Lossy Encryption from General Assumptions

Brett Hemenway* and Rafail Ostrovsky†

July 26, 2009

Abstract

In this paper, we present new, general constructions of lossy encryption schemes. Applying the results of [Hof08] and [BHY09], we obtain general cryptosystems secure against a Selective Opening Adversary (SOA). Although it was recognized almost twenty years ago that SOA security was important, it was not until the recent breakthrough works of Hofheinz [Hof08] and Bellare, Hofheinz and Yilek [BHY09] that any progress was made on this fundamental problem.

The Selective Opening problem is as follows: suppose an adversary receives n commitments (or encryptions) of (possibly) correlated messages, and now the adversary can choose $n/2$ of the messages, and receive decommitments (or decryptions *and* the randomness used to encrypt them). Do the unopened commitments (encryptions) remain secure? A protocol which achieves this type of security is called *secure against a Selective Opening Adversary (SOA)*. This question arises naturally in the context of Byzantine Agreement and Secure Multiparty Computation, where an active adversary is able to eavesdrop on all the wires, and then choose a subset of players to corrupt. Unfortunately, the traditional definitions of security (IND-CPA, IND-CCA) do not guarantee security in this setting. In this paper:

- We formally define *re-randomizable* encryption and show that *every* re-randomizable encryption scheme gives rise to efficient encryptions secure against a selective opening adversary. (Very informally, an encryption is re-randomizable, if given any ciphertext, there is an efficient way to map it to an almost uniform re-encryption of the same underlying message).
- We formally define *re-randomizable* one-way functions and show that *every* re-randomizable one-way function family gives rise to efficient commitments secure against a Selective Opening Adversary.
- Applying our constructions to the known cryptosystems of El-Gamal, Paillier, and Goldwasser and Micali, we obtain selective opening secure commitments and encryptions from the Decisional Diffie-Hellman (DDH), Decisional Composite Residuosity (DCR) and Quadratic Residuosity (QR) assumptions, that are either simpler or more efficient than existing constructions of Bellare, Hofheinz and Yilek.
- We show that Statistically-Hiding 2-round Oblivious Transfer (OT) implies Lossy Encryption. Combining this with known results immediately gives the following new results
  - Private Information Retrieval implies Lossy Encryption, and hence selective opening secure encryption.
  - Homomorphic Encryption implies Lossy Encryption, and hence selective opening secure encryption.
- Applying our general results to the Paillier Cryptosystem we demonstrate the first cryptosystem to achieve Semantic Selective Opening security from the DCR assumption.
- We define the notion of indistinguishability-based adaptive chosen ciphertext security (CCA-2) in the selective opening setting, and describe the first encryption scheme which is CCA-2 secure and simultaneously SOA-secure, relative to this definition.

*E-mail: bretth@math.ucla.edu

†E-mail: rafail@cs.ucla.edu

- We define the notion of simulation-based adaptive chosen ciphertext security (CCA-2) in the selective opening setting, and describe the first encryption scheme which is CCA-2 secure and simultaneously SOA-secure, relative to this definition.

Keywords: Public Key Encryption, Commitment, Selective Opening, Homomorphic Encryption, Chosen Ciphertext Security, Lossy Encryption

1 Introduction

In Byzantine agreement, and more generally in secure multiparty computation, it is often assumed that all parties are connected to each other via private channels. In practice, these private channels are implemented using a public-key cryptosystem. An adaptive adversary in an MPC setting, however, has very different powers than an adversary in an IND-CPA or IND-CCA game. In particular, an adaptive MPC adversary may view all the encryptions sent in a given round, and then choose to corrupt a certain fraction of the players, thus revealing the decryptions of those players' messages *and the randomness used to encrypt them*. A natural question is whether the messages sent from the uncorrupted players remain secure. If the messages (and randomness) of all the players are chosen independently, then security in this setting follows immediately from the IND-CPA security of the underlying encryption. If, however, the messages are not chosen independently, the security does not immediately follow from the IND-CPA (or even IND-CCA) security of the underlying scheme. In fact, although this problem was first investigated over twenty years ago, it remains an open question whether IND-CPA (or IND-CCA) security implies this *Selective Opening* security.

A similar question may be asked regarded in terms of commitments as well. Suppose an adversary is allowed to see commitments to a number of related messages, the adversary may then choose a subset of the commitments for the challenger to decommit. Does this reveal any information about the unopened commitments? This question has applications to concurrent zero-knowledge proofs.

2 Previous Work

There have been many attempts to design encryption protocols that can be used to implement secure multiparty computation against an adaptive adversary. The first protocols by Beaver and Haber [BH92] required interaction between the sender and receiver, required erasure and were fairly inefficient. The first non-interactive protocol was given by Canetti, Feige, Goldreich and Naor in [CFGN96]. In [CFGN96] the authors defined a new primitive called Non-Committing Encryption, and gave an example of such a scheme based on the RSA assumption. In [Bea97], Beaver extended the work of [CFGN96], and created adaptively secure key exchange under the Diffie-Hellman assumption. In subsequent work Damgård and Nielsen improved the efficiency of the schemes of Canetti et al. and Beaver, they were also able to obtain Non-Committing Encryption based on one-way trapdoor functions with invertible sampling. In [CHK05], Canetti, Halevi and Katz presented a Non-Committing encryption protocols with evolving keys.

In [CDNO97], Canetti, Dwork, Naor and Ostrovsky extended the notion of Non-Committing Encryption to a new protocol which they called Deniable Encryption. In Non-Committing Encryption schemes there is a simulator, which can generate non-committing ciphertexts, and later open them to any desired message, while in Deniable Encryption, valid encryptions generated by the sender and receiver can later be opened to any desired message. The power of this primitive made it relatively difficult to realize, and Canetti et al. were only able to obtain modest examples of Deniable Encryption and left it as an open question whether fully deniable schemes could be created.

The notions of security against an adaptive adversary can also be applied to commitments. In fact, according to [DNRS03] the necessity of adaptively-secure commitments was realized by 1985. Despite its utility, until recently there have been relatively few papers that directly address the question of commitments secure against a Selective Opening Adversary (SOA). The work of Dwork, Naor, Reingold and Stockmeyer [DNRS03] was the first to explicitly address the problem. In [DNRS03], Dwork et al.

showed that non-interactive SOA-secure commitments can be used to create a 3-round zero-knowledge proof system for NP with negligible soundness error, and they gave constructions of a weak form of SOA-secure commitments, but leave open the question of whether general SOA-secure commitments exist.

The question of SOA-secure commitments was put on firm foundations by Hofheinz [Hof08] and Bellare, Hofheinz and Yilek in [BHY09]. In [BHY09], Bellare et al. distinguished between simulation-based and indistinguishability-based definitions of security, and gave a number of constructions and black-box separations. In particular, Hofheinz showed that in the simulation-based setting, non-interactive SOA-secure commitments cannot be realized in a black-box manner from standard cryptographic assumptions, but if interaction is allowed, they can be created from one-way permutations in a non-black-box manner. In the indistinguishability-based setting, they showed that any statistically-hiding scheme achieves this level of security, but that there is a black-box separation between perfectly-binding SOA-secure commitments and most standard cryptographic assumptions. Our results in the Selective Opening setting build on the breakthrough results of [BHY09].

3 Our Contributions

In this paper, we primarily consider encryptions secure against a selective opening adversary. In particular, we formalize the notion of **Re-Randomizable** Public-Key Encryption and we show that re-randomizable encryption implies Lossy Encryption as defined in [PVW08], and expanded in [BHY09]. Combining this with the recent result of Bellare, Hofheinz and Yilek [BHY09] showing that Lossy Encryption is IND-SO-ENC secure, we have an efficient construction of IND-SO-ENC secure encryption from any re-randomizable encryption (which generalizes and extends previous results). Furthermore, these constructions retain the efficiency of the underlying re-randomizable encryption protocol.

Applying our results to the Paillier Cryptosystem, we obtain a cryptosystem which attains a strong, simulation-based form of semantic security under selective openings (SEM-SO-ENC security). This is the first construction of this type from the Decisional Composite Residuosity (DCR) assumption, and the most efficient known construction of SEM-SO-ENC secure encryption.

We go on to show that statistically-hiding $\binom{2}{1}$-OT implies lossy encryption. Combining this with the results of [PVW08], we recognize that this "new" primitive, lossy encryption, is essentially just a different way to view the well known primitive statistically-hiding $\binom{2}{1}$-OT. Applying the reductions in [BHY09] to this result, yields constructions of SOA secure encryption from both PIR and Homomorphic Encryption.

These results show that the primitives Lossy Encryption and Selective Opening Secure Encryption, which are fairly new and not very well-studied primitives are in fact implied by many well-known primitives i.e. Re-randomizable encryption, PIR, Homomorphic Encryption and statistically-hiding $\binom{2}{1}$-OT.

Prior to this work the only known general[1] constructions of lossy encryption were from Lossy Trapdoor Functions. Our results show that they are implied by many seemingly weaker primitives. The full relationship can be seen below.

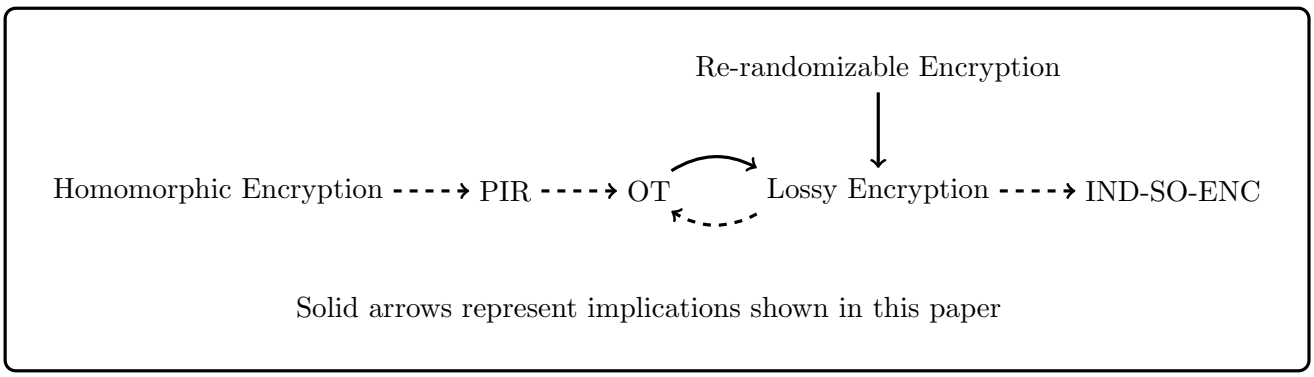

[1] i.e. not based on specific number-theoretic assumptions

Finally, we present a definition of security against a chosen ciphertext (CCA-2) attack in the selective opening setting (in both the indistinguishability and simulation-based models) and create the first public-key cryptosystems that satisfy these strengthened forms of security. We note that our constructions are completely orthogonal to the recent work of Prabhakaran and Rosulek [PR07] creating RCCA Encryption. In their work, they create encryptions which satisfy a version of security against a chosen-ciphertext attack, while remaining re-randomizable. In this work, we use re-randomizable (CPA secure) encryption to create Selective Opening secure encryption, and then use Selective Opening secure encryption (and other tools) to create a cryptosystem that retains its Selective Opening security against an adaptive chosen ciphertext attack.

