the adversary's view in this simulation is identical to the one in the game. Overall, \mathcal{D} outputs 1 in this case with noticeable probability. If \mathcal{D} 's input y^* , on the other hand, is truly random, then the probability of returning 1 drops to negligible by assumption. But then \mathcal{D} would successfully distinguish the two cases, contradicting the pseudorandomness of \mathcal{H} .

We conclude that \mathcal{A} 's probability of asking r^* to G if we replace $G(r^*)$ and $H(K, s^*)$ by random elements ω^*, u^* must remain noticeable. However, since r^* is information-theoretically hidden in the challenge ciphertext with these substitutions, the probability can only be 2^{-k_0} . This yields a contradiction to our initial assumption about \mathcal{A} asking G about r^* in $\text{Game}_{\mathcal{A}, b}^0$ with noticeable probability. As discussed, this implies that D also works in $\text{Game}_{\mathcal{A}, b}^1$ with overwhelming probability.

COMPARING GAMES ZERO, ONE AND TWO. By the previous considerations about D 's behavior in $\text{Game}_{\mathcal{A}, b}^1$ adversary asks G about r^* with negligible probability in $\text{Game}_{\mathcal{A}, b}^0$ only. Given this all three games are identical from \mathcal{A} 's viewpoint.

We conclude that the encryption scheme is IND-CCA2 in the RO model. \blacksquare

E Proof of Theorem 4.1

Let \mathcal{A} be an attacker on the $\$$ NM-CPA property and $R_{\mathcal{A}}$ be a relation. Once more we look at a sequence of games $\text{Game}_{\mathcal{A}, b}^i(k)$ for $i = 0, 1, 2, 3$ and bit $b = 0, 1$ where $\text{Game}_{\mathcal{A}, R_{\mathcal{A}}, b}^0(k)$ describes \mathcal{A} 's attack in experiment $\text{Exp}_{\mathcal{AS}, \mathcal{A}}^{\$nm-cpa-b}$ for uniform messages distributions. The games are described formally in Figure E. Informally, the differences between the games is as follows:

- As mentioned before, $\text{Game}_{\mathcal{A}, R_{\mathcal{A}}, b}^i(k)$ measures \mathcal{A} 's success probability in scenario $\text{Exp}_{\mathcal{AS}, \mathcal{A}, R_{\mathcal{A}}}^{\$nm-cpa-b}$. That is, for $b = 0$ the adversary gets a ciphertext $f(s^*) || \gamma^* || t^*$ of message M^* with $\gamma^* = \text{lsb}_{k_1}(\mathcal{G}(\text{KG}, r^*))$ and $t^* = H(\text{KH}, s^* || \gamma^*) \oplus r^*$ and tries to find a ciphertext $f(s) || \gamma || t$ of a related message M . For $b = 1$ the adversary sees a ciphertext of an independent message M' instead.
- In the next experiment $\text{Game}_{\mathcal{A}, R_{\mathcal{A}}, b}^1(k)$ we replace the computation of $H_{\text{KH}}(s^* || \gamma^*)$ in the challenge ciphertext by an evaluation for an independent $s' || \gamma'$. That is, the challenge ciphertext is of the form $f(s^*) || \gamma^* || H_{\text{KH}}(s' || \gamma') \oplus r^*$.
- In $\text{Game}_{\mathcal{A}, R_{\mathcal{A}}, b}^2(k)$ we then substitute the value $H_{\text{KH}}(s' || \gamma')$ by the evaluation of a truly random element u' such that the challenge ciphertext equals $f(s^*) || \gamma^* || u' \oplus r^*$.
- We finally replace the value $\mathcal{G}_{\text{KG}}(r^*)$ in the challenge ciphertext by a random element v^* .

Proceeding from $\text{Game}_{\mathcal{A},R_{\mathcal{A}},b}^i(k)$ to $\text{Game}_{\mathcal{A},R_{\mathcal{A}},b}^{i+1}(k)$ for fixed bit b we show that \mathcal{A} 's success probability cannot change noticeably. But in the final game the challenge ciphertext is independent of the message and the experiments $\text{Game}_{\mathcal{A},R_{\mathcal{A}},0}^3(k)$ and $\text{Game}_{\mathcal{A},R_{\mathcal{A}},1}^3(k)$ generate the same output distribution. We conclude that the initial games $\text{Game}_{\mathcal{A},R_{\mathcal{A}},0}^0(k)$ and $\text{Game}_{\mathcal{A},R_{\mathcal{A}},1}^0(k)$ must be negligibly close.

