An Uninstantiable Random-Oracle-Model Scheme for a Hybrid-Encryption Problem

Mihir Bellare, Alexandra Boldyreva and Adriana Palacio

Dept. 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 uninstantiable, 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. These results extend our understanding 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 can arise in domains where RO schemes are commonly designed.

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 [6], 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. Let us begin by reviewing previous work and then explain our results.

1.1 Previous Work

Let us call a RO-model scheme *uninstantiable*, 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 the scheme meets the goal in question.

Canetti, Goldreich and Halevi [8] provided the first examples of uninstantiable schemes, the goals in question being IND-CPA-secure asymmetric encryption and digital signatures secure against chosen-message attacks. Further examples followed: Nielsen [19] presented an uninstantiable RO-model scheme for the goal of non-interactive, non-committing encryption [7], and Goldwasser and Taumann [17] showed the existence of a 3-move protocol which, when collapsed via a RO as per the Fiat-Shamir heuristic [14], yields an uninstantiable RO-model signature scheme.

The results of [8] indicate that it is possible for the RO paradigm to fail to yield secure "real-world" schemes. The example schemes provided by [8], 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 [18].) The same is true of the scheme of [17]. In contrast, the scheme of [19] 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 uninstantiable 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 the full version of this paper [3].) IND-CCA preservation, however, is actually a weaker requirement on an asymmetric encryption scheme than IND-CCA security itself. In fact, since the messages to be encrypted using the asymmetric scheme are randomly-chosen symmetric keys, the encryption itself need not even be randomized. 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 [9, 10, 15, 20].

1.3 The Hash ElGamal Scheme and its Security

It is easy to see that the ElGamal encryption scheme [13] 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 [15] 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

$$\mathsf{AE}^{G,H}((q, g, X), K) \;=\; (g^{H(K)}\,,\, G(X^{H(K)}) \oplus K)\,, \qquad (1)$$

[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.

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, 4].

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 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.

We follow this with Theorem 2, however, which says that the Hash ElGamal scheme is uninstantiable. 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

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 contrived, but not particularly complex. 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 [15, 2, 4].) 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 uninstantiable thus points to the difficulty of being able to distinguish uninstantiable RO-model schemes from ones that at least *may* be securely instantiable, even in the context of natural and practical goals.

[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.

1.5 Generalizations

In the full version of the paper [3] 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. Note that ciphertext verifiability contradicts IND-CPA security, but it need not prevent a scheme from being IND-CCA preserving, since the latter notion considers the use of the asymmetric scheme only for the encryption of messages that are chosen at random.

In [3] we prove 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. However, as we also prove in [3], 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. This generalizes Theorem 2 because any instantiation of the Hash ElGamal scheme is key-verifiable and ciphertext-verifiable, and hence cannot be IND-CCA-preserving.

In [3] we lift 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.

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 [9, 10, 15, 20, 12, 21] 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 [9, 10, 15, 20] 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, 4] of which Hash ElGamal is a special case, are uninstantiable.

2 Definitions

NOTATION AND CONVENTIONS. If S is a randomized algorithm, then $[S(x, y, \ldots)]$ denotes the set of all points having positive probability of being output by S on inputs $x, y, \ldots$. 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 [6].

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

We define security following [5] and addressing the possibility of the symmetric scheme being in the RO model. Let $\mathrm{LR}(M_0, M_1, b) = M_b$ if M_0, M_1 are strings of equal length, and $\perp$ otherwise. Associate to SS, an adversary $\boldsymbol{S}$, and $k \in \mathbb{N}$, the following experiment.

Experiment $\mathbf{Exp}^{\text{ind-cca}}_{\mathsf{SS},\boldsymbol{S}}(k)$
  Randomly choose RO $R_s\colon \{0,1\}^* \to \{0,1\}$
  $K \xleftarrow{\$} \mathsf{SK}^{R_s}(1^k)$; $b \xleftarrow{\$} \{0,1\}$
  Run $\boldsymbol{S}$ with input 1^k and oracles $\mathsf{SE}^{R_s}(K, \mathrm{LR}(\cdot,\cdot,b))$, $\mathsf{SD}^{R_s}(K,\cdot)$, R_s
  Let d denote the output of $\boldsymbol{S}$
  If $d = b$ then return 1 else return 0.