4 Notation

If $f : X \to Y$ is a function, for any $Z \subset X$, we let $f(Z) = \{f(x) : x \in Z\}$.

If A is a PPT machine, then we use $a \leftarrow A$ to denote running the machine A and obtaining an output, where a is distributed according to the internal randomness of A. For a PPT machine A, we use $\mathsf{coins}(A)$ to denote the distribution of the internal randomness of A. So the distributions $\{a \leftarrow A\}$ and $\{r \leftarrow \mathsf{coins}(A) : a = A(r)\}$ are identical. If R is a set, we use $r \leftarrow R$ to denote sampling uniformly from R.

If X and Y are families of distributions indexed by a security parameter λ, we use $X \approx_s Y$ to mean the distributions X and Y are statistically close, i.e. for all polynomials p and sufficiently large λ we have

$$\sum_x |\Pr[X = x] - \Pr[Y = x]| < \frac{1}{p(\lambda)},$$

We use $X \approx_c Y$ to mean X and Y are computationally close, i.e. for all PPT adversaries A, for all polynomials p, then for all sufficiently large λ,

$$\left|\Pr[A^X = 1] - \Pr[A^Y = 1]\right| < \frac{1}{p(\lambda)}.$$

5 Re-randomizable Encryption

In many cryptosystems, given a ciphertext c, and a public-key it is possible to re-randomize the ciphertext to a new ciphertext c', such that c and c' are valid encryptions of the same plaintext, but they are statistically independent. Formally, we call a Public Key Cryptosystem given by algorithms (G, E, D) *re-randomizable* (RRPKC) if

- (G, E, D) is semantically-secure in the standard sense (IND-CPA).
- There is an efficient function ReRand such that if r' is chosen uniformly from $\mathsf{coins}(\mathsf{ReRand})$, and r_0 are chosen uniformly from $\mathsf{coins}(E)$, then the distributions

  $$\{r_0 \leftarrow \mathsf{coins}(E) : E(pk, m, r_0)\} \approx_s \{r' \leftarrow \mathsf{coins}(\mathsf{ReRand}) : \mathsf{ReRand}(E(pk, m, r_1), r')\}$$

  for all public keys pk and messages m, and randomness r_1.

We note that this definition of re-randomizable encryption provides a statistical re-randomization. It is possible to define re-randomizable encryption which satisfies perfect re-randomization (stronger) or computational re-randomization (weaker). Such definitions already exist in the literature (see for example [PR07],[Gro04],[JJS04],[CKN03]). Our constructions require statistical re-randomization, and do not go through under a computational re-randomization assumption.

There are many known examples of re-randomizable encryption. For example, if (G, E, D) is *homomorphic*, i.e. $E(pk, m_0, r_0) \cdot E(pk, m_1, r_1) = E(pk, m_0 + m_1, r^*)$, we can re-randomize by taking $\mathsf{ReRand}(pk, c, r') = c \cdot E(pk, 0, r')$. For all known homomorphic cryptosystems, (e.g. El-Gamal, Paillier, Damgård-Jurik, Goldwasser-Micali) we obtain re-randomizable encryption with this definition of ReRand.

We note that since re-randomization does not require any kind of group structure on the plaintext space, or any method for combining ciphertexts, it appears to be a weaker primitive than homomorphic encryption. It is not, however, implied by homomorphic encryption. See Appendix B for a more thorough discussion of the relationship between these primitives.

6 Selective Opening Secure Encryption

6.1 Preliminaries

Here we present a definition of encryption secure against a Selective Opening Adversary (this was originally formalized in [BHY09]).

We define two games, a real and an ideal game which should be indistinguishable to any efficient adversary. The key point to notice is that the adversary receives *both* the messages and the randomness for his selection. This mirrors the fact that an adaptive MPC adversary learns the entire history of the corrupted players (i.e. there are no secure erasures). If the adversary receives only the messages this would reduce to standard CPA security.

Definition 1. (Indistinguishability under selective openings/IND-SO-ENC).

Let (G, E, D) be a Public Key Cryptosystem (PKC), we say that (G, E, D) is indistinguishable under selective openings (IND-SO-ENC secure) if for every PPT message distribution M and every PPT adversary A, we have that

$$\left|\Pr\left[A^{\mathsf{ind\text{-}so\text{-}real}} = 1\right] - \Pr\left[A^{\mathsf{ind\text{-}so\text{-}ideal}} = 1\right]\right| < \nu$$

for some negligible function ν, and where the games $\mathsf{ind\text{-}so\text{-}real}$ and $\mathsf{ind\text{-}so\text{-}ideal}$ are defined as follows

IND-SO-ENC (Real)

- $(m_1, \ldots, m_n) \leftarrow M$
- $r_1, \ldots, r_n \leftarrow \mathsf{coins}(E)$
- $I \leftarrow A((E(m_1, r_i), \ldots, E(m_n, r_n))$
- $b \leftarrow A(((m_i, r_i))_{i \in I}, (m_1, \ldots, m_n))$

IND-SO-ENC (Ideal)

- $(m_1, \ldots, m_n) \leftarrow M$
- $r_1, \ldots, r_n \leftarrow \mathsf{coins}(E)$
- $I \leftarrow A((E(m_1, r_i), \ldots, E(m_n, r_n))$
- $(m'_1, \ldots, m'_n) \leftarrow M|M_I$
- $b \leftarrow A((m_i, r_i))_{i \in I}, (m'_1, \ldots, m'_n))$

More explicitly, in the real game,

- The challenger samples messages $(m_1, \ldots, m_n) \leftarrow M$, from the joint message distribution.
- The challenger generates randomness $r_1, \ldots, r_n \leftarrow \mathsf{coins}(E)$.
- The challenger sends $(E(m_1, r_1), \ldots, E(m_n, r_n)$ to A.
- The adversary A responds with a subset $I \subset \{1, \ldots, n\}$, with $|I| = n/2$.
- The challenger reveals *both* m_i and r_i for $i \in I$.

- The challenger sends $(m_1, \ldots, m_n)$ to the adversary.
- The adversary outputs a bit b.

In the ideal game,

- The challenger samples messages $(m_1, \ldots, m_n) \leftarrow M$, from the joint message distribution.
- The challenger generates randomness $r_1, \ldots, r_n \leftarrow \mathsf{coins}(E)$.
- The challenger sends $(E(m_1, r_1), \ldots, E(m_n, r_n))$ to A.
- The adversary A responds with a subset $I \subset \{1, \ldots, n\}$, with $|I| = n/2$.
- The challenger reveals *both* m_i and r_i for $i \in I$.
- The challenger samples a new vector $m' \leftarrow M|M_I$, from M conditioned on the fact that $m_i = m'_i$ for $i \in I$, and sends M' to A.
- The adversary outputs a bit b.

We emphasize that the challenger reveals both the messages m_i and the randomness r_i for the selected messages. If the challenger only revealed the messages m_i, this type of security would follow immediately from IND-CPA security.

7 Lossy Encryption

In [PVW08], Peikert, Vaikuntanathan and Waters defined Dual-Mode Encryption, a type of cryptosystem with two types public-keys, injective keys on which the cryptosystem behaves normally and "lossy" or "messy" keys on which the system loses information about the plaintext. In particular they require that the encryptions of any two plaintexts under a lossy key yield distributions that are statistically close, yet injective and lossy keys remain computationally indistinguishable.

In [BHY09] Bellare, Hofheinz and Yilek define *Lossy Encryption*, expanding on the definitions of Dual-Mode Encryption in [PVW08], and Meaningful/Meaningless Encryption in [KN08]. At a high level, a 'lossy' (or 'messy' in the terminology of [PVW08]) cryptosystem is one which has two types of public keys which specify two different modes of operation. In the normal mode, encryption is injective, while in the lossy (or 'messy') mode, the ciphertexts generated by the encryption algorithm are independent of the plaintext. We also require that no efficient adversary can distinguish normal keys from lossy keys. In [BHY09], they also require openability, which basically allows the decryptor to decrypt a ciphertext generated from a lossy key to *any* plaintext.

Definition 2. Formally, an *lossy public-key encryption scheme* is a tuple $(G_{\mathsf{inj}}, G_{\mathsf{lossy}}, E, D)$ of polynomial-time algorithms such that

- $G_{\mathsf{inj}}(1^\lambda)$ outputs keys (pk, sk), keys generated by G_{inj} are called *injective keys*.
- $G_{\mathsf{lossy}}(1^\lambda)$ outputs keys $(pk_{\mathsf{lossy}}, sk_{\mathsf{lossy}})$, keys generated by G_{lossy} are called *lossy keys*.

Additionally, the algorithms must satisfy the following properties:

1. *Correctness on injective keys.* For all $x \in X$,

$$\Pr\left[(pk, sk) \leftarrow G_{\mathsf{inj}}(1^\lambda); r \leftarrow \mathsf{coins}(E) : D(sk, E(pk, x, r)) = x\right] = 1.$$

2. *Indistinguishability of keys.* This basically says that the the pk in lossy mode and injective mode are computationally indistinguishable. Specifically, if $\text{proj} : (pk, sk) \mapsto pk$ is the projection map, then the two distributions

$$\{\text{proj}(G_{\mathsf{inj}}(1^\lambda))\} \approx_c \{\text{proj}(G_{\mathsf{lossy}}(1^\lambda))\}$$

3. *Lossiness of lossy keys.* If $(pk_{\mathsf{lossy}}, sk_{\mathsf{lossy}}) \leftarrow G_{\mathsf{lossy}}$, then for all $x_0, x_1 \in X$, the two distributions $E(pk_{\mathsf{lossy}}, x_0, R)$ and $E(pk_{\mathsf{lossy}}, x_1, R)$ are statistically close, i.e. the statistical distance is negligible in λ.

