Standard Security Does Not Imply Security against Selective-Opening

Mihir Bellare[1], Rafael Dowsley[1], Brent Waters[2], and Scott Yilek[3]

[1] Department of Computer Science & Engineering, University of California San Diego
http://cseweb.ucsd.edu/~mihir, http://cseweb.ucsd.edu/~rdowsley
[2] Department of Computer Science, University of Texas at Austin
http://www.cs.utexas.edu/~bwaters
[3] Department of Computer and Information Sciences, University of St. Thomas
http://personal.stthomas.edu/yile5901

Abstract. We show that no commitment scheme that is hiding and binding according to the standard definition is semantically-secure under selective opening attack (SOA), resolving a long-standing and fundamental open question about the power of SOAs. We also obtain the first examples of IND-CPA encryption schemes that are not secure under SOA, both for sender corruptions where encryption coins are revealed and receiver corruptions where decryption keys are revealed. These results assume only the existence of collision-resistant hash functions.

1 Introduction

A *commitment scheme* $\mathcal{E}$ can be applied to a message m and coins r to (deterministically) produce a commitment $c \leftarrow \mathcal{E}(m; r)$ that is sent to a receiver. The sender can later "open" the commitment by providing m, r and the receiver checks that $\mathcal{E}(m; r) = c$. The first security requirement, often called hiding, is formalized as IND-CPA, namely an adversary knowing m_0, m_1 and $\mathcal{E}(m_b; r)$ for random b, r has negligible advantage in computing challenge bit b. The second requirement, binding, asks that it be hard for an adversary to produce r_0, r_1 and *distinct* m_0, m_1 such that $\mathcal{E}(m_0; r_0) = \mathcal{E}(m_1; r_1) \neq \perp$. Let us refer to a commitment scheme as HB-secure (Hiding and Binding) if it satisfies both these properties. HB-security is the standard requirement and HB-secure commitment schemes are a fundamental tool in cryptography in general and in protocol design in particular. HB-secure commitment implies PRGs [31], PRFs [21] and ZK proofs for **NP** [24].

Suppose there are n committers, the i-th computing its commitment $\mathbf{c}[i] \leftarrow \mathcal{E}(\mathbf{m}[i]; \mathbf{r}[i])$ to its message $\mathbf{m}[i]$ using coins $\mathbf{r}[i]$, the coins of different committers being of course not only random but also independent of each other. An adversary computes, as a function of the vector $\mathbf{c}$ of commitments, a subset $I \subseteq \{1, \ldots, n\}$ of the senders, and obtains the corresponding openings, namely $\langle \mathbf{m}[i] : i \in I \rangle$ and $\langle \mathbf{r}[i] : i \in I \rangle$. This is called a selective opening attack (SOA). We say that $\mathcal{E}$ is SOA-secure if privacy of the un-opened messages is preserved, meaning the adversary, after its SOA, cannot learn anything about $\langle \mathbf{m}[i] : i \notin I \rangle$

D. Pointcheval and T. Johansson (Eds.): EUROCRYPT 2012, LNCS 7237, pp. 645–662, 2012.

other than it would from possession of $\langle \mathbf{m}[i] : i \in I \rangle$. (That is, the coins are unhelpful.) SOAs arise quite naturally in multi-party cryptographic protocols and SOA-security is desirable in many such settings.

A fundamental question that was posed in this area is whether (standard) HB-security implies SOA-security, meaning, is a HB-secure commitment scheme also SOA-secure? So far, the question has received neither a positive nor a negative answer. Intuitively, the answer would appear to be "yes," for how could the coins accompanying the opened messages help, beyond the opened messages themselves, in revealing something about the un-opened messages? Yet attempts to prove SOA-security of a commitment scheme based on its HB-security have failed. But attempts to find a counter-example have also failed. We do not have a single example, even artificial, of a HB-secure commitment scheme that is demonstrably not SOA-secure. This situation has vexed and intrigued cryptographers for many years and been the subject or inspiration for much work [12,19,13,35,3,20,29,7,28].

This paper answers this long-standing open question. We show that the answer is negative. We give an example of a HB-secure commitment scheme which we prove is not SOA-secure. In fact our result is much stronger. It shows that *no* HB-secure commitment scheme is SOA-secure. Given any HB-secure commitment scheme, we present an attack showing it is not SOA-secure. Before going on to our results on encryption let us expand on this result on commitment including its implications and its relation to previous work.

SOA-secure commitment. Dwork, Naor, Reingold and Stockmeyer (DNRS) [19] gave a definition of SOA-secure commitments, henceforth referred to as SS-SOA, that captures semantic security for relations via a simulation-based formalization. Suitable for applications and widely accepted as the right definition, SS-SOA is what we use in our results. We show that no HB-secure commitment scheme is SS-SOA-secure by presenting, for any given HB-secure commitment scheme $\mathcal{E}$, an adversary for which we prove that there is no successful simulator. We do *not* assume the simulation is blackbox. The only assumption made is the existence of a collision-resistant (CR) hash function.

This general result rules out SS-SOA security for particular schemes. For example, a widely employed way to commit to $m \in \mathbb{Z}_p$ is by picking $r \in \mathbb{Z}_p$ at random and returning $\mathcal{E}(m; r) = g^m h^r \in \mathbb{G}$ where g, h are generators of a group $\mathbb{G}$ of prime order p [36]. This scheme is binding if the DL problem is hard in $\mathbb{G}$ and it is unconditionally hiding. Our results imply that it is not SS-SOA secure. They yield a specific attack, in the form of an adversary for which there is no simulator. Since CR hash functions exist if DL is hard, one does not even need extra assumptions. We stress that this is just an example; our result rules out SS-SOA security for *all* HB-secure schemes.

Implications for IND-SOA-CRS. An indistinguishability-based definition of SOA-secure commitment is given in [3,29]. It only applies when the message vector $\mathbf{m}$ is drawn from what's called a "conditionally re-samplable (CRS) distribution," and accordingly we denote it IND-SOA-CRS. This definition is of limited

use in applications because message distributions there are often not CRS, but for CRS distributions the definition is intuitively compelling and sound.

Letting SS-SOA-CRS denote the restriction of SS-SOA to CRS distributions, [3,29] had noted that SS-SOA-CRS implies IND-SOA-CRS and asked whether the converse was true. We settle this question in the negative, showing that SS-SOA-CRS is strictly stronger. We arrive at this separation by combining two facts. First, the message distribution underlying our negative result is CRS, meaning we say that there does not exist a HB-secure commitment scheme that is SS-SOA-CRS, not just SS-SOA. Second, it is known that there does exist a HB-secure commitment scheme that is IND-SOA-CRS [3,29].

Hofheinz [3,29] shows that any commitment scheme that is *statistically* hiding and binding is IND-SOA-CRS. This positive result does not contradict our result, because, as we have just seen (indeed, invoking this positive result to do so), IND-SOA-CRS is a strictly weaker requirement than SS-SOA or SS-SOA-CRS. A question that still remains open is whether HB-security implies IND-SOA-CRS security.

Message distribution. It has been suggested that the difficulty in showing that HB-security implies SS-SOA is that the messages in the vector $\mathbf{m}$ may be related to each other. Our results imply that although showing HB-security implies SS-SOA-security is not just hard but impossible, it is not for this reason. We have already noted that our negative result holds for a message distribution that is CRS. In fact, the message distribution is uniform, meaning the messages in the vector are uniformly and independently distributed strings. Even for this uniform distribution, no HB-secure commitment scheme is SS-SOA secure. This may at first glance appear to contradict known results, for DNRS [19] showed that HB-security implied SOA-security for independently distributed messages. The difference is that they only showed this for what they called semantic security for functions, a notion implied by, but not known to imply their main notion of semantic security for relations that we call SS-SOA. Thus, not only is there no contradiction, but our results settle an open question from [19]. Namely we show that their result does not extend to SS-SOA and also that SS-SOA is strictly stronger than semantic security for functions.

Random oracles. Our result holds in the standard model and in the non-programmable random oracle (RO) model [32]. (In the latter the simulator is given oracle access to the RO and cannot define it.) In the standard (programmable) RO model [5], where the simulator can define the RO, our result is not true: there *do* exist HB-secure schemes that are SS-SOA secure. As an example, commitment scheme $\mathcal{E}^H(m;r) = H(m;r)$, where H is the RO, is HB-secure in the non-programmable RO. Our results show it is not SS-SOA in this model. However, it can be easily shown SS-SOA in the programmable RO model. Consequently, our results yield another separation between the programmable and non-programmable RO models complementing that of [32].

Previous negative results. Hofheinz [3,29] shows that no HB-secure scheme can be proven SS-SOA secure via blackbox reduction to "standard" assumptions.

(A "standard" assumption as defined in [17,3,29] is one specified by a certain type of game.) However, it might still be possible to prove that a particular HB-secure scheme was SS-SOA in some ad hoc and non-blackbox way. The blackbox separation does not yield a single example of an HB-secure scheme that is not SS-SOA secure, let alone show, as we do, that all HB-secure schemes fail to be SS-SOA secure.

INTERACTION. Our result applies to non-interactive commitment schemes. When commitment involves an interactive protocol between sender and receiver the corresponding claim is not true. There *does* exist an interactive HB and SS-SOA secure commitment scheme. Specifically, Hofheinz [3,29] presents a particular construction of such a scheme based on one-way permutations. Further results on interactive SOA-secure commitment are [39,34].

SOA-SECURE ENCRYPTION FOR SENDER CORRUPTIONS. Turning now to encryption, consider a setting with n senders and one receiver, the latter having public encryption key ek. Sender i picks random coins $\mathbf{r}[i]$, encrypts its message $\mathbf{m}[i]$ via $\mathbf{c}[i] \leftarrow \mathcal{E}(ek, \mathbf{m}[i]; \mathbf{r}[i])$, and sends ciphertext $\mathbf{c}[i]$ to the receiver. The adversary selects, as a function of $\mathbf{c}$, a set $I \subseteq \{1, \ldots, n\}$ of the senders and corrupts them, obtaining their messages $\langle \mathbf{m}[i] : i \in I \rangle$ and coins $\langle \mathbf{r}[i] : i \in I \rangle$. As before, we say that $\mathcal{E}$ is SOA-secure if privacy of the un-opened messages is preserved. An SS-SOA definition analogous to the one for commitment was given in [3,8].

The standard and accepted security condition for encryption since [26] is of course IND-CPA. SOA-security was identified upon realizing that it is necessary to implement the assumed-secure channels in multi-party secure computation protocols like those of [9,14]. The central open question was whether or not IND-CPA implies SS-SOA. Neither a proof showing the implication is true, nor a counter-example showing it is false, had been given. We show that IND-CPA does not imply SS-SOA by exhibiting a large class of IND-CPA encryption schemes that we prove are not SS-SOA. The class includes many natural and existing schemes.

DNRS [19] had pointed out that the obstacle to proving that IND-CPA implies SS-SOA is that most encryption schemes are "committing." Our results provide formal support for this intuition. We formalize a notion of binding-security for encryption. Our result is that no binding encryption scheme is SS-SOA secure. As with commitment, it holds when the distribution on messages is uniform.

The existence of a decryption algorithm corresponding to the encryption algorithm means that for any ek created by honest key-generation, there do not exist r_0, r_1 and distinct m_0, m_1 such that $\mathcal{E}(ek, m_0; r_0) = \mathcal{E}(ek, m_1; r_1)$. Binding strengthens this condition to also hold when ek is adversarially chosen, while also relaxing it from unconditional to computational. It is thus a quite natural condition and is met by many schemes.

Inability to show that IND-CPA implies SS-SOA led to the search for specific SS-SOA secure encryption schemes. Non-commiting encryption [12] yields a solution when the number of bits encrypted is bounded by the length of the public key. The first full solution was based on lossy encryption [3,8]. Deniable

encryption [11] was used to obtain further solutions [20,7]. More lossy-encryption based solutions appear in [28]. In all these solutions, the encryption scheme is *not* binding. Our results show that this is necessary to achieve SS-SOA security.

SOA-security has so far been viewed as a theoretical rather than practical issue because even if there was no proof that IND-CPA implies SS-SOA, there were no attacks on standard, practical schemes such as ElGamal. Our results change this situation for they show that ElGamal and other practical schemes are not SS-SOA secure. Thus, the above-mentioned schemes that achieve SS-SOA in more involved ways are necessary if we want SS-SOA security.

IND-CCA doesn't help: The Cramer-Shoup scheme [15] meets our definition of binding and is thus not SS-SOA secure. As with commitment, our results imply that IND-SOA-CRS security is *strictly* weaker than SS-SOA-CRS security, answering an open question from [3,8]. Subsequent to our work, the relations between different notions of SOA-security under sender corruptions were further clarified in [10] but whether there exist schemes that are IND-CPA but not IND-SOA-CRS secure remains open.

SOA-secure encryption for receiver corruptions. In a dual of the above setting, there are n receivers and one sender, receiver i having public encryption key $\mathbf{ek}[i]$ and secret decryption key $\mathbf{dk}[i]$. For each i the sender picks random coins $\mathbf{r}[i]$, encrypts message $\mathbf{m}[i]$ via $\mathbf{c}[i] \leftarrow \mathcal{E}(\mathbf{ek}[i], \mathbf{m}[i]; \mathbf{r}[i])$, and sends ciphertext $\mathbf{c}[i]$ to receiver i. The adversary selects, as a function of $\mathbf{c}$, a set $I \subseteq \{1, \ldots, n\}$ of the receivers and corrupts them, obtaining not only the messages $\langle \mathbf{m}[i] : i \in I \rangle$ but also the decryption keys $\langle \mathbf{dk}[i] : i \in I \rangle$. As usual, we say that $\mathcal{E}$ is SOA-secure if privacy of the un-opened messages is preserved. An SS-SOA definition analogous to the ones for commitment and sender-corruptions in encryption is given in Section 5.

The status and issues are analogous to what we have seen above, namely that it has been open whether IND-CPA security implies SS-SOA for receiver corruptions, neither a proof nor a counter-example ever being given. We settle this with the first counter-examples. We define a notion of decryption verifiability for encryption that can be seen as a weak form of robustness [1]. It asks that there is an algorithm $\mathcal{W}$ such that it is hard to find ek, dk_0, dk_1 and distinct m_0, m_1 such that $\mathcal{W}(ek, dk_0, m_0)$ and $\mathcal{W}(ek, dk_1, m_1)$ both accept. We show that no IND-CPA and decryption-verifiable encryption scheme is SS-SOA secure. Standard encryption schemes like ElGamal are decryption verifiable (even though they are not robust) so our result continues to rule out SS-SOA security for many natural schemes.

Non-committing encryption [12] yields an SS-SOA scheme secure for receiver corruptions when the number of bits encrypted is bounded by the length of the secret key. Nielsen [32] showed that any non-committing encryption scheme has keys larger than the total number of message bits it can securely encrypt. This result is not known to extend to SS-SOA, meaning the existence of an SS-SOA scheme for receiver corruptions without this restriction is open. Our results do not rule out such a full solution but indicate that the scheme must not be decryption-verifiable.

2 Technical Approach

We provide a high-level description of our approach, focusing for simplicity on commitment schemes and the claim that no HB-secure commitment scheme is SS-SOA secure. We then discuss extensions and variants of our results.

THE DEFINITION. Let $\mathcal{E}$ be a commitment scheme. To compact notation, we extend it to vector inputs by letting $\mathcal{E}(\mathbf{m};\mathbf{r})$ be the vector whose i-th component is $\mathcal{E}(\mathbf{m}[i];\mathbf{r}[i])$. Let $\mathcal{M}$ be a message sampler that outputs a vector $\mathbf{m}$ of messages and let R be a relation. Adversary A, given ciphertext vector $\mathbf{c} = \mathcal{E}(\mathbf{m};\mathbf{r})$ will corrupt a subset I of the senders, get their messages and coins, and output a value w. It is said to win if $\mathsf{R}(\mathbf{m}, I, w)$ is true. The simulator, given no ciphertexts, can also corrupt a subset I of senders but gets back only the corresponding messages, and outputs a value w. It too is said to win if $\mathsf{R}(\mathbf{m}, I, w)$ is true. Security requires that for every $\mathcal{M}, \mathsf{R}$ and adversary A there is a simulator S such that S wins with about the same probability as A. DNRS [19, Sec 7.1] require this to be true even for any auxiliary input a given initially to A and also to S. See Section 4 for a formal definition.

THE ATTACK. Let $\mathcal{E}$ be any, given HB-secure commitment scheme. We construct $\mathcal{M}, \mathsf{R}, A$ for which we prove there is no simulator. We let $\mathcal{M}$ output $n = 2h$ randomly and independently distributed messages, each of length ℓ. Our adversary A applies to the vector $\mathbf{c} = \mathcal{E}(\mathbf{m};\mathbf{r})$ of commitments a hash function H to get back an h-bit string $b[1] \ldots b[h]$ and then corrupts the set of indices $I = \{2j - 1 + b[j] \ : \ 1 \le j \le h\}$ to get back $\langle \mathbf{m}[i] : i \in I\rangle$ and $\langle \mathbf{r}[i] : i \in I\rangle$. Its output w consists of $\mathbf{c}$ and $\langle \mathbf{r}[i] : i \in I\rangle$. We define R, on inputs $\mathbf{m}, I$ and w, to check two constraints. The *opening constraint* is that $\mathcal{E}(\mathbf{m}[i];\mathbf{r}[i]) = \mathbf{c}[i]$ for all $i \in I$. The *hash constraint* is that $I = \{2j - 1 + b[j] \ : \ 1 \le j \le h\}$ for $b[1] \ldots b[h] = H(\mathbf{c})$. A detailed description of A and R is in Fig. 3.