We say that adversary $\boldsymbol{S}$ is legitimate if it never queries $\mathsf{SD}^{R_s}(K,\cdot)$ with a ciphertext previously returned by $\mathsf{SE}^{R_s}(K, \mathrm{LR}(\cdot,\cdot,b))$. Symmetric encryption scheme SS is said to be IND-CCA secure if the function

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

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

ASYMMETRIC ENCRYPTION. An asymmetric encryption scheme $\mathsf{AS} = (\mathsf{AK}, \mathsf{AE}, \mathsf{AD})$ is specified by three polynomial-time algorithms: via $(pk, sk) \xleftarrow{\$} \mathsf{AK}(1^k)$ one can generate keys; via $C \xleftarrow{\$} \mathsf{AE}(pk, K)$ one can encrypt a message $K \in \{0,1\}^k$; and via $K \leftarrow \mathsf{AD}(sk, C)$ 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 $\mathsf{AD}(sk, \mathsf{AE}(pk, K)) = K$ for all $(pk, sk) \in [\mathsf{AK}(1^k)]$ 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 [3].

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 $(\mathsf{AS}, \mathsf{SS})$ consisting of an asymmetric encryption scheme $\mathsf{AS} = (\mathsf{AK}, \mathsf{AE}, \mathsf{AD})$ and a symmetric encryption scheme $\mathsf{SS} = (\mathsf{SK}, \mathsf{SE}, \mathsf{SD})$. We associate to $(\mathsf{AS}, \mathsf{SS})$, a *hybrid adversary* $\boldsymbol{H}$, and $k \in \mathbb{N}$, the following experiment.

Experiment $\mathbf{Exp}^{\text{ind-cca}}_{\mathsf{AS},\mathsf{SS},\boldsymbol{H}}(k)$
  Randomly choose RO R: $\{0,1\}^* \to \{0,1\}$
  Define ROs $R_s(\cdot) = R(0\|\cdot)$ and $R_a(\cdot) = R(1\|\cdot)$
  $(pk, sk) \xleftarrow{\$} \mathsf{AK}^{R_a}(1^k)$; $K \xleftarrow{\$} \mathsf{SK}^{R_s}(1^k)$; $b \xleftarrow{\$} \{0,1\}$
  $C_a \xleftarrow{\$} \mathsf{AE}^{R_a}(pk, K)$
  Run $\boldsymbol{H}$ with inputs pk, C_a and
  oracles $\mathsf{SE}^{R_s}(K, \mathrm{LR}(\cdot,\cdot,b))$, $\mathsf{SD}^{R_s}(K,\cdot)$, $\mathsf{AD}^{R_a}(sk,\cdot)$, R
  Let d denote the output of $\boldsymbol{H}$
  If $d = b$ then return 1 else return 0.

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

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

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

Finally, we say that an asymmetric encryption scheme AS is *IND-CCA preserving* if the mm-hybrid encryption scheme $(\mathsf{AS}, \mathsf{SS})$ 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 is secure.

Remark 1. Suppose we have two RO-model schemes, and are composing them, or executing them in a common context. (Above, this is happening with the asymmetric encryption scheme and the symmetric encryption scheme.) We claim that, in this case, the ROs of the two schemes should be chosen independently. (This does not mean that we need to assume two RO oracles are given. The formal model always provides just one RO. But one can easily derive several independent ROs from a single one, as we did above.) The correctness of this principle of independent instantiation of ROs in a common context can be seen in many ways. First, it is easy to come up with an example of a pair of secure RO-model schemes that, when composed, yield an insecure one if the ROs in the two schemes are defined to be the same. Second, one can reason by analogy with the way we need to choose keys in composing primitives. For example, suppose we have a MAC and symmetric encryption scheme, each individually secure. If we use them to construct an authenticated-encryption scheme, we should use different keys for the MAC and the symmetric encryption scheme. (There is no reason to think otherwise that the composition will be secure.) The principle, for ROs, is exactly the same. They are just like keys provided to primitives.

The existence of IND-CCA-preserving asymmetric encryption schemes is easy to establish since, as we show in [3], 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 [13] 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

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

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

is said to be hard for CG if the function

$$\mathsf{Adv}^{\mathrm{cdh}}_{\mathsf{CG},\boldsymbol{C}}(k) \;=\; \Pr\left[\,(q,g) \xleftarrow{\$} \mathsf{CG}(1^k)\,;\; x,y \xleftarrow{\$} \mathbb{Z}_q \;:\; \boldsymbol{C}(q,g,g^x,g^y) = g^{xy}\,\right]$$