4. *Openability.* If $(pk_{\mathsf{lossy}}, sk_{\mathsf{lossy}}) \leftarrow G_{\mathsf{lossy}}$, and $r \leftarrow \mathsf{coins}(E)$, then for all $x_0, x_1 \in X$ with all but negligible probability, there exists an $r' \in \mathsf{coins}(E)$, such that $E(pk_{\mathsf{lossy}}, x_0, r) = E(pk_{\mathsf{lossy}}, x_1, r')$. While this is a statistical property that follows immediately from property (3), it is convenient, to state it explicitly, and to rephrase it in terms of an algorithm. We require that with all but negligible probability there is an (unbounded) algorithm opener that can open a lossy ciphertext to *any* plaintext.

Although the Openability property is implied by property (3), it is useful to include it explicitly because it simplifies the exposition somewhat. It also generalizes nicely, and in [BHY09] they show that if the algorithm opener is efficient, then the encryption scheme is actually SEM-SO-ENC secure (instead of only IND-SO-ENC).

We do not explicitly assume, that the scheme is IND-CPA secure, and in fact, the semantic security of the scheme follows from the indistinguishability of keys and the lossiness of the lossy keys, since for any $x_0, x_1 \in X$,

$$E(\pi(G_{\mathsf{inj}}(1^\lambda)), x_0, R) \approx_c E(\pi(G_{\mathsf{lossy}}(1^\lambda)), x_0, R)) \approx_s E(\pi(G_{\mathsf{lossy}}(1^\lambda)), x_1, R) \approx_c E(\pi(G_{\mathsf{inj}}(1^\lambda)), x_1, R).$$

In [BHY09] it was shown that Lossy Encryption can be constructed in a straightforward manner from Lossy-Trapdoor Functions, in fact, they simply observe that the CPA-secure system given in [PW08] is a Lossy Encryption.

Next, they showed

Theorem 1. Lossy Encryption is IND-SO-ENC secure.

Proof. This is proven in [BHY09].
For a review of the proof, see the proof of Theorem 5 in Appendix A.2. □

Thus to create IND-SO-ENC secure encryptions, it suffices to construct Lossy Encryption.

7.1 Re-randomizable Encryption Implies Lossy Encryption

Our first result gives a simple and efficient method for creating Lossy Encryption from Re-randomizable encryption.
Let (G, E, D) be a re-randomizable public-key cryptosystem, and we create Lossy Encryption $(\bar{G}_{\mathsf{inj}}, \bar{G}_{\mathsf{lossy}}, \bar{E}, \bar{D})$ as follows:

- $\bar{G}_{\mathsf{inj}}(1^\lambda)$ runs $\bar{G}(1^\lambda)$, generating a pair (pk, sk). Then G_{inj} picks $r_0, r_1 \leftarrow \mathsf{coins}(E)$, and generates $e_0 = E(pk, 0, r_0)$, $e_1 = E(pk, 1, r_1)$. $\bar{G}_{\mathsf{inj}}$ returns $(\bar{pk}, \bar{sk}) = ((pk, e_0, e_1), sk)$.
- $\bar{G}_{\mathsf{lossy}}(1^\lambda)$ runs $\bar{G}(1^\lambda)$, generating a pair (pk, sk). Then G_{lossy} picks $r_0, r_1 \leftarrow \mathsf{coins}(E)$, and generates $e_0 = E(pk, 0, r_0)$, $e_1 = E(pk, 0, r_1)$. $\bar{G}_{\mathsf{lossy}}$ returns $(\bar{pk}, \bar{sk}) = ((pk, e_0, e_1), sk)$.
- $\bar{E}(\bar{pk}, b, r') = \mathsf{ReRand}(pk, e_b, r')$ for $b \in \{0, 1\}$.
- $\bar{D}(\bar{sk}, c)$, simply outputs $D(sk, c)$.

To see that this is a lossy encryption we notice that under an injective key it is clearly injective by the correctness of the decryption algorithm D, while in lossy mode, it will be statistically lossy by the properties of the ReRand function. The proof that this Lossy Encryption is straightforward and we check the details here.

1. *Correctness on injective keys.* This follows immediately from the correctness of E.
2. *Indistinguishability of keys.* This follows immediately from the IND-CPA security of (G, E, D).
3. *Lossiness of lossy keys.* Notice that under a lossy public-key $\bar{pk}$, e_0 and e_1 are both encryptions of zero, so $\bar{E}(\bar{pk}, b, r)$ will also be an encryption of zero for $b \in \{0, 1\}$. By the properties of ReRand, we have that the distributions $\{\bar{E}(\bar{pk}, 0, r)\}$ and $\{\bar{E}(\bar{pk}, 1, r)\}$ are statistically close, which is exactly what is required for a key to be "lossy".
4. *Openability.* Under a lossy public-key, $\bar{E}(\bar{pk}, b, r') = \mathsf{ReRand}(E(pk, 0, r_b), r')$. Since r' is chosen uniformly from $\mathsf{coins}(\mathsf{ReRand})$, the properties of ReRand guarantee that the distributions $\mathsf{ReRand}(E(pk, 0, r_b), r')$ and $\mathsf{ReRand}(E(pk, 0, r_{1-b}, r''))$ are statistically close. That there exists an r'' such that $\mathsf{ReRand}(E(pk, 0, r_b), r') = \mathsf{ReRand}(E(pk, 0, r_{1-b}), r'')$ then follows from lemma 1.

Lemma 1. If R is a random variable, and $f : R \to X$, $g : R \to Y$ and

$$\sum_{z \in X \cup Y} \Pr\left[r \leftarrow R : f(r) = z\right] - \Pr\left[r \leftarrow R : g(r) = z\right] = \nu,$$

then

$$\Pr\left[r \leftarrow R : \nexists r' \in R \text{ such that } f(r) = g(r')\right] < \nu.$$

Proof. It suffices to notice that

$$\begin{aligned}
\nu &= \sum_{z \in X \cup Y} \Pr\left[r \leftarrow R : f(r) = z\right] - \Pr\left[r \leftarrow R : g(r) = z\right] \\
&\geq \sum_{z \in X \setminus Y} \Pr\left[r \leftarrow R : f(r) = z\right] - \Pr\left[r \leftarrow R : g(r) = z\right] \\
&= \Pr\left[r \leftarrow R : \nexists r' \in R \text{ such that } f(r) = g(r')\right]
\end{aligned}$$

□

It is clear that the same construction also gives a perfectly-binding SOA secure commitment scheme (with trusted setup). If our goal is only to construct SOA secure commitments, we do not need Re-randomizable encryption, and a weaker primitive suffices. In Appendix A, we define Re-randomizable One-Way Functions and show that these imply SOA secure commitments. While both these constructions require trusted setup, in a sense that is inevitable since it was shown in [BHY09] that perfectly-binding SOA secure commitments without trusted setup cannot be created in a black-box manner from any primitive with a game-based definition of security.

8 Oblivious Transfer

We briefly recall the definition of honest-receiver two round statistically-hiding $\binom{2}{1}$-OT.

oblivious transfer is a protocol between two parties, a sender Sen and a receiver $\mathsf{Rec} = (\mathsf{Rec}_q, \mathsf{Rec}_r)$. the sender Sen has two two strings s_1, s_2, and the receiver has a bit b. the receiver generates a query q and sends q to the sender. the sender evaluates $\mathsf{q}(s_1, s_2, r)$, and sends the result to the receiver.

- **Correctness:**
  For all $s_1, s_2 \in \{0,1\}^k$, for all $b \in \{0,1\}$,

  $$\Pr[(\mathsf{q}, sk) \leftarrow \mathsf{Rec}_q(1^\lambda, b); r \leftarrow \mathsf{Sen}(\mathsf{q}, s_1, s_2) : \mathsf{Rec}_r(sk, r) = s_b] \geq 1 - \nu(\lambda).$$

  For some negligible function ν.

- **Receiver Privacy:**
  The distributions

  $$\{(\mathsf{q}, sk) \leftarrow \mathsf{Rec}_q(1^\lambda, 0) : \mathsf{q}\} \approx_c \{(\mathsf{q}, sk) \leftarrow \mathsf{Rec}_q(1^\lambda, 1) : \mathsf{q}\}$$

  are computationally indistinguishable, where the probability is taken over the internal randomness of Rec_q.

- **Sender Privacy:**
  The distributions

  $$\{(\mathsf{q}, sk) \leftarrow \mathsf{Rec}_q(1^\lambda, b); \mathsf{r} \leftarrow \mathsf{Sen}(\mathsf{q}, s_1, s_2) : \mathsf{r}\} \approx_s \{(\mathsf{q}, sk) \leftarrow \mathsf{Rec}_q(, 1^\lambda, b); \mathsf{r} \leftarrow \mathsf{Sen}(\mathsf{q}, s_1', s_2') : \mathsf{r}\}$$

  for all s_1', s_2' with $s_b = s_b'$, where the randomness is taken over the internal randomness of Rec_q and Sen.

9 Statistically-Hiding $\binom{2}{1}$-OT Implies Lossy Encryption

Let $(\mathsf{Sen}, \mathsf{Rec})$ be a two round honest-receiver statistically-hiding $\binom{2}{1}$-OT.
We construct a lossy encryption as follows:

- **Key Generation:**
  Define $g(1^\lambda, inj) = \mathsf{Rec}(1^\lambda, 0)$, define $g(1^\lambda, lossy) = \mathsf{Rec}(1^\lambda, 1)$, so $pk = \mathsf{q}$, and $sk = sk$.

- **Encryption:**
  Define $e(pk, m, (r, r^*)) = \mathsf{Sen}(\mathsf{q}, m, r; r^*)$, where r^* is the randomness used in $\mathsf{Sen}(\mathsf{q}, m, r)$.

- **Decryption:**
  If we are in injective mode, then we may define $d(sk, r) = \mathsf{Rec}_r(sk, r)$.

Now, we must show that g, e, d forms a lossy encryption.

- **Correctness on Injective Keys:**
  This follows immediately from the correctness of the oblivious transfer.

- **Indistinguishability of Keys:**
  This follows immediately from the receiver privacy of the oblivious transfer.

- **Lossiness of Lossy Keys:**
  This will follow from the statistical sender privacy of the oblivious transfer.

  If the cryptosystem is in lossy mode, the sender privacy of the OT says that for all m_0, m_1

  $$\{\mathsf{Sen}(\mathsf{q}, m_0, r)\} \approx_s \{\mathsf{Sen}(\mathsf{q}, m_1, r)\},$$

  where the distribution is taken over the internal randomness of Sen.

  Now, if we view the randomness of Sen as an explicit input to Sen (as we do in encryption), then we have that for all m_0, m_1 and r,

  $$\Delta(\mathsf{Sen}(\mathsf{q}, m_0, r; \cdot), \mathsf{Sen}(\mathsf{q}, m_1, r); \cdot) < \nu,$$

where we view the distribution as over the internal randomness of Sen. Applying lemma 2, we have

$$\Delta(\mathsf{Sen}(\mathsf{q}, m_0, \cdot); \cdot), \mathsf{Sen}(\mathsf{q}, m_1, \cdot); \cdot) \leq \nu.$$

where the distribution ranges over the uniform choice of r, and the internal randomness of Sen. But this is exactly what we require.

Lemma 2. Let X, Y, Z be random variables, and suppose

$$\Delta(X, Y|Z = z) < \epsilon,$$

for all z, then $\Delta(X, Y) < \epsilon$.

Proof.

$$\begin{aligned}
\Delta(X, Y) &= \sum_a |\Pr(X = a) - \Pr(Y = a)| \\
&= \sum_a \sum_z |\Pr(X = a, Z = z) - \Pr(Y = a, Z = z)| \\
&= \sum_a \sum_z |\Pr(X = a|Z = z) - \Pr(Y = a|Z = z)| \Pr(z = z) \\
&= \sum_z \Pr(Z = z) \sum_a |\Pr(X = a|Z = z) - \Pr(Y = a|Z = z)| \\
&= \sum_z \Pr(Z = z)\Delta(X, Y|Z = z) \\
&\leq \epsilon \sum_z \Pr(Z = z) \\
&= \epsilon.
\end{aligned}$$

□

Applying the results of [CMO00] which show that Single-Server Private Information Retrieval (PIR) implies Statistically-Hiding OT, we have

Corollary 1. Single-Server PIR implies Lossy-Encryption.

Since Homomorphic Encryption is known to imply PIR [KO97],[Man98],[IKO05] we have

Corollary 2. Homomorphic Encryption implies Lossy-Encryption.

Applying Theorem 1, we have

Corollary 3. Statistically-Hiding 2-round honest-player $\binom{2}{1}$-OT implies IND-SO-ENC secure encryption.
Single-Server PIR implies IND-SO-ENC secure encryption.
Homomorphic Encryption implies IND-SO-ENC secure encryption.

10 A Simulation-Based Definition of Security

While we have focused on an indistinguishability-based definition of security for commitments and encryptions, it is also possible to give a simulation-based definition. Roughly, this says that anything an adversary can learn by playing the Selective Opening game with the challenger can be efficiently simulated by a simulator that sees only I and $(m_i)_{i \in I}$, and never sees the ciphertexts at all. This is called SEM-SO-ENC security. While it appears that the simulation-based definition offers a stronger form of security

than the indistinguishability-based definition, in fact, this remains unknown. It does seem, however, that protocols satisfying SEM-SO-ENC security are harder to construct.

In [BHY09] it was shown that if a lossy encryption scheme has an efficient algorithm opener that can "open" a lossy ciphertext to a desired plaintext, then the scheme is already SEM-SO-ENC secure. Since our constructions of Re-randomizable encryptions give rise to lossy encryptions, to create SEM-SO-ENC security from known assumptions, it suffices to check which re-randomizable encryptions have an efficient opener algorithm.

When we instantiate our encryption scheme with the Paillier Cryptosystem, or the Goldwasser-Micali cryptosystem, the factorization of the modulus N allows us to devise an efficient opening algorithm. That the Goldwasser-Micali scheme is SEM-SO-ENC secure was already recognized in [BHY09], however instantiating our re-randomizable encryption with the Paillier Cryptosystem gives rise to the first SEM-SO-ENC secure cryptosystem under the Decisional Composite Residuosity (DCR) assumption.

10.1 Simulation-Based Security

While we have mostly focused on an indistinguishability-based notion of security under selecting openings, in [BHY09], Hofheinz et al. also formalized a simulation-based notion of security under selective openings. Their simulation-based definition of security intuitively seems stronger than the indistinguishability-based definition, however, it remains unknown whether SEM-SO-ENC implies IND-SO-ENC.

Definition 3. (Semantic Security under selective openings/SEM-SO-ENC).

Let Enc be a Public Key Cryptosystem (PKC), we say that Enc is simulatable under selective openings (SEM-SO-ENC secure) if for every PPT message distribution M, every PPT adversary A, and every PPT relation $\mathcal{R}$, there exists an efficient simulator $S = (S_1, S_2)$ such that we have that

$$\left|\Pr\left[A^{\mathsf{sem\text{-}so\text{-}real}} = 1\right] - \Pr\left[A^{\mathsf{sem\text{-}so\text{-}ideal}} = 1\right]\right| < \nu$$

for some negligible function ν, and where the games sem-so-real and sem-so-ideal are defined as follows

SEM-SO-ENC (Real)

- $(m_1, \ldots, m_n) \leftarrow M$
- $r_1, \ldots, r_n \leftarrow \mathsf{coins}(E)$
- $I \leftarrow A((E(m_1, r_i), \ldots, E(m_n, r_n))$
- $w \leftarrow A(((m_i, r_i))_{i \in I})$
- Output $\mathcal{R}(m, w)$.

SEM-SO-ENC (Ideal)

- $(m_1, \ldots, m_n) \leftarrow M$.
- $(I, st) \leftarrow S_1(1^\lambda)$.
- $w \leftarrow S_2(st, \{m_i\}_{i \in I})$.
- Output $\mathcal{R}(m, w)$.

More explicitly, in the real game,

- The challenger samples messages $(m_1, \ldots, m_n) \leftarrow M$, from the joint message distribution.
- The challenger generates randomness $r_1, \ldots, r_n \leftarrow \mathsf{coins}(E)$.
- The challenger sends $(E(m_1, r_1), \ldots, E(m_n, r_n)$ to A.
- The adversary A responds with a subset $I \subset \{1, \ldots, n\}$, with $|I| = n/2$.
- The challenger reveals *both* m_i and r_i for $i \in I$.

- The adversary outputs a string w.
- The value of the game is $\mathcal{R}(m, w)$.

In the ideal game,

- The challenger samples messages $(m_1, \ldots, m_n) \leftarrow M$, from the joint message distribution.
- Without seeing any encryptions, the simulator chooses a subset I, and some state information st.
- Without seeing any randomness, after seeing the messages $\{m_i\}_{i \in I}$, and the state information, the simulator outputs a string w.
- The value of the game is $\mathcal{R}(m, w)$.

In [BHY09], Hofheinz, Bellare and Yilek, proved that a lossy encryption scheme, with an *efficient* opener procedure are SEM-SO-ENC secure.

Definition 4. A *lossy public-key encryption scheme with efficient opening* is a tuple $(G_{\mathsf{inj}}, G_{\mathsf{lossy}}, E, D)$ satisfying Definition 2, with the additional property that the algorithm opener is efficient, i.e.

- *Openability.* There is an *efficient* algorithm opener, such that if $(pk_{\mathsf{lossy}}, sk_{\mathsf{lossy}}) \leftarrow G_{\mathsf{lossy}}$, and $r \leftarrow \mathsf{coins}(E)$, then for all $x_0, x_1 \in X$ with all but negligible probability, $r' \leftarrow \mathsf{opener}(pk_{\mathsf{lossy}}, E(pk_{\mathsf{lossy}}, x_0, r))$, and $E(pk_{\mathsf{lossy}}, x_1, r')$.

Theorem 2. Lossy Encryption with efficient opening is SEM-SO-ENC secure.

Proof. This is Theorem 2 in [BHY09]. The proof is straightforward, and we only sketch it here.

We proceed in a series of games.

- G_0 is the real SEM-SO-ENC experiment.
- G_1 is the same as G_0, except the adversary is given a lossy public key, instead of a real public key.
- G_2 instead of giving the adversary the real randomness $\{r_i\}_{i \in I}$, the Challenger uses the efficient opener procedure to generate valid randomness.
- G_3 instead of giving the adversary encryptions of m_i, the adversary is given encryptions of a dummy message δ, but the adversary is still given openings to actual messages $\{m_i\}_{i \in I}$ obtained from the opener procedure.

Now, the simulator can simulate G_3 with the adversary. The simulator generates a lossy key pair, and encrypts a sequence of dummy messages and forwards the encryptions to A. The adversary, A, replies with a set I, which S forwards to the challenger. Then S uses the efficient opener procedure to open the selected messages for A. At which point A outputs a string w, and S outputs the same string. Since the outputs of A in G_0 and G_3 are computationally close, the outputs of S, and A in the real and ideal experiments will also be computationally close. □

10.2 Selective Opening Security From the Decisional Composite Residuosity Assumption

Here we give an overview of our construction when applied to the Paillier Cryptosystem (a review of the details of the Paillier Cryptosystem can be found in Appendix C).

By defining $\mathsf{ReRand}(c, r) = c \cdot E(pk, 0, r) \mod N^2$, we obtain IND-SO-ENC secure encryptions through our general construction in 7.1.

It was already known how to build IND-SO-ENC from DCR, since Peikert and Waters [PW08], and Boldyreva, Fehr and O'Neill showed how to build Lossy-Trapdoor Functions from DCR, and Bellare,

Hofheinz and Yilek showed that Lossy-Trapdoor Functions imply IND-SO secure encryptions. We note, however, that our constructions are significantly more efficient than those that follow from [PW08], and somewhat more efficient than those that follow from [BFO08].

While the results of [BHY09] imply that IND-SO-ENC secure encryptions follow from DCR, the question of SEM-SO-ENC secure encryptions was left open, indeed, the only previous construction of SEM-SO-ENC secure encryptions were given in [BHY09] and based on the Quadratic Residuosity Assumption (QR). By instantiating our scheme in 7.1 with the Paillier (or Damgård-Jurik) cryptosystem, we observe that the function opener is efficient, and hence the results of [BHY09] show that the resulting encryption scheme achieves SEM-SO-ENC security.

To see this, recall that $E(pk, m, r) = c^m r^N \mod N^2$, where, in lossy mode, c is an Nth power. Thus, the algorithm opener, on input $e = r_1^N$ and some target message m must find $r' \in \mathbb{Z}/N\mathbb{Z}$ such that $c^m (r')^N = e$. If we write $c = r_0^N$, then opener must find a solution to

$$(r')^N = \left(\frac{r_1}{r_0^m}\right)^N.$$

So the efficiency of opener reduces to the efficiency of taking Nth roots modulo N^2. But this is easily done if the factorization of N is known, since we can set $d = N^{-1} \mod \phi(N)$, and then taking Nth roots, is equivalent to exponentiating modulo N, i.e.

$$(r^N)^d = r^{Nd} = r \mod N.$$

Thus we immediately get a SEM-SO-COM secure encryption protocol from the DCR assumption. Thus we arrive at

Corollary 4. Under the Decisional Composite Residuosity assumption (DCR), the system described in §7.1 is SEM-SO-ENC secure.

Since the Paillier cryptosystem (and the Damgård-Jurik extension), have smaller ciphertext expansion than the Goldwasser-Micali cryptosystem (which only encrypts bits), we arrive at a more efficient system than the only known SEM-SO-ENC secure cryptosystem.

11 Chosen Ciphertext Security

11.1 Definitions

It has long been recognized that if an adversary is given access to a decryption oracle, many cryptosystems may become insecure. The notion of Chosen-Ciphertext Security ([NY90],[RS91],[DDN91]) was created to address this issue, and since then there have been many schemes that achieve this level of security. The attacks of Bleichenbacher on RSA PKCS#1 [Ble98] emphasized the practical importance of security against Chosen-Ciphertext Attacks (CCA).

The need for Selective Opening Security was first recognized in the setting of Multi-Party Computation (MPC), where an active MPC adversary can view all the ciphertexts sent in a current round, and then choose a subset of senders to corrupt. It is natural to imagine an MPC adversary, who, in addition to corrupting a subset of senders, can also mount a Chosen-Ciphertext Attack against the receiver. It is easy to see that the schemes proposed (based on re-randomizable, or homomorphic encryption) become trivially insecure in this setting.

In this section, we show how to extend the notion of a Chosen Ciphertext Attack to the selective opening setting. As in the standard Selective-Opening setting, we can define security in two different ways, either by indistinguishability, or by simulatability. We will give definitions of security as well as constructions for both settings.

11.2 Chosen Ciphertext Security: Indistinguishability

We begin with the indistinguishability-based definition. We define two games, a real game (ind-cca2-real) and an ideal game (ind-cca2-ideal). In both games, the challenger runs the key-generation algorithm to generate a public-key secret-key pair, and sends the public-key to the adversary. The adversary is then allowed to adaptively make two types of queries.

- **Selective Opening Query:** The adversary A chooses a message distribution M, and sends a description of M to the challenger. The challenger samples $(m_1, \ldots, m_n) \leftarrow M$, and generates

$$(c_1, \ldots, c_n) = (E(pk, m_1, r_1), \ldots, E(pk, m_n, r_n)).$$

  The challenger sends $(c_1, \ldots, c_n)$ to the adversary, and the adversary chooses a subset $I \subset [n]$, with $|I| = n/2$, and sends I to the challenger. The challenger then sends $\{(m_i, r_i)\}_{i \in I}$ to the Adversary. We call the ciphertexts $c_1, \ldots, c_n$ *target ciphertexts*.

  - In the real game, the challenger then sends $\{m_j\}_{j \notin I}$ to the adversary.
  - In the ideal game, the challenger resamples $(m'_1, \ldots, m'_n) \leftarrow M|_{M_I}$, and sends $\{m'_j\}_{j \notin I}$ to the adversary.

- **Decryption Queries:** The adversary A chooses a ciphertext c that has never appeared as a target ciphertext, and sends c to the challenger. If c is a valid ciphertext (i.e. $D(c) \neq \perp$) then the challenger responds with $m = D(c)$.

After adaptively making polynomially many queries, with at most one of them being a Selective Opening Query, the adversary outputs a bit b.

Definition 5. (IND-SO-CCA2) A public key encryption scheme E is called IND-SO-CCA2 secure, if, for all PPT adversaries A, A's output in the real game is negligibly different from its output in the ideal game, i.e.

$$\left|\Pr[A^{\mathsf{ind\text{-}cca2\text{-}real}} = 1] - \Pr[A^{\mathsf{ind\text{-}cca2\text{-}ideal}} = 1]\right| < \nu.$$

For some negligible function ν.

We remark that if the adversary is not allowed to make decryption queries, this reduces to IND-SO-ENC security.

Our construction of an IND-SO-CCA2 secure cryptosystem requires some basic tools, outlined below.

11.3 Strongly Unforgeable Signatures

A signature scheme is a triple of PPT algorithms $(\mathsf{G}, \mathsf{Sign}, \mathsf{Ver})$ such that

- The algorithm G takes a security parameter λ, and returns a verification key and a signing key.

$$(vk, sk) \leftarrow \mathsf{G}(1^\lambda).$$

- The algorithm Sign takes a message m and the signing key, and produces a signature sig.

$$\mathsf{sig} \leftarrow \mathsf{Sign}(m, sk).$$

- The algorithm Ver takes a verification key, a message, and a signature, and returns a bit b.

$$b \leftarrow \mathsf{Ver}(vk, m, \mathsf{sig}).$$

We require

- **Completeness:** For all m

$$\Pr\left[(vk, sk) \leftarrow \mathsf{G}; \mathsf{sig} \leftarrow \mathsf{Sign}(m, sk); \mathsf{Ver}(vk, m, \mathsf{sig}) = 1\right] = 1.$$

- **Strongly Unforgeable:** For all PPT adversaries A,

$$\Pr\left[(vk, sk) \leftarrow \mathsf{G};\ (m, \mathsf{sig}') \leftarrow A^{\mathsf{Sign}(\cdot, sk)}(vk) : \right.$$
$$\left. \mathsf{Ver}(vk, m, \mathsf{sig}') = 1 \text{ and } \mathsf{sig}' \text{ was never the output of } \mathsf{Sign}(\cdot, sk)\right]$$

  If we restrict A to make at most one oracle query to $\mathsf{Sign}(\cdot, sk)$ we say that $(\mathsf{G}, \mathsf{Sign}, \mathsf{Ver})$ is a one-time strongly unforgeable signature scheme.

11.4 Unduplicatable Set Selection

Unduplicatable set selection was used implicitly in [NY90] and [CIO98], and formalized in [Sah99]. The description below is essentially that of [Sah99].

The goal of unduplicatable set selection is to create a mapping from $\mathfrak{g} : \{0,1\}^k \to B$ such that for all distinct $a^1, \ldots, a^n, a^{n+1} \in \{0,1\}^k$,

$$\mathfrak{g}(a^{n+1}) \not\subset \bigcup_{i=1}^{n} \mathfrak{g}(a^i).$$

In [Sah99], Sahai gives a simple general construction based on polynomials which we recall here. Let $\ell = 2^{\lceil \log_2 2nk \rceil}$, so $\ell > 2nk$, and let $Y = \mathbb{F}_\ell \times \mathbb{F}_\ell$, and $B \subset \mathcal{P}(Y)$. To each $a \in \{0,1\}^k$ we may associate a polynomial

$$f_a(x) = a_0 + a_1 x + \cdots a_{k-1} x^{k-1} \in \mathbb{F}_\ell[x].$$

Then if we set

$$\mathfrak{g}(a) = \{(t, f_a(t)) : t \in \mathbb{F}_\ell\} \subset Y.$$

Now, $|\mathfrak{g}(a)| = \ell$, and if $a \neq a'$, we have $|\mathfrak{g}(a) \cap \mathfrak{g}(a')| \leq k - 1$. Thus

$$\begin{aligned}
\left|\mathfrak{g}(a^{n+1}) \setminus \bigcup_{i=1}^{n} \mathfrak{g}(a^i)\right| &= \left|\mathfrak{g}(a^{n+1}) \setminus \bigcup_{i=1}^{n} \mathfrak{g}(a^{n+1}) \cap \mathfrak{g}(a^i)\right| \\
&\geq \left|\mathfrak{g}(a^{n+1})\right| - \sum_{i=1}^{n} \left|\mathfrak{g}(a^{n+1}) \cap \mathfrak{g}(a^i)\right| \\
&\geq \ell - n(k-1) \\
&\geq \frac{\ell}{2}.
\end{aligned}$$

We call $\mathfrak{g}$ an (n, k, ℓ) unduplicatable set selector.

11.5 Lossy Trapdoor Functions

Lossy Trapdoor Functions were first defined in [PW08], and we review the definition here.

A tuple $(S_{\mathrm{LTDF}}, F, F^{-1})$ of PPT algorithms is called a family of (d, k)-Lossy Trapdoor Functions if the following properties hold:

- **Sampling Injective Functions:** $S_{\mathrm{LTDF}}(1^\lambda, 1)$ outputs s, t where s is a function index, and t its trapdoor. We require that $F(s, \cdot)$ is an injective deterministic function on $\{0,1\}^d$, and $F^{-1}(t, F(s, x)) = x$ for all x.
- **Sampling Lossy Functions:** $S_{\mathrm{LTDF}}(1^\lambda, 0)$ outputs $(s, \perp)$ where s is a function index and $F(s, \cdot)$ is a function on $\{0,1\}^d$, where the image of $F(s, \cdot)$ has size at most 2^{d-k}.

- **Indistinguishability:** The first outputs of $S_{\text{LTDF}}(1^\lambda, 0)$ and $S_{\text{LTDF}}(1^\lambda, 1)$ are computationally indistinguishable.

Along with Lossy Trapdoor Functions, we can define All But One (ABO) Functions. Essentially, these are lossy trapdoor functions, except instead of having two branches (a lossy branch and an injective branch) they have many branches, all but one of which are injective. A tuple $(S_{\text{ABO}}, G, G^{-1})$ of PPT algorithms is called a family of (d, k)-ABO Functions if the following properties hold:

- **Sampling with a given Lossy Branch:** For $b^* \in \mathcal{B}$, $S_{\text{ABO}}(1^\lambda, b^*)$ outputs s, t where s is a function index, and t its trapdoor. We require that for any $b \neq b^*$, $G(s, b, \cdot)$ is an injective deterministic function on $\{0,1\}^d$, and $G^{-1}(t, b, G(s, b, x)) = x$ for all x.

  Additionally, the image $G(s, b^*, \cdot)$ has size at most 2^{d-k}.

- **Hidden Lossy Branch:** For any $b_0^*, b_1^* \in \mathcal{B}$, the first outputs of $S_{\text{ABO}}(1^\lambda, b_0^*)$ and $S_{\text{ABO}}(1^\lambda, b_1^*)$ are computationally indistinguishable.

11.6 All-But-n Functions

In the original Peikert Waters construction, they require an All-But-One family of functions, so that the single challenge ciphertext in the CCA2 game can be evaluated on the lossy branch. Since the IND-SO-CCA security game has n challenge ciphertexts, and we generalize the Peikert-Waters construction to an All-But-n (ABN) Functions, where all the branches except the specified ones are injective.

- **Sampling with a given Lossy Set:** For $I \subset \mathcal{B}$, $S_{\text{ABN}}(1^\lambda, I)$ outputs s, t where s is a function index, and t its trapdoor. We require that for any $b \in \mathcal{B} \setminus I$, $G(s, b, \cdot)$ is an injective deterministic function on $\{0,1\}^d$, and $G^{-1}(t, b, G(s, b, x)) = x$ for all x.

  Additionally, for $b \in I$, the image $G(s, b, \cdot)$ has size at most 2^{d-k}.

- **Hidden Lossy Branches:** For any $b_0^*, b_1^* \in \mathcal{B}$, the first outputs of $S_{\text{ABN}}(1^\lambda, b_0^*)$ and $S_{\text{ABN}}(1^\lambda, b_1^*)$ are computationally indistinguishable.

We can construct a family of ABM Functions given any sufficiently lossy family of LTDFs as follows. Given a set $I \subset \mathcal{B}$, we create an unduplicatable set selector $\mathfrak{g} : \mathcal{B} \to \hat{\mathcal{B}}$. For each $\hat{b} \in \hat{\mathcal{B}}$, we will associate a Lossy Trapdoor Function. Let $\hat{I} = \bigcup_{i \in I} \mathfrak{g}(i)$. For each $\hat{\imath} \in \hat{I}$, we will set create a LTDF in lossy mode, and for each $\hat{b} \in \hat{\mathcal{B}} \setminus \hat{I}$, we will associate a LTDF in injective mode.

- **Sampling with a given Lossy Set:** Create an $(n, \lceil \log |\mathcal{B}| \rceil)$ unduplicatable set selector $\mathfrak{g}$. Suppose $\mathcal{B} \subset \{0,1\}^v$, then the construction outlined above produces $\mathfrak{g}$ which maps $\{0,1\}^v$ to subsets of $\mathbb{F}_\ell \times \mathbb{F}_\ell$, where $\ell = 2^{\lceil \log_2 2nv \rceil}$. For each element in $\mathbb{F}_\ell \times \mathbb{F}_\ell$, we will associate a Lossy Trapdoor Function. Let $\hat{I} = \bigcup_{i \in I} \mathfrak{g}(i) \subset \mathbb{F}_\ell \times \mathbb{F}_\ell$. For each $y \in \hat{I}$ let F_y be an LTDF in lossy mode, and for each $y \in \mathbb{F}_\ell \times \mathbb{F}_\ell \setminus \hat{I}$, let F_y be an LTDF in injective mode.

  Now, define $G(b, x) = (F_{y_1}(x), \ldots, F_{y_\ell}(x))_{y_i \in \mathfrak{g}(b)}$.

Notice that if any of the functions F_y are injective, then G is also injective, and if the image size of F in lossy mode is 2^r, then the images size of G on a lossy branch is $2^{r\ell}$. Finally, we notice that the lossy set is hidden by the indistinguishability of modes of the LTDF.

11.7 IND-SO-CCA Construction

We now give a method for constructing IND-SO-CCA secure encryption from LTDFs. The construction (and proof) mimics that of [PW08].

Let $(S_{\text{LTDF}}, F, F^{-1})$ be a family of (d, k) Lossy-Trapdoor Functions, and let $(S_{\text{ABN}}, G, G^{-1})$ be a family of (d, k') all-but-n functions with branch set $\{0,1\}^v$ where v is the length of a verification key for our one-time signature scheme. We require that $2d - k - k' \leq t - \kappa$, for $\kappa = \kappa(t) = \omega(\log t)$. Let $\mathcal{H}$ be a pairwise independent hash family from $\{0,1\}^d \to \{0,1\}^\ell$, with $0 < \ell < \kappa - 2\log(1/\nu)$, for some negligible $\nu = \nu(\lambda)$. The message space will $\{0,1\}^\ell$.

- **KeyGen:**
  Generate a hash function $h \leftarrow \mathcal{H}$, and

  $$(s,t) \leftarrow S_{\text{LTDF}}(1^\lambda, inj), \quad (s',t') \leftarrow S_{\text{ABN}}(1^\lambda, \{0,1,\ldots,n-1\}).$$

  The public key will be (s, s', h) and the secret key will be (t, t').

- **Encryption:**
  Generate $r \leftarrow \{0,1\}^t$, and $r^{sig} \leftarrow \mathsf{coins}(\mathsf{G})$, and generate keys for a one-time signature using randomness r^{sig}

  $$(vk, sk) = \mathsf{G}(r^{sig}).$$

  For a message m, output the ciphertext

  $$c = (vk, F(s,x), G(s,vk,x), h(x) \oplus m, \mathsf{sig}),$$

  where $\mathsf{sig} = \mathsf{Sign}(sk, F(s,x), G(s,vk,x), h(x) \oplus m)$.

- **Decryption:**
  Given a ciphertext $c = (vk, c_1, c_2, c_3, \mathsf{sig})$, first check that $\mathsf{Ver}(vk, (c_1, c_2, c_3), \mathsf{sig}) = 1$, otherwise output $\perp$.
  Then let $x = F^{-1}(t, c_1)$, and check whether $G(s, vk, x) = c_2$, if not, output $\perp$.

  Finally, output $m = c_3 \oplus h(x)$.

Theorem 3. The algorithms described above form an IND-SO-CCA2 secure cryptosystem.

Proof. The correctness of the scheme is clear, so we focus on the security.

We prove security through a sequence of games.

Let Game_0 be the real IND-SO-CCA2 game.

Let Game_1 be identical to Game_0 except that we generate $\{(vk_i, sk_i)\}_{i=1}^n$, to be used in the challenge query during key-generation.

Let Game_2 be identical to Game_1 except we modify the decryption algorithm to output $\perp$ on ciphertexts of the form $c = (vk, c_1, c_2, c_3, \mathsf{sig})$, if $vk \in \{vk_i\}_{i=1}^n$.

Let Game_3 be identical to Game_2 except that we set the lossy branches of the All-But-n function G to be $\{vk_1, \ldots, vk_n\}$.

Let Game_4 be identical to Game_3 except that in the decryption algorithm we use G^{-1} to decrypt instead of F^{-1}, i.e. we set $x = G^{-1}(t', vk, c_2)$ instead of $x = F^{-1}(t, c_1)$.

Let Game_5 be identical to Game_4 except that we replace the injective function with a lossy one, i.e. during key-generation we generate $(s, \perp) \leftarrow S_{\text{LTDF}}(1^\lambda, lossy)$, instead of $(s,t) \leftarrow S_{\text{LTDF}}(1^\lambda, inj)$.

- Clearly the adversary's view in Game_0 and Game_1 are identical.

- The only way the adversary's view can be different in Game_2 than in Game_1 is if the adversary successfully generates a signature using one of keys in the set $\{vk_i\}$. But this can only happen with negligible probability by the strong unforgeability of the signature scheme.
- The indistinguishability of Game_2 and Game_3 follows from the indistinguishability of All-But-n functions with different lossy branches.
- The adversary's views in Game_3 and Game_4 are identical, because the adversary can never make a decryption query on a lossy-branch of G.
- The indistinguishability of Game_4 and Game_5 follows from the indistinguishability of modes of Lossy-Trapdoor Functions.

Now, if we can show that an adversary's probability of success in Game_5 is negligible we will be done. To do this, we follow the the proof that Lossy Encryption is Selective Opening secure. See [BHY09] or Theorem 5. The key observation is that in Game_5 the challenge ciphertexts are *statistically* independent of the underlying messages. We begin by showing that this is, in fact, the case.

Now, $F(s,\cdot)$ and $G(s', vk_i, \cdot)$ are lossy functions with image sizes at most 2^{d-k} and $2^{d-k'}$ respectively for each $i \in [n]$. Thus the function $x \mapsto (F(s,x), G(s', vk_i, x))$ takes on at most $2^{2d-k-k'} \leq 2^{d-\kappa}$ values. Now by Lemma 2.1 of [PW08], the average min-entropy is bounded below

$$\tilde{H}_\infty(x|c_1, c_2, s, s') \geq H_\infty(x|s, s') - (d-\kappa) = t - (d-\kappa) = \kappa.$$

Since $\ell \leq \kappa - 2\log(1/\nu)$, by Lemma 2.2 of [PW08], we have

$$\Delta((c_1, c_2, h, h(x)), (c_1, c_2, h, r')) \leq \nu.$$

Now, we can incorporate the ideas of Theorem 5. Since the challenge ciphertexts are statistically independent of the underlying plaintexts, there is a (possibly inefficient)[2] algorithm opener, which, given (vk, c_1, c_2, c_3, m) outputs x such that $F(s,x) = c_1$, $G(s, vk, x) = c_2$, and $h(x) \oplus m = c_3$. If no such x exists, opener outputs $\perp$ (the statistical closeness guarantees that this happens with probability at most ν).

Now, let us imagine a new series of games.

Let Game_{5_0} be identical to Game_5, except that the challenge ciphertexts are opened using the output of opener, instead of the actual randomness used by the challenger.

Now, for $j \in [n]$, let Game_{5_j} be identical to Game_{5_0} except that for $i \leq j$, the challenge ciphertexts will be

$$(E(pk, \delta, r_1), \ldots, E(pk, \delta, r_j), E(pk, m_{j+1}, r_{j+1}), \ldots, E(pk, m_n, r_n))$$

So, the only difference between the Game_{5_j} and $\text{Game}_{5_{j-1}}$ is whether the jth encryption in the challenger ciphertext is an encryption of δ or m_j. Since these two distributions are *statistically* close, even an *unbounded* adversary has a negligible chance of distinguishing them. Thus by the triangle inequality, an unbounded adversary has a negligible probability of distinguishing Game_{5_0} from Game_{5_n}.

But Game_{5_n} is identical in both the ind-cca2-real and ind-cca2-ideal games, so an adversary has at most a negligible probability of distinguishing the two worlds.

□

12 Chosen Ciphertext Security: Simulatability

- **Selective Opening Query:** The adversary A chooses a message distribution M, and sends a description of M to the challenger. The challenger samples $(m_1, \ldots, m_n) \leftarrow M$, and generates

$$(c_1, \ldots, c_n) = (E(pk, m_1, r_1), \ldots, E(pk, m_n, r_n)).$$

[2]The algorithm opener is inefficient for the DDH construction in [PW08], but is efficient for the DCR construction in [BFO08]. When the algorithme opener is efficient this construction will achieve SEM-SO-CCA2 security.

The challenger sends $(c_1, \ldots, c_n)$ to the adversary, and the adversary chooses a subset $I \subset [n]$, with $|I| = n/2$, and sends I to the challenger. The challenger then sends $\{(m_i, r_i)\}_{i \in I}$ to the Adversary. We call the ciphertexts $c_1, \ldots, c_n$ *target ciphertexts*.

The challenger then sends $\{m_j\}_{j \notin I}$ to the adversary.

- **Decryption Queries:** The adversary A chooses a ciphertext c that has never appeared as a target ciphertext, and sends c to the challenger. If c is a valid ciphertext (i.e. $D(c) \neq \perp$) then the challenger responds with $m = D(c)$.

After adaptively making polynomially many queries, with at most one of them being a Selective Opening Query, the adversary outputs w, and the value of the game is $\mathcal{R}(m, w)$.

In the ideal game, the challenger samples $(m_1, \ldots, m_n) \leftarrow M$.

- The simulator chooses a subset, $I \leftarrow S_1$.
- The simulator views the chosen messages and outputs a w, $w \leftarrow S_2(\{m_i\}_{i \in I})$.

The value of the game is $\mathcal{R}(m, w)$.

Definition 6. (SEM-SO-CCA2) A public key encryption scheme E is called SEM-SO-CCA2 secure, if, for all PPT adversaries A, and all PPT relations $\mathcal{R}$, there exists a simulator $S = (S_1, S_2)$ such that the values of the real and ideal games are identical with all but negligible probability, i.e.

$$\Pr[\mathsf{sem\text{-}cca2\text{-}real} \neq \mathsf{sem\text{-}cca2\text{-}ideal}] \leq \nu.$$

For some negligible function ν.

We remark that if the adversary is not allowed to make decryption queries, this reduces to SEM-SO-ENC security.

12.1 Non-Interactive Zero Knowledge

The most successful technique in constructing systems secure against an adaptive chosen ciphertext attack has been the Naor-Yung paradigm [NY90]. Roughly, the idea is to encrypt the message twice and include a Non-Interactive Zero Knowledge (NIZK) proof that both encryptions encrypt the same plaintext. The proof of security then uses the simulator for the NIZK to simulate the proof for the challenge ciphertext. This method has since been refined in [DDN91],[Sah99],[SCO$^+$01], and [Lin06] (among others).

Our constructions of IND-SO-CCA2 encryption follow the general structure of the Naor-Yung paradigm [NY90], however, the selective opening of the encryption query poses new challenges. In particular, if we naïvely try to follow the Naor-Yung paradigm, we immediately encounter difficulties because our challenger must reveal the messages and randomness for half of the ciphertexts in the challenge. This will immediately reveal to the adversary that the proofs were simulated. It requires new ideas to overcome this difficulty.

We now give a brief definition of the properties of a Non-Interactive Zero Knowledge Proof of Knowledge with Honest Prover State-Reconstruction (originally defined and constructed in [GOS06]).

Let $\mathcal{R}$ be an efficiently computable binary relation and let $L = \{x : \exists w \text{ such that } (x, w) \in \mathcal{R}\}$. We refer to L as a language, x as a statement, and w as a witness.

A Non-Interactive Proof System for L is a triple of PPT algorithms $(\mathsf{CRSgen}, \mathsf{Prover}, \mathsf{Verifier})$ such that

- $\sigma \leftarrow \mathsf{CRSgen}(1^\lambda)$.
  Generates a Common Reference String.
- $\pi \leftarrow \mathsf{Prover}(CRS, x, w)$.
  On inputs x, and a witness w for x, such that $\mathcal{R}(x, w) = 1$, the Prover outputs a proof π.

- $b \leftarrow \mathsf{Verifier}(CRS, x, p)$.
  On inputs x and a purported proof π, $\mathsf{Verifier}$ outputs a bit b.

Definition 7. A triple of algorithms is called a Non-Interactive Zero Knowledge Proof of Knowledge if

- **Completeness:** For all adversaries A,

  $$\Pr\left[\sigma \leftarrow \mathsf{CRSgen}(1^\lambda); (x,w) \leftarrow A(\sigma); \pi \leftarrow \mathsf{Prover}(\sigma,x,w) : \mathsf{Verifier}(\sigma,x,\pi) = 1 \text{ if } (x,w) \in R\right] > 1-\nu,$$

  For some negligible function ν.

- **Soundness:** For all adversaries A,

  $$\Pr\left[\sigma \leftarrow \mathsf{CRSgen}(1^\lambda); (x,\pi) \leftarrow A(\sigma) : \mathsf{Verifier}(\sigma,x,\pi) = 0 \text{ if } x \notin L\right] > 1-\nu.$$

- **Knowledge Extraction:** There exists an extractor $\mathsf{Ext} = (\mathsf{Ext}_1, \mathsf{Ext}_2)$ such that for all adversaries A

  $$\left|\Pr\left[\sigma \leftarrow \mathsf{CRSgen}(1^\lambda) : A(\sigma) = 1\right] - \Pr\left[(\sigma,\tau) \leftarrow \mathsf{Ext}_1(1^\lambda) : A(\sigma) = 1\right]\right| < \nu$$

  and

  $$\Pr\left[(\sigma,\tau) \leftarrow \mathsf{Ext}_1(1^\lambda); (x,\pi) \leftarrow A(\sigma); w \leftarrow \mathsf{Ext}_2(\sigma,\tau,x,\pi) : \mathsf{Verifier}(\sigma,x,\pi) = 0 \text{ or } (x,w) \in \mathcal{R}\right] > 1-\nu$$

  For some negligible function ν.

- **Zero-Knowledge:** There exists a simulator $S = (S_1, S_2)$, such that for all adversaries A,

  $$\left|\Pr\left[\sigma \leftarrow \mathsf{CRSgen}(1^\lambda) : A^{P(\sigma,\cdot,\cdot)}(\sigma) = 1\right] - \Pr\left[(\sigma,\tau) \leftarrow S_1(1^\lambda) : A^{S'(\sigma,\tau,\cdot,\cdot)}(\sigma) = 1\right]\right| < \nu,$$

  where S' is defined

  $$S' = \begin{cases} S_2(\sigma,\tau,x) & \text{if } (x,w) \in \mathcal{R}, \\ \perp & \text{otherwise.} \end{cases}$$

- **Honest-Prover State Reconstruction:** There exists a simulator $S = (S_1, S_2, S_2)$ such that for all adversaries A

  $$\left|\Pr\left[\sigma \leftarrow \mathsf{CRSgen}(1^\lambda); A^{PR(\sigma,\cdot,\cdot)}(\sigma) = 1\right] - \Pr\left[(\sigma,\tau) \leftarrow S_1(1^\lambda) : A^{SR(\sigma,\tau,\cdot,\cdot)}(\sigma) = 1\right]\right| < \nu,$$

  where $PR(\sigma,x,w)$ samples $r \leftarrow \mathsf{coins}(\mathsf{Prover})$, and sets $\pi = \mathsf{Prover}(\sigma,x,w,r)$, and returns (π, r), and where SR samples $r^* \leftarrow \mathsf{coins}(S_2)$, and sets $\pi' = S_2(\sigma,\tau,x,r^*)$, finally PR sets $r' \leftarrow S_3(\sigma,\tau,x,w,r^*)$ and returns (π', r'). Both oracles output $\perp$ if $(x,w) \notin \mathcal{R}$.

12.2 A SEM-SO-CCA2 Construction

Along with the NIZK Proofs of Knowledge with Honest Prover State Reconstruction, our construction relies on a number of common cryptographic tools. We will also require a strongly unforgeable one-time signature scheme. In our game, a single encryption query is actually n separate encryptions, so we will require an unduplicatable set selector $\mathfrak{g}$ for sets of size n. Finally, we will require an IND-SO-ENC secure cryptosystem.

Let (G_{so}, E, D) be an efficiently openable lossy cryptosystem (by Theorem 2 it will also be SEM-SO-ENC secure). Let $(\mathsf{G}, \mathsf{Sign}, \mathsf{Ver})$ be a strongly unforgeable one-time signature scheme. Let $\mathfrak{g}$ be an (n, λ) unduplicatable set selector, and let $\ell = |g(0^\lambda)|$, and $L = \mathfrak{g}(\{0,1\}^\lambda)$.

Let $(\mathsf{CRSgen}, \mathsf{Prover}, \mathsf{Verifier})$ be a Noninteractive Zero Knowledge Proof of Knowledge with Honest Prover State Reconstruction for the language given by the relation $((e_0,e_1),(m,r_0,r_1)) \in \mathcal{R}$ if $e_0 = E(m,r_0)$ and $e_1 = E(m,r_1)$.

Our scheme works as follows

- **KeyGen:**

$$(pk_0, sk_0) \leftarrow G_{so}(1^\lambda),\ (pk_1, sk_1) \leftarrow G_{so}(1^\lambda),\ \text{and } (\sigma_i, \tau_i) \leftarrow \mathsf{Ext}_1(1^\lambda) \text{ for } i \in L$$

Set

$$pk = (pk_0, pk_1, \{\sigma_i\}_{i\in L}) \quad \text{and} \quad sk = (sk_0, sk_1, \{\tau_i\}_{i\in L}).$$

- **Encryption:** Pick

$$r^{sig} \leftarrow \mathsf{coins}(\mathsf{Sign}),\ r_0 \leftarrow \mathsf{coins}(E),\ r_1 \leftarrow \mathsf{coins}(E),\ r_i^{nizk} \leftarrow \mathsf{coins}(\mathsf{Prover}) \text{ for } i = 1, \ldots, \ell.$$

Generate keys for a one-time signature using randomness r^{sig}.

$$(vk, sk) = \mathsf{G}(r^{sig}).$$

For a message m, calculate

$$e_0 = E(pk_0, m, r_0), \quad e_1 = E(pk_1, m, r_1)$$

set $w = (m, r_0, r_1)$.

$$\overline{\pi} = (\pi_1, \ldots, \pi_\ell) = (\mathsf{Prover}(\sigma_i, (e_0, e_1), w))_{i\in\mathfrak{g}(vk)}$$

using randomness r_i^{nizk} in the ith iteration of Prover. Set

$$\mathsf{sig} = \mathsf{Sign}(e_0, e_1, \overline{\pi}),$$

output the ciphertext

$$c = (vk, e_0, e_1, \overline{\pi}, \mathsf{sig})$$

- **Decryption:** Given a ciphertext

$$c = (vk, e_0, e_1, \overline{\pi}, \mathsf{sig})$$

Check that

$$\mathsf{Ver}(vk, (e_0, e_1, \overline{\pi})) = 1,$$

otherwise return $\perp$. For $i \in \mathfrak{g}(vk)$, check that

$$\mathsf{Verifier}(\sigma_i, (e_0, e_1), \pi_i) = 1,$$

otherwise return $\perp$.

Pick a random $i \in \mathfrak{g}(vk)$ and use the Extractor Ext_2 to recover the witness (m, r_0, r_1), i.e.

$$(m, r_0, r_1) \leftarrow \mathsf{Ext}_2(\sigma_i, \tau_i, (e_0, e_1), \pi_i)$$

return m.

Theorem 4. This scheme is SEM-SO-CCA2 secure.

Proof. We will show how to use an adversary in the sem-cca2-real game to construct a simulator for the sem-cca2-ideal game.

To do this, we begin by considering a series of games

- Game_0: This is the sem-cca2-real game.

- Game$_1$: Pick the verification key $((vk^{chal,1}, sk^{chal,1}), \ldots, (vk^{chal,n}, sk^{chal,n}))$ to be used in the challenge ciphertexts during parameter generation.
- Game$_2$: Generate the Common Reference Strings by

$$\sigma_i = \begin{cases} S_1(1^\lambda) & \text{if } i \in \mathfrak{g}(vk^{chal,j}) \text{ for some } j \in [n] \\ \mathsf{Ext}_1(1^\lambda) & \text{otherwise.} \end{cases}$$

  In decryption, we now use $i \notin \mathfrak{g}(vk)$ to recover (m, r_0, r_1).
- Game$_3$: When generating the target ciphertexts, ignore the witness and generate the "proof"

$$\overline{\pi} = \{\pi_i\}_{i \in \mathfrak{g}(vk)} = \{S_2(\sigma_i, \tau_i(e_0, e_1), r_i^*)\}_{i \in \mathfrak{g}(vk)}$$

  when the adversary asks for the decryption and randomness of a subset of the target ciphertexts, use the State Reconstructor to generate

$$r_i \leftarrow S_3(\sigma_i, \tau_i, (e_0, e_1), (m, r_0, r_1, r_i^*)),$$

  and return these r_i instead of the r_i^* that were actually used.
- Game$_4$: When generating the target ciphertexts, generate them all as the encryption of a dummy message δ and when the adversary asks for the decryption and randomness of a subset of the target ciphertexts, use the efficient openability of (G_{so}, E, D) to generate $\{r_i\}_{i \in I}$. Then proceed as in Game$_3$.

Let W_i be the distribution of the adversary's output in game i. Clearly $W_0 = W_1$, since from the adversary's point of view they are identical. To show that W_1 and W_2 are only negligibly different, notice that by the strong unforgeability of $(\mathsf{G}, \mathsf{Sign}, \mathsf{Ver})$, the adversary can never ask for the decryption of a ciphertext signed with vk, so by the unduplicatability of $\mathfrak{g}$, there will always be at least one valid proof generated with an extractable CRS. Now, it's easy to see that any PPT adversary that can distinguish between Game 2 and Game 1 can be used to distinguish the CRS generated by the extraction simulator Ext_1, and Honest Prover Reconstruction simulator S_1 (really $n\ell$ such simulators), but if

$$\left|\Pr\left[(\sigma, \tau) \leftarrow S_1(1^\lambda) : A^{SR(\sigma,\tau,\cdot,\cdot)}(\sigma) = 1\right] - \Pr\left[(\sigma, \tau) \leftarrow S_1(1^\lambda) : A^{S'(\sigma,\tau,\cdot,\cdot)}(\sigma) = 1\right]\right| > \epsilon,$$

the either

$$\left|\Pr\left[\sigma \leftarrow \mathsf{CRSgen}(1^\lambda) : A^{P(\sigma,\cdot,\cdot)}(\sigma) = 1\right] - \Pr\left[(\sigma, \tau) \leftarrow S_1(1^\lambda) : A^{S'(\sigma,\tau,\cdot,\cdot)}(\sigma) = 1\right]\right| > \frac{\epsilon}{2},$$

or

$$\left|\Pr\left[\sigma \leftarrow \mathsf{CRSgen}(1^\lambda); A^{PR(\sigma,\cdot,\cdot)}(\sigma) = 1\right] - \Pr\left[(\sigma, \tau) \leftarrow S_1(1^\lambda) : A^{SR(\sigma,\tau,\cdot,\cdot)}(\sigma) = 1\right]\right| > \frac{\epsilon}{2}.$$

Since these are both negligible by the definition of our NIZK, the difference between W_1 and W_2 is negligible.

To see that the difference between W_2 and W_3 is negligible, we notice that we can immediately transform an adversary that distinguishes Game$_2$ from Game$_3$ into an adversary that breaks the indistinguishability of the Honest Prover State Reconstruction simulator, losing a factor of $n\ell$ (because we are making $n\ell$ comparisons).

Thus we have shown that the value of the Game$_0$ run against an efficient adversary A, will be computationally indistinguishable from the value of Game$_3$ when run against A. Now, we show how to use Game$_3$ to build a simulator for the sem-cca2-ideal game.

Specifically the simulator, will run Game$_4$ with A. When A asks for a subset I, the simulator will ask for openings of the same subset I. Using the received messages $\{m_i\}_{i \in I}$, the simulator will run the efficient opening procedure of (G_{so}, E, D), to generate $\{r_i\}_{i \in I}$. The simulator then proceeds as in Game$_4$,

i.e. the simulator uses the State Reconstructor to generate randomness for the proofs using the witnesses $\{(m_i, r_i)\}_{i \in I}$, and answering further decryption queries as in Game_3. Finally, when A outputs w, the simulator will output the same w. Since the output of A in Game_4 is indistinguishable from the output of A in the sem-cca2-real game, the output of the simulator will be indistinguishable from the output of A in the sem-cca2-real game.

□

A similar argument shows that this construction will be IND-SO-CCA2 if the underlying encryption scheme is IND-SO-ENC instead of SEM-SO-ENC secure.

13 Conclusion

We have shown that re-randomizable encryption implies IND-SO-ENC secure encryptions. In the process we have shown that re-randomizable encryption implies Lossy Encryption, which is interesting in its own right. These constructions are relatively simple and retain the efficiency of the underlying re-randomizable encryption protocol. Our constructions can be applied to known cryptosystems, and immediately yields simple and efficient IND-SO-COM secure commitments and IND-SO-ENC secure encryptions from the Decisional Diffie-Hellman (DDH), Decisional-Composite Residuosity (DCR) and Quadratic Residuosity (QR) assumptions.

Applying our general construction to the Paillier Cryptosystem yields the first construction of SEM-SO-ENC secure encryptions from the DCR assumption, and this construction is the most efficient that is currently known.

We have shown that Statistically-Hiding $\binom{2}{1}$-OT implies lossy encryption, which, when combined with known results, implies that both PIR and Homomorphic Encryption imply IND-SO-ENC secure encryptions.

We formalized Chosen Ciphertext security in the selective opening setting in both the indistinguishability and simulation-based settings, and gave a general construction based existing primitives.

We note, however, that both the indistinguishability and simulation-based CCA secure constructions suffer from the drawback that an upper bound on the the number of ciphertexts in the challenge query (n) must be known in advance. Although the standard (not CCA) SOA-secure constructions don't suffer from this restriction, it seems very difficult to construct a CCA secure construction whose parameters are independent of n. We note, however, that it is very simple to do in the Random Oracle model.[3]

Acknowledgements:

We thank Yuval Ishai for suggesting a connection between Oblivious Transfer and Lossy Encryption.

References

[Bea97] Donald Beaver. Plug and play encryption. In *CRYPTO '97*, pages 75–89, London, UK, 1997. Springer-Verlag.

[BFO08] Alexandra Boldyreva, Serge Fehr, and Adam O'Neill. On notions of security for deterministic encryption, and efficient constructions without random oracles. In David Wagner, editor, *CRYPTO '08*, volume 5157 of *Lecture Notes in Computer Science*, pages 335–359. Springer, 2008.

[BH92] Donald Beaver and Stuart Haber. cryptographic protocols provably secure against dynamic adversaries. In *EUROCRYPT '92*, number 658 in Lecture Notes in Computer Science, pages 307–323. Springer-Verlag, 1992.

[3] If f is a LTDF, and h is a pairwise independent hash, then setting $E(m, r) = (f(r), h(r) \oplus m, \mathcal{RO}(m, r))$, and copying the proof in [BR93], gives one such general construction.

[BHY09] Mihir Bellare, Dennis Hofheinz, and Scott Yilek. Possibility and impossibility results for encryption and commitment secure under selective opening. In *EUROCRYