

An Un-Instantiable Random-Oracle-Model Scheme for a Hybrid-Encryption Problem

MIHIR BELLARE

ALEXANDRA BOLDYREVA

ADRIANA PALACIO

Department of Computer Science & Engineering
University of California, San Diego
9500 Gilman Drive, La Jolla, CA 92093-0114, USA.
E-mail: {mihir, aboldyre, apalacio}@cs.ucsd.edu
URL: <http://www.cse.ucsd.edu/users/{mihir, aboldyre, apalacio}>

Abstract

We present a simple, natural random oracle (RO) model scheme, for a practical goal, that is un-instantiable, meaning is proven in the RO model to meet its goal yet admits *no* standard-model instantiation that meets this goal. The goal in question is *IND-CCA-preserving asymmetric encryption* which formally captures security of the most common practical usage of asymmetric encryption, namely to transport a symmetric key in such a way that symmetric encryption under the latter remains secure. The scheme is an ElGamal variant, called Hash ElGamal, that resembles numerous existing RO-model schemes, and on the surface shows no evidence of its anomalous properties.

We believe these results deepen our understanding of the nature and extent of the gap between the standard and RO models, and bring concerns raised by previous work closer to practice by indicating that the problem of RO-model schemes admitting no secure instantiation is a very real one that can and does arise in domains where RO schemes are commonly designed.

Keywords: Random-Oracle Model, asymmetric encryption, symmetric encryption, chosen-ciphertext attacks.

1 Introduction

A random oracle (RO) model scheme is one whose algorithms have oracle access to a random function. (Its security is evaluated with respect to an adversary with oracle access to the same function.) An “instantiation” of such a scheme is the standard-model scheme obtained by replacing this function with a member of a polynomial-time computable family of functions, described by a short key. (The security of the scheme is evaluated with respect to an adversary given the same key.) In the random-oracle paradigm, as enunciated by Bellare and Rogaway [5], one first designs and proves secure a scheme in the RO model, and then instantiates it to get a (hopefully still secure) standard-model scheme.

The RO model has proven quite popular and there are now numerous practical schemes designed and proven secure in this model. But the important issue of whether such schemes can be securely instantiated, and, if so, how, remains less clear.

This paper adds to existing concerns in this regard via a new example of an “un-instantiable” RO-model scheme. The interest of this example, compared to previous ones, is that the scheme is simple and natural, resembling RO-model schemes typically being designed, and the goal is a practical one, related to popular targets for RO-model schemes. This indicates that the problem of RO-model schemes admitting no secure instantiation is a very real one that can and does arise in domains where RO schemes are commonly designed.

We also prove a more general result that identifies a goal achievable in the RO model but not the standard model.

Below we begin with some background and then describe our contributions in more detail.

1.1 Previous work

Let us call a RO-model scheme *un-instantiable*, with respect to some underlying cryptographic goal, if the scheme can be proven to meet this goal in the random-oracle model, but *no* instantiation of this scheme meets the goal in question. Canetti, Goldreich and Halevi [7] provided the first examples of un-instantiable schemes, the goals in question being IND-CPA-secure asymmetric encryption and digital signatures secure against chosen-message attacks. Further examples followed: Nielsen [18] presented an un-instantiable RO-model scheme for the goal of non-interactive, non-committing encryption [6], and Goldwasser and Taumann [16] showed the existence of a 3-move protocol which, when collapsed via a RO as per the Fiat-Shamir heuristic [13], yields an un-instantiable RO-model signature scheme.

The results of [7] indicate that it is possible for the RO paradigm to fail to yield secure “real-world” schemes. The example schemes provided by [7], however, are complex and contrived ones that do not resemble the kinds of RO schemes typically being designed. (Their schemes are designed to return the secret key depending on the result of some test applied to an output of the oracle, and they use diagonalization and CS proofs [17].) The same is true of the scheme of [16]. In contrast, the scheme of [18] is simple, but the goal, namely non-interactive, non-committing encryption, is somewhat distant from ones that are common practical targets of RO-model designs. Accordingly, based on existing work, one might be tempted to think that “in practice,” or when confined to “natural” schemes for practical problems commonly being targeted by RO-scheme designers, the RO paradigm is sound.

This paper suggests that even this might not always be true. For a practical cryptographic goal, we present an un-instantiable RO-model scheme that is simple and natural, closely resembling the types of schemes being designed in this domain. We begin below by discussing the goal, which we call IND-CCA-preserving asymmetric encryption and which arises in the domain of hybrid encryption.

1.2 IND-CCA-preserving asymmetric encryption

In practice, the most common usage of asymmetric encryption is to transport a symmetric key that is later used for symmetric encryption of the actual data. The notion of an asymmetric encryption scheme AS being

IND-CCA-preserving, that we introduce, captures the security attribute that AS must possess in order to render this usage of AS secure. We now elaborate.

Encryption, in practice, largely employs the “hybrid” paradigm. The version of this paradigm that we consider here is quite general. In a first phase, the sender picks at random a “session” key K for a symmetric encryption scheme, encrypts K asymmetrically under the receiver’s public key to get a ciphertext C_a , and transfers C_a to the receiver. In a second phase, it can encrypt messages of its choice symmetrically under K and transfer the corresponding ciphertexts to the receiver. We call this multi-message (mm) hybrid encryption.¹

A choice of an asymmetric encryption scheme AS and a symmetric encryption scheme SS gives rise to a particular mm-hybrid scheme. We introduce in Section 2 a definition of the IND-CCA security of this mm-hybrid scheme which captures the privacy of the encrypted messages even in the presence of an adversary allowed chosen-ciphertext attacks on both component schemes and allowed to choose the messages to be encrypted adaptively and as a function of the asymmetric ciphertext, denoted C_a above, that transports the symmetric key.

Now let us say that an asymmetric encryption scheme AS is *IND-CCA preserving* if the mm-hybrid associated to AS and symmetric encryption scheme SS is IND-CCA secure for *every* IND-CCA secure SS. This notion of security for an asymmetric encryption scheme captures the security attribute of its being able to securely transport a session key for the purpose of mm-hybrid encryption. The goal we consider is IND-CCA-preserving asymmetric encryption.

It is easy to see that any IND-CCA-secure asymmetric encryption scheme is IND-CCA preserving. (For completeness, this is proved in Appendix C.1.) IND-CCA preservation, however, is actually a weaker requirement on an asymmetric encryption scheme than IND-CCA security itself, hence there might be IND-CCA-preserving asymmetric encryption schemes that are simpler and more efficient than IND-CCA-secure ones. In particular, it is natural to seek an efficient IND-CCA-preserving scheme in the RO model along the lines of existing hybrid encryption schemes such as those of [8, 9, 14, 19].

1.3 The Hash ElGamal scheme and its security

It is easy to see that the ElGamal encryption scheme [12] is not IND-CCA preserving. An effort to strengthen it to be IND-CCA preserving lead us to a variant that we call the Hash ElGamal scheme. It uses the idea underlying the Fujisaki-Okamoto [14] transformation, namely to encrypt under the original (ElGamal) scheme using coins obtained by applying a random oracle H to the message. Specifically, encryption of a message K under public key (q, g, X) in the Hash ElGamal scheme is given by

$$\text{AE}^{G,H}((q, g, X), K) = (g^{H(K)}, G(X^{H(K)}) \oplus K), \quad (1)$$

where G, H are random oracles, $q, 2q + 1$ are primes, g is a generator of the order q cyclic subgroup of \mathbb{Z}_{2q+1}^* , and the secret key is (q, g, x) where $g^x = X$. Decryption is performed in the natural way as detailed in Figure 1.

The Hash ElGamal scheme is very much like practical RO-model schemes presented in the literature. In fact, it is a particular case of an asymmetric encryption scheme proposed by Baek, Lee and Kim [2, 3].

We note that the Hash ElGamal asymmetric encryption scheme is not IND-CCA secure, or even IND-CPA secure, in particular because the encryption algorithm is deterministic. But Theorem 3.1 guarantees that the Hash ElGamal asymmetric encryption scheme is IND-CCA-preserving in the RO model, if the Computational Diffie-Hellman (CDH) problem is hard in the underlying group.

¹ The term multi-message refers to the fact that multiple messages may be encrypted, in the second phase, under the same session key. The main reason for using such a hybrid paradigm, as opposed to directly encrypting the data asymmetrically under the receiver’s public key, is that the number-theoretic operations underlying popular asymmetric encryption schemes are computationally more expensive than the block-cipher operations underlying symmetric encryption schemes, so hybrid encryption brings significant performance gains.

We follow this with Theorem 4.1, however, which says that the Hash ElGamal scheme is un-instantiable. In other words, the standard-model asymmetric encryption scheme obtained by instantiating the RO-model Hash ElGamal scheme is not IND-CCA preserving, regardless of the choice of instantiating functions.² (We allow these to be drawn from any family of polynomial-time computable functions.)

1.4 A closer look

We remark that in our proof that the Hash ElGamal scheme is IND-CCA preserving in the RO model, we do not “program” the random oracle(s), meaning that we only need to be able to “watch” the adversary’s oracle queries. (That is, Theorem 3.1 holds in what Nielsen calls the “non-programmable” RO model [18].) This indicates that the security of the scheme relies on relatively weak properties of the RO, and, since we show the scheme is nonetheless un-instantiable, strengthens the impact of our result.

As noted above, we show that no instantiation of the Hash ElGamal scheme is IND-CCA-preserving. The way we establish this is the following. We let AS be some (any) instantiation of the Hash ElGamal scheme. Then, we construct a particular IND-CCA-secure symmetric encryption scheme SS such that the mm-hybrid associated to AS and SS is not IND-CCA secure. The latter is proven by presenting an explicit attack on the mm-hybrid. We clarify that the symmetric scheme SS constructed in this proof is not a natural one. (It is not a particularly complex one either.) We do not view this as subtracting much from the value of our result, which lies rather in the nature of the Hash ElGamal scheme itself and the practicality of the underlying goal.

What we suggest is interesting about the result is that the Hash ElGamal scheme, on the surface, seems innocuous enough. It does not seem to be making any “peculiar” use of its random oracle that would lead us to think it is “wrong.” (Indeed, it uses random oracles in ways they have been used previously, in particular by [14, 2, 3].) The scheme is simple, efficient, and similar to other RO-model schemes out there. In addition, we contend that the definition of IND-CCA-preserving asymmetric encryption is natural and captures a practical requirement. The fact that the Hash ElGamal scheme is un-instantiable thus points to the difficulty of being able to distinguish un-instantiable RO-model schemes from ones that at least *may* be securely instantiable, even in the context of natural and practical goals.

1.5 Generalizations

Next we provide some results that generalize the above. We consider the class of IND-CCA-preserving asymmetric encryption schemes that possess a pair of properties that we call *key verifiability* and *ciphertext verifiability*. Key verifiability means there is a way to recognize valid public keys in polynomial time. Ciphertext verifiability means there is a polynomial-time procedure to determine whether a given ciphertext is an encryption of a given message under a given valid public key.

Theorem B.2 points out that the goal of key-verifiable, ciphertext-verifiable IND-CCA-preserving asymmetric encryption is achievable in the RO model, by the Hash El Gamal scheme in particular, assuming the CDH problem is hard in the underlying group. Theorem B.3, however, says that this goal is not achievable in the standard model. In other words, there exist RO-model schemes meeting this goal, but there exist no standard-model schemes meeting it. Theorem B.3 generalizes Theorem 4.1 because any instantiation of the Hash ElGamal is key-verifiable and ciphertext-verifiable, and hence cannot be IND-CCA-preserving.

Theorem B.3 lifts our results from being about a particular scheme to being about a primitive, or class of schemes. The generalization also helps better understand what aspects of the Hash ElGamal scheme lead to its admitting no IND-CCA-preserving instantiation. In particular, we see that this is not due to some “peculiar” use of random oracles but rather due to some simply stated properties of the resulting asymmetric encryption scheme itself.

² This result is based on the assumption that one-way functions exist (equivalently, IND-CCA-secure symmetric encryption schemes exist), since, otherwise, by default, *any* asymmetric encryption scheme is IND-CCA preserving, and, indeed, the entire mm-hybrid encryption problem we are considering is vacuous. This assumption is made implicitly in all results in this paper.

1.6 Related work

In the cryptographic community, the term “hybrid encryption” seems to be used quite broadly, to refer to a variety of goals or methods in which symmetric and asymmetric primitives are combined to achieve privacy. We have considered one goal in this domain, namely mm-hybrid encryption. We now discuss related work that has considered other goals or problems in this domain.

Works such as [8, 9, 14, 19, 11, 20] provide designs of IND-CCA-secure asymmetric encryption schemes that are referred to as “hybrid encryption schemes” because they combine the use of asymmetric and symmetric primitives. (Possible goals of such designs include gaining efficiency, increasing the size of the message space, or reducing the assumptions that must be made on the asymmetric component in order to guarantee the IND-CCA security of the construction.) The schemes of [8, 9, 14, 19] are in the RO model and, although addressing a different goal, form an important backdrop for our work because the Hash ElGamal scheme is based on similar techniques and usage of random oracles. We stress, however, that we have no reason to believe that any of these schemes, or that of [2, 3] of which Hash ElGamal is a special case, are un-instantiable.

2 Definitions

NOTATION AND CONVENTIONS. If S is a randomized algorithm, then $[S(x, y, \dots)]$ denotes the set of all points having positive probability of being output by S on inputs x, y, \dots . If x is a binary string, then $|x|$ denotes its length, and if $n \geq 1$ is an integer, then $|n|$ denotes the length of its binary encoding, meaning the unique integer ℓ such that $2^{\ell-1} \leq n < 2^\ell$. The string-concatenation operator is denoted “||”.

Formal definitions in the RO model provide as an oracle, to the algorithms and the adversary, a single random function R mapping $\{0, 1\}^*$ to $\{0, 1\}$. Schemes might, however, use and refer to multiple random functions of different domains and ranges. These can be derived from R via standard means [5].

SYMMETRIC ENCRYPTION. A symmetric encryption scheme $SS = (SK, SE, SD)$ is specified by three polynomial-time algorithms: via $K \xleftarrow{\$} SK(1^k)$ one can generate a key; via $C \xleftarrow{\$} SE(K, M)$ one can encrypt a message $M \in \{0, 1\}^*$; and via $M \leftarrow SD(K, C)$ one can decrypt a ciphertext C . It is required that $SD(K, SE(K, M)) = M$ for all $K \in [SK(1^k)]$ and all $M \in \{0, 1\}^*$. We assume (without loss of generality) that $[SK(1^k)] \subseteq \{0, 1\}^k$. In the RO model, all algorithms have access to the RO.

We define security following [4] and addressing the possibility of the symmetric scheme being in the RO model. Associate to SS , an adversary \mathcal{S} , and $k \in \mathbb{N}$, the following experiment:

Experiment $\mathbf{Exp}_{SS, \mathcal{S}}^{\text{ind-cca}}(k)$

Randomly choose RO $R_s: \{0, 1\}^* \rightarrow \{0, 1\}$

$K \xleftarrow{\$} SK^{R_s}(1^k)$; $b \xleftarrow{\$} \{0, 1\}$

Run \mathcal{S} with input 1^k and oracles $SE^{R_s}(K, LR(\cdot, \cdot, b))$, $SD^{R_s}(K, \cdot)$, R_s

Let d denote the output of \mathcal{S}

If $d = b$ then return 1 else return 0.

Above, $LR(M_0, M_1, b) = M_b$ if M_0, M_1 are strings of equal length, and \perp otherwise. We say that adversary \mathcal{S} is legitimate if it never queries $SD^{R_s}(K, \cdot)$ with a ciphertext previously returned by $SE^{R_s}(K, LR(\cdot, \cdot, b))$. Symmetric encryption scheme SS is said to be IND-CCA secure if the function

$$\text{Adv}_{SS, \mathcal{S}}^{\text{ind-cca}}(k) = 2 \cdot \Pr \left[\mathbf{Exp}_{SS, \mathcal{S}}^{\text{ind-cca}}(k) = 1 \right] - 1$$

is negligible for all legitimate polynomial-time adversaries \mathcal{S} .

ASYMMETRIC ENCRYPTION. An asymmetric encryption scheme $AS = (AK, AE, AD)$ is specified by three polynomial-time algorithms: via $(pk, sk) \xleftarrow{\$} AK(1^k)$ one can generate keys; via $C \xleftarrow{\$} AE(pk, K)$ one can

encrypt a message $K \in \{0, 1\}^k$; and via $K \leftarrow \text{AD}(sk, C)$ one can decrypt a ciphertext C . (We denote the message by K because we will set it to a key for a symmetric encryption scheme.) It is required that $\text{AD}(sk, \text{AE}(pk, K)) = K$ for all $(pk, sk) \in [\text{AK}(1^k)]$ and all $K \in \{0, 1\}^k$. In the RO model, all algorithms have access to the RO.

Discussions and peripheral results in this paper sometimes refer to standard notions of security for such schemes like IND-CPA and IND-CCA, but these are not required for the main results and accordingly are not defined here but recalled in Appendix C.1.

IND-CCA-PRESERVING ASYMMETRIC ENCRYPTION. We provide the formal definitions first and explanations later. A *multi-message hybrid (mm-hybrid) encryption scheme* is simply a pair (AS, SS) consisting of an asymmetric encryption scheme $\text{AS} = (\text{AK}, \text{AE}, \text{AD})$ and a symmetric encryption scheme $\text{SS} = (\text{SK}, \text{SE}, \text{SD})$. We associate to (AS, SS) , a *hybrid adversary* \mathbf{H} , and $k \in \mathbb{N}$, the following experiment:

Experiment $\mathbf{Exp}_{\text{AS}, \text{SS}, \mathbf{H}}^{\text{ind-cca}}(k)$

Randomly choose RO $R: \{0, 1\}^* \rightarrow \{0, 1\}$

Define ROs $R_s(\cdot) = R(0\|\cdot)$ and $R_a(\cdot) = R(1\|\cdot)$

$(pk, sk) \xleftarrow{\$} \text{AK}^{R_a}(1^k)$; $K \xleftarrow{\$} \text{SK}^{R_s}(1^k)$; $b \xleftarrow{\$} \{0, 1\}$

$C_a \xleftarrow{\$} \text{AE}^{R_a}(pk, K)$

Run \mathbf{H} with inputs pk, C_a and oracles $\text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, b))$, $\text{SD}^{R_s}(K, \cdot)$, $\text{AD}^{R_a}(sk, \cdot)$, R

Let d denote the output of \mathbf{H}

If $d = b$ then return 1 else return 0.

We say that adversary \mathbf{H} is legitimate if it does not query $\text{SD}^{R_s}(K, \cdot)$ on a ciphertext previously returned by $\text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, b))$, and it does not query $\text{AD}^{R_a}(sk, \cdot)$ on C_a . Mm-hybrid encryption scheme (AS, SS) is said to be IND-CCA secure if the function

$$\text{Adv}_{\text{AS}, \text{SS}, \mathbf{H}}^{\text{ind-cca}}(k) = 2 \cdot \Pr \left[\mathbf{Exp}_{\text{AS}, \text{SS}, \mathbf{H}}^{\text{ind-cca}}(k) = 1 \right] - 1$$

is negligible for all legitimate polynomial-time adversaries \mathbf{H} .

Finally, we say that an asymmetric encryption scheme AS is *IND-CCA preserving* if the mm-hybrid encryption scheme (AS, SS) is IND-CCA secure for *all* IND-CCA-secure symmetric encryption schemes SS . Here, the set of symmetric encryption schemes over which we quantify includes RO-model ones if AS is a RO-model scheme, and includes only standard-model ones if AS is a standard-model scheme.

Let us now explain the ideas behind these formalisms. Recall that we are modelling the security of the following two-phase scenario: in phase one, the sender picks a key K for symmetric encryption, asymmetrically encrypts it under the receiver's public key to get a ciphertext C_a , and sends C_a to the receiver; in phase two, the sender symmetrically encrypts messages of its choice under K and transmits the resulting ciphertexts to the receiver. The definition above captures the requirement of privacy of the symmetrically encrypted data under a chosen-ciphertext attack. Privacy is formalized in terms of indistinguishability via left-or-right oracles, and the chosen-ciphertext attack is formalized via the adversary's access to decryption oracles for *both* the symmetric and asymmetric schemes. The legitimacy requirement, as usual, disallows decryption queries on challenge ciphertexts since they would lead to trivial adversary victory. The experiment reflects the possibility that SS and AS are RO-model schemes by picking random oracles for their encryption and decryption algorithms. The standard model is the special case where the algorithms of the schemes do not refer to any oracles, and thus the definition above covers security in both models. The notion of AS being IND-CCA preserving reflects a valuable pragmatic requirement, namely that one may use, in conjunction with AS , any symmetric encryption scheme and be guaranteed security of the mm-hybrid under the minimal assumption that the symmetric scheme itself was secure.

$\text{AK}(1^k)$ $(q, g) \stackrel{\$}{\leftarrow} \text{CG}(1^k)$ $x \stackrel{\$}{\leftarrow} \mathbb{Z}_q$ $X \leftarrow g^x$ Return $((q, g, X), (q, g, x))$	$\text{AE}^{G,H}((q, g, X), K)$ $y \leftarrow H(K)$ $Y \leftarrow g^y$ $T \leftarrow G(X^y)$ $W \leftarrow T \oplus K$ Return (Y, W)	$\text{AD}^{G,H}((q, g, x), (Y, W))$ $T \leftarrow G(Y^x)$ $K \leftarrow T \oplus W$ If $g^{H(K)} = Y$ then Return K else Return \perp EndIf
--	---	---

Figure 1: Algorithms of the RO-model asymmetric encryption scheme $\text{HEG}[\text{CG}] = (\text{AK}, \text{AE}, \text{AD})$ associated to cyclic-group generator CG . Here $G: \langle g \rangle \rightarrow \{0, 1\}^k$ and $H: \{0, 1\}^k \rightarrow \mathbb{Z}_q$ are random oracles.

Note that in $\text{Exp}_{\text{AS}, \text{SS}, \text{H}}^{\text{ind-cca}}$ the random oracles for AS and SS are defined via the main random oracle R in such a way that they are *independent* random functions of $\{0, 1\}^*$ to $\{0, 1\}$. The correctness of the principle of independently choosing random oracles of different RO-model schemes in a common context (as opposed to having them use the same oracle) should hopefully be clear. This independence turns out to be important in relation to Theorem 4.1, but for reasons that are not apparent until we see the proof of this theorem, and, accordingly, we postpone further discussion until Remark 4.4.

The existence of IND-CCA-preserving asymmetric encryption schemes is easy to establish since, as Theorem C.1 indicates, any IND-CCA-secure asymmetric encryption scheme is IND-CCA preserving. The interesting question is to find IND-CCA-preserving asymmetric encryption schemes that are more efficient than existing IND-CCA-secure asymmetric encryption schemes. Hash El Gamal is one such scheme.

3 The HEG scheme and its security in the RO model

In this section we introduce a variant of the ElGamal encryption scheme of [12] that, although not IND-CCA secure, is IND-CCA preserving in the RO model under a standard assumption. In Section 4, we will show that this scheme admits no IND-CCA-preserving instantiation.

PRELIMINARIES. A *cyclic-group generator* is a randomized, polynomial-time algorithm CG which on input 1^k outputs a pair (q, g) , where q is a prime such that $p = 2q + 1$ is also a prime, g is a generator of the cyclic, order q subgroup $\langle g \rangle$ of \mathbb{Z}_p^* , and $|p| = k$. Recall that the Computational Diffie-Hellman (CDH) problem is said to be hard for CG if the function

$$\text{Adv}_{\text{CG}, \mathcal{C}}^{\text{cdh}}(k) = \Pr \left[(q, g) \stackrel{\$}{\leftarrow} \text{CG}(1^k); x, y \stackrel{\$}{\leftarrow} \mathbb{Z}_q : \mathcal{C}(q, g, g^x, g^y) = g^{xy} \right]$$

is negligible for all polynomial-time *cdh adversaries* \mathcal{C} .

SCHEME AND RESULT STATEMENT. To any cyclic-group generator CG we associate the RO-model asymmetric encryption scheme $\text{HEG}[\text{CG}] = (\text{AK}, \text{AE}, \text{AD})$ whose constituent algorithms are depicted in Figure 1. (The scheme makes reference to two ROs, namely $G: \langle g \rangle \rightarrow \{0, 1\}^k$ and $H: \{0, 1\}^k \rightarrow \mathbb{Z}_q$, while the formal definition of an asymmetric encryption scheme provides a single RO $R: \{0, 1\}^* \rightarrow \{0, 1\}$, but G, H may be implemented via R in standard ways [5].) We call this variant of the ElGamal encryption scheme the *Hash ElGamal* encryption scheme associated to CG . Our result about its security in the RO model is the following:

Theorem 3.1 If the CDH problem is hard for cyclic-group generator CG , then the associated Hash ElGamal asymmetric encryption scheme $\text{HEG}[\text{CG}]$ is IND-CCA preserving in the RO model. \blacksquare

For the definition of what it means to be IND-CCA preserving, we refer the reader to Section 2.

REMARKS. We note that the encryption algorithm AE of HEG[CG] is deterministic. For this reason alone, HEG[CG] is not an IND-CCA secure, or even IND-CPA secure, asymmetric encryption scheme. Nonetheless, Theorem 3.1 says that it is IND-CCA preserving as long as the CDH problem is hard for CG. This is not a contradiction. Very roughly, the reason HEG[CG] can preserve IND-CCA while not itself being even IND-CPA is that the former notion considers the use of the scheme only for the encryption of messages that are symmetric keys, which (as long as the associated symmetric encryption scheme is secure) have relatively high entropy, and the entropy in these messages compensates for the lack of any introduced by AE. We add that previous work [8, 9, 14, 19] has shown that in the RO model, relatively weak asymmetric components suffice to ensure strong security properties of the hybrid based on them. Thus, it is not surprising that, although HEG[CG] is not secure with respect to standard measures like IND-CPA and IND-CCA, it is secure enough to permit its use for transport of a symmetric encryption key as indicated by Theorem 3.1.

The full proof of Theorem 3.1 is in Appendix A. Below we provide an intuitive overview that highlights the main areas of novelty.

PROOF SETUP. Let $AS = \text{HEG}[CG]$ and let AK, AE, AD denote its constituent algorithms. Let $SS = (SK, SE, SD)$ be any IND-CCA-secure symmetric encryption scheme. We need to show that (AS, SS) is an IND-CCA-secure mm-hybrid encryption scheme.

Let H be a polynomial-time hybrid adversary attacking (AS, SS) . We will construct polynomial-time adversaries S and C such that

$$\text{Adv}_{AS,SS,H}^{\text{ind-cca}}(k) \leq \text{poly}(k) \cdot \text{poly}\left(\text{Adv}_{SS,S}^{\text{ind-cca}}(k), \text{Adv}_{CG,C}^{\text{cdh}}(k)\right) + \frac{\text{poly}(k)}{2^k}. \quad (2)$$

Since SS is assumed IND-CCA secure and the CDH problem is hard for CG , the advantage functions related to S and C above are negligible, and thus so is the advantage function related to H . To complete the proof, we need to specify adversaries S, C for which Equation (2) is true. Below we let GH be the event that there is a time at which g^{xy} is queried to G but K has not been queried to H ; HG the event that there is a time at which K is queried to H but g^{xy} has not been queried to G ; and $\text{Succ}(H)$ the event that H is successful at guessing the value of its challenge bit b . We will construct C so that

$$\Pr[\text{GH}] \leq \text{poly}(k) \cdot \text{Adv}_{CG,C}^{\text{cdh}}(k) + \frac{\text{poly}(k)}{2^k},$$

and we will construct S so that

$$\Pr[\text{HG} \vee (\text{Succ}(H) \wedge \neg\text{GH} \wedge \neg\text{HG})] \leq \text{Adv}_{SS,S}^{\text{ind-cca}}(k) + \frac{\text{poly}(k)}{2^k}. \quad (3)$$

Equation (2) follows.

THE ADVERSARIES. The design of C relies mostly on standard techniques, and so we leave it to Appendix A. We turn to S . The latter gets input 1^k and oracles $SE^{R_s}(K, \text{LR}(\cdot, \cdot, b))$, $SD^{R_s}(K, \cdot)$, R_s , and begins with the initializations

$$((q, g, X), (q, g, x)) \stackrel{\$}{\leftarrow} \text{AK}(1^k); y \stackrel{\$}{\leftarrow} \mathbb{Z}_q; Y \leftarrow g^y; W \stackrel{\$}{\leftarrow} \{0, 1\}^k; C_a \leftarrow (Y, W), \quad (4)$$

and then runs H on inputs $(q, g, X), C_a$, itself responding to the oracle queries of the latter. Its aim is to do this in such a way that the key K underlying S 's oracles plays the role of the quantity of the same name for H . Eventually, it will output what H outputs. The difficulty faced by this adversary is that H might query K to H . (Other oracle queries are dealt with in standard ways.) In that case, H expects to be returned y . (And it cannot be fooled since, knowing $Y = g^y$, it can verify whether or not the value returned is y .) The difficulty for S is not that it does not know the right answer (via Equation (4), it actually knows y), but rather that it is not clear how it would know that a query being made to H equals the key K underlying its oracles, so that it would know *when* to return y as the answer to a query to H .

In order to “detect” when query K is made, we would, ideally, like a test that can be performed on a value L , accepting if $L = K$ and rejecting otherwise. However, it is not hard to see that, in general, such a

test does not exist.³ Instead, we introduce a test that has a weaker property and show that it suffices for us.

Our test `KeyTest` takes input L and has access to S 's $SE^{R_s}(K, LR(\cdot, \cdot, b))$ oracle. It returns a pair (dec, gs) such that: (1) If $L = K$ then $(\text{dec}, \text{gs}) = (1, b)$, meaning in this case it correctly computes the challenge bit b , and (2) If $L \neq K$ then, with overwhelming probability, either $\text{dec} = 0$ (the test is saying $L \neq K$) or $(\text{dec}, \text{gs}) = (1, b)$ (the test is saying it does not know whether or not $L = K$, but it has successfully calculated the challenge bit anyway). With `KeyTest` in hand, S can answer a query L made to H as follows. It runs $(\text{dec}, \text{gs}) \stackrel{\$}{\leftarrow} \text{KeyTest}(L)$. If $\text{dec} = 0$, it can safely assume $L \neq K$ and return a random answer, while if $\text{dec} = 1$, it can output gs as its guess to challenge bit b and halt.

A precise description and analysis of `KeyTest` are in Appendix A, but we briefly sketch the ideas here. The algorithm has two phases. In the first phase, it repeatedly tests whether or not

$$SD^{R_s}(L, SE^{R_s}(K, LR(T_0, T_0, b))) = T_0 \quad \text{and} \quad SD^{R_s}(L, SE^{R_s}(K, LR(T_1, T_1, b))) = T_1,$$

where T_0, T_1 are some distinct ‘‘test’’ messages. If any of these checks fails, it knows that $L \neq K$ and returns $(0, 0)$. (However, the checks can succeed with high probability even if $L \neq K$.) In the next phase, it repeatedly computes $SD^{R_s}(L, SE^{R_s}(K, LR(T_0, T_1, b)))$ and, if *all* these computations yield T_{gs} for some bit gs , it returns $(1, \text{gs})$. The analysis shows that, conditional on the first phase not returning $(0, 0)$, the bit gs from the second stage equals b with overwhelming probability.

A subtle point arises with relation to the test. Recall that H is making queries to $SD^{R_s}(K, \cdot)$. S will answer these via its own oracle of the same name. Now, consider the event that H queries to $SD^{R_s}(K, \cdot)$ a ciphertext C generated in some execution of `KeyTest`. If S calls $SD^{R_s}(K, C)$ to obtain the answer, it would immediately become an illegitimate adversary and thus forgo its advantage, since C is a result of a call to $SE^{R_s}(K, LR(\cdot, \cdot, b))$ made by S via subroutine `KeyTest`. There are a few ways around this, and the one we use is to choose the initial ‘‘test’’ messages randomly so that H has low probability of being able to query a ciphertext C generated in some execution of `KeyTest`.

This is all put together in Appendix A to show that Equation (3) holds.

We note that one might consider an alternative solution to S 's problem of wanting to ‘‘detect’’ query K to H . Namely, reply to queries to H at random, then, after H terminates, pick one such query L at random, decrypt a challenge ciphertext via L , and use that to predict the challenge bit. Unfortunately, even though $L = K$ with probability $1/\text{poly}(k)$, the advantage over one-half obtained by S via the strategy just outlined could be negligible because the wrong answers from the wrong random choices could overwhelm the right answer that arises when K is chosen.

4 Un-instantiability of the Hash ElGamal scheme

In this section we show (cf. Theorem 4.1) that the RO-model Hash ElGamal scheme admits no IND-CCA-preserving instantiation. Below we begin by detailing what we mean by instantiation of a RO-model asymmetric encryption scheme. This will refer to a RO-model scheme which, as per the formal definitions in Section 2, uses a single random oracle mapping $\{0, 1\}^*$ to $\{0, 1\}$.

INSTANTIATING RO-MODEL ASYMMETRIC ENCRYPTION SCHEMES. A *poly-time family of functions* \overline{F} associates to security parameter $k \in \mathbb{N}$ and key $fk \in \{0, 1\}^{\text{FKL}(k)}$ a map $\overline{F}_{fk}^k: \{0, 1\}^* \rightarrow \{0, 1\}$. The *key length* FKL of the family of functions is a polynomial in k . We require that there exist a polynomial t such that $\overline{F}_{fk}^k(x)$ is computable in $t(k + |x|)$ time for all $k \in \mathbb{N}$, $fk \in \{0, 1\}^{\text{FKL}(k)}$ and $x \in \{0, 1\}^*$.

³ Suppose, for example, that algorithms SE, SD only depend on the first half of the bits of their k -bit key. This is consistent with their being IND-CCA secure (in the sense that, if there exists an IND-CCA-secure symmetric encryption scheme, there also exists one with this property), but now, any test has probability at most $2^{-k/2}$ of being able to differentiate between K and a key $L \neq K$ that agrees with K in its first half.

$\overline{AK}(1^k)$ $fk \xleftarrow{s} \{0, 1\}^{\text{FKL}(k)}$ $(pk, sk) \xleftarrow{s} AK^{\overline{F}_{fk}^k}(1^k)$ Return $((pk, fk), (sk, fk))$	$\overline{AE}((pk, fk), K)$ $C \xleftarrow{s} AE^{\overline{F}_{fk}^k}(pk, K)$ Return C	$\overline{AD}((sk, fk), C)$ $K \leftarrow AD^{\overline{F}_{fk}^k}(sk, C)$ Return K
--	--	--

Figure 2: Algorithms of the standard-model asymmetric encryption scheme $\overline{AS} = (\overline{AK}, \overline{AE}, \overline{AD})$ obtained by instantiating RO-model asymmetric encryption scheme $AS = (AK, AE, AD)$ via poly-time family of functions \overline{F} .

An *instantiation* of a RO-model asymmetric encryption scheme $AS = (AK, AE, AD)$ via family \overline{F} is the standard-model asymmetric encryption scheme $\overline{AS} = (\overline{AK}, \overline{AE}, \overline{AD})$ whose constituent algorithms are illustrated in Figure 2. As these indicate, the public and secret keys of the original scheme are enhanced to also include a key fk specifying the function \overline{F}_{fk}^k , and calls to the random oracle are then replaced by evaluations of this function in all algorithms.

THE UN-INSTANTIABILITY RESULT. The formal statement of the result is the following:

Theorem 4.1 Let $\text{HEG}[\text{CG}] = (AK, AE, AD)$ be the RO-model Hash ElGamal scheme associated to a cyclic-group generator CG. Let $\overline{\text{HEG}}[\text{CG}] = (\overline{AK}, \overline{AE}, \overline{AD})$ be *any* instantiation of $\text{HEG}[\text{CG}]$ via a poly-time family of functions. Then $\overline{\text{HEG}}[\text{CG}]$ is *not* IND-CCA preserving. ■

PROOF OF THEOREM 4.1. We will construct an IND-CCA-secure symmetric encryption scheme SS such that the mm-hybrid encryption scheme $(\overline{\text{HEG}}[\text{CG}], SS)$ is not IND-CCA secure. This proves the theorem.

Let us say that a value pk is an (AS, k) -valid public key if there exists a value sk such that $(pk, sk) \in [AK(1^k)]$. We first define two polynomial-time algorithms VfPK and VfCtxt which are used by SS .

Algorithm VfPK , which we call a *key verifier*, takes inputs 1^k and pk , and outputs 1 if and only if pk is a $(\overline{\text{HEG}}[\text{CG}], k)$ -valid public key. The algorithm works by parsing pk as (q, g, X, fk) , where $fk \in \{0, 1\}^{\text{FKL}}$, and then returning 1 if and only if q and $2q+1$ are primes, g is a generator of the order q cyclic subgroup $\langle g \rangle$ of \mathbb{Z}_{2q+1}^* , $|2q+1| = k$, and $X \in \langle g \rangle$. This algorithm can be implemented in polynomial-time based on standard facts from computational number theory, and even deterministically, given the existence of polynomial-time primality tests [1]. We omit the details.

Algorithm VfCtxt , which we call a *ciphertext verifier*, takes inputs $1^k, pk, K, C$, where pk is a $(\overline{\text{HEG}}[\text{CG}], k)$ -valid public key and $K \in \{0, 1\}^k$. It runs $\overline{AE}(pk, K)$ and outputs 1 if the result is C , and 0 otherwise. In other words, VfCtxt verifies whether C is indeed an encryption of message K under the given public key pk . This is possible because the encryption algorithm AE of $\text{HEG}[\text{CG}]$ (cf. Figure 1), and hence the encryption algorithm \overline{AE} of $\overline{\text{HEG}}[\text{CG}]$, is deterministic.