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

Scheme and result statement. To any cyclic-group generator CG we associate the RO-model asymmetric encryption scheme $\mathsf{HEG}[\mathsf{CG}] = (\mathsf{AK},\mathsf{AE},\mathsf{AD})$ whose constituent algorithms are depicted in Figure 1. (The scheme makes reference to two ROs, namely $G\colon \langle g\rangle \to \{0,1\}^k$ and $H\colon \{0,1\}^k \to \mathbb{Z}_q$, while the formal definition of an asymmetric encryption scheme provides a single RO $R\colon \{0,1\}^* \to \{0,1\}$, but G,H may be implemented via R in standard ways [6].) 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 1. *If the CDH problem is hard for cyclic-group generator* CG*, then the associated Hash ElGamal asymmetric encryption scheme* $\mathsf{HEG}[\mathsf{CG}]$ *is IND-CCA preserving in the RO model.*

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 $\mathsf{HEG}[\mathsf{CG}]$ is deterministic. For this reason alone, $\mathsf{HEG}[\mathsf{CG}]$ is not an IND-CCA secure, or even IND-CPA secure, asymmetric encryption scheme. Nonetheless, Theorem 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 $\mathsf{HEG}[\mathsf{CG}]$ 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 [9, 10, 15, 20] 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 $\mathsf{HEG}[\mathsf{CG}]$ 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 1.

The full proof of Theorem 1 is in [3]. Below we provide an intuitive overview that highlights the main areas of novelty.

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

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

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

Since SS is assumed IND-CCA secure and the CDH problem is hard for CG, the advantage functions related to $\boldsymbol{S}$ and $\boldsymbol{C}$ above are negligible, and thus so is the advantage function related to $\boldsymbol{H}$. To complete the proof, we need to specify adversaries $\boldsymbol{S}, \boldsymbol{C}$ for which Equation (2) is true.

Consider $\mathbf{Exp}^{\text{ind-cca}}_{\mathsf{AS},\mathsf{SS},\boldsymbol{H}}(k)$. Let (q, g, X) be the public key and (q, g, x) the secret key chosen, where $X = g^x$. Let $C_a = (Y, W)$ where $Y = g^y$. Let K denote the symmetric encryption key chosen. 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 $\mathsf{Succ}(\boldsymbol{H})$ the event that $\boldsymbol{H}$ is successful at guessing the value of its challenge bit b. We will construct $\boldsymbol{C}$ so that

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

and we will construct $\boldsymbol{S}$ so that

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

Equation (2) follows.

THE ADVERSARIES. The design of $\boldsymbol{C}$ relies mostly on standard techniques, and so we leave it to [3]. We turn to $\boldsymbol{S}$. The latter gets input 1^k and oracles $\mathsf{SE}^{R_s}(K, \text{LR}(\cdot,\cdot,b))$, $\mathsf{SD}^{R_s}(K,\cdot)$, R_s, and begins with the initializations

$$\begin{aligned} &((q, g, X), (q, g, x)) \xleftarrow{\$} \mathsf{AK}(1^k)\,; \\ &y \xleftarrow{\$} \mathbb{Z}_q\,;\; Y \leftarrow g^y\,;\; W \xleftarrow{\$} \{0,1\}^k\,;\; C_a \leftarrow (Y, W)\,. \end{aligned} \quad (4)$$

It then runs $\boldsymbol{H}$ 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 $\boldsymbol{S}$'s oracles plays the role of the quantity of the same name for $\boldsymbol{H}$. Eventually, it will output what $\boldsymbol{H}$ outputs. The difficulty faced by this adversary is that $\boldsymbol{H}$ might query K to H. (Other oracle queries are dealt with in standard ways.) In that case, $\boldsymbol{H}$ 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

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

A precise description and analysis of $\mathsf{KeyTest}$ are in [3], but we briefly sketch the ideas here. The algorithm has two phases. In the first phase, it repeatedly tests whether or not

$$\begin{aligned} \mathsf{SD}^{R_s}(L, \mathsf{SE}^{R_s}(K, \mathrm{LR}(T_0, T_0, b))) &= T_0 \text{ and} \\ \mathsf{SD}^{R_s}(L, \mathsf{SE}^{R_s}(K, \mathrm{LR}(T_1, T_1, b))) &= T_1 \ , \end{aligned}$$

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 $\mathsf{SD}^{R_s}(L, \mathsf{SE}^{R_s}(K, \mathrm{LR}(T_0, T_1, b)))$ and, if *all* these computations yield T_{gs} for some bit gs, it returns $(1, \mathrm{gs})$. 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 $\boldsymbol{H}$ is making queries to $\mathsf{SD}^{R_s}(K, \cdot)$. $\boldsymbol{S}$ will answer these via its own oracle of the same name. Now, consider the event that $\boldsymbol{H}$ queries to $\mathsf{SD}^{R_s}(K, \cdot)$ a ciphertext C generated in some execution of $\mathsf{KeyTest}$. If $\boldsymbol{S}$ calls $\mathsf{SD}^{R_s}(K, C)$ 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 $\mathsf{SE}^{R_s}(K, \mathrm{LR}(\cdot,\cdot,b))$ made by $\boldsymbol{S}$ via subroutine $\mathsf{KeyTest}$. There are a few ways around this, and the one we use is to choose the initial "test" messages randomly so that $\boldsymbol{H}$ has low probability of being able to query a ciphertext C generated in some execution of $\mathsf{KeyTest}$.

[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.

$\overline{\mathsf{AK}}(1^k)$	$\overline{\mathsf{AE}}(\overline{pk}, K)$	$\overline{\mathsf{AD}}(\overline{sk}, C)$
$fk \xleftarrow{\$} \{0,1\}^{\mathrm{FKL}(k)}$	Parse $\overline{pk}$ as (pk, fk)	Parse $\overline{sk}$ as (sk, fk)
$(pk, sk) \xleftarrow{\$} \mathsf{AK}^{\overline{F}^k(fk,\cdot)}(1^k)$	$C \xleftarrow{\$} \mathsf{AE}^{\overline{F}^k(fk,\cdot)}(pk, K)$	$K \leftarrow \mathsf{AD}^{\overline{F}^k(fk,\cdot)}(sk, C)$
Return $((pk, fk), (sk, fk))$	Return C	Return K

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

We note that one might consider an alternative solution to $\boldsymbol{S}$'s problem of wanting to "detect" query K to H. Namely, reply to queries to H at random, then, after $\boldsymbol{H}$ 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 $1/\operatorname{poly}(k)$, the advantage over one-half obtained by $\boldsymbol{S}$ 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.

We provide all the details and justify Equation (2) in [3].

4 Uninstantiability of the Hash ElGamal scheme

In this section we show (cf. Theorem 2) 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 $fk \in \{0,1\}^{\mathrm{FKL}(k)}$ a map $\overline{F}^k(fk, \cdot)\colon \{0,1\}^* \to \{0,1\}$. The *key length* FKL of the family of functions is a polynomial in k. We require that there exist a polynomial t such that $\overline{F}^k(fk, x)$ is computable in $t(k + |x|)$ time for all $k \in \mathbb{N}$, $fk \in \{0,1\}^{\mathrm{FKL}(k)}$ and $x \in \{0,1\}^*$.

An *instantiation* of a RO-model asymmetric encryption scheme $\mathsf{AS} = (\mathsf{AK}, \mathsf{AE}, \mathsf{AD})$ via family $\overline{F}$ is the standard-model asymmetric encryption scheme $\overline{\mathsf{AS}} = (\overline{\mathsf{AK}}, \overline{\mathsf{AE}}, \overline{\mathsf{AD}})$ 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 $\overline{F}^k(fk, \cdot)$, and calls to the random oracle are then replaced by evaluations of this function in all algorithms.

The uninstantiability result. The formal statement of the result is the following.

Theorem 2. *Let* $\mathsf{HEG}[\mathsf{CG}] = (\mathsf{AK}, \mathsf{AE}, \mathsf{AD})$ *be the RO-model Hash ElGamal scheme associated to a cyclic-group generator* CG. *Let* $\overline{\mathsf{HEG}}[\mathsf{CG}] = (\overline{\mathsf{AK}}, \overline{\mathsf{AE}}, \overline{\mathsf{AD}})$ *be any instantiation of* $\mathsf{HEG}[\mathsf{CG}]$ *via a poly-time family of functions. Then* $\overline{\mathsf{HEG}}[\mathsf{CG}]$ *is not IND-CCA preserving.*

Proof of Theorem 2. Let $\overline{F}$ be the poly-time family of functions used in $\overline{\mathsf{HEG}}[\mathsf{CG}]$ to replace the random oracle. We will construct an IND-CCA-secure symmetric encryption scheme SS such that the mm-hybrid encryption scheme $(\overline{\mathsf{HEG}}[\mathsf{CG}], \mathsf{SS})$ is not IND-CCA secure. This proves the theorem.