<p>Experiment $\text{Game}_{\mathcal{A},R_{\mathcal{A}},b}^0(k)$: $(pk, sk) \xleftarrow{\\$} \mathcal{K}(1^k)$ Compute ciphertext $C^* \gamma^* t^*$: Pick $M^* \xleftarrow{\\$} \mathcal{M}$ if $b = 0$ then $M' \leftarrow M^*$ else $M' \xleftarrow{\\$} \mathcal{M}$ Pick $r^* \xleftarrow{\\$} \{0, 1\}^{k_0}$ Set $s^* \gamma^* \leftarrow \mathsf{G}_{\text{KG}}(r^*) \oplus M' 0^{k_1}$ Compute $C^* \leftarrow f(s^*)$ Compute $t^* \leftarrow \mathsf{H}(\text{KH}, s^* \gamma^*) \oplus r^*$ $C \gamma t \xleftarrow{\\$} \mathcal{A}(C^* \gamma^* t^*, pk)$ $M \leftarrow \mathcal{D}(sk, C \gamma t)$ return $R(M^*, M)$</p>	<p>Experiment $\text{Game}_{\mathcal{A},b}^1(k)$: $(pk, sk) \xleftarrow{\\$} \mathcal{K}(1^k)$ Compute ciphertext $C^* \gamma^* t^*$: Pick $M^* \xleftarrow{\\$} \mathcal{M}$ if $b = 0$ then $M' \leftarrow M^*$ else $M' \xleftarrow{\\$} \mathcal{M}$ Pick $r^* \xleftarrow{\\$} \{0, 1\}^{k_0}$ Set $s^* \gamma^* \leftarrow \mathsf{G}(\text{KG}, r^*) \oplus M' 0^{k_1}$ Compute $C^* \leftarrow f(s^*)$ Pick $s' \gamma' \xleftarrow{\\$} \{0, 1\}^k$ Compute $t^* \leftarrow \mathsf{H}(\text{KH}, s' \gamma')$ $\oplus r^*$ $C \gamma t \xleftarrow{\\$} \mathcal{A}(C^* \gamma^* t^*, pk)$ $M \leftarrow \mathcal{D}(sk, C \gamma t)$ return $R(M^*, M)$</p>
<p>Experiment $\text{Game}_{\mathcal{A},b}^2(k)$: $(pk, sk) \xleftarrow{\\$} \mathcal{K}(1^k)$ Compute ciphertext $C^* \gamma^* t^*$: Pick $M^* \xleftarrow{\\$} \mathcal{M}$ if $b = 0$ then $M' \leftarrow M^*$ else $M' \xleftarrow{\\$} \mathcal{M}$ Pick $r^* \xleftarrow{\\$} \{0, 1\}^{k_0}$ Set $s^* \gamma^* \leftarrow \mathsf{G}(\text{KG}, r^*) \oplus M' 0^{k_1}$ Compute $C^* \leftarrow f(s^*)$ Pick $\gamma' \xleftarrow{\\$} \{0, 1\}^{k_1}$ and $u' \xleftarrow{\\$} \{0, 1\}^{k_0}$ Compute $t^* \leftarrow u' \oplus r^*$ $C \gamma t \xleftarrow{\\$} \mathcal{A}(C^* \gamma^* t^*, pk)$ $M \leftarrow \mathcal{D}(sk, C \gamma t)$ return $R(M^*, M)$</p>	<p>Experiment $\text{Game}_{\mathcal{A},b}^3(k)$: $(pk, sk) \xleftarrow{\\$} \mathcal{K}(1^k)$ Compute ciphertext $C^* \gamma^* t^*$: Pick $M^* \xleftarrow{\\$} \mathcal{M}$ if $b = 0$ then $M' \leftarrow M^*$ else $M' \xleftarrow{\\$} \mathcal{M}$ Pick $r^* \xleftarrow{\\$} \{0, 1\}^{k_0}$ Pick $v^* \xleftarrow{\\$} \{0, 1\}^{k-k_0}$ Set $s^* \gamma^* \leftarrow v^* \oplus M' 0^{k_1}$ Compute $C^* \leftarrow f(s^*)$ Pick $\gamma' \xleftarrow{\\$} \{0, 1\}^{k_1}$ and $u' \xleftarrow{\\$} \{0, 1\}^{k_0}$ Compute $t^* \leftarrow u' \oplus r^*$ $C \gamma t \xleftarrow{\\$} \mathcal{A}(C^* \gamma^* t^*, pk)$ $M \leftarrow \mathcal{D}(sk, C \gamma t)$ return $R(M^*, M)$</p>

Figure 4: Games in the Proof of Theorem 4.1: Shaded areas indicate the differences between the games. It is always assumed that \mathcal{M} is the uniform distribution and that \mathcal{A} returns a ciphertext different from the challenge ciphertext.

COMPARING GAMES ZERO AND ONE. Recall that the difference between experiments $\text{Game}_{\mathcal{A},b}^0(k)$ and $\text{Game}_{\mathcal{A},b}^1(k)$ is that we replace the value $\mathsf{H}_{\text{KH}}(s^* || \gamma^*)$ by $\mathsf{H}_{\text{KH}}(s' || \gamma')$ for an independent $s' || \gamma'$. We show that the non-malleability of \mathcal{H} guarantees that the output of $\text{Game}_{\mathcal{A},R_{\mathcal{A}},b}^0(k)$ and $\text{Game}_{\mathcal{A},R_{\mathcal{A}},b}^1(k)$ are indistinguishable.

Consider an adversary \mathcal{B}_b for fixed bit b attacking the non-malleability property of \mathcal{H} , and a relation $R_{\mathcal{B}}$, based on the pre-defined relation $R_{\mathcal{A}}$. The relation $R_{\mathcal{B}}$ is defined as follows. It start by sampling keys $\text{KG}, \text{KG}^{-1}$ and outputs them as part of the description, together with $R_{\mathcal{A}}$. For input $s^* || \gamma^*, s || \gamma$ relation $R_{\mathcal{B}}$ first recovers r^*, r from γ^*, γ with the help of KG^{-1} , it then computes $M^* || 0^{k_1} = s^* || \gamma^* \oplus \mathsf{G}(\text{KG}, r^*)$ as well

as $M||z = s||\gamma \oplus \mathsf{G}(\mathsf{KG}, r)$ and outputs 1 iff $z = 0^{k_1}$ and $R_{\mathcal{A}}(M^*, M) = 1$.

Algorithm \mathcal{B}_b gets as input a tuple $(K, f, f(s^*)||\gamma^*, y^*)$ where either $y^* = \mathsf{H}_{\mathsf{KH}}(s^*||\gamma^*)$ or $y^* = \mathsf{H}_{\mathsf{KH}}(s' || \gamma')$ for independent $s' || \gamma'$. Then \mathcal{B}_b runs a black-box simulation of \mathcal{A} by supplementing a key KG of the trapdoor pseudorandom generator \mathcal{G} and storing KG^{-1} , taking them from the description of $R_{\mathcal{B}}$, and starting \mathcal{A} on the public key $pk = (f, \mathsf{KG}, \mathsf{KH})$. For generating the challenge ciphertext algorithm \mathcal{B}_b computes r^* from its input γ^* via the trapdoor procedure TdG of \mathcal{G} . This yields a suitable r^* and \mathcal{B}_b returns $f(s^*)||\gamma^*||t^*$ for $t^* = y^* \oplus r^*$ to \mathcal{A} . Note that s^* here is random as it is in the \mathcal{A} 's attack in the game (because there M^* or M' are uniformly distributed). When \mathcal{A} finally outputs a ciphertext $f(s)||\gamma||t$ then \mathcal{B}_b reconstructs r from γ via TdG (if the ciphertext is valid then this works again). \mathcal{B}_b then outputs $f(s)||\gamma, t \oplus r$.

For the analysis note that $r \neq r^*$ implies $\gamma = \mathsf{lsb}_{k_1}(\mathsf{G}(\mathsf{KG}, r)) \neq \gamma^* = \mathsf{lsb}_{k_1}(\mathsf{G}(\mathsf{KG}, r^*))$ with overwhelming probability; else \mathcal{B}_b 's values r, r^* would contradict the near-collision resistance in a straightforward way (where one would use f^{-1} to decrypt the adversarial ciphertext such that knowledge of KG^{-1} is not necessary). But then the pre-images $s||\gamma$ and $s^*||\gamma^*$ must be different, as required for a success in the non-malleability attack. Similarly, if $r = r^*$ then we must have $s \neq s^*$ else the ciphertext returned by \mathcal{A} would be equal to the challenge ciphertext.