The simulator gets no ciphertexts. It must corrupt some set I of indices to get back $\langle \mathbf{m}[i] : i \in I\rangle$. Now it must create a ciphertext vector $\mathbf{c}$ and a list $\langle \mathbf{r}[i] : i \in I\rangle$ of coins to output as w to R, and to satisfy the latter it must satisfy both constraints. Intuitively, the simulator faces a Catch-22. It is helpful for the intuition to think of H as a random oracle. The simulator could first pick I in some way, get $\langle \mathbf{m}[i] : i \in I\rangle$ from its oracle, and compute $\mathbf{c}$ and $\langle \mathbf{r}[i] : i \in I\rangle$ to satisfy the opening constraint. But it is unlikely, given only $\mathrm{poly}(\cdot)$ queries to H, to satisfy the hash constraint. On the other hand it could pick some $\mathbf{c}$, define I to satisfy the hash constraint, and get $\langle \mathbf{m}[i] : i \in I\rangle$ from its oracle. But now it would have a hard time satisfying the opening constraint *because the commitment scheme is binding.*

This intuition that the simulator's task is hard is, however, not a proof that a simulator does not exist. Furthermore, the intuition relies on the hash function being a random oracle and we only want to assume collision-resistance. Our proof takes an arbitrary simulator and proves that the probability that it makes the relation true is small unless it finds a hash collision or violates binding. The proof involves backing up the simulator, feeding it different, random responses

to its corruption query, and applying a Reset Lemma analogous to that of [4]. We do not assume the simulation is blackbox. See Theorem 2.

Related work. The strategy of specifying challenges by a hash of commitments arose first in showing failure of parallel-repetition to preserve zero-knowledge [22,23]. The model, goals and techniques are however quite different. Also in [23] the simulator is assumed to make only blackbox calls to the adversary (verifier) and we make no such assumption, and they use a pairwise independent hash rather than a CR one. We point out that although the seed of our technique can be traced back 20 years it was not noted until now that it could be of use in settling the long-standing open question of whether HB-secure commitments are SS-SOA-secure.

Adaptive security. Our definition of SS-SOA, following [19,3,7] is one-shot, meaning the adversary gets all the ciphertexts at once and performs all its corruptions in parallel. A definition where the adversary can make adaptive ciphertext-creation and corruption requests is more suitable for applications. But our result is negative so using a restricted adversary only makes it stronger. (We are saying there is an attack with a one-shot adversary so certainly there is an attack with an adaptive adversary.)

The flip side is that if the adversary is allowed to be adaptive, so is the simulator. Our theorems only consider (and rule out) one-shot simulators for simplicity, but the proofs can be extended to also rule out adaptive simulators. We discuss briefly how to do this following the proof of Theorem 2.

Auxiliary inputs. As indicated above, the definition of DNRS [19] that we use allows both the adversary and simulator to get an auxiliary input, denoted "z" in [19, Sec 7.1]. The simplest and most basic form of our result exploits the auxiliary input to store the key describing the CR hash function. (If the simulator can pick this key the function will not be CR.)

Auxiliary inputs model history. They were introduced in the context of zero-knowledge by Goldreich and Oren [25] who showed that in their presence ZK had natural and desirable composability properties absent under the original definition of [27]. They have since become standard in zero-knowledge and also in simulation-based definitions in other contexts [18,19] to provide composability. Their inclusion in the SS-SOA definition of commitment by DNRS [19] was thus correct and justified and we put them to good use.

Later definitions [3,29] however appear to have dropped the auxiliary inputs. Although this appears to be only for notational simplicity (modern works on ZK also often drop auxiliary inputs since it is well understood how to extend the definition to include them) it does raise an interesting technical question, namely what negative results can we prove without auxiliary inputs?

A simple solution is to use one of the messages as a key. The adversary would corrupt the corresponding party to get this key, thereby defining the hash function, and then proceed as above. This however makes the adversary adaptive, and while this is still a significant result, we ask whether anything can be shown for one-shot adversaries without using auxiliary inputs.

This turns out to be technically challenging. The difficulty is that the simulator can control the hash key. In [2] we present a construction relying on a new primitive we call an encrypted hash scheme (EHS). The idea is that there is an underlying core hash function whose keys are messages and an encrypted hash function whose keys are ciphertexts. We show how to build an EHS based on DDH.

We remark that from a practical perspective these distinctions are moot since hash functions like SHA-256 are keyless. Also, it is possible to work theoretically with keyless hash functions [38]. But in classical asymptotic theoretical cryptography, hash functions are keyed and we were interested in results in this setting.

3 Preliminaries

Notation and conventions. If $n \in \mathbb{N}$ then let 1^n denote the string of n ones and $[n]$ the set $\{1, \ldots, n\}$. The empty string is denoted by ε. By $a \,\|\, b$ we denote the concatenation of strings a, b. If a is tuple then $(a_1, \ldots, a_n) \leftarrow a$ means we parse a into its constituents. We use boldface letters for vectors. If $\mathbf{x}$ is a vector then we let $|\mathbf{x}|$ denote the number of components of $\mathbf{x}$ and for $1 \le i \le |\mathbf{x}|$ we let $\mathbf{x}[i]$ denote its i-th component. For a set $I \subseteq [|\mathbf{x}|]$ we let $\mathbf{x}[I]$ be the $|\mathbf{x}|$-vector whose i-th component is $\mathbf{x}[i]$ if $i \in I$ and $\perp$ otherwise. We let $\perp_n$ denote the n-vector all of whose components are $\perp$. We define the Embedding subroutine Emb to take 1^n, $I \subseteq [n]$, a $|I|$-vector $\mathbf{x}^*$ and a n-vector $\overline{\mathbf{x}}$ and return the n-vector that consists of $\overline{\mathbf{x}}$ with $\mathbf{x}^*$ embedded in the positions indexed by I. More precisely,

Subroutine $\mathrm{Emb}(1^n, I, \mathbf{x}^*, \overline{\mathbf{x}})$
$j \leftarrow 0$; For $i = 1, \ldots, n$ do If $i \in I$ then $j \leftarrow j + 1$; $\overline{\mathbf{x}}[i] \leftarrow \mathbf{x}^*[j]$
Return $\overline{\mathbf{x}}$.

All algorithms are randomized, unless otherwise specified as being deterministic. We use the abbreviation PT for polynomial-time. If A is an algorithm then $y \leftarrow A(x_1, \ldots, x_n; r)$ represents the act of running the algorithm A with inputs $x_1, \ldots, x_n$ and coins r to get an output y and $y \leftarrow\!\!\$\, A(x_1, \ldots, x_n)$ represents the act of picking r at random and letting $y \leftarrow A(x_1, \ldots, x_n; r)$. By $[A(x_1, \ldots, x_n)]$ we denote the set of all y for which there exists r such that $y = A(x_1, \ldots, x_n; r)$.

Games. We use the language of code-based game-playing [6]. A game (see Fig. 1 for examples) has an Initialize procedure, procedures to respond to adversary oracle queries, and a Finalize procedure. A game G is executed with an adversary A and security parameter λ as follows. A is given input 1^λ and can then call game procedures. Its first oracle query must be $\textsc{Initialize}(1^\lambda)$ and its last oracle query must be to Finalize, and it must make exactly one query to each of these oracles. In between it can query the other procedures as oracles as it wishes. The output of Finalize, denoted $\mathrm{G}^A(\lambda)$, is called the output of the game, and we let "$\mathrm{G}^A(\lambda)$" denote the event that this game output takes value true.

$\underline{\textsc{Initialize}(1^\lambda)}$	
$b \leftarrow\!\!\$\ \{0,1\}\ ;\ \pi \leftarrow\!\!\$\ \mathcal{P}(1^\lambda)$	$\underline{\textsc{Initialize}(1^\lambda)}$
$(ek, dk) \leftarrow\!\!\$\ \mathcal{K}(\pi)$	$\pi \leftarrow\!\!\$\ \mathcal{P}(1^\lambda)$
Return (π, ek)	Return π
$\underline{\textsc{LR}(m_0, m_1)}$	$\underline{\textsc{Finalize}(ek, c, m_0, m_1, r_0, r_1)}$
$c \leftarrow\!\!\$\ \mathcal{E}(1^\lambda, \pi, ek, m_b)$	$d_0 \leftarrow \mathcal{V}(1^\lambda, \pi, ek, c, m_0, r_0)$
Return c	$d_1 \leftarrow \mathcal{V}(1^\lambda, \pi, ek, c, m_1, r_1)$
$\underline{\textsc{Finalize}(b')}$	Return $(d_0 \wedge d_1 \wedge (m_0 \neq m_1))$
Return $(b' = b)$	

Fig. 1. Game IND_{Π} (left) and game BIND_{Π} (right) defining, respectively, IND-CPA privacy and binding security of CE scheme $\Pi = (\mathcal{P}, \mathcal{K}, \mathcal{E}, \mathcal{V})$