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

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

Let $\mathsf{SS}' = (\mathsf{SK}', \mathsf{SE}', \mathsf{SD}')$ 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 $\mathsf{VfCtxt}_{\overline{F}}$. We use the notation $\langle(\cdot,\cdot)\rangle$ to denote an injective, polynomial-time computable encoding of pairs of strings as strings such that given $\langle(M_1, M_2)\rangle$, M_1 and M_2 can be recovered in polynomial time. If s is a string and $a \le b$ are integers then $s[a \ldots b]$ denotes the string consisting of bit positions a through b of s. The algorithms constituting $\mathsf{SS} = (\mathsf{SK}, \mathsf{SE}, \mathsf{SD})$ are depicted in Figure 3. To conclude the proof, we need only establish the following propositions.

Proposition 1. *Symmetric encryption scheme* SS *is IND-CCA secure.*

Proposition 2. *Multi-message hybrid encryption scheme* $(\overline{\mathsf{HEG}}[\mathsf{CG}], \mathsf{SS})$ *is not IND-CCA secure.*

Proof (Proposition 1). Let us first provide some intuition. Note that on input M, encryption algorithm $\mathsf{SE}(K_1' || K_2, \cdot)$ uses the encryption algorithm SE' of an

```
SK(1^k)
  K' <$- SK'(1^⌈k/2⌉)
  K_2 <$- {0,1}^⌊k/2⌋
  Return K'||K_2
```

```
SE(K, M)
  k ← |K|
  K' ← K[1 ... ⌈k/2⌉]
  K_2 ← K[1 + ⌈k/2⌉ ... k]
  C' ← SE'(K', M)
  Parse M as ⟨(M_1, M_2)⟩
  If the parsing fails then
      Return C'||1 EndIf
  p ← VfPK(1^k, M_1)
  c ← VfCtxt_F̄(1^k, M_1, K, M_2)
  If (p = 1 and c = 1) then
      Return C'||0
  else Return C'||1 EndIf
```

```
SD(K, C)
  k ← |K|
  K' ← K[1 ... ⌈k/2⌉]
  K_2 ← K[1 + ⌈k/2⌉ ... k]
  Parse C as C'||d,
      where d ∈ {0,1}
  M' ← SD'(K', C')
  Parse M' as ⟨(M_1, M_2)⟩
  If the parsing fails then
      If d = 1 then Return M'
      else Return ⊥ EndIf
  p ← VfPK(1^k, M_1)
  c ← VfCtxt_F̄(1^k, M_1, K, M_2)
  If (d = 0 and p = 1 and c = 1)
  then Return M' EndIf
  If (d = 1 and (p ≠ 1 or c ≠ 1))
  then Return M' EndIf
  Return ⊥
```

Fig. 3. Algorithms of the symmetric encryption scheme $\mathsf{SS} = (\mathsf{SK}, \mathsf{SE}, \mathsf{SD})$ for the proof of Theorem 2. Above, $\langle (M_1, M_2) \rangle$ denotes an encoding of the pair of strings (M_1, M_2) as a string.

IND-CCA-secure scheme to compute $C' \xleftarrow{\$} \mathsf{SE}'(K_1', M)$ 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 $\mathsf{VfPK}, \mathsf{VfCtxt}_{\overline{F}}$ indicate that M_1 is a $(\overline{\mathsf{HEG}}[\mathsf{CG}], k)$-valid public key and $M_2 \in [\overline{\mathsf{AE}}(M_1, K_1'||K_2)]$. The decryption algorithm $\mathsf{SD}(K_1'||K_2, \cdot)$ on input $C'||d$, where d is a bit, computes $M' \leftarrow \mathsf{SD}'(K_1', C')$ 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 $\mathsf{SE}(K, \mathrm{LR}(\cdot, \cdot, b))$ 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 $\overline{\mathsf{AE}}$ 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 $\boldsymbol{S}$ be a legitimate polynomial-time adversary attacking SS. We will construct a legitimate polynomial-time adversary $\boldsymbol{S}'$ such that

$$\mathsf{Adv}^{\text{ind-cca}}_{\mathsf{SS},\boldsymbol{S}}(k) \;\leq\; \mathsf{Adv}^{\text{ind-cca}}_{\mathsf{SS}',\boldsymbol{S}'}(\lceil k/2 \rceil) + \frac{O(Q(k))}{2^{\lfloor k/2 \rfloor}}\,, \tag{5}$$

where Q is a polynomial upper bounding the total number of queries made by $\boldsymbol{S}$ to its different oracles. Since SS' is assumed IND-CCA secure, the advantage function associated to $\boldsymbol{S}'$ above is negligible, and thus so is the advantage func-

tion associated to $\boldsymbol{S}$. To complete the proof, we need to specify adversary $\boldsymbol{S}'$ and prove Equation (5).