Let $SS' = (SK', SE', SD')$ be any standard-model IND-CCA-secure symmetric encryption scheme. (Recall an implicit assumption is that some such scheme exists, since otherwise *all* asymmetric encryptions schemes are by default IND-CCA preserving and the entire problem we are considering is moot.) The construction of SS is in terms of SS' and algorithms VfPK and VfCtxt . We use the notation $\langle\langle \cdot, \cdot \rangle\rangle$ to denote an injective, polynomial-time computable encoding of pairs of strings as strings such that given $\langle\langle M_1, M_2 \rangle\rangle$, M_1 and M_2 can be recovered in polynomial time. If s is a string and $a \leq b$ are integers then $s[a \dots b]$ denotes the string consisting of bit positions a through b of s . The algorithms constituting $SS = (SK, SE, SD)$ are depicted in Figure 3. To conclude the proof, we need only establish the following:

Claim 4.2 Symmetric encryption scheme SS is IND-CCA secure. ■

$\text{SK}(1^k)$ $K' \xleftarrow{\$} \text{SK}'(1^{\lceil k/2 \rceil})$ $K_2 \xleftarrow{\$} \{0, 1\}^{\lceil k/2 \rceil}$ Return $K' K_2$	$\text{SE}(K, M)$ $k \leftarrow K $ $K' \leftarrow K[1 \dots \lceil k/2 \rceil]$ $K_2 \leftarrow K[1 + \lceil k/2 \rceil \dots k]$ $C' \leftarrow \text{SE}'(K', M)$ Parse M as $\langle (M_1, M_2) \rangle$ If the parsing fails then Return $C' 1$ EndIf $p \leftarrow \text{VfPK}(1^k, M_1)$ $c \leftarrow \text{VfCtxt}(1^k, M_1, K, M_2)$ If $(p = 1 \text{ and } c = 1)$ then Return $C' 0$ else Return $C' 1$ EndIf	$\text{SD}(K, C)$ $k \leftarrow K $ $K' \leftarrow K[1 \dots \lceil k/2 \rceil]$ $K_2 \leftarrow K[1 + \lceil k/2 \rceil \dots k]$ Parse C as $C' d$, where $d \in \{0, 1\}$ $M' \leftarrow \text{SD}'(K', C')$ Parse M' as $\langle (M_1, M_2) \rangle$ If the parsing fails then If $d = 1$ then Return M' else Return \perp EndIf $p \leftarrow \text{VfPK}(1^k, M_1)$ $c \leftarrow \text{VfCtxt}(1^k, M_1, K, M_2)$ If $(d = 0 \text{ and } p = 1 \text{ and } c = 1)$ then Return M' EndIf If $(d = 1 \text{ and } (p \neq 1 \text{ or } c \neq 1))$ then Return M' EndIf Return \perp
--	--	---

Figure 3: Algorithms of the symmetric encryption scheme $\text{SS} = (\text{SK}, \text{SE}, \text{SD})$ for the proof of Theorem 4.1. Above, $\langle (M_1, M_2) \rangle$ denotes an encoding of the pair of strings (M_1, M_2) as a string.

Claim 4.3 Multi-message hybrid encryption scheme $(\overline{\text{HEG}}[\text{CG}], \text{SS})$ is not IND-CCA secure. ■

Proof of Claim 4.2: Let us first provide some intuition. Note that on input M , encryption algorithm $\text{SE}(K'_1 || K_2, \cdot)$ uses the encryption algorithm SE' of an IND-CCA-secure scheme to compute $C' \xleftarrow{\$} \text{SE}'(K'_1, M)$ and outputs $C' || 0$ or $C' || 1$, depending on whether M has some “special” form or not. The ciphertext ends with 0 if M parses as a pair (M_1, M_2) such that algorithms VfPK , VfCtxt indicate that M_1 is a $(\overline{\text{HEG}}[\text{CG}], k)$ -valid public key and $M_2 \in [\overline{\text{AE}}(M_1, K'_1 || K_2)]$. The decryption algorithm $\text{SD}(K'_1 || K_2, \cdot)$ on input $C' || d$, where d is a bit, computes $M' \leftarrow \text{SD}'(K'_1, C')$ and returns M' only if either M' is of the special form and $d = 0$, or M' is not of this form and $d = 1$. Therefore, an obvious strategy for an adversary against SS is to query its oracle $\text{SE}(K, \text{LR}(\cdot, \cdot, b))$ on a pair of messages such that one of them is of this special form and the other is not. Using the unique decryptability of $\overline{\text{AE}}$ and the fact that K_2 is chosen at random, independently from the adversary’s view, we show that it cannot find such queries except with negligible probability. Moreover, we show that any strategy for the adversary can be employed by an attacker against scheme SS' to win its game. Details follow.

Let \mathcal{S} be a legitimate polynomial-time adversary attacking SS . We will construct a legitimate polynomial-time adversary \mathcal{S}' such that

$$\text{Adv}_{\text{SS}, \mathcal{S}}^{\text{ind-cca}}(k) \leq \text{Adv}_{\text{SS}', \mathcal{S}'}^{\text{ind-cca}}(\lceil k/2 \rceil) + \frac{O(Q(k))}{2^{\lceil k/2 \rceil}}, \quad (5)$$

where Q is a polynomial upper bounding the total number of queries made by \mathcal{S} to its different oracles. Since SS' is assumed IND-CCA secure, the advantage function associated to \mathcal{S}' above is negligible, and thus so is the advantage function associated to \mathcal{S} . To complete the proof, we need to specify adversary \mathcal{S}' and prove Equation (5).

Adversary \mathcal{S}' is given input $1^{\lceil k/2 \rceil}$ and has access to oracles $\text{SE}'(K'_1, \text{LR}(\cdot, \cdot, b))$ and $\text{SD}'(K'_1, \cdot)$. Its goal is to guess the bit b . It runs \mathcal{S} on input 1^k . In this process, \mathcal{S} will query its two oracles $\text{SE}(K, \text{LR}(\cdot, \cdot, b))$ and $\text{SD}(K, \cdot)$. To answer a query to the first of these oracles, \mathcal{S}' forwards the query to its oracle

$SE'(K'_1, LR(\cdot, \cdot, b))$, appends 1 to the oracle's reply and returns the result to \mathcal{S} . To answer a query to the second oracle, \mathcal{S}' checks the last bit of the query. If it is 0, \mathcal{S}' returns \perp to \mathcal{S} . Otherwise, it removes the last bit, forwards the result to its oracle $SD'(K'_1, \cdot)$, and returns the answer to \mathcal{S} . When \mathcal{S} outputs its guess b' , \mathcal{S}' returns b' .

We now analyze \mathcal{S}' . Consider the experiment in which \mathcal{S}' attacks SS' . We define the following events.

- Succ(\mathcal{S}') : \mathcal{S}' is successful, meaning its output equals the challenge bit b
- BadE : \mathcal{S} makes a query to oracle $SE(K, LR(\cdot, \cdot, b))$ in which one of the messages can be parsed as $\langle (M_1, M_2) \rangle$ such that M_1 is a $(\overline{HEG}[CG], k)$ -valid public key and $M_2 \in \overline{AE}(M_1, K)$
- BadD : \mathcal{S} makes a query to oracle $SD(K, \cdot)$ that can be parsed as $C' || d$, where d is a bit, such that $SD'(K'_1, C') = \langle (M_1, M_2) \rangle$, where M_1 is a $(\overline{HEG}[CG], k)$ -valid public key and $M_2 \in \overline{AE}(M_1, K)$

For the experiment in which \mathcal{S} attacks SS , we define the following event.

- Succ(\mathcal{S}) : \mathcal{S} is successful, meaning its output equals the challenge bit b

We claim that if events BadE and BadD do not occur, then \mathcal{S}' simulates perfectly the environment provided to \mathcal{S} in its attack against SS . First, note that answers to queries to oracle $SE(K, LR(\cdot, \cdot, b))$ can only be off by the last bit. In the absence of the “bad” events, each ciphertext returned to \mathcal{S} as a reply to a query to oracle $SE(K, LR(\cdot, \cdot, b))$ has 1 as the last bit. This is also the case in \mathcal{S}' 's real attack. If \mathcal{S} queries $SD(K, \cdot)$ with a ciphertext $C' || 0$, assuming events BadE and BadD do not occur, \mathcal{S}' gives \mathcal{S} the response it would get in the real attack, namely \perp . Since \mathcal{S} is legitimate, if it queries oracle $SD(K, \cdot)$ with a ciphertext $C' || 1$, then C' must not have previously been returned by oracle $SE'(K'_1, LR(\cdot, \cdot, b))$. Thus \mathcal{S}' can legitimately make query C' to its oracle $SD'(K'_1, \cdot)$. If M is the response, then, assuming that events BadE and BadD do not occur, the answer \mathcal{S} expects is exactly M . Therefore,

$$\begin{aligned} \Pr [\text{Succ}(\mathcal{S}')] &\geq \Pr [\text{Succ}(\mathcal{S}') \mid \neg \text{BadE} \wedge \neg \text{BadD}] - \Pr [\text{BadE} \vee \text{BadD}] \\ &\geq \Pr [\text{Succ}(\mathcal{S})] - \Pr [\text{BadE} \vee \text{BadD}] . \end{aligned}$$

We now provide an upper bound for the probability of event $\text{BadE} \vee \text{BadD}$. Let $q_e(k)$ and $q_d(k)$ be the number of queries \mathcal{S} makes to oracles $SE(K, LR(\cdot, \cdot, b))$ and $SD(K, \cdot)$, respectively, on input 1^k . We observe that if M_1 is a $(\overline{HEG}[CG], k)$ -valid public key, then for any $M_2 \in \{0, 1\}^*$, there exists a unique $K' \in [\text{SK}(1^k)]$ such that $M_2 \in \overline{AE}(M_1, K')$. Recall that the key for oracles $SE(K, LR(\cdot, \cdot, b))$ and $SD(K, \cdot)$ is $K = K'_1 || K_2$, where K_2 is chosen uniformly at random from $\{0, 1\}^{\lfloor k/2 \rfloor}$ and is independent from \mathcal{S}' 's view. Therefore, for any query made by \mathcal{S} to oracle $SE(K, LR(\cdot, \cdot, b))$, the probability that one of the messages in the query parses as $\langle (M_1, M_2) \rangle$ such that M_1 is a $(\overline{HEG}[CG], k)$ -valid public key and $M_2 \in \overline{AE}(M_1, K)$ is at most $2/2^{\lfloor k/2 \rfloor}$. Similarly, for any query $C' || d$, where d is a bit, made by \mathcal{S} to oracle $SD(K, \cdot)$, the probability that $SD'(K'_1, C') = M'$, where M' parses as $\langle (M_1, M_2) \rangle$, M_1 is a $(\overline{HEG}[CG], k)$ -valid public key and $M_2 \in \overline{AE}(M_1, K)$ is at most $1/2^{\lfloor k/2 \rfloor}$. Therefore,

$$\Pr [\text{BadE} \vee \text{BadD}] \leq \frac{2q_e(k) + q_d(k)}{2^{\lfloor k/2 \rfloor}} \leq \frac{2 \cdot Q(k)}{2^{\lfloor k/2 \rfloor}} ,$$

where $Q(k) = q_e(k) + q_d(k)$. Hence

$$\begin{aligned} \text{Adv}_{SS', \mathcal{S}'}^{\text{ind-cca}}(\lceil k/2 \rceil) &= 2 \cdot \Pr [\text{Succ}(\mathcal{S}')] - 1 \geq 2 \cdot \left(\Pr [\text{Succ}(\mathcal{S})] - \frac{O(Q(k))}{2^{\lfloor k/2 \rfloor}} \right) - 1 \\ &= \text{Adv}_{SS, \mathcal{S}}^{\text{ind-cca}}(k) - \frac{O(Q(k))}{2^{\lfloor k/2 \rfloor}} . \end{aligned}$$

Rearranging terms gives Equation (5). \blacksquare

Proof of Claim 4.3: We define a hybrid adversary \mathbf{H} attacking $(\overline{\text{HEG}}[\text{CG}], \text{SS})$. \mathbf{H} is given inputs $pk = (q, g, X, fk)$ and C_a and has access to oracles $\text{SE}(K, \text{LR}(\cdot, \cdot, b))$, $\text{SD}(K, \cdot)$, and $\text{AD}(sk, \cdot)$, where $sk = (q, g, x, fk)$. Its goal is to guess the challenge bit b . By the definition of experiment $\text{Exp}_{\overline{\text{HEG}}[\text{CG}], \text{SS}, \mathbf{H}}^{\text{ind-cca}}(k)$, pk is a $(\overline{\text{HEG}}[\text{CG}], k)$ -valid public key and $C_a \in [\overline{\text{AE}}(pk, K)]$. Therefore, $\langle (pk, C_a) \rangle$ is a message which, when encrypted with $\text{SE}(K, \cdot)$, yields a ciphertext that has last bit 0. We observe that for any string C chosen at random from $\{0, 1\}^{|C_a|} \setminus \{C_a\}$, the probability that $K = \text{AD}(sk, C)$ is 0 (since $\text{AE}(pk, K) = C_a$ and AE is deterministic), i.e., the probability that $C \in [\overline{\text{AE}}(pk, K)]$ is 0. Hence $\langle (pk, C) \rangle$ is a message which, when encrypted with $\text{SE}(K, \cdot)$, yields a ciphertext that has last bit 1. (If $C \notin [\overline{\text{AE}}(pk, K)]$, then the last bit will be 1.) Thus, adversary \mathbf{H} can construct two messages for which it can guess with probability 1 the last bit of the corresponding ciphertext. Using this information it can then guess the challenge bit. Details follow.

Adversary \mathbf{H} chooses C at random from $\{0, 1\}^{|C_a|} \setminus \{C_a\}$, makes a query $\langle (pk, C_a) \rangle, \langle (pk, C) \rangle$ to oracle $\text{SE}(K, \text{LR}(\cdot, \cdot, b))$, parses the response as $C' || d$, where d is a bit, and returns d . The running time of \mathbf{H} is clearly polynomial in k . We claim that $\text{Adv}_{\overline{\text{HEG}}[\text{CG}], \text{SS}, \mathbf{H}}^{\text{ind-cca}}(k) = 1$. To prove this, we consider the event

$\text{Succ}(\mathbf{H})$: \mathbf{H} is successful, meaning its output equals the challenge bit b

If challenge bit b is 0, then the response to \mathbf{H} 's query is a ciphertext that has last bit 0. If bit b is 1, then the response is a ciphertext that has last bit 1. Thus

$$\Pr[\text{Succ}(\mathbf{H})] = \frac{1}{2} + \frac{1}{2} = 1.$$

Hence

$$\text{Adv}_{\overline{\text{HEG}}[\text{CG}], \text{SS}, \mathbf{H}}^{\text{ind-cca}}(k) = 2 \cdot \Pr[\text{Succ}(\mathbf{H})] - 1 = 1,$$

as desired. \blacksquare

Notice that the adversary constructed in the proof of Claim 4.3 does not make any queries to its oracles $\text{SD}(K, \cdot)$ and $\text{AD}(sk, \cdot)$.

Remark 4.4 An interesting question at this point may be why the proof of Theorem 4.1 fails for the RO-model Hash ElGamal scheme $\text{HEG}[\text{CG}]$ associated to a cyclic-group generator CG —it must, since otherwise Theorem 3.1 would be contradicted—but succeeds for any instantiation of this scheme. The answer is that symmetric encryption scheme SS , depicted in Figure 3 runs a ciphertext verifier VfCtxt for the asymmetric encryption scheme in question. In the case of the RO-model scheme $\text{HEG}[\text{CG}]$, any ciphertext verifier must query random oracles G and H . But as we clarified in Section 2, SS does not have access to these oracles (although it might have access to its own, independently chosen oracle R_s), and so cannot run such a ciphertext verifier. The adversary of course does have access to G, H , but has no way to “pass” these objects to the encryption algorithm of the symmetric encryption scheme. On the other hand, in the instantiated scheme, the keys describing the functions instantiating the random oracles may be passed by the adversary to the encryption algorithm of SS in the form of a message containing the public key, giving SS the ability to run the ciphertext verifier.

This might lead one to ask why SS does not have oracle access to G, H . The answer is that in general, when several RO-model schemes are being executed in a common context, their random oracles should be chosen independently of each other. In this case, RO-model schemes SS and AS are being executed in a common context in $\text{Exp}_{\text{AS}, \text{SS}, \mathbf{H}}^{\text{ind-cca}}(k)$ which, correctly, chose their random oracles independently, and thus the algorithms of SS do not have access to the oracle of AS . \blacksquare

As we discussed in Section 1, in Appendix B we provide a more general impossibility result.

References

- [1] M. AGARWAL, N. SAXENA AND N. KAYAL, “PRIMES is in P,” *Preprint. Available at <http://www.cse.iitk.ac.in/news/primality.html>*, August 6, 2002.
- [2] J. BAEK, B. LEE AND K. KIM, “Secure length-saving ElGamal encryption under the Computational Diffie-Hellman assumption,” *Proceedings of the Fifth Australasian Conference on Information Security and Privacy–ACISP 2000*, Lecture Notes in Computer Science Vol. 1841, E. Dawson, A. Clark and C. Boyd ed., Springer-Verlag, 2000.
- [3] J. BAEK, B. LEE AND K. KIM, “Provably secure length-saving public-key encryption scheme under the computational Diffie-Hellman assumption,” *ETRI Journal*, 22(4), 2000.
- [4] M. BELLARE, A. DESAI, E. JOKIPII AND P. ROGAWAY, “A concrete security treatment of symmetric encryption: Analysis of the DES modes of operation,” *Proceedings of the 38th Symposium on Foundations of Computer Science*, IEEE, 1997.
- [5] M. BELLARE AND P. ROGAWAY, Random oracles are practical: a paradigm for designing efficient protocols, *First ACM Conference on Computer and Communications Security*, ACM, 1993.
- [6] R. CANETTI, U. FEIGE, O. GOLDREICH AND M. NAOR, “Adaptively secure multi-party computation,” *Proceedings of the 28th Annual Symposium on the Theory of Computing*, ACM, 1996.
- [7] R. CANETTI, O. GOLDREICH, S. HALEVI, “The random oracle methodology, revisited,” *Proceedings of the 30th Annual Symposium on the Theory of Computing*, ACM, 1998.
- [8] J.-S. CORON, H. HANDSCHUH, M. JOYE, P. PAILLIER, D. POINTCHEVAL, C. TYMEN, “GEM: A Generic Chosen-Ciphertext Secure Encryption Method”, *Topics in Cryptology – CT-RSA ’02*, Lecture Notes in Computer Science Vol. 2271 , B. Preneel ed., Springer-Verlag, 2002.
- [9] J.-S. CORON, H. HANDSCHUH, M. JOYE, P. PAILLIER, D. POINTCHEVAL, C. TYMEN, “Optimal Chosen-Ciphertext Secure Encryption of Arbitrary-Length Messages,” *Proceedings of the Fifth International workshop on practice and theory in Public Key Cryptography (PKC’02)*, Lecture Notes in Computer Science Vol. 1431, D. Naccache and P. Paillier eds., Springer-Verlag, 2002.
- [10] R. CRAMER AND V. SHOUP, “A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack,” *Advances in Cryptology – CRYPTO ’98*, Lecture Notes in Computer Science Vol. 1462, H. Krawczyk ed., Springer-Verlag, 1998.
- [11] R. CRAMER AND V. SHOUP, “Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack,” IACR ePrint archive Record 2001/108, 2001, <http://eprint.iacr.org/>.
- [12] T. ELGAMAL, “A public key cryptosystem and signature scheme based on discrete logarithms,” *IEEE Transactions on Information Theory*, vol 31, 1985.
- [13] A. FIAT AND A. SHAMIR, “How to prove yourself: practical solutions to identification and signature problems,” *Advances in Cryptology – CRYPTO ’86*, Lecture Notes in Computer Science Vol. 263, A. Odlyzko ed., Springer-Verlag, 1986.
- [14] E. FUJISAKI, T. OKAMOTO, “Secure Integration of Asymmetric and Symmetric Encryption Schemes,” *Advances in Cryptology – CRYPTO ’99*, Lecture Notes in Computer Science Vol. 1666, M. Wiener ed., Springer-Verlag, 1999.
- [15] S. GOLDWASSER AND S. MICALI, “Probabilistic encryption,” *Journal of Computer and System Science*, Vol. 28, 1984, pp. 270–299.
- [16] S. GOLDWASSER AND Y. TAUMANN, “On the (in)security of the Fiat-Shamir paradigm,” *Proceedings of the 44th Symposium on Foundations of Computer Science*, IEEE, 2003.
- [17] S. MICALI, “Computationally sound proofs,” *SIAM Journal on Computing*, Vol. 30, No. 4, 2000, pp. 1253-1298.

```

Subroutine GSim( $Z$ )
  If GT[ $Z$ ] is not defined then GT[ $Z$ ]  $\stackrel{\$}{\leftarrow}$   $\{0, 1\}^k$  EndIf
  Return GT[ $Z$ ]

```

```

Subroutine HSim( $L$ )
  If HT[ $L$ ] is defined then return it as the answer EndIf
  (dec, gs)  $\stackrel{\$}{\leftarrow}$  KeyTest( $L$ ); HT[ $L$ ]  $\stackrel{\$}{\leftarrow}$   $\mathbb{Z}_q$ 
  If dec = 0 then return HT[ $L$ ] as the answer EndIf
  If dec = 1 then output gs (as a guess to the value of challenge bit  $b$ ) and halt EndIf

```

```

Subroutine KeyTest( $L$ )
  dec  $\leftarrow$  1
  For  $i = 1, \dots, k$  do
     $C_0^i[L] \stackrel{\$}{\leftarrow} SE^{R_s}(K, LR(T_0, T_0, b))$ ; If  $SD^{R_s}(L, C_0^i[L]) \neq T_0$  then dec  $\leftarrow$  0 EndIf
     $C_1^i[L] \stackrel{\$}{\leftarrow} SE^{R_s}(K, LR(T_1, T_1, b))$ ; If  $SD^{R_s}(L, C_1^i[L]) \neq T_1$  then dec  $\leftarrow$  0 EndIf
  EndFor
  If dec = 0 return (0, 0) EndIf
  For  $i = 1, \dots, k$  do  $C^i[L] \stackrel{\$}{\leftarrow} SE^{R_s}(K, LR(T_0, T_1, b))$ ;  $T^i \leftarrow SD^{R_s}(L, C^i[L])$  EndFor
  If  $T^1 = T^2 = \dots = T^k = T_0$  then return (1, 0) EndIf
  If  $T^1 = T^2 = \dots = T^k = T_1$  then return (1, 1) EndIf
  Return (0, 0)

```

Figure 4: Subroutines defined by \mathbf{S} and used to simulate \mathbf{H} 's oracles.

-
- [18] J. B. NIELSEN “Separating Random Oracle Proofs from Complexity Theoretic Proofs: The Non-committing Encryption Case,” *Advances in Cryptology – CRYPTO ’02*, Lecture Notes in Computer Science Vol. 2442, M. Yung ed., Springer-Verlag, 2002.
 - [19] T. OKAMOTO AND D. POINTCHEVAL “REACT: Rapid Enhanced-security Asymmetric Cryptosystem Transform,” *Topics in Cryptology – CT-RSA ’01*, Lecture Notes in Computer Science Vol. 2020, D. Naccache ed., Springer-Verlag, 2001.
 - [20] V. SHOUP, “A proposal for an ISO standard for public key encryption”, IACR ePrint archive Record 2001/112, 2001, <http://eprint.iacr.org/>.

A Proof of Theorem 3.1

We explained the ideas behind this proof in Section 3. Here we provide the full adversary constructions and analyses.

PROOF SETUP. Let \mathbf{H} be a polynomial-time hybrid adversary attacking (AS, SS). We will construct polynomial-time adversaries \mathbf{S} and \mathbf{C} such that

$$\text{Adv}_{\text{AS,SS},\mathbf{H}}^{\text{ind-cca}}(k) \leq \text{Adv}_{\text{SS},\mathbf{S}}^{\text{ind-cca}}(k) + O(Q(k)) \cdot \text{Adv}_{\text{CG},\mathbf{C}}^{\text{cdh}}(k) + \frac{O(Q(k)^2)}{2^k}, \quad (6)$$

where $Q(k)$ is a polynomial upper bounding the number of queries made by \mathbf{H} to the G and H oracles. (This includes queries made directly by \mathbf{H} and those made indirectly as a consequence of \mathbf{H} 's queries to its $\text{AD}^{G,H}((q, g, x), \cdot)$ oracle.) Since SS is assumed IND-CCA secure and the CDH problem is hard for CG, the advantage functions related to \mathbf{S} and \mathbf{C} above are negligible, and thus so is the advantage function related to \mathbf{H} . To complete the proof, we need to specify the adversaries \mathbf{S} , \mathbf{C} and prove Equation (6).

```

Subroutine GSim( $Z$ )
  If  $\text{GT}[Z]$  is not defined then  $\text{GT}[Z] \stackrel{\$}{\leftarrow} \{0, 1\}^k$  EndIf
  Return  $\text{GT}[Z]$ 

```

```

Subroutine HSim( $L$ )
  If  $\text{HT}[L]$  is not defined then  $\text{HT}[L] \stackrel{\$}{\leftarrow} \mathbb{Z}_q$  EndIf
  Return  $\text{HT}[L]$ 

```

```

Subroutine ADSim( $Y', W'$ )
  If there is no  $L$  such that  $g^{\text{HT}[L]} = Y'$  then return  $\perp$  EndIf
  Let  $L$  be such that  $g^{\text{HT}[L]} = Y'$ 
   $Z' \leftarrow X^{\text{HT}[L]}$ ;  $T' \leftarrow \text{GSim}(Z')$ ;  $K' \leftarrow T' \oplus W'$ ; Return  $K'$ 

```

Figure 5: Subroutines defined by \mathbf{C} and used to simulate \mathbf{H} 's oracles.

DESCRIPTION OF \mathbf{S} . Adversary \mathbf{S} is given input 1^k and has access to oracles $\text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, b))$, $\text{SD}^{R_s}(K, \cdot)$, and R_s . Its goal is to guess the bit b . It begins with the following initializations:

$$((q, g, X), (q, g, x)) \stackrel{\$}{\leftarrow} \text{AK}(1^k); y \stackrel{\$}{\leftarrow} \mathbb{Z}_q; Y \leftarrow g^y; W \stackrel{\$}{\leftarrow} \{0, 1\}^k; C_a \leftarrow (Y, W);$$

$$T_0 \stackrel{\$}{\leftarrow} \{0, 1\}^k; T_1 \stackrel{\$}{\leftarrow} \{0, 1\}^k - \{T_0\}.$$

Then it runs \mathbf{H} on inputs public key (q, g, X) and ciphertext C_a . In the process \mathbf{H} will query its oracles

$$G, H, R_s, \text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, b)), \text{SD}^{R_s}(K, \cdot), \text{AD}^{G, H}((q, g, x), \cdot). \quad (7)$$

\mathbf{S} will answer these queries. To that end, it defines the subroutines shown in Figure 4. It answers a query Z to G by running $\text{GSim}(Z)$ and returning the answer to \mathbf{H} . It answers a query L to H by running $\text{HSim}(L)$ and returning the answer to \mathbf{H} . It answers queries to the $\text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, b))$ and R_s oracles via its own oracles of the same name. It answers each query C to the $\text{SD}^{R_s}(K, \cdot)$ oracle using its own decryption oracle, unless there exist i, j and L such that L was queried to H and either $C = C_j^i[L]$ or $C = C^i[L]$. In that case, \mathbf{S} aborts. Since \mathbf{S} possesses the secret key (q, g, x) , it can answer queries to $\text{AD}^{G, H}((q, g, x), \cdot)$ by performing the computation of the decryption algorithm, replacing calls that the latter makes to G or H by calls to the relevant subroutines just mentioned. If \mathbf{H} runs to completion (\mathbf{S} can output its guess as to the value of b , and halt, before this), then \mathbf{S} outputs whatever \mathbf{H} outputs.