Furthermore, if \mathcal{A} 's ciphertext is valid then \mathcal{B}_b generates a valid pair $f(s)||\gamma, \mathsf{H}_{\mathsf{KH}}(s||\gamma)$ from \mathcal{A} 's output. Hence, the probability that \mathcal{B}_b 's attack returns 1 for given $y^* = \mathsf{H}_{\mathsf{KH}}(s^*||\gamma^*)$ is negligibly close to the probability that $\mathsf{Game}_{\mathcal{A},b}^0(k)$ yields 1. Analogously it follows that \mathcal{B}_b 's attack gives 1 for $y^* = \mathsf{H}_{\mathsf{KH}}(s' || \gamma')$ is negligibly close to the one in $\mathsf{Game}_{\mathcal{A},b}^1(k)$. Hence, by the non-malleability both probabilities in the games are close.

COMPARING GAMES ONE AND TWO. The difference between the two games is now that $\mathsf{H}_{\mathsf{KH}}(s' || \gamma')$ is replaced by a truly random string u' . We claim again that the output probabilities of the two games are not significantly affected by this. For this consider (for fixed bit b) a distinguisher \mathcal{D}_b against the pseudorandomness of \mathcal{H} (with respect to $\mathsf{hint}(s' || \gamma') = \gamma'$). Algorithm \mathcal{D}_b gets as input a pair $(\mathsf{KH}, y', \gamma')$ where γ' is random and y' is either $y' = \mathsf{H}(\mathsf{KH}, s' || \gamma')$ for random s' , or $y' = u'$ is truly random. Algorithm \mathcal{D}_b generates $(f, f^{-1}) \leftarrow F$ and $(\mathsf{KG}, \mathsf{KG}^{-1}) \leftarrow \mathsf{KGenG}(1^k)$ and starts a black-box simulation of \mathcal{A} .

To generate the challenge ciphertext \mathcal{D}_b first samples a uniform message $M^* \xleftarrow{\$} \mathcal{M}$ and sets $M' \leftarrow M^*$ if $b = 0$, or samples another $M' \xleftarrow{\$} \mathcal{M}$ if $b = 1$. Algorithm \mathcal{D}_b next computes $r^* \xleftarrow{\$} \{0, 1\}^{k_0}$ from its input γ' via the generator's extraction procedure TdG , as well as $s^* || \gamma' \leftarrow \mathsf{G}(\mathsf{KG}, r^*) \oplus M' || 0^{k_1}$. It returns $f(s^*)||\gamma' || y' \oplus r^*$ to \mathcal{A} . When \mathcal{A} finally outputs a ciphertext $f(s)||\gamma||t$ then \mathcal{D}_b uses f^{-1} to decrypt the ciphertext to M and returns 1 if and only if $R_{\mathcal{A}}(M^*, M) = 1$.

It is now easy to see that for $y' = \mathsf{H}(\mathsf{KH}, s' || \gamma')$ algorithm \mathcal{D}_b returns 1 with the same probability that $\mathsf{Game}_{\mathcal{A},b}^1(k)$ yields 1. On the other hand, for random $y' = u'$ distinguisher \mathcal{D}_b outputs 1 with the same probability that $\mathsf{Game}_{\mathcal{A},b}^2(k) = 1$. Hence, both games must be indistinguishable.

COMPARING GAMES TWO AND THREE. The indistinguishability of the two games follows again by a distinguisher against the pseudorandomness, this time, however, against \mathcal{G} . We omit a formal description since the construction of the algorithm is similar to the previous case and is straightforward noting that the distribution of $t^* = u' \oplus r^*$ in $\mathsf{Game}_{\mathcal{A},b}^2(k)$ can be replaced by a random u' instead.