CE Schemes. We introduce CE (Committing Encryption) schemes as a way to unify commitment and encryption schemes under a single syntax and avoid duplicating similar definitions and results for the two cases. A *CE scheme* $\Pi = (\mathcal{P}, \mathcal{K}, \mathcal{E}, \mathcal{V})$ is specified by four PT algorithms. Via $\pi \leftarrow\!\!\$\ \mathcal{P}(1^\lambda)$ the parameter-generation algorithm $\mathcal{P}$ generates system parameters such as a description of a group. Via $(ek, dk) \leftarrow\!\!\$\ \mathcal{K}(\pi)$ the key-generation algorithm $\mathcal{K}$ generates an encryption key ek and decryption key dk. Via $c \leftarrow \mathcal{E}(1^\lambda, \pi, ek, m; r)$ the encryption algorithm deterministically maps a message m and coins $r \in \{0,1\}^{\rho(\lambda)}$ to a ciphertext $c \in \{0,1\}^* \cup \{\perp\}$ where $\rho\colon\ \mathbb{N} \to \mathbb{N}$ is the *randomness length* associated to Π and $c \neq \perp$ iff $|m| = \ell(\lambda)$ where $\ell\colon\ \mathbb{N} \to \mathbb{N}$ is the *message length* associated to Π. Via $d \leftarrow \mathcal{V}(1^\lambda, \pi, ek, c, m, r)$, deterministic verification algorithm $\mathcal{V}$ returns true or false. We require that $\mathcal{V}(1^\lambda, \pi, ek, \mathcal{E}(1^\lambda, \pi, ek, m; r), m, r) = \mathsf{true}$ for all $\lambda \in \mathbb{N}$, all $\pi \in [\mathcal{P}(1^\lambda)]$, all $(ek, dk) \in [\mathcal{K}(\pi)]$, all $r \in \{0,1\}^{\rho(\lambda)}$ and all $m \in \{0,1\}^*$ such that $\mathcal{E}(1^\lambda, \pi, ek, m; r) \neq \perp$. We say that the verification algorithm $\mathcal{V}$ is *canonical* if $\mathcal{V}(1^\lambda, \pi, ek, c, m, r)$ returns the boolean $(\mathcal{E}(1^\lambda, \pi, ek, m; r) = c \neq \perp)$.

Game IND_{Π} of Fig. 1 captures the standard notion of indistinguishability under chosen-plaintext attack (IND-CPA) [26] and serves to define privacy for CE schemes. The adversary is allowed only one LR query and the messages m_0, m_1 involved must be of the same length. Game BIND_{Π} captures binding security. For adversaries A, B we let

$$\mathbf{Adv}^{\mathrm{indcpa}}_{\Pi,A}(\lambda) \;=\; 2\Pr[\mathrm{IND}^A_{\Pi}(\lambda)] - 1 \qquad \text{and} \qquad \mathbf{Adv}^{\mathrm{bind}}_{\Pi,B}(\lambda) \;=\; \Pr[\mathrm{BIND}^B_{\Pi}(\lambda)]\,.$$

We say that Π is IND-CPA secure if $\mathbf{Adv}^{\mathrm{indcpa}}_{\Pi,A}(\cdot)$ is negligible for all PT A, *binding* if $\mathbf{Adv}^{\mathrm{bind}}_{\Pi,B}(\cdot)$ is negligible for all PT B and *perfectly binding* if $\mathbf{Adv}^{\mathrm{bind}}_{\Pi,B}(\cdot) = 0$ for all (not necessarily PT) B.

Discussion. Commitment and encryption schemes can be recovered as special cases of CE schemes as follows. We say that Π is a *commitment scheme* if $\mathcal{K}$ always returns $(\varepsilon, \varepsilon)$. We see that our two security requirements capture the standard hiding and binding properties. In Section 1 we had simplified by assuming the verification algorithm is canonical and there were no parameters but here we are more general. We say that $\mathcal{D}$ is a decryption algorithm for CE scheme Π if

$\mathcal{D}(1^\lambda, \pi, dk, \mathcal{E}(1^\lambda, \pi, ek, m; r)) = m$ for all $\lambda \in \mathbb{N}$, all $\pi \in [\mathcal{P}(1^\lambda)]$, all $(ek, dk) \in [\mathcal{K}(\pi)]$, all $r \in \{0,1\}^{\rho(\lambda)}$ and all $m \in \{0,1\}^*$ such that $\mathcal{E}(1^\lambda, \pi, ek, m; r) \neq \bot$. We say that Π *admits decryption* if it has a PT decryption algorithm and in that case we say Π is an *encryption scheme*. IND-CPA is then, of course, the standard privacy goal.

Typical encryption schemes are perfectly binding under canonical verification with some added checks. For example, the ElGamal encryption scheme over a order-p group $\mathbb{G}$ with generator g (these quantities in the parameters) is binding under a verification algorithm that performs the re-encryption check and then also checks that quantities that should be in $\mathbb{G}$ or $\mathbb{Z}_p$ really are. RSA-based schemes can be made binding by requiring the encryption exponent to be a prime larger than the modulus.

Lossy encryption schemes [3,30,37] are not binding because the adversary could provide a lossy encryption key and, under this, be able to generate encryption collisions. Non-commiting [12,16] and deniable [11,33] encryption schemes are intentionally not binding. These types of encryption schemes have been shown to have SOA security. Our results show that the lack of binding was necessary for their success at this task.

Hash Functions. A hash function $\Gamma = (\mathcal{A}, \mathcal{H})$ with associated output length $h\colon \mathbb{N} \to \mathbb{N}$ is a tuple of PT algorithms. Via $a \leftarrow\!\!\$\, \mathcal{A}(1^\lambda)$ the key-generation algorithm $\mathcal{A}$ produces a key a. Via $y \leftarrow \mathcal{H}(a, x)$ the deterministic hashing algorithm $\mathcal{H}$ produces the $h(\lambda)$-bit hash of a string x under key a. Collision-resistance is defined via game CR_Γ whose Initialize(1^λ) procedure returns $a \leftarrow\!\!\$\, \mathcal{A}(1^\lambda)$ and whose Finalize procedure on input (x, x') returns $(x \neq x') \wedge (\mathcal{H}(a, x) = \mathcal{H}(a, x'))$. There are no other procedures. The advantage of an adversary C is defined by $\mathbf{Adv}^{\mathrm{cr}}_{\Gamma,C}(\lambda) = \Pr\left[\mathrm{CR}^C_\Gamma(\lambda)\right]$. We say that Γ is collision-resistant (CR) if $\mathbf{Adv}^{\mathrm{cr}}_{\Gamma,C}(\cdot)$ is negligible for every PT C. The following says that CR hash functions must have super-logarithmic output length and will be useful later:

Proposition 1. Let $\Gamma = (\mathcal{A}, \mathcal{H})$ be a hash function with associated output length $h\colon \mathbb{N} \to \mathbb{N}$. If Γ is collision-resistant then the function $2^{-h(\cdot)}$ is negligible.

4 SOA-C Insecurity of CE Schemes

Here we show that no CE-scheme that is binding is SOA-C secure. This implies that no HB-secure commitment scheme is SOA-secure and that no binding IND-CPA encryption scheme is SOA-secure under sender corruptions. In [2] we establish similar results for SOA-K to show that no robust IND-CPA encryption scheme is SOA-secure for receiver corruptions.

SOA-C Security. A *relation* is a PT algorithm with boolean output. A *message sampler* is a PT algorithm $\mathcal{M}$ taking input 1^λ and a string α and returning a vector over $\{0,1\}^*$. There must exist a function $n\colon \mathbb{N} \to \mathbb{N}$ (called the number of messages) and a function $\ell\colon \mathbb{N} \times \{0,1\}^* \times \mathbb{N} \to \mathbb{N}$ (called the message length) such that $|\mathbf{m}| = n(\lambda)$ and $|\mathbf{m}[i]| = \ell(\lambda, \alpha, i)$ for all $\mathbf{m} \in [\mathcal{M}(1^\lambda, \alpha)]$ and all $i \in [n]$. An *auxiliary-input generator* is a PT algorithm.