DESCRIPTION OF \mathbf{C} . Adversary \mathbf{C} is given inputs q, g, X, Y , where $X, Y \in \langle g \rangle$ have been chosen uniformly at random. Its goal is to compute g^{xy} where $g^x = X$ and $g^y = Y$. Let $k \leftarrow |\langle 2q + 1 \rangle|$. \mathbf{C} begins with the following initializations:

$$K \stackrel{\$}{\leftarrow} \text{SK}(1^k); b \stackrel{\$}{\leftarrow} \{0, 1\}; W \stackrel{\$}{\leftarrow} \{0, 1\}^k; C_a \leftarrow (Y, W).$$

Then it runs \mathbf{H} on inputs public key (q, g, X) and ciphertext C_a . In the process \mathbf{H} will query the oracles listed in Equation (7). \mathbf{C} will answer these queries. Queries to R_s are simulated the standard way, by returning a random value for each new query and the previously returned value for each repeated query. To simulate the rest of the oracles it defines the subroutines shown in Figure 5. It answers a query Z to G by running $\text{GSim}(Z)$ and returning the answer to \mathbf{H} . It answers a query L to H by running $\text{HSim}(L)$ and returning the answer to \mathbf{H} . Since it possesses K and b , it can answer queries to the $\text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, b))$ or $\text{SD}^{R_s}(K, \cdot)$ oracles by simply performing the relevant computation and returning the answer. It answers a query (Y', W') to $\text{AD}^{G, H}((q, g, x), \cdot)$ by running $\text{ADSim}(Y', W')$ and returning the answer. When \mathbf{H} has terminated, \mathbf{C} picks Z at random from the set $\{Z : \text{GT}[Z] \text{ is defined} \}$ and outputs Z .

ANALYSIS. For the analysis, define the following experiment:

$$\mathbf{Exp}_{\mathbf{CG},\mathbf{C}}^{\text{cdh}}(k) : (q, g) \xleftarrow{\$} \mathbf{CG}(1^k); x, y \xleftarrow{\$} \mathbb{Z}_q; Z \leftarrow \mathbf{C}(q, g, g^x, g^y)$$

If $Z = g^{xy}$ then return 1 else return 0

We let $\Pr_{\mathbf{C}}[\cdot]$, $\Pr_{\mathbf{S}}[\cdot]$, and $\Pr_{\mathbf{H}}[\cdot]$ denote the probabilities in experiments $\mathbf{Exp}_{\mathbf{CG},\mathbf{C}}^{\text{cdh}}(k)$, $\mathbf{Exp}_{\mathbf{SS},\mathbf{S}}^{\text{ind-cca}}(k)$, and $\mathbf{Exp}_{\mathbf{AS},\mathbf{SS},\mathbf{H}}^{\text{ind-cca}}(k)$, respectively.

Let $((q, g, X), (q, g, x)) \in [\mathbf{AK}(1^k)]$ and $K \in [\mathbf{SK}(1^k)]$. We define the following events relating to \mathbf{H} 's execution on inputs public key (q, g, X) and ciphertext $C_a = (Y, W)$ where $g^y = Y$. These events are defined in any of the three experiments we are considering:

- GH : There exists a time at which g^{xy} is queried to G but K has not been queried to H
- HG : There exists a time at which K has been queried to H but g^{xy} has not been queried to G
- Succ(\mathbf{H}) : \mathbf{H} is successful, meaning its output equals the challenge bit b .

We clarify that the queries referred to above include both direct and indirect queries of \mathbf{H} , but, in the case of $\mathbf{Exp}_{\mathbf{AS},\mathbf{SS},\mathbf{H}}^{\text{ind-cca}}(k)$, they do *not* include the queries to G and H made by the computation $C_a \leftarrow \mathbf{AE}^{G,H}((q, g, X), \cdot)$ that initializes the experiment. (We are only considering queries to G, H resulting from the execution of \mathbf{H} .) The main claims related to the analysis are:

$$\Pr_{\mathbf{H}}[\text{HG} \vee (\text{Succ}(\mathbf{H}) \wedge \neg\text{HG} \wedge \neg\text{GH})] \leq \Pr_{\mathbf{S}}[\mathbf{Exp}_{\mathbf{SS},\mathbf{S}}^{\text{ind-cca}}(k) = 1] + \frac{O(Q(k))}{2^k} \quad (8)$$

$$\Pr_{\mathbf{H}}[\text{GH}] \leq Q(k) \cdot \Pr_{\mathbf{C}}[\mathbf{Exp}_{\mathbf{CG},\mathbf{C}}^{\text{cdh}}(k) = 1] + \frac{O(Q(k)^2)}{2^k}. \quad (9)$$

Let us see how these enable us to conclude the proof, and then return to prove them. We have:

$$\begin{aligned} & \frac{1}{2} \cdot \text{Adv}_{\mathbf{AS},\mathbf{SS},\mathbf{H}}^{\text{ind-cca}}(k) + \frac{1}{2} \\ &= \Pr_{\mathbf{H}}[\mathbf{Exp}_{\mathbf{AS},\mathbf{SS},\mathbf{H}}^{\text{ind-cca}}(k) = 1] \\ &= \Pr_{\mathbf{H}}[\text{Succ}(\mathbf{H})] \\ &= \Pr_{\mathbf{H}}[(\text{Succ}(\mathbf{H}) \wedge \text{HG}) \vee (\text{Succ}(\mathbf{H}) \wedge \neg\text{HG} \wedge \neg\text{GH})] + \Pr_{\mathbf{H}}[\text{Succ}(\mathbf{H}) \wedge \text{GH}] \\ &\leq \Pr_{\mathbf{H}}[\text{HG} \vee (\text{Succ}(\mathbf{H}) \wedge \neg\text{HG} \wedge \neg\text{GH})] + \Pr_{\mathbf{H}}[\text{GH}] \\ &\leq \Pr_{\mathbf{S}}[\mathbf{Exp}_{\mathbf{SS},\mathbf{S}}^{\text{ind-cca}}(k) = 1] + \frac{O(Q(k))}{2^k} + Q(k) \cdot \Pr_{\mathbf{C}}[\mathbf{Exp}_{\mathbf{CG},\mathbf{C}}^{\text{cdh}}(k) = 1] + \frac{O(Q(k)^2)}{2^k} \\ &= \frac{1}{2} \cdot \text{Adv}_{\mathbf{SS},\mathbf{S}}^{\text{ind-cca}}(k) + \frac{1}{2} + Q(k) \cdot \text{Adv}_{\mathbf{CG},\mathbf{C}}^{\text{cdh}}(k) + \frac{O(Q(k)^2)}{2^k}. \end{aligned}$$

Re-arranging terms and simplifying, we get Equation (6). To complete the proof, we must establish Equations (8) and (9).

PROOF OF EQUATION (8). An important ingredient in this proof is the following lemma that characterizes what Subroutine KeyTest accomplishes:

Lemma A.1 If $L = K$ then $\text{KeyTest}(L)$ returns $(1, b)$, while if $L \neq K$ then

$$\Pr\left[(\text{dec}, \text{gs}) \xleftarrow{\$} \text{KeyTest}(L) : (\text{dec}, \text{gs}) = (1, 1 - b)\right] \leq 4^{-k}. \quad \blacksquare$$

In other words, if $L \neq K$, then with high probability either the test indicates this by returning $\text{dec} = 0$ or it successfully computes the value of the challenge bit b . Above, the probability is over the coin tosses made by the $\text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, b))$ oracle called in KeyTest , with K and b fixed.

Proof of Lemma A.1: The fact that $\text{KeyTest}(L)$ returns $(1, b)$ when $L = K$ is a consequence merely of the unique decryptability of SS , namely the fact that for all $K \in [\text{SK}(1^k)]$ and all $M \in \{0, 1\}^*$ we have $\text{SD}^{R_s}(K, \text{SE}^{R_s}(K, M)) = M$ with probability one, the probability being over the coin tosses of SE .

Now assume $L \neq K$. Let $\Pr[\cdot]$ denote the probability taken over the coin tosses of $\text{SE}^{R_s}(K, \cdot)$, with K fixed. Let

$$\begin{aligned} P_0 &= \Pr[\text{SD}^{R_s}(L, \text{SE}^{R_s}(K, T_0)) = T_0] \quad \text{and} \\ P_1 &= \Pr[\text{SD}^{R_s}(L, \text{SE}^{R_s}(K, T_1)) = T_1]. \end{aligned}$$

The probability that $\text{dec} = 1$ at the end of the first For loop in subroutine KeyTest is $P_0^k P_1^k$ and the probability that $T^1 = \dots = T^k = T_{1-b}$ is at most $(1 - P_b)^k$. So we have

$$\begin{aligned} \Pr\left[(\text{dec}, \text{gs}) \stackrel{\$}{\leftarrow} \text{KeyTest}(L) : (\text{dec}, \text{gs}) = (1, 1 - b)\right] &= P_0^k P_1^k \cdot (1 - P_b)^k \\ &\leq P_b^k \cdot (1 - P_b)^k \\ &= [P_b(1 - P_b)]^k \\ &\leq 4^{-k}. \end{aligned}$$

The last line is true because the function $f: [0, 1] \rightarrow \mathbb{R}$ defined by $f(x) = x(1 - x)$ attains its maximum at $x = 1/2$ and the value of this maximum is $1/4$. This concludes the proof. \blacksquare

Returning to the proof of Equation (8), we define the following events in $\mathbf{Exp}_{\text{SS}, \mathcal{S}}^{\text{ind-cca}}(k)$:

- FailTest** : There exists $L \neq K$ such that L was queried to H and $\text{KeyTest}(L)$ returned $(1, 1 - b)$ in subroutine $\text{HSim}(L)$
- Illegit** : There exist i, j and L such that L was queried to H and either $C_j^i[L]$ or $C^i[L]$ was queried by H to $\text{SD}^{R_s}(K, \cdot)$.

We obtain Equation (8) as shown below. Justifications follow the formulas:

$$\begin{aligned} &\Pr_{\mathbf{H}}[\text{HG} \vee (\text{Succ}(\mathbf{H}) \wedge \neg \text{HG} \wedge \neg \text{GH})] \\ &\leq \Pr_{\mathcal{S}}[\text{HG} \vee (\text{Succ}(\mathbf{H}) \wedge \neg \text{HG} \wedge \neg \text{GH}) \mid \neg \text{FailTest}] + \Pr_{\mathcal{S}}[\text{FailTest}] \end{aligned} \quad (10)$$

$$\leq \Pr_{\mathcal{S}}[\mathbf{Exp}_{\text{SS}, \mathcal{S}}^{\text{ind-cca}}(k) = 1] + \Pr_{\mathcal{S}}[\text{Illegit}] + \Pr_{\mathcal{S}}[\text{FailTest}] \quad (11)$$

$$\begin{aligned} &\leq \Pr_{\mathcal{S}}[\mathbf{Exp}_{\text{SS}, \mathcal{S}}^{\text{ind-cca}}(k) = 1] + \Pr_{\mathcal{S}}[\text{Illegit} \mid \neg \text{FailTest}] + 2 \cdot \Pr_{\mathcal{S}}[\text{FailTest}] \\ &\leq \Pr_{\mathcal{S}}[\mathbf{Exp}_{\text{SS}, \mathcal{S}}^{\text{ind-cca}}(k) = 1] + \frac{O(Q(k))}{2^k}. \end{aligned} \quad (12)$$

To justify Equation (10), observe that if event FailTest does not happen, then the simulation of \mathbf{H} done by \mathcal{S} is correct. (If HG occurs, then prior to this g^{xy} was not a query to G , so the simulation of the G oracle is correct. If $\neg \text{HG} \wedge \neg \text{GH}$ occurs, then also g^{xy} was not a query to G , so the simulation of the G oracle is correct. If FailTest does not occur, then the replies to queries to H are correct.)

To justify Equation (11), first note that if event HG occurs, then the $L = K$ case of Lemma A.1 tells us that \mathcal{S} halts with correct output. On the other hand, if neither HG nor GH occur, then \mathcal{S} halts with correct output as long as \mathbf{H} does. But $\mathbf{Exp}_{\text{SS}, \mathcal{S}}^{\text{ind-cca}}(k)$ can still fail to return 1 because \mathcal{S} aborted due to the occurrence

of Illegit. (When the latter occurs, \mathcal{S} aborts to avoid calling its oracle $\text{SD}^{R_s}(K, \cdot)$ on a ciphertext returned by its $\text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, b))$ oracle.)