This proves that the scheme is $\$$ NM-CPA. \blacksquare

F Proof of Theorem 5.2

Let \mathcal{A} be a probabilistic polynomial-time algorithm attacking NM-CPA security of the hybrid scheme AS' . Below we prove the more general result saying that the scheme is still NM-CPA if the adversary outputs a vector of ciphertexts C_1, C_2, \dots . For this we consider again a sequence of games, formally described in Figure F. such that the starting game corresponds to \mathcal{A} 's attack scenario in the experiment $\mathbf{Exp}_{\mathsf{AS}', \mathcal{A}}^{\text{nm-cpa-b}}(k)$:

- $\mathsf{Game}_{\mathcal{A},b}^0$ describes the attack on the hybrid encryption scheme AS' when the challenge message M^* ($b = 0$) or an independent message M' ($b = 1$) is encrypted.
- $\mathsf{Game}_{\mathcal{A},b}^1$ describes the game where we let the random oracle G return the undefined symbol \perp if queried about r^* encrypted in the challenge ciphertext. We denote this oracle by $G - \{r^*\}$.

- $\text{Game}_{\mathcal{A},b}^2$ describes the game where we set the experiment's output to 0 if some asymmetric part $C_{\text{asym},i}$ of adversary's ciphertext is different from the part C_{asym}^* in the challenge ciphertext, yet encrypts the same value $r_i = r^*$.

We show that proceeding from $\text{Game}_{\mathcal{A},b}^i(k)$ to $\text{Game}_{\mathcal{A},b}^{i+1}(k)$ for fixed bit b the adversary's success probability cannot change noticeably. We finally discuss that the output distribution in $\text{Game}_{\mathcal{A},0}^2$ and $\text{Game}_{\mathcal{A},1}^2$ are indistinguishable, proving that the initial games $\text{Game}_{\mathcal{A},0}^0(k)$ and $\text{Game}_{\mathcal{A},1}^0(k)$ must be negligibly close.

<p>Experiment $\text{Game}_{\mathcal{A},b}^0(k)$: $(pk, sk) \xleftarrow{\\$} \mathcal{K}(1^k)$ $(\mathcal{M}, \text{state}) \xleftarrow{\\$} \mathcal{A}^G(pk)$ Pick $M^* \xleftarrow{\\$} \mathcal{M}$ if $b = 0$ then $M' \leftarrow M^*$ else $M' \xleftarrow{\\$} \mathcal{M}$ Compute ciphertext $C^* = (C_{\text{asym}}^*, C_{\text{sym}}^*)$: Pick $r^* \xleftarrow{\\$} \{0, 1\}^k$ Compute $C_{\text{asym}}^* \xleftarrow{\\$} \mathcal{E}_{\text{asym}}(pk, r^*)$ Compute $C_{\text{sym}}^* \xleftarrow{\\$} \mathcal{E}_{\text{sym}}(G(r^*), M')$ $(R, \mathbf{C}) \xleftarrow{\\$} \mathcal{A}^G(C^*, \text{state})$ $\mathbf{M} \leftarrow \mathcal{D}(sk, \mathbf{C})$ return $R(M^*, \mathbf{M})$</p>	<p>Experiment $\text{Game}_{\mathcal{A},b}^1(k)$: $(pk, sk) \xleftarrow{\\$} \mathcal{K}(1^k)$ Pick $r^* \xleftarrow{\\$} \{0, 1\}^k$ $(\mathcal{M}, \text{state}) \xleftarrow{\\$} \mathcal{A}^{G - \{r^*\}}(pk)$ Pick $M^* \xleftarrow{\\$} \mathcal{M}$ if $b = 0$ then $M' \leftarrow M^*$ else $M' \xleftarrow{\\$} \mathcal{M}$ Compute ciphertext $C^* = (C_{\text{asym}}^*, C_{\text{sym}}^*)$: Compute $C_{\text{asym}}^* \xleftarrow{\\$} \mathcal{E}_{\text{asym}}(pk, r^*)$ Compute $C_{\text{sym}}^* \xleftarrow{\\$} \mathcal{E}_{\text{sym}}(G(r^*), M')$ $(R, \mathbf{C}) \xleftarrow{\\$} \mathcal{A}^{G - \{r^*\}}(C^*, \text{state})$ $\mathbf{M} \leftarrow \mathcal{D}(sk, \mathbf{C})$ return $R(M^*, \mathbf{M})$</p>
<p>Experiment $\text{Game}_{\mathcal{A},b}^2(k)$: $(pk, sk) \xleftarrow{\\$} \mathcal{K}(1^k)$ Pick $r^* \xleftarrow{\\$} \{0, 1\}^k$ $(\mathcal{M}, \text{state}) \xleftarrow{\\$} \mathcal{A}^{G - \{r^*\}}(pk)$ Pick $M^* \xleftarrow{\\$} \mathcal{M}$ if $b = 0$ then $M' \leftarrow M^*$ else $M' \xleftarrow{\\$} \mathcal{M}$ Compute ciphertext $C^* = (C_{\text{asym}}^*, C_{\text{sym}}^*)$: Compute $C_{\text{asym}}^* \xleftarrow{\\$} \mathcal{E}_{\text{asym}}(pk, r^*)$ Compute $C_{\text{sym}}^* \xleftarrow{\\$} \mathcal{E}_{\text{sym}}(G(r^*), M')$ $(R, \mathbf{C}) \xleftarrow{\\$} \mathcal{A}^{G - \{r^*\}}(C^*, \text{state})$ $\mathbf{M} \leftarrow \mathcal{D}(sk, \mathbf{C})$ and let $r_i \leftarrow \mathcal{D}(sk, C_{\text{asym},i})$ if $\exists i : C_{\text{asym},i} \neq C_{\text{asym}}^*$ and $r_i = r^*$ then return 0 else return $R(M^*, \mathbf{M})$</p>	