Adversary $\boldsymbol{S}'$ is given input $1^{\lceil k/2 \rceil}$ and has access to oracles $\mathsf{SE}'(K_1', \mathrm{LR}(\cdot,\cdot,b))$ and $\mathsf{SD}'(K_1', \cdot)$. Its goal is to guess the bit b. It runs $\boldsymbol{S}$ on input 1^k. In this process, $\boldsymbol{S}$ will query its two oracles $\mathsf{SE}(K, \mathrm{LR}(\cdot,\cdot,b))$ and $\mathsf{SD}(K,\cdot)$. To answer a query to the first of these oracles, $\boldsymbol{S}'$ forwards the query to its oracle $\mathsf{SE}'(K_1', \mathrm{LR}(\cdot,\cdot,b))$, appends 1 to the oracle's reply and returns the result to $\boldsymbol{S}$. To answer a query to the second oracle, $\boldsymbol{S}'$ checks the last bit of the query. If it is 0, $\boldsymbol{S}'$ returns $\perp$ to $\boldsymbol{S}$. Otherwise, it removes the last bit, forwards the result to its oracle $\mathsf{SD}'(K_1', \cdot)$, and returns the answer to $\boldsymbol{S}$. When $\boldsymbol{S}$ outputs its guess b', $\boldsymbol{S}'$ returns b'.

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

$\mathsf{Succ}(\boldsymbol{S}')$	:	$\boldsymbol{S}'$ is successful, meaning its output equals the challenge bit b
BadE	:	$\boldsymbol{S}$ makes a query to oracle $\mathsf{SE}(K, \mathrm{LR}(\cdot,\cdot,b))$ in which one of the messages can be parsed as $\langle (M_1, M_2) \rangle$ such that M_1 is a $(\overline{\mathsf{HEG}}[\mathsf{CG}], k)$-valid public key and $M_2 \in [\overline{\mathsf{AE}}(M_1, K)]$
BadD	:	$\boldsymbol{S}$ makes a query to oracle $\mathsf{SD}(K,\cdot)$ that can be parsed as $C'\|\|d$, where d is a bit, such that $\mathsf{SD}'(K_1', C') = \langle (M_1, M_2) \rangle$, where M_1 is a $(\overline{\mathsf{HEG}}[\mathsf{CG}], k)$-valid public key and $M_2 \in [\overline{\mathsf{AE}}(M_1, K)]$

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

$\mathsf{Succ}(\boldsymbol{S})$	:	$\boldsymbol{S}$ is successful, meaning its output equals the challenge bit b

We claim that if events BadE and BadD do not occur, then $\boldsymbol{S}'$ simulates perfectly the environment provided to $\boldsymbol{S}$ in its attack against SS. First, note that answers to queries to oracle $\mathsf{SE}(K, \mathrm{LR}(\cdot,\cdot,b))$ can only be off by the last bit. In the absence of the "bad" events, each ciphertext returned to $\boldsymbol{S}$ as a reply to a query to oracle $\mathsf{SE}(K, \mathrm{LR}(\cdot,\cdot,b))$ has 1 as the last bit. This is also the case in $\boldsymbol{S}$'s real attack. If $\boldsymbol{S}$ queries $\mathsf{SD}(K,\cdot)$ with a ciphertext $C'||0$, assuming events BadE and BadD do not occur, $\boldsymbol{S}'$ gives $\boldsymbol{S}$ the response it would get in the real attack, namely $\perp$. Since $\boldsymbol{S}$ is legitimate, if it queries oracle $\mathsf{SD}(K,\cdot)$ with a ciphertext $C'||1$, then C' must not have previously been returned by oracle $\mathsf{SE}'(K_1', \mathrm{LR}(\cdot,\cdot,b))$. Thus $\boldsymbol{S}'$ can legitimately make query C' to its oracle $\mathsf{SD}'(K_1', \cdot)$. If M is the response, then, assuming that events BadE and BadD do not occur, the answer $\boldsymbol{S}$ expects is exactly M. Therefore,