To justify Equation (12), first note that Lemma A.1 together with the fact that the total number of queries is at most $Q(k)$ implies that $\Pr_{\mathcal{S}}[\text{FailTest}] \leq Q(k)/4^k$. Next, we observe that if FailTest does not occur, then \mathbf{H} gets no information about T_0, T_1 other than that they are random distinct k -bit strings. The unique decryptability of SS then tells us that $\Pr_{\mathcal{S}}[\text{Illegit} \mid \neg\text{FailTest}]$ is bounded above by the probability of guessing either T_0 or T_1 in $Q(k)$ tries, and this is $O(Q(k)/2^k)$.

PROOF OF EQUATION (9). We define the following event in $\text{Exp}_{\text{CG}, \mathcal{C}}^{\text{cdh}}(k)$:

FailDec : There exist times $t_1 < t_2$ and Y', W', L such that all the following hold:

- query (Y', W') was made to $\text{AD}^{G, H}((q, g, x), \cdot)$ at time t_1 and $\text{ADSim}(Y', W')$ returned \perp
- query L was made to H at time t_2
- $g^{\text{HT}[L]} = Y'$.

The answers provided by $\text{ADSim}(\cdot, \cdot)$ are correct exactly when this event does not occur. Furthermore, if there is a time at which query g^{xy} to G occurs and GH is true, then query K to H has not occurred at this time, and thus the answers to queries to H have been correct. Hence

$$\Pr_{\mathcal{C}}[\text{Exp}_{\text{CG}, \mathcal{C}}^{\text{cdh}}(k) = 1] \geq \frac{\Pr_{\mathbf{H}}[\text{GH}] - \Pr_{\mathcal{C}}[\text{FailDec}]}{Q(k)}.$$

Re-arranging, we get

$$\Pr_{\mathbf{H}}[\text{GH}] \leq Q(k) \cdot \Pr_{\mathcal{C}}[\text{Exp}_{\text{CG}, \mathcal{C}}^{\text{cdh}}(k) = 1] + \Pr_{\mathcal{C}}[\text{FailDec}]. \quad (13)$$

At any point in time, a query L to H has probability at most ℓ/q of making FailDec happen, where ℓ is the number of queries that have been made to $\text{AD}^{G, H}((q, g, x), \cdot)$ at this time. Recall that $k = \lfloor (2q + 1) \rfloor$ and thus $q \geq 2^{k-2}$. Putting these observations together we get

$$\Pr_{\mathcal{C}}[\text{FailDec}] \leq \frac{Q(k)^2}{q} \leq \frac{Q(k)^2}{2^{k-2}} = \frac{O(Q(k)^2)}{2^k}.$$

Putting this together with Equation (13) completes the proof of Equation (9).

B A generalization

In this section, we identify a subclass of IND-CCA-preserving asymmetric encryption schemes that we call key-verifiable, ciphertext-verifiable IND-CCA-preserving asymmetric encryption schemes. We show that such schemes exist in the RO model, but do not exist in the standard model. We then discuss how this generalizes our results about the Hash El Gamal scheme. We begin by defining the two properties mentioned above, namely, key verifiability and ciphertext verifiability.

Let $\text{AS} = (\text{AK}, \text{AE}, \text{AD})$ be an asymmetric encryption scheme. We say that AS is *key verifiable* if there exists a polynomial-time, possibly randomized algorithm VfPK (called the *key verifier*) and a negligible function ν (called the *error probability* of VfPK) such that VfPK($1^k, pk$) returns 1 with probability at least $1 - \nu(k)$ if pk is an (AS, k) -valid public key, and returns 1 with probability at most $\nu(k)$ otherwise.

We say that asymmetric encryption scheme $\text{AS} = (\text{AK}, \text{AE}, \text{AD})$ is *ciphertext verifiable* if there exists a polynomial-time, possibly randomized algorithm VfCtxt (called the *ciphertext verifier*) and a negligible function ν (called the *error probability* of VfCtxt) such that, if VfCtxt is run on inputs $1^k, pk, K, C$, where pk is an (AS, k) -valid public key and $K \in \{0, 1\}^k$, then VfCtxt returns 1 with probability at least $1 - \nu(k)$ if $C \in [\text{AE}(pk, K)]$, and returns 1 with probability at most $\nu(k)$ otherwise. If AE or AD access a random oracle, then VfCtxt is given access to the same random oracle.

The following will be used later:

Proposition B.1 Suppose AS is a RO-model asymmetric encryption scheme that is both key verifiable and ciphertext verifiable. Let \overline{AS} be any instantiation of AS via a poly-time family of functions. Then \overline{AS} is also both key verifiable and ciphertext verifiable. ■

Proof of Proposition B.1: Let $VfPK$ and $VfCtxt$ be a key verifier and a ciphertext verifier for AS , respectively. Let \overline{F} be the poly-time family of functions used in \overline{AS} to replace the random oracle. Recall that a public key of \overline{AS} contains a public key pk of AS and also a key fk specifying an instance of \overline{F} . We define algorithms \overline{VfPK} and \overline{VfCtxt} .

On inputs $1^k, s$, \overline{VfPK} attempts to parse s as a pair (pk, fk) . If it fails, it returns 0. Otherwise, it runs $VfPK(1^k, pk)$. If the result is 0, it returns 0. Otherwise, it verifies that $fk \in \{0, 1\}^{FKL(k)}$. If so, it returns 1, if not it returns 0. Clearly, \overline{VfPK} is a key verifier for \overline{AS} .

\overline{VfCtxt} is identical to $VfCtxt$ except that the random oracle is replaced with the same instance of \overline{F} used in \overline{AS} to replace the oracle. ■

We now observe that, in the RO model, there exist key-verifiable, ciphertext-verifiable IND-CCA-preserving asymmetric encryption schemes, meaning the goal of key-verifiable, ciphertext-verifiable asymmetric encryption is achievable in this model.

Theorem B.2 Suppose there exists a cyclic-group generator for which the CDH problem is hard. Then there exists a key-verifiable, ciphertext-verifiable RO-model asymmetric encryption scheme that is IND-CCA-preserving in the RO model. ■

Proof of Theorem B.2: If the CDH problem is hard for cyclic-group generator CG , then Theorem 3.1 guarantees that the associated Hash ElGamal asymmetric encryption scheme $HEG[CG]$, defined in Section 3, is IND-CCA preserving in the RO model. The proof of Theorem 3.1 defines a key verifier $VfPK$ and a ciphertext verifier $VfCtxt$ for $HEG[CG]$, each having error probability 0. ■

Next, we show that in the standard model, there do *not* exist key-verifiable, ciphertext-verifiable IND-CCA-preserving asymmetric encryption schemes, meaning the goal of key-verifiable, ciphertext-verifiable asymmetric encryption is *not* achievable in this model.

Theorem B.3 Let \overline{AS} be a standard-model asymmetric encryption scheme that is both key verifiable and ciphertext verifiable. Then \overline{AS} is *not* IND-CCA preserving. ■

Theorem B.3 is proved below. We first state and prove our final result.

Theorem B.4 Let AS be a RO-model asymmetric encryption scheme that is both key verifiable and ciphertext verifiable. Let \overline{AS} be any instantiation of AS via a poly-time family of functions. Then \overline{AS} is *not* IND-CCA preserving. ■

Proof of Theorem B.4: \overline{AS} is a standard-model asymmetric encryption scheme. Proposition B.1 implies that it inherits the key verifiability and ciphertext verifiability of AS . Theorem B.3 then implies that it is *not* IND-CCA preserving. ■

Note that Theorem B.4 implies Theorem 4.1 because Hash El Gamal scheme is a RO model scheme that is key verifiable and ciphertext verifiable. Theorem B.4 is, however, more general, and shows that the un-instantiability of the Hash El Gamal scheme arises not due to some “peculiar” use of random oracles, but due to the fact that the scheme possesses the properties of key verifiability and ciphertext verifiability.

PROOF OF THEOREM B.3. The proof is almost identical to the proof of Theorem 4.1. Accordingly, we use the same notation and the previous results, and only indicate the differences. Let $VfPK$ and $VfCtxt$ be a key

verifier and a ciphertext verifier for \overline{AS} , respectively. The main difference is that now VfPK and VfCtxt can be randomized algorithms with non-zero error probabilities.

We present an IND-CCA-secure symmetric encryption scheme SS such that the mm-hybrid encryption scheme (\overline{AS}, SS) is not IND-CCA secure. This proves the theorem.

Let $SS' = (SK', SE', SD')$ be any standard-model IND-CCA-secure symmetric encryption scheme. The construction of SS is in terms of SS' and algorithms VfPK and VfCtxt, and is exactly as in the proof of Theorem 4.1. See Figure 3. To conclude the proof, we need only establish the following:

Claim B.5 Symmetric encryption scheme SS is IND-CCA secure. ■

Claim B.6 Multi-message hybrid encryption scheme (\overline{AS}, SS) is not IND-CCA secure. ■

Proof of Claim B.5: Let \mathcal{S} be a legitimate polynomial-time adversary attacking SS . We will construct a legitimate polynomial-time adversary \mathcal{S}' such that

$$\text{Adv}_{SS, \mathcal{S}}^{\text{ind-cca}}(k) \leq \text{Adv}_{SS', \mathcal{S}'}^{\text{ind-cca}}(\lceil k/2 \rceil) + \frac{O(Q(k))}{2^{\lfloor k/2 \rfloor}} + O(Q(k)) \cdot \nu(k), \quad (14)$$

where Q is a polynomial upper bounding the total number of queries made by \mathcal{S} to its different oracles, and ν is a negligible function related to the error probabilities of algorithms VfPK and VfCtxt. Note that the last term is the only difference with Equation (5). Since SS' is assumed IND-CCA secure, the advantage function associated to \mathcal{S}' above is negligible, and thus so is the advantage function associated to \mathcal{S} . To complete the proof, we need to specify adversary \mathcal{S}' and prove Equation (14).

Adversary \mathcal{S}' is identical to the adversary in the proof of Claim 4.2. The analysis of \mathcal{S}' is similar, but we need to take into account the possibility that algorithms VfPK and VfCtxt err. For this reason, for the experiment in which \mathcal{S} attacks SS , we define the following additional event.

Crct : Every time algorithms VfPK and VfCtxt are invoked, they return the correct value

We claim that if events BadE and BadD do not occur, then \mathcal{S}' simulates perfectly the environment provided to \mathcal{S} in its attack against SS when algorithms VfPK and VfCtxt never err. First, note that answers to queries to oracle $SE^{R_s}(K, \text{LR}(\cdot, \cdot, b))$ can only be off by the last bit. In the absence of the “bad” events, each ciphertext returned to \mathcal{S} as a reply to a query to oracle $SE^{R_s}(K, \text{LR}(\cdot, \cdot, b))$ has 1 as the last bit. This is also the case in \mathcal{S}' 's real attack when algorithms VfPK and VfCtxt are always correct. If \mathcal{S} queries $SD(K, \cdot)$ with a ciphertext $C' || 0$, assuming events BadE and BadD do not occur, \mathcal{S}' gives \mathcal{S} the response it would get in the real attack when algorithms VfPK and VfCtxt are always correct, namely \perp . Since \mathcal{S} is legitimate, if it queries oracle $SD(K, \cdot)$ with a ciphertext $C' || 1$, then C' must not have previously been returned by oracle $SE'(K'_1, \text{LR}(\cdot, \cdot, b))$. Thus \mathcal{S}' can legitimately make query C' to its oracle $SD'(K'_1, \cdot)$. If M is the response, then, assuming that events BadE and BadD do not occur, the answer \mathcal{S} expects when algorithms VfPK and VfCtxt are always correct is exactly M . Therefore,

$$\begin{aligned} \Pr [\text{Succ}(\mathcal{S}')] &\geq \Pr [\text{Succ}(\mathcal{S}') \mid \neg \text{BadE} \wedge \neg \text{BadD}] - \Pr [\text{BadE} \vee \text{BadD}] \\ &\geq \Pr [\text{Succ}(\mathcal{S}) \mid \text{Crct}] - \Pr [\text{BadE} \vee \text{BadD}] \\ &\geq \Pr [\text{Succ}(\mathcal{S})] - \Pr [\neg \text{Crct}] - \Pr [\text{BadE} \vee \text{BadD}]. \end{aligned}$$

We now provide an upper bound for the probability of event $\neg \text{Crct}$. (The bound for $\text{BadE} \vee \text{BadD}$ is identical to the one in the proof of Claim 4.2.) Let $q_e(k)$ and $q_d(k)$ be the number of queries \mathcal{S} makes to oracles $SE^{R_s}(K, \text{LR}(\cdot, \cdot, b))$ and $SD(K, \cdot)$, respectively, on input 1^k . Let ν_1 be the error probability of key verifier VfPK, and ν_2 the error probability of ciphertext verifier VfCtxt. Then

$$\Pr [\neg \text{Crct}] \leq q_e(k) \cdot (\nu_1(k) + \nu_2(k)) + q_d(k) \cdot (\nu_1(k) + \nu_2(k)) = Q(k) \cdot \nu(k),$$

where $Q(k) = q_e(k) + q_d(k)$ and $\nu(k) = \nu_1(k) + \nu_2(k)$.