Figure 5: Games in the Proof of Theorem 5.2: Shaded areas indicate the differences between the games. It is always assumed that the support of \mathcal{M} consists of equal-length messages and that $C^* \notin \mathbf{C}$.

COMPARING GAMES ZERO AND ONE. Assume that the probability that \mathcal{A} queries G about r^* in $\text{Game}_{\mathcal{A},b}^0$ was noticeable. Then we construct a successful attacker \mathcal{B}_b (for fixed bit b) on the one-wayness of the asymmetric scheme. That is, \mathcal{B}_b is given as input (pk, C_{asym}^*) where $C_{\text{asym}}^* = \mathcal{E}(pk, r^*)$ for an unknown random string r^* , and the goal is to find r^* . For this, \mathcal{B}_b runs a black-box simulation of \mathcal{A} of $\text{Game}_{\mathcal{A},b}^0$ as described next.

Algorithm \mathcal{B}_b first picks a random integer j between 1 and the maximum number of G -oracle queries \mathcal{A} makes. \mathcal{B}_b will stop the simulation when \mathcal{A} makes the j -th query r to G and will output r . For the simulation

\mathcal{B}_b emulates \mathcal{A} 's oracle access with standard list techniques during the simulation. At the beginning \mathcal{B}_b lets \mathcal{A} for input pk output a distribution \mathcal{M} in the first phase. \mathcal{B} picks $M^* \xleftarrow{\$} \mathcal{M}$ and sets $M' \leftarrow M^*$ if $b = 0$ or picks an independent M' if $b = 1$. It selects a random key $K^* \xleftarrow{\$} \mathcal{EK}_{\text{sym}}(1^k)$, encrypts $C_{\text{sym}}^* \xleftarrow{\$} \mathcal{E}_{\text{sym}}(K^*, M')$ and returns $C^* = (C_{\text{asym}}^*, C_{\text{sym}}^*)$ to \mathcal{A} . Algorithm \mathcal{B} stops with output r if \mathcal{A} makes the j -th query r to G (or \mathcal{B}_b outputs \perp if \mathcal{A} stops prematurely).

Since the simulation is perfect from \mathcal{A} 's viewpoint up to the point where \mathcal{A} makes the query r^* to G , we conclude that \mathcal{B}_b returns the pre-image r^* of C_{asym}^* with noticeable probability. This, however, contradicts the one-wayness of the encryption scheme.

COMPARING GAMES ONE AND TWO. Recall that the difference between the two games is that the experiment in $\text{Game}_{\mathcal{A},b}^2$ now returns 0 if there exists an i such that $C_{\text{asym},i} \neq C_{\text{asym}}^*$ but they both encrypt the same value $r_i = r^*$. Assume that the probability that \mathcal{A} outputs such a vector of ciphertexts in $\text{Game}_{\mathcal{A},b}^2$ was noticeable. In particular, we can then consider an adversary \mathcal{B}_b against the non-malleability of the asymmetric scheme, i.e., given a challenge (pk, C_{asym}^*) of an unknown random string r^* this adversary tries to find a ciphertext C_{asym} of r for the equality relation $R(r^*, r) = 1$ iff $r = r^*$. Clearly, R can be defined before seeing the generator's key.

The construction of \mathcal{B}_b is analogously to the previous case (against the one-wayness) and \mathcal{B}_b too builds the symmetric ciphertext part by picking an independent K^* and letting $C_{\text{sym}}^* \xleftarrow{\$} \mathcal{E}_{\text{sym}}(K^*, M')$. When \mathcal{A} finally outputs $(C_{\text{asym},i}, C_{\text{sym},i})$ for $i = 1, 2, \dots$ we let \mathcal{B}_b pick one of these ciphertexts at random, say, $(C_{\text{asym},j}, C_{\text{sym},j})$. Algorithm \mathcal{B}_b then returns $C_{\text{asym},j}$.

Since the oracle $G - \{r^*\}$ leaves the value for r^* unspecified we can assume that $G(r^*) = K^*$. Hence the simulation above is perfect and \mathcal{A} outputs some $C_{\text{asym},j} \neq C_{\text{asym}}^*$ with $r_j = r^*$ with noticeable probability in the simulation, and hence \mathcal{B}_b also satisfies the equality relation with noticeable probability. By the non-malleability of the asymmetric scheme for random strings \mathcal{B}_b 's success must remain noticeable if it is given input (pk, C_{asym}^*) for an independent string r' instead of r^* . But then the probability that \mathcal{B}_b 's output $C_{\text{asym},j}$ contains a string r_j such that $R(r^*, r_j) = 1$ for the information-theoretically hidden string r^* is at most $\#\mathbf{C} \cdot 2^{-k}$, contradicting our initial assumption about \mathcal{A} 's success probability.

COMPARING GAMES TWO FOR $b = 0, 1$. It remains to show that the output distributions of experiments $\text{Game}_{\mathcal{A},b}^2$ for $b = 0, 1$ are negligibly close. This follows from the IND-CCA2 security of the symmetric scheme and the fact that for all i , either $C_{\text{asym},i} = C_{\text{asym}}^*$ or $r_i \neq r^*$.

Suppose that the experiments differ with noticeable probability. Consider the following attacker \mathcal{B} on the IND-CCA2 security of the symmetric scheme. This algorithm is given a ciphertext C_{sym}^* under an unknown random key K^* for known messages M^* or M' (depending on b) and is also allowed to make decryption queries for different ciphertext. \mathcal{B} invokes a black-box simulation of \mathcal{A} by selecting the public encryption key and the secret decryption key of the asymmetric scheme and simulating the random oracle. To prepare the challenge ciphertext \mathcal{B} uses the given C_{sym}^* and augments the asymmetric part of the challenge ciphertext as in the game.

When \mathcal{A} eventually outputs a value β for the relation R (recall that β can be different from \perp) and a vector \mathbf{C} of ciphertexts $(C_{\text{asym},i}, C_{\text{sym},i})$ algorithm \mathcal{B} decrypts each symmetric ciphertext part as follows:

- If $C_{\text{asym},i} = C_{\text{asym}}^*$ then we must have $C_{\text{sym},i} \neq C_{\text{sym}}^*$ and \mathcal{B} decrypts to some (possibly invalid) message M_i by submitting $C_{\text{sym},i}$ to the decryption oracle of the symmetric scheme.
- If $C_{\text{asym},i} \neq C_{\text{asym}}^*$ then we must have $r_i \neq r^*$ by exclusion of the case in $\text{Game}_{\mathcal{A},b}^2$. This means that \mathcal{B} first decrypts the value r_i in \mathcal{A} 's asymmetric ciphertext part and then looks up the value $K_i = G(r_i)$ in the list for simulating the random oracle G (or \mathcal{B} picks a new random value if unspecified so far). \mathcal{B} finally uses this key K_i to decrypt $C_{\text{sym},i}$ to M_i .

Algorithm \mathcal{B} outputs $R(M^*, \mathbf{M})$. Since the simulation is perfect this yields a noticeable advantage in predicting b by assumption. Hence, for this case the two experiments differ only insignificantly.

The claim of the theorem now follows. \blacksquare