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

We now provide an upper bound for the probability of event $\mathsf{BadE} \vee \mathsf{BadD}$. Let $q_e(k)$ and $q_d(k)$ be the number of queries $\boldsymbol{S}$ makes to oracles $\mathsf{SE}(K, \mathrm{LR}(\cdot,\cdot,b))$ and $\mathsf{SD}(K,\cdot)$, respectively, on input 1^k. We observe that if M_1 is a $(\overline{\mathsf{HEG}}[\mathsf{CG}], k)$-valid public key, then for any $M_2 \in \{0,1\}^*$, there exists a unique $K' \in [\mathsf{SK}(1^k)]$ such that $M_2 \in [\overline{\mathsf{AE}}(M_1, K')]$. Recall that the key for oracles $\mathsf{SE}(K, \mathrm{LR}(\cdot,\cdot,b))$ and $\mathsf{SD}(K,\cdot)$ 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 $\boldsymbol{S}$'s view. Therefore, for any query made by $\boldsymbol{S}$ to oracle $\mathsf{SE}(K, \mathrm{LR}(\cdot,\cdot,b))$, the probability that one of the messages in the query parses as $\langle (M_1, M_2) \rangle$ such that M_1 is a $(\overline{\mathsf{HEG}}[\mathsf{CG}], k)$-valid public key and $M_2 \in [\overline{\mathsf{AE}}(M_1, K)]$ is at most $2/2^{\lfloor k/2 \rfloor}$. Similarly, for any query $C'||d$, where d is a bit, made by $\boldsymbol{S}$ to oracle $\mathsf{SD}(K,\cdot)$, the probability that $\mathsf{SD}'(K_1', C') = M'$, where M' parses as $\langle (M_1, M_2) \rangle$, M_1 is a $(\overline{\mathsf{HEG}}[\mathsf{CG}], k)$-valid public key and $M_2 \in [\overline{\mathsf{AE}}(M_1, K)]$ is at most $1/2^{\lfloor k/2 \rfloor}$. Therefore,

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

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

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

Rearranging terms gives Equation (5).

Proof (Proposition 2). We define a hybrid adversary $\boldsymbol{H}$ attacking $(\overline{\mathsf{HEG}}[\mathsf{CG}], \mathsf{SS})$. $\boldsymbol{H}$ is given inputs $\overline{pk} = ((q, g, X), fk)$ and C_a and has access to oracles $\mathsf{SE}(K, \mathrm{LR}(\cdot,\cdot,b))$, $\mathsf{SD}(K,\cdot)$, and $\overline{\mathsf{AD}}(\overline{sk},\cdot)$, where $\overline{sk} = ((q, g, x), fk)$. Its goal is to guess the challenge bit b. By the definition of experiment $\mathbf{Exp}^{\text{ind-cca}}_{\overline{\mathsf{HEG}}[\mathsf{CG}],\mathsf{SS},\boldsymbol{H}}(k)$, $\overline{pk}$ is a $(\overline{\mathsf{HEG}}[\mathsf{CG}], k)$-valid public key and $C_a \in [\overline{\mathsf{AE}}(\overline{pk}, K)]$. Therefore, $\langle (\overline{pk}, C_a) \rangle$ is a message which, when encrypted with $\mathsf{SE}(K,\cdot)$, 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 $K = \overline{\mathsf{AD}}(\overline{sk}, C)$ is 0 (since $\overline{\mathsf{AE}}(\overline{pk}, K) = C_a$ and $\overline{\mathsf{AE}}$ is deterministic), i.e., the probability that $C \in [\overline{\mathsf{AE}}(\overline{pk}, K)]$ is 0. Hence $\langle (\overline{pk}, C) \rangle$ is a message which, when encrypted with $\mathsf{SE}(K,\cdot)$, yields a ciphertext that has last bit 1. (If $C \notin [\overline{\mathsf{AE}}(\overline{pk}, K)]$, then the last bit will be 1.) Thus, adversary $\boldsymbol{H}$ 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 $\boldsymbol{H}$ chooses C at random from $\{0,1\}^{|C_a|} \setminus \{C_a\}$, makes a query $\langle (\overline{pk}, C_a) \rangle, \langle (\overline{pk}, C) \rangle$ to oracle $\mathsf{SE}(K, \mathrm{LR}(\cdot,\cdot,b))$, parses the response as $C'||d$, where d is a bit, and returns d. The running time of $\boldsymbol{H}$ is clearly polynomial in k. We claim that $\mathsf{Adv}^{\text{ind-cca}}_{\overline{\mathsf{HEG}}[\mathsf{CG}],\mathsf{SS},\boldsymbol{H}}(k) = 1$. To prove this, we consider the event

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

If challenge bit b is 0, then the response to $\boldsymbol{H}$'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

$$\Pr[\,\mathsf{Succ}(\boldsymbol{H})\,] \;=\; \frac{1}{2} + \frac{1}{2} \;=\; 1\,.$$

Hence

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

as desired.

Notice that the adversary constructed in the proof of Proposition 2 does not make any queries to its oracles $\mathsf{SD}(K, \cdot)$ and $\overline{\mathsf{AD}}(\overline{sk}, \cdot)$.

Remark 2. An interesting question at this point may be why the proof of Theorem 2 fails for the RO-model Hash ElGamal scheme $\mathsf{HEG}[\mathsf{CG}]$ associated to a cyclic-group generator CG —it must, since otherwise Theorem 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 $\mathsf{VfCtxt}_{\overline{F}}$ for the asymmetric encryption scheme in question. In the case of the RO-model scheme $\mathsf{HEG}[\mathsf{CG}]$, 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. This is explained in Remark 1.

As we discussed in Section 1, in [3] we provide a more general impossibility result.

5 Acknowledgements

We thank the anonymous referees for their comments. Mihir Bellare and Alexandra Boldyreva were supported in part by NSF grants CCR-0098123, ANR-0129617 and CCR-0208842, and an IBM Faculty Partnership Development Award. Adriana Palacio is supported by an NSF Graduate Research Fellowship.

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*, LNCS Vol. 1841, E. Dawson, A. Clark and C. Boyd ed., Springer-Verlag, 2000.
3. M. Bellare, A. Boldyreva and A. Palacio, "An Uninstantiable Random-Oracle-Model Scheme for a Hybrid-Encryption Problem," Full version of this paper. Available at `http://www-cse.ucsd.edu/users/mihir/`.
4. 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.

5. 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.
6. 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.
7. R. Canetti, U. Feige, O. Goldreich and M. Naor, "Adaptively secure multiparty computation," *Proceedings of the 28th Annual Symposium on the Theory of Computing*, ACM, 1996.
8. R. Canetti, O. Goldreich, S. Halevi, "The random oracle methodology, revisited," *Proceedings of the 30th Annual Symposium on the Theory of Computing*, ACM, 1998.
9. 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 2002*, LNCS Vol. 2271 , B. Preneel ed., Springer-Verlag, 2002.
10. 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 2002*, LNCS Vol. 1431, D. Naccache and P. Paillier eds., Springer-Verlag, 2002.
11. R. Cramer and V. Shoup, "A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack," *Advances in Cryptology – CRYPTO 1998*, LNCS Vol. 1462, H. Krawczyk ed., Springer-Verlag, 1998.
12. 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/`.
13. T. ElGamal, "A public key cryptosystem and signature scheme based on discrete logarithms," *IEEE Transactions on Information Theory*, Vol. 31, 1985.
14. A. Fiat and A. Shamir, "How to prove yourself: practical solutions to identification and signature problems," *Advances in Cryptology – CRYPTO 1986*, LNCS Vol. 263, A. Odlyzko ed., Springer-Verlag, 1986.
15. E. Fujisaki, T. Okamoto, "Secure Integration of Asymmetric and Symmetric Encryption Schemes," *Advances in Cryptology – CRYPTO 1999*, LNCS Vol. 1666, M. Wiener ed., Springer-Verlag, 1999.
16. S. Goldwasser and S. Micali, "Probabilistic encryption," *Journal of Computer and System Science*, Vol. 28, 1984, pp. 270–299.
17. 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.
18. S. Micali, "Computationally sound proofs," *SIAM Journal on Computing*, Vol. 30, No. 4, 2000, pp. 1253-1298.
19. J. B. Nielsen "Separating Random Oracle Proofs from Complexity Theoretic Proofs: The Non-committing Encryption Case," *Advances in Cryptology – CRYPTO 2002*, LNCS Vol. 2442 , M. Yung ed., Springer-Verlag, 2002.
20. T. Okamoto and D. Pointcheval "REACT: Rapid Enhanced-security Asymmetric Cryptosystem Transform," *Topics in Cryptology – CT-RSA 2001*, LNCS Vol. 2020, D. Naccache ed., Springer-Verlag, 2001.
21. V. Shoup, "A proposal for an ISO standard for public key encryption", IACR ePrint archive, record 2001/112, 2001, `http://eprint.iacr.org/`.