Hence

$$\begin{aligned} \text{Adv}_{\text{SS}, \text{S}'}^{\text{ind-cca}}(\lceil k/2 \rceil) &= 2 \cdot \Pr[\text{Succ}(\mathbf{S}')] - 1 \geq 2 \cdot \left(\Pr[\text{Succ}(\mathbf{S})] - Q(k) \cdot \nu(k) - \frac{O(Q(k))}{2^{\lceil k/2 \rceil}} \right) - 1 \\ &= \text{Adv}_{\text{SS}, \mathbf{S}}^{\text{ind-cca}}(k) - O(Q(k)) \cdot \nu(k) - \frac{O(Q(k))}{2^{\lceil k/2 \rceil}}. \end{aligned}$$

Rearranging terms gives Equation (14). \blacksquare

Proof of Claim B.6: We define a hybrid adversary \mathbf{H} attacking $(\overline{\text{AS}}, \text{SS})$ exactly as in the proof of Claim 4.3. We claim that $\text{Adv}_{\text{AS}, \text{SS}, \mathbf{H}}^{\text{ind-cca}}(k) \geq 1 - 2^{-k} - \nu(k)$, where ν is a negligible function related to the error probabilities of algorithms VfPK and VfCtxt. The analysis is similar to the one in the proof of Claim 4.3, but we need to take into account the additional event

Crct : Every time algorithms VfPK and VfCtxt are invoked, they return the correct value,

and the possibility that $K = \text{AD}(sk, C)$, where C is chosen at random from $\{0, 1\}^{|C_a|} \setminus \{C_a\}$. The latter can happen with probability at most 2^{-k} . I.e., the probability that $C \in [\overline{\text{AE}}(pk, K)]$ is at most 2^{-k} . Hence $\langle (pk, C) \rangle$ is a message which, when encrypted with $\text{SE}(K, \cdot)$, yields a ciphertext that with overwhelming probability has last bit 1. (If $C \notin [\overline{\text{AE}}(pk, K)]$, then the last bit will be 1.) Assume that event Crct occurs. If challenge bit b is 0, then the response to \mathbf{H} 's query is a ciphertext that has last bit 0. If bit b is 1, then with probability at least $1 - 2^{-k}$, the response is a ciphertext that has last bit 1. Thus

$$\Pr[\text{Succ}(\mathbf{H})] \geq \Pr[\text{Succ}(\mathbf{H}) \mid \text{Crct}] - \Pr[\neg \text{Crct}] \geq \frac{1}{2} \cdot \left(1 - \frac{1}{2^k} \right) + \frac{1}{2} - \Pr[\neg \text{Crct}]$$

If ν_1 is the error probability of key verifier VfPK, and ν_2 is the error probability of ciphertext verifier VfCtxt, then $\Pr[\neg \text{Crct}] \leq \nu_1(k) + \nu_2(k)$. Hence

$$\text{Adv}_{\text{AS}, \text{SS}, \mathbf{H}}^{\text{ind-cca}}(k) = 2 \cdot \Pr[\text{Succ}(\mathbf{H})] - 1 \geq 1 - 2^{-k} - 2 \cdot (\nu_1(k) + \nu_2(k)) = 1 - 2^{-k} - \nu(k),$$

where $\nu(k) = 2 \cdot (\nu_1(k) + \nu_2(k))$. \blacksquare

C Any IND-CCA-secure scheme is IND-CCA preserving

We remarked in Section 1.2 that any asymmetric encryption scheme that is IND-CCA secure is also IND-CCA preserving. (The interesting thing about the Hash ElGamal scheme is that it is not IND-CCA secure but is still IND-CCA preserving.) For completeness, we state and prove this formally here. We begin by recalling the definition of IND-CCA security of an asymmetric encryption scheme.

DEFINITION. This follows [4]. Associate to AS, an adversary \mathbf{A} , and $k \in \mathbb{N}$, the following experiment:

Experiment $\text{Exp}_{\text{AS}, \mathbf{A}}^{\text{ind-cca}}(k)$

Randomly choose RO $R_a: \{0, 1\}^* \rightarrow \{0, 1\}$

$(pk, sk) \xleftarrow{\$} \text{AK}^{R_a}(1^k)$; $b \xleftarrow{\$} \{0, 1\}$

Run \mathbf{A} with input $1^k, pk$ and oracles $\text{AE}^{R_a}(pk, \text{LR}(\cdot, \cdot, b))$, $\text{AD}^{R_a}(sk, \cdot)$, R_a

Let d denote the output of \mathbf{A}

If $d = b$ then return 1 else return 0.

We say that adversary \mathbf{A} is legitimate if it never queries $\text{AD}^{R_a}(sk, \cdot)$ with a ciphertext previously returned by $\text{AE}^{R_s}(pk, \text{LR}(\cdot, \cdot, b))$. Asymmetric encryption scheme AS is said to be IND-CCA secure if the function

$$\text{Adv}_{\text{AS}, \mathbf{A}}^{\text{ind-cca}}(k) = 2 \cdot \Pr \left[\mathbf{Exp}_{\text{AS}, \mathbf{A}}^{\text{ind-cca}}(k) = 1 \right] - 1$$

is negligible for all legitimate polynomial-time adversaries \mathbf{A} . IND-CPA security is defined similarly, except the adversary is not given access to oracle $\text{AD}^{R_a}(sk, \cdot)$.

RESULT. The following holds in both the standard and the RO models.

Theorem C.1 Let AS be an IND-CCA-secure asymmetric encryption scheme. Then AS is IND-CCA preserving. \blacksquare

Proof of Theorem C.1: Let $\text{AS} = (\text{AK}, \text{AE}, \text{AD})$ be an IND-CCA-secure asymmetric encryption scheme and let $\text{SS} = (\text{SK}, \text{SE}, \text{SD})$ be an IND-CCA-secure symmetric encryption scheme. We will show that for any polynomial-time legitimate hybrid adversary \mathbf{H} attacking mm-hybrid encryption scheme (AS, SS) there exist polynomial-time legitimate adversaries \mathbf{A} and \mathbf{S} such that for any $k \in \mathbb{N}$

$$\text{Adv}_{\text{AS}, \text{SS}, \mathbf{H}}^{\text{ind-cca}}(k) \leq 2\text{Adv}_{\text{AS}, \mathbf{A}}^{\text{ind-cca}}(k) + \text{Adv}_{\text{SS}, \mathbf{S}}^{\text{ind-cca}}(k). \quad (15)$$

Since AS and SS are assumed IND-CCA secure, the advantage functions related to \mathbf{A} and \mathbf{S} above are negligible, and thus so is the advantage function related to \mathbf{H} . To complete the proof, we need to specify the adversaries \mathbf{A} , \mathbf{S} and prove Equation (15).

We first associate to (AS, SS) , \mathbf{H} , and $k \in \mathbb{N}$, the following experiments, for $i \in \{1, 2, 3, 4\}$:

Experiment $\mathbf{Exp}_{\text{AS}, \text{SS}, \mathbf{H}}^i(k)$

Randomly choose RO $R: \{0, 1\}^* \rightarrow \{0, 1\}$

Define ROs $R_s(\cdot) = R(0\|\cdot)$ and $R_a(\cdot) = R(1\|\cdot)$

$(pk, sk) \xleftarrow{\$} \text{AK}^{R_a}(1^k)$; $K \xleftarrow{\$} \text{SK}^{R_s}(1^k)$; $K' \xleftarrow{\$} \text{SK}^{R_s}(1^k)$

If $i = 1$ or $i = 4$ then $C_a \xleftarrow{\$} \text{AE}^{R_a}(pk, K)$ else $C_a \xleftarrow{\$} \text{AE}^{R_a}(pk, K')$ EndIf

If $i = 1$ or $i = 2$ then run \mathbf{H} with inputs pk, C_a and oracles

$\text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, 0))$, $\text{SD}^{R_s}(K, \cdot)$, $\text{AD}^{R_a}(sk, \cdot)$, R

Else run \mathbf{H} with inputs pk, C_a and oracles

$\text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, 1))$, $\text{SD}^{R_s}(K, \cdot)$, $\text{AD}^{R_a}(sk, \cdot)$, R

EndIf

Let d denote the output of \mathbf{H}

Return d .

For $i \in \{1, 2, 3, 4\}$, let P_i denote the probability that $\mathbf{Exp}_{\text{AS}, \text{SS}, \mathbf{H}}^i(k)$ returns 1. It is easy to see that

$$\text{Adv}_{\text{AS}, \text{SS}, \mathbf{H}}^{\text{ind-cca}}(k) = P_4 - P_1 = (P_4 - P_3) + (P_3 - P_2) + (P_2 - P_1).$$

We will show that there exist legitimate polynomial-time adversaries \mathbf{A}' , \mathbf{S} , and \mathbf{A}'' such that

$$P_4 - P_3 = \text{Adv}_{\text{AS}, \mathbf{A}'}^{\text{ind-cca}}(k), \quad P_3 - P_2 = \text{Adv}_{\text{SS}, \mathbf{S}}^{\text{ind-cca}}(k), \quad \text{and} \quad P_2 - P_1 = \text{Adv}_{\text{AS}, \mathbf{A}''}^{\text{ind-cca}}(k). \quad (16)$$

We obtain Equation (15) from the above by setting $\mathbf{A} = \mathbf{A}'$ if $\text{Adv}_{\text{AS}, \mathbf{A}'}^{\text{ind-cca}}(k) \geq \text{Adv}_{\text{AS}, \mathbf{A}''}^{\text{ind-cca}}(k)$, and $\mathbf{A} = \mathbf{A}''$ otherwise. We now define adversaries \mathbf{A}' , \mathbf{A}'' , \mathbf{S} and prove Equation (16).

DESCRIPTION OF \mathbf{A}' . Adversary \mathbf{A}' is given inputs $1^k, pk$ and has access to oracles $\text{AE}^{R_a}(pk, \text{LR}(\cdot, \cdot, b))$, $\text{AD}^{R_a}(sk, \cdot)$, and R_a . Its goal is to guess the bit b . It begins with the following initializations:

$K \xleftarrow{\$} \{0, 1\}^k$; $K' \xleftarrow{\$} \{0, 1\}^k$; Make query (K', K) to $\text{AE}^{R_a}(pk, \text{LR}(\cdot, \cdot, b))$, and let C_a be the response

Then it runs \mathbf{H} on inputs public key pk and ciphertext C_a . In the process \mathbf{H} will query its oracles

$$R_a, R_s, \text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, b)), \text{SD}^{R_s}(K, \cdot), \text{AD}^{R_a}(sk, \cdot). \quad (17)$$

\mathbf{A}' will answer these queries. Queries to R_s are simulated the standard way, by returning a random value for each new query and the previously returned value for each repeated query. \mathbf{A}' answers queries to the $\text{AD}^{R_a}(sk, \cdot)$ and R_a oracles via its own oracles of the same name. Since it possesses K , it can answer queries to $\text{SD}^{R_s}(K, \cdot)$ by simply performing the computation of the decryption algorithm, replacing calls that the latter makes to R_s by the above-mentioned simulation, and returning the answer. \mathbf{A}' answers queries to $\text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, b))$ by using K to simulate oracle $\text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, 1))$. When \mathbf{H} halts and outputs d , \mathbf{A}' outputs d .

DESCRIPTION OF \mathbf{A}'' . Adversary \mathbf{A}'' is identical to adversary \mathbf{A}' , except that it makes query (K, K') to oracle $\text{AE}^{R_a}(pk, \text{LR}(\cdot, \cdot, b))$ and it answers queries to $\text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, b))$ by using K to simulate oracle $\text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, 0))$.

DESCRIPTION OF \mathbf{S} . Adversary \mathbf{S} is given input 1^k and has access to oracles $\text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, b))$, $\text{SD}^{R_s}(K, \cdot)$, and R_s . Its goal is to guess the bit b . It begins with the following initializations:

$$K' \xleftarrow{\$} \{0, 1\}^k; (pk, sk) \xleftarrow{\$} \text{AK}^{R_a}(1^k); C_a \xleftarrow{\$} \text{AE}_{pk}(K')$$

Then it runs \mathbf{H} on inputs public key pk and ciphertext C_a . In the process \mathbf{H} will query the oracles listed in Equation (17). \mathbf{S} will answer these queries. Queries to R_a are simulated the standard way, by returning a random value for each new query and the previously returned value for each repeated query. \mathbf{S} answers queries to the $\text{SE}^{R_s}(K, \text{LR}(\cdot, \cdot, b))$, $\text{SD}^{R_s}(K, \cdot)$, and R_s oracles via its own oracles of the same name. Since it possesses the secret key sk , it can answer queries to $\text{AD}^{R_a}(sk, \cdot)$ by simply performing the computation of the decryption algorithm, replacing calls that the latter makes to R_a by the above-mentioned simulation, and returning the answer. When \mathbf{H} halts and outputs d , \mathbf{S} outputs d .

ANALYSIS. Clearly, if \mathbf{H} is polynomial-time and legitimate, so are \mathbf{A}' , \mathbf{A}'' , and \mathbf{S} . It is easy to see that Equation (16) holds. ■