$\underline{\textsc{Initialize}(1^\lambda)}$
$\pi \leftarrow\!\!\$\ \mathcal{P}(1^\lambda)\ ;\ a \leftarrow \mathcal{A}(1^\lambda)\ ;\ (ek, dk) \leftarrow\!\!\$\ \mathcal{K}(\pi)$
Return (a, π, ek)

$\underline{\textsc{Enc}(\alpha)}$
$\mathbf{m} \leftarrow\!\!\$\ \mathcal{M}(1^\lambda, \alpha)$
For $i = 1, \ldots, n(\lambda)$ do
  $\mathbf{r}[i] \leftarrow\!\!\$\ \{0,1\}^{\rho(\lambda)}\ ;\ \mathbf{c}[i] \leftarrow \mathcal{E}(1^\lambda, \pi, ek, \mathbf{m}[i]; \mathbf{r}[i])$
Return $\mathbf{c}$

$\underline{\textsc{Corrupt}(I)}$
Return $\mathbf{m}[I], \mathbf{r}[I]$

$\underline{\textsc{Finalize}(w)}$
Return $\mathsf{R}(1^\lambda, a, \pi, \mathbf{m}, \alpha, I, w)$

$\underline{\textsc{Initialize}(1^\lambda)}$
$\pi \leftarrow\!\!\$\ \mathcal{P}(1^\lambda)\ ;\ a \leftarrow \mathcal{A}(1^\lambda)$
Return (a, π)

$\underline{\textsc{Msg}(\alpha)}$
$\mathbf{m} \leftarrow\!\!\$\ \mathcal{M}(1^\lambda, \alpha)$

$\underline{\textsc{Corrupt}(I)}$
Return $\mathbf{m}[I]$

$\underline{\textsc{Finalize}(w)}$
Return $\mathsf{R}(1^\lambda, a, \pi, \mathbf{m}, \alpha, I, w)$

Fig. 2. Game $\mathrm{RSOAC}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A}}$ capturing the real-world SOA-C attack to be mounted by an adversary (left) and game $\mathrm{SSOAC}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A}}$ capturing the simulated-world SOA-C attack to be mounted by a simulator (right)

Let $\Pi = (\mathcal{P}, \mathcal{K}, \mathcal{E}, \mathcal{V})$ be a CE-scheme, R a relation, $\mathcal{M}$ a message sampler and $\mathcal{A}$ an auxiliary-input generator. We define SOA-C security via the games of Fig. 2. "Real" game $\mathrm{RSOAC}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A}}$ will be executed with an adversary A. An soa-c adversary's (mandatory, starting) $\textsc{Initialize}(1^\lambda)$ call results in its being returned an auxiliary input, parameters, and an encryption key, the latter corresponding to the single receiver modeled here. The adversary is then required to make exactly one $\textsc{Enc}(\alpha)$ call. This results in production of a message vector whose encryption is provided to the adversary. Now the adversary is required to make exactly one $\textsc{Corrupt}(I)$ call to get back the messages *and coins* corresponding to the senders named in the set $I \subseteq [n(\lambda)]$. It then calls $\textsc{Finalize}$ with some value w of its choice and wins if the relation returns true on the inputs shown. A soa-c simulator S runs with the simulator game $\mathrm{SSOAC}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A}}$ and gets back only auxiliary input and parameters from its $\textsc{Initialize}(1^\lambda)$ call, there being no encryption key in its world. It is then required to make exactly one $\textsc{Msg}(\alpha)$ call resulting in creation of a message vector but the simulator is returned nothing related to it. It must then make its $\textsc{Corrupt}(I)$ and $\textsc{Finalize}(w)$ calls like the adversary and wins under the same conditions. The soa-c-advantage of an soa-c-adversary A with respect to CE-scheme Π, message sampler $\mathcal{M}$, relation R, auxiliary input generator $\mathcal{A}$ and soa-c simulator S is defined by

$$\mathbf{Adv}^{\text{soa-c}}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A},A,S}(\lambda) \;=\; \Pr\left[\,\mathrm{RSOAC}^{A}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A}}(\lambda)\,\right] - \Pr\left[\,\mathrm{SSOAC}^{S}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A}}(\lambda)\,\right].$$

We say that Π is $(\mathcal{M},\mathcal{A})$-SOA-C-secure if for every PT R and every PT soa-c adversary A there exists a PT soa-c simulator S such that $\mathbf{Adv}^{\text{soa-c}}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A},A,S}(\cdot)$ is negligible. We say that Π is SOA-C-secure if it is $(\mathcal{M},\mathcal{A})$-SOA-C-secure for every PT $\mathcal{M},\mathcal{A}$.

Result. The following implies that any binding CE-scheme is not SOA-C-secure.

Theorem 2. *Let $\Pi = (\mathcal{P}, \mathcal{K}, \mathcal{E}, \mathcal{V})$ be a binding CE-scheme with message length ℓ: $\mathbb{N} \to \mathbb{N}$. Let $\Gamma = (\mathcal{A}, \mathcal{H})$ be a collision-resistant hash function with associated output length h: $\mathbb{N} \to \mathbb{N}$. Let $n(\cdot) = 2h(\cdot)$ and let $\mathcal{M}$ be the message sampler that on input $1^\lambda, \alpha$ (ignores α and) returns a $n(\lambda)$-vector whose components are uniformly and independently distributed over $\{0,1\}^{\ell(\lambda)}$. Then there exists a PT soa-c adversary A and a PT relation R such that for all PT simulators S there is a negligible function ν such that $\mathbf{Adv}^{\text{soa-c}}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A},A,S}(\lambda) \geq 1 - \nu(\lambda)$ for all $\lambda \in \mathbb{N}$.* ∎

Thus, Π is not $(\mathcal{M}, \mathcal{A})$-SOA-C-secure and hence cannot be SOA-C-secure. Moreover, this is true when the distribution on messages is uniform. These claims would only require $\mathbf{Adv}^{\text{soa-c}}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A},A,S}(\cdot)$ in the theorem to be non-negligible, but we show more, namely that it is almost one. Note that ℓ is arbitrary and could even be $\ell(\cdot) = 1$, meaning we rule out SOA-C-security even for bit-commitment and encryption of 1-bit messages. The proof will make use of the following variant of the Reset Lemma of [4].

Lemma 3. *Let $V = \{V_\lambda\}_{\lambda \in \mathbb{N}}$ be a collection of non-empty sets. Let P_1, P_2 be algorithms, the second with boolean output. The single-execution acceptance probability $\mathbf{AP}_1(P_1, P_2, V, \lambda)$ is defined as the probability that $d = \mathsf{true}$ in the single execution experiment $\overline{St} \leftarrow\!\!\$\, P_1(1^\lambda)$; $\mathbf{m}^* \leftarrow\!\!\$\, V_\lambda$; $d \leftarrow\!\!\$\, P_2(\overline{St}, \mathbf{m}^*)$. The double-execution acceptance probability $\mathbf{AP}_2(P_1, P_2, V, \lambda)$ is defined as the probability that $d_1 = d_2 = \mathsf{true}$ and $\mathbf{m}_0^* \neq \mathbf{m}_1^*$ in the double execution experiment $\overline{St} \leftarrow\!\!\$\, P_1(1^\lambda); \mathbf{m}_0^*, \mathbf{m}_1^* \leftarrow\!\!\$\, V_\lambda; d_0 \leftarrow\!\!\$\, P_2(\overline{St}, \mathbf{m}_0^*); d_1 \leftarrow\!\!\$\, P_2(\overline{St}, \mathbf{m}_1^*)$. Then $\mathbf{AP}_1(P_1, P_2, V, \lambda) \leq 1/|V_\lambda| + \sqrt{\mathbf{AP}_2(P_1, P_2, V, \lambda)}$ for all $\lambda \in \mathbb{N}$.* ∎

The two executions in the double-execution experiment are not independent because $\overline{St}$ is the same for both, which is why the lemma is not trivial.

Proof (Lemma 3). Let $\delta = 1/|V_\lambda|$. Let $\mathsf{X}(\omega) = \Pr[d = \mathsf{true}]$ in the experiment $\overline{St} \leftarrow P_1(1^\lambda; \omega)$; $\mathbf{m}^* \leftarrow\!\!\$\, V_\lambda$; $d \leftarrow\!\!\$\, P_2(\overline{St}, \mathbf{m}^*)$. So $\mathbf{E}[\mathsf{X}] = \mathbf{AP}_1(P_1, P_2, V, \lambda)$ where the expectation is over the coins ω of P_1. Let $a_1 = \mathbf{AP}_1(P_1, P_2, V, \lambda)$ and $a_2 = \mathbf{AP}_2(P_1, P_2, V, \lambda)$. Then

$$a_2 \;\geq\; \mathbf{E}[\mathsf{X}(\mathsf{X} - \delta)] \;=\; \mathbf{E}[\mathsf{X}^2] - \delta \cdot \mathbf{E}[\mathsf{X}] \;\geq\; \mathbf{E}[\mathsf{X}]^2 - \delta \cdot \mathbf{E}[\mathsf{X}] \;=\; a_1^2 - \delta \cdot a_1$$

where the third step above is by Jensen's inequality. Now $a_1^2 - \delta \cdot a_1 = (a_1 - \delta/2)^2 - \delta^2/4$ so

$$a_1 \;\leq\; \delta/2 + \sqrt{a_2 + \delta^2/4} \;\leq\; \delta/2 + \sqrt{a_2} + \sqrt{\delta^2/4} \;=\; \delta + \sqrt{a_2}$$

which yields the lemma. ∎

Proof (Theorem 2). The adversary A and relation R are depicted in Fig. 3. Let S be any PT soa-c simulator. In the real game the adversary always makes the relation return true hence

$$\mathbf{Adv}^{\text{soa-c}}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A},A,S}(\lambda) \;=\; 1 - \Pr\left[\,\mathrm{SSOAC}^S_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A}}(\lambda)\,\right].$$

Adversary $A(1^\lambda)$	Relation $\mathsf{R}(1^\lambda, a, \pi, \mathbf{m}, \alpha, I, w)$
$(a, \pi, ek) \leftarrow$ INITIALIZE(1^λ)	If $\alpha \neq \varepsilon$ then return false
$\mathbf{c} \leftarrow$ ENC(ε)	$(ek, \mathbf{c}, \overline{\mathbf{r}}) \leftarrow w$; $b[1] \ldots b[h(\lambda)] \leftarrow \mathcal{H}(a, ek \,\|\, \mathbf{c})$
$b[1] \ldots b[h(\lambda)] \leftarrow \mathcal{H}(a, ek \,\|\, \mathbf{c})$	If $(I \neq \{2j - 1 + b[j] \,:\, 1 \leq j \leq h(\lambda)\})$ then
$I \leftarrow \{2j - 1 + b[j] \,:\, 1 \leq j \leq h(\lambda)\}$	return false
$(\overline{\mathbf{m}}, \overline{\mathbf{r}}) \leftarrow$ CORRUPT(I)	If $\lvert\mathbf{c}\rvert \neq n(\lambda)$ or $\lvert\overline{\mathbf{r}}\rvert \neq n(\lambda)$ then return false
$w \leftarrow (ek, \mathbf{c}, \overline{\mathbf{r}})$	For all $i \in I$ do
FINALIZE(w)	If $\mathcal{V}(1^\lambda, \pi, ek, \mathbf{c}[i], \mathbf{m}[i], \overline{\mathbf{r}}[i]) =$ false then
	return false
	Return true

Fig. 3. Adversary A and relation R for the proof of Theorem 2

We will construct a binding-adversary B and cr-adversary C such that

$$\Pr\left[\,\mathrm{SSOAC}^{S}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A}}(\lambda)\,\right] \;\leq\; 2^{-h(\lambda)\ell(\lambda)} + \sqrt{\mathbf{Adv}^{\mathrm{cr}}_{\Gamma,C}(\lambda) + \mathbf{Adv}^{\mathrm{bind}}_{\Pi,B}(\lambda)}\,. \qquad (1)$$

The assumptions that Γ is collision-resistant, Π is binding, together with Proposition 1, imply that the RHS of Eq. (1) is negligible, which proves the theorem. It remains to construct B and C. Given S we can define sub-algorithms S_1, S_2 such that S can be written in terms of S_1, S_2 as follows:

Simulator $S(1^\lambda)$
$(a, \pi) \leftarrow$ INITIALIZE(1^λ) ; MSG$(1^\lambda, \varepsilon)$; $(St, I) \leftarrow\!\!\$\; S_1(a, \pi)$
$\overline{\mathbf{m}} \leftarrow$ CORRUPT(I) ; $w \leftarrow S_2(St, \overline{\mathbf{m}})$; FINALIZE(w)

We clarify that we are not defining S; the latter is given and arbitrary. Rather, *any* S has the form above for *some* S_1, S_2 that can be determined given S. Specifically, S_1 runs S until S makes its CORRUPT(I) query, returning I along with the current state St of S. Then S_2, given the response $\overline{\mathbf{m}}$ to the query, feeds it back to S and continues executing S from St. By having S_1 put all S's coins in St we can assume S_2 is deterministic. We may assume wlog that $|I|$ is always $h(\lambda)$ and that the argument α in S's MSG call is ε since otherwise R rejects. We now define adversary B. The embedding subroutine Emb it calls and the notation $\perp_{n(\lambda)}$ were defined in Section 3:

Adversary $B(1^\lambda)$
$\pi \leftarrow$ INITIALIZE(1^λ) ; $a \leftarrow\!\!\$\; \mathcal{A}(1^\lambda)$; $(St, I) \leftarrow\!\!\$\; S_1(a, \pi)$
$\mathbf{m}^*_0, \mathbf{m}^*_1 \leftarrow\!\!\$\; (\{0,1\}^{\ell(\lambda)})^{h(\lambda)}$
$\overline{\mathbf{m}}_0 \leftarrow \mathrm{Emb}(1^{n(\lambda)}, I, \mathbf{m}^*_0, \perp_{n(\lambda)})$; $\overline{\mathbf{m}}_1 \leftarrow \mathrm{Emb}(1^{n(\lambda)}, I, \mathbf{m}^*_1, \perp_{n(\lambda)})$
$w_0 \leftarrow S_2(St, \overline{\mathbf{m}}_0)$; $(ek_0, \mathbf{c}_0, \overline{\mathbf{r}}_0) \leftarrow w_0$
$w_1 \leftarrow S_2(St, \overline{\mathbf{m}}_1)$; $(ek_1, \mathbf{c}_1, \overline{\mathbf{r}}_1) \leftarrow w_1$; $t \leftarrow\!\!\$\; I$
For all $i \in I$ do If $\overline{\mathbf{m}}_0[i] \neq \overline{\mathbf{m}}_1[i]$ then $t \leftarrow i$
FINALIZE$(ek_0, \mathbf{c}_0[t], \overline{\mathbf{m}}_0[t], \overline{\mathbf{m}}_1[t], \overline{\mathbf{r}}_0[t], \overline{\mathbf{r}}_1[t])$

Adversary B is running S to get its CORRUPT query I and then, by backing it up, providing two different responses. Adversary C has a similar strategy, only deviating in how the final values are used:

$$\underline{\text{Adversary } C(1^\lambda)}$$
$$a \leftarrow \text{INITIALIZE}(1^\lambda)\ ;\ \pi \leftarrow\!\!\$\ \mathcal{P}(1^\lambda)\ ;\ (St, I) \leftarrow\!\!\$\ S_1(a, \pi)$$
$$\mathbf{m}_0^*, \mathbf{m}_1^* \leftarrow\!\!\$\ (\{0,1\}^{\ell(\lambda)})^{h(\lambda)}$$
$$\overline{\mathbf{m}}_0 \leftarrow \text{Emb}(1^{n(\lambda)}, I, \mathbf{m}_0^*, \perp_{n(\lambda)})\ ;\ \overline{\mathbf{m}}_1 \leftarrow \text{Emb}(1^{n(\lambda)}, I, \mathbf{m}_1^*, \perp_{n(\lambda)})$$
$$w_0 \leftarrow S_2(St, \overline{\mathbf{m}}_0)\ ;\ (ek_0, \mathbf{c}_0, \overline{\mathbf{r}}_0) \leftarrow w_0$$
$$w_1 \leftarrow S_2(St, \overline{\mathbf{m}}_1)\ ;\ (ek_1, \mathbf{c}_1, \overline{\mathbf{r}}_1) \leftarrow w_1$$
$$\text{FINALIZE}((ek_0 \,\|\, \mathbf{c}_0, ek_1 \,\|\, \mathbf{c}_1)).$$

The analysis will use Lemma 3. Let $V_\lambda = (\{0,1\}^{\ell(\lambda)})^{h(\lambda)}$ and $V = \{V_\lambda\}_{\lambda \in \mathbb{N}}$. Define P_1, P_2 via:

$$\underline{\text{Algorithm } P_1(1^\lambda)}$$
$$\pi \leftarrow \mathcal{P}(1^\lambda)\ ;\ a \leftarrow\!\!\$\ \mathcal{A}(1^\lambda)\ ;\ (St, I) \leftarrow\!\!\$\ S_1(a, \pi)$$
$$\mathbf{m} \leftarrow\!\!\$\ (\{0,1\}^{\ell(\lambda)})^{n(\lambda)}\ ;\ \overline{St} \leftarrow (1^\lambda, a, \pi, \mathbf{m}, I, St)$$
$$\text{Return } \overline{St}$$

$$\underline{\text{Algorithm } P_2(\overline{St}, \mathbf{m}^*)}$$
$$(1^\lambda, a, \pi, \mathbf{m}, I, St) \leftarrow \overline{St}$$
$$\overline{\mathbf{m}} \leftarrow \text{Emb}(1^{n(\lambda)}, I, \mathbf{m}^*, \perp_{n(\lambda)})$$
$$w \leftarrow S_2(St, \overline{\mathbf{m}})$$
$$\mathbf{m} \leftarrow \text{Emb}(1^{n(\lambda)}, I, \mathbf{m}^*, \mathbf{m})$$
$$\text{Return } \mathsf{R}(1^\lambda, a, \pi, \mathbf{m}, \varepsilon, I, w)$$

Above the argument $\mathbf{m}^*$ to P_2 is drawn from V_λ. Now

$$\begin{aligned}\Pr\left[\,\text{SSOAC}^S_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A}}(\lambda)\,\right] &= \mathbf{AP}_1(P_1, P_2, V, \lambda) \\ &\le 2^{-h(\lambda)\ell(\lambda)} + \sqrt{\mathbf{AP}_2(P_1, P_2, V, \lambda)} \qquad (2)\end{aligned}$$

Above the equality is from the definitions and the inequality is by Lemma 3. Finally we claim that

$$\mathbf{AP}_2(P_1, P_2, V, \lambda) \ \le\ \mathbf{Adv}^{\text{cr}}_{\Gamma,C}(\lambda) + \mathbf{Adv}^{\text{bind}}_{\Pi,B}(\lambda)\ . \qquad (3)$$

Eqs. (2) and (3) imply Eq. (1) and conclude the proof. We now justify Eq. (3). To do so it is helpful to write down the double-execution experiment underlying $\mathbf{AP}_2(P_1, P_2, V, \lambda)$:

$$\pi \leftarrow \mathcal{P}(1^\lambda)\ ;\ a \leftarrow\!\!\$\ \mathcal{A}(1^\lambda)\ ;\ (St, I) \leftarrow\!\!\$\ S_1(a, \pi)\ ;\ \mathbf{m} \leftarrow\!\!\$\ (\{0,1\}^{\ell(\lambda)})^{n(\lambda)}$$
$$\mathbf{m}_0^*, \mathbf{m}_1^* \leftarrow\!\!\$\ (\{0,1\}^{\ell(\lambda)})^{h(\lambda)}$$
$$\overline{\mathbf{m}}_0 \leftarrow \text{Emb}(1^{n(\lambda)}, I, \mathbf{m}_0^*, \perp_{n(\lambda)})\ ;\ \overline{\mathbf{m}}_1 \leftarrow \text{Emb}(1^{n(\lambda)}, I, \mathbf{m}_1^*, \perp_{n(\lambda)})$$
$$w_0 \leftarrow S_2(St, \overline{\mathbf{m}}_0)\ ;\ w_1 \leftarrow S_2(St, \overline{\mathbf{m}}_1)\ ;\ (ek_0, \mathbf{c}_0, \overline{\mathbf{r}}_0) \leftarrow w_0\ ;\ (ek_1, \mathbf{c}_1, \overline{\mathbf{r}}_1) \leftarrow w_1$$
$$\mathbf{m}_0 \leftarrow \text{Emb}(1^{n(\lambda)}, I, \mathbf{m}_0^*, \mathbf{m})\ ;\ \mathbf{m}_1 \leftarrow \text{Emb}(1^{n(\lambda)}, I, \mathbf{m}_1^*, \mathbf{m})$$
$$\text{Return } \mathsf{R}(1^\lambda, a, \pi, \mathbf{m}_0, \varepsilon, I, w_0) \wedge \mathsf{R}(1^\lambda, a, \pi, \mathbf{m}_1, \varepsilon, I, w_1) \wedge (\mathbf{m}_0^* \neq \mathbf{m}_1^*).$$

Assume this experiment returns true. By definition of R it must be that $I = \{2j - 1 + b_0[j] \ :\ 1 \le j \le h(\lambda)\}$ where $b_0[1] \ldots b_0[h(\lambda)] = \mathcal{H}(a, ek_0 \,\|\, \mathbf{c}_0)$ and also $I = \{2j - 1 + b_1[j] : 1 \le j \le h(\lambda)\}$ where $b_1[1] \ldots b_1[h(\lambda)] = \mathcal{H}(a, ek_1 \,\|\, \mathbf{c}_1)$. However, I is the same in both cases, so we must have $\mathcal{H}(a, ek_0 \,\|\, \mathbf{c}_0) = \mathcal{H}(a, ek_1 \,\|\, \mathbf{c}_1)$, meaning we have a hash collision. This means that C succeeds unless $ek_0 \,\|\, \mathbf{c}_0 = ek_1 \,\|\, \mathbf{c}_1$. But we now argue that in the latter case, B succeeds. We know $\mathbf{m}_0^* \neq \mathbf{m}_1^*$ so there is some $t \in I$ such that $\overline{\mathbf{m}}_0[t] \neq \overline{\mathbf{m}}_1[t]$. The definition of R implies that $\mathcal{V}(1^\lambda, \pi, ek_0, \mathbf{c}_0[t], \overline{\mathbf{m}}_0[t], \overline{\mathbf{r}}_0[t]) = \mathsf{true}$ and also $\mathcal{V}(1^\lambda, \pi, ek_1, \mathbf{c}_1[t], \overline{\mathbf{m}}_1[t], \overline{\mathbf{r}}_1[t]) = \mathsf{true}$. But since $ek_0 \,\|\, \mathbf{c}_0 = ek_1 \,\|\, \mathbf{c}_1$ we have $\mathcal{V}(1^\lambda, \pi, ek_0, \mathbf{c}_0[t], \overline{\mathbf{m}}_0[t], \overline{\mathbf{r}}_0[t]) = \mathsf{true}$ and also $\mathcal{V}(1^\lambda, \pi, ek_0, \mathbf{c}_0[t], \overline{\mathbf{m}}_1[t], \overline{\mathbf{r}}_1[t]) = \mathsf{true}$ with $\overline{\mathbf{m}}_0[t] \neq \overline{\mathbf{m}}_1[t]$ so B wins. ∎

Extensions, applications and remarks. The SOA-C definition could be weakened by allowing the simulator's corruptions to be adaptive, meaning S is allowed multiple queries to procedure Corrupt that now would take input $i \in [n(\lambda)]$ and return $\mathbf{m}[i]$. The proof strategy of Theorem 2 no longer works but can be extended to also rule out adaptive simulators. We would back S up to its last Corrupt query and give a new response only to this query. We would now require $\ell(\cdot)$ to be super-logarithmic so that collisions are rare on single messages. We omit the details.

Theorem 2 applies to all commitment schemes since they are binding by definition. Not all encryption schemes are binding, but many popular ones are. For example, the ElGamal scheme is binding. The Cramer-Shoup scheme [15] is also binding, showing that IND-CCA is not a panacea against SOAs.

Our model allows a scheme to have system parameters π that effectively function as auxiliary input. This means the simulator cannot modify them. This is not necessary but merely makes the results more general. If one wishes to view commitment, as in DNRS [19], as having no parameters, just restrict attention to schemes where π is always 1^λ. Our result applies to these as a special case.

5 SOA-K Insecurity of Encryption Schemes

Here we show that no decryption-verifiable IND-CPA encryption scheme is SOA-secure for receiver corruptions.

SOA-K Security. This is the dual of SOA-C where there are multiple receivers and a single sender rather than a single receiver and multiple senders, and corruptions reveal decryption keys rather than coins. The definition uses games $\mathrm{RSOAK}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A}}$ and $\mathrm{SSOAK}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A}}$ of Fig. 4. The soa-k-advantage of an soa-k-adversary A with respect to the encryption scheme Π, message sampler $\mathcal{M}$, relation R, auxiliary input generator $\mathcal{A}$ and soa-k simulator S is defined by

$$\mathbf{Adv}^{\text{soa-k}}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A},A,S}(\lambda) \;=\; \Pr\left[\,\mathrm{RSOAK}^{A}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A}}\,\right] - \Pr\left[\,\mathrm{SSOAK}^{S}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A}}\,\right].$$

We say that Π is $(\mathcal{M},\mathcal{A})$-SOA-K-secure if for every PT R and every PT soa-k adversary A there exists a PT soa-k simulator S such that $\mathbf{Adv}^{\text{soa-k}}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A},A,S}(\cdot)$ is negligible. We say that Π is SOA-K-secure if it is $(\mathcal{M},\mathcal{A})$-SOA-K-secure for every PT $\mathcal{M},\mathcal{A}$.

Result. The following implies that any decryption-verifiable encryption scheme is not SOA-K-secure. Decryption-verifiable encryption schemes are defined in [2] and include many common schemes. The proof is in [2].

Theorem 4. *Let $\Pi = (\mathcal{P},\mathcal{K},\mathcal{E},\mathcal{V})$ be a decryption-verifiable encryption scheme with decryption verifier $\mathcal{W}$ and message length $\ell\colon \mathbb{N} \to \mathbb{N}$. Let $\Gamma = (\mathcal{A},\mathcal{H})$ be a collision-resistant hash function with associated output length $h\colon \mathbb{N} \to \mathbb{N}$. Let $n(\cdot) = 2h(\cdot)$ and let $\mathcal{M}$ be the message sampler that on input $1^\lambda, \alpha$ (ignores α and) returns a $n(\lambda)$-vector whose components are uniformly and independently*

$\underline{\text{INITIALIZE}(1^\lambda)}$
$\pi \xleftarrow{\$} \mathcal{P}(1^\lambda)$; $a \leftarrow \mathcal{A}(1^\lambda)$
For $i = 1, \ldots, n(\lambda)$ do $(\mathbf{ek}[i], \mathbf{dk}[i]) \xleftarrow{\$} \mathcal{K}(\pi)$
Return $(a, \pi, \mathbf{ek})$

$\underline{\text{ENC}(\alpha)}$
$\mathbf{m} \xleftarrow{\$} \mathcal{M}(1^\lambda, \alpha)$
For $i = 1, \ldots, n(\lambda)$ do
  $\mathbf{r}[i] \xleftarrow{\$} \{0,1\}^{\rho(\lambda)}$; $\mathbf{c}[i] \leftarrow \mathcal{E}(1^\lambda, \pi, \mathbf{ek}[i], \mathbf{m}[i]; \mathbf{r}[i])$
Return $\mathbf{c}$

$\underline{\text{CORRUPT}(I)}$
Return $\mathbf{m}[I], \mathbf{dk}[I]$

$\underline{\text{FINALIZE}(w)}$
Return $\mathsf{R}(1^\lambda, a, \pi, \mathbf{m}, \alpha, I, w)$

$\underline{\text{INITIALIZE}(1^\lambda)}$
$\pi \xleftarrow{\$} \mathcal{P}(1^\lambda)$; $a \leftarrow \mathcal{A}(1^\lambda)$
Return (a, π)

$\underline{\text{MSG}(\alpha)}$
$\mathbf{m} \xleftarrow{\$} \mathcal{M}(1^\lambda, \alpha)$

$\underline{\text{CORRUPT}(I)}$
Return $\mathbf{m}[I]$

$\underline{\text{FINALIZE}(w)}$
Return $\mathsf{R}(1^\lambda, a, \pi, \mathbf{m}, \alpha, I, w)$

Fig. 4. Game $\text{RSOAK}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A}}$ capturing the real-world SOA-K attack to be mounted by an adversary (left) and game $\text{SSOAK}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A}}$ capturing the simulated-world SOA-K attack to be mounted by a simulator (right)

distributed over $\{0,1\}^{\ell(\lambda)}$. *Then there exists a PT soa-k adversary* A *and a PT relation* R *such that for all PT soa-k simulators* S *there is a negligible function* ν *such that* $\mathbf{Adv}^{\text{soa-k}}_{\Pi,\mathcal{M},\mathsf{R},\mathcal{A},A,S}(\lambda) \geq 1 - \nu(\lambda)$ *for all* $\lambda \in \mathbb{N}$. ∎

Acknowledgments. The first two authors were supported in part by NSF grants CNS-0627779 and CCF-0915675. The third author was supported in part by NSF grants CNS-0915361 and CNS-0952692, AFOSR grant FA9550-08-1-0352, DARPA PROCEED, DARPA N11AP20006, a Google Faculty Research award, the Alfred P. Sloan Fellowship, a Microsoft Faculty Fellowship, and a Packard Foundation Fellowship.

References

1. Abdalla, M., Bellare, M., Neven, G.: Robust Encryption. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 480–497. Springer, Heidelberg (2010)
2. Bellare, M., Dowsley, R., Waters, B., Yilek, S.: Standard security does not imply security against selective-opening. Cryptology ePrint Archive, Report 2011/581 (2011), Full version of this abstract, `http://eprint.iacr.org/`
3. Bellare, M., Hofheinz, D., Yilek, S.: Possibility and Impossibility Results for Encryption and Commitment Secure under Selective Opening. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 1–35. Springer, Heidelberg (2009)
4. Bellare, M., Palacio, A.: GQ and Schnorr Identification Schemes: Proofs of Security against Impersonation under Active and Concurrent Attacks. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 162–177. Springer, Heidelberg (2002)
5. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Ashby, V. (ed.) ACM CCS 1993, pp. 62–73. ACM Press (November 1993)
6. Bellare, M., Rogaway, P.: The Security of Triple Encryption and a Framework for Code-Based Game-Playing Proofs. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 409–426. Springer, Heidelberg (2006)

7. Bellare, M., Waters, B., Yilek, S.: Identity-Based Encryption Secure against Selective Opening Attack. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 235–252. Springer, Heidelberg (2011)
8. Bellare, M., Yilek, S.: Encryption schemes secure under selective opening attack. Cryptology ePrint Archive, Report 2009/101 (2009), http://eprint.iacr.org/
9. Ben-Or, M., Goldwasser, S., Wigderson, A.: Completeness theorems for noncryptographic fault-tolerant distributed computations. In: 20th ACM STOC, pp. 1–10. ACM Press (May 1988)
10. Böhl, F., Hofheinz, D., Kraschewski, D.: On definitions of selective opening security. Cryptology ePrint Archive, Report 2011/678 (2011), http://eprint.iacr.org/
11. Canetti, R., Dwork, C., Naor, M., Ostrovsky, R.: Deniable Encryption. In: Kaliski Jr., B.S. (ed.) CRYPTO 1997. LNCS, vol. 1294, pp. 90–104. Springer, Heidelberg (1997)
12. Canetti, R., Feige, U., Goldreich, O., Naor, M.: Adaptively secure multi-party computation. In: 28th ACM STOC, pp. 639–648. ACM Press (May 1996)
13. Canetti, R., Halevi, S., Katz, J.: Adaptively-Secure, Non-interactive Public-Key Encryption. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 150–168. Springer, Heidelberg (2005)
14. Chaum, D., Crépeau, C., Damgård, I.: Multiparty unconditionally secure protocols. In: 20th ACM STOC, pp. 11–19. ACM Press (May 1988)
15. Cramer, R., Shoup, V.: Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. SIAM Journal on Computing 33(1), 167–226 (2003)
16. Damgård, I.B., Nielsen, J.B.: Improved Non-committing Encryption Schemes Based on a General Complexity Assumption. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 432–450. Springer, Heidelberg (2000)
17. Dodis, Y., Oliveira, R., Pietrzak, K.: On the Generic Insecurity of the Full Domain Hash. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 449–466. Springer, Heidelberg (2005)
18. Dolev, D., Dwork, C., Naor, M.: Nonmalleable cryptography. SIAM Journal on Computing 30(2), 391–437 (2000)
19. Dwork, C., Naor, M., Reingold, O., Stockmeyer, L.J.: Magic functions. Journal of the ACM 50(6), 852–921 (2003)
20. Fehr, S., Hofheinz, D., Kiltz, E., Wee, H.: Encryption Schemes Secure against Chosen-Ciphertext Selective Opening Attacks. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 381–402. Springer, Heidelberg (2010)
21. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. Journal of the ACM 33, 792–807 (1986)
22. Goldreich, O., Krawczyk, H.: On the Composition of Zero-Knowledge Proof Systems. In: Paterson, M. (ed.) ICALP 1990. LNCS, vol. 443, pp. 268–282. Springer, Heidelberg (1990)
23. Goldreich, O., Krawczyk, H.: On the composition of zero-knowledge proof systems. SIAM Journal on Computing 25(1), 169–192 (1996)
24. Goldreich, O., Micali, S., Wigderson, A.: Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. Journal of the ACM 38(3), 691–729 (1991)
25. Goldreich, O., Oren, Y.: Definitions and properties of zero-knowledge proof systems. Journal of Cryptology 7(1), 1–32 (1994)
26. Goldwasser, S., Micali, S.: Probabilistic encryption. Journal of Computer and System Sciences 28(2), 270–299 (1984)

27. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems. SIAM Journal on Computing 18(1), 186–208 (1989)
28. Hemenway, B., Libert, B., Ostrovsky, R., Vergnaud, D.: Lossy Encryption: Constructions from General Assumptions and Efficient Selective Opening Chosen Ciphertext Security. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 70–88. Springer, Heidelberg (2011)
29. Hofheinz, D.: Possibility and impossibility results for selective decommitments. Journal of Cryptology 24(3), 470–516 (2011)
30. Kol, G., Naor, M.: Cryptography and Game Theory: Designing Protocols for Exchanging Information. In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 320–339. Springer, Heidelberg (2008)
31. Naor, M.: Bit commitment using pseudorandomness. Journal of Cryptology 4(2), 151–158 (1991)
32. Nielsen, J.B.: Separating Random Oracle Proofs from Complexity Theoretic Proofs: The Non-committing Encryption Case. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 111–126. Springer, Heidelberg (2002)
33. O'Neill, A., Peikert, C., Waters, B.: Bi-Deniable Public-Key Encryption. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 525–542. Springer, Heidelberg (2011)
34. Ostrovsky, R., Rao, V., Scafuro, A., Visconti, I.: Revisiting lower and upper bounds for selective decommitments. Cryptology ePrint Archive, Report 2011/536 (2011), `http://eprint.iacr.org/`
35. Panjwani, S.: Tackling Adaptive Corruptions in Multicast Encryption Protocols. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 21–40. Springer, Heidelberg (2007)
36. Pedersen, T.P.: Non-interactive and Information-Theoretic Secure Verifiable Secret Sharing. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, pp. 129–140. Springer, Heidelberg (1992)
37. Peikert, C., Vaikuntanathan, V., Waters, B.: A Framework for Efficient and Composable Oblivious Transfer. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 554–571. Springer, Heidelberg (2008)
38. Rogaway, P.: Formalizing Human Ignorance. In: Nguyên, P.Q. (ed.) VIETCRYPT 2006. LNCS, vol. 4341, pp. 211–228. Springer, Heidelberg (2006)
39. Xiao, D.: (Nearly) Round-Optimal Black-Box Constructions of Commitments Secure against Selective Opening Attacks. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 541–558. Springer, Heidelberg (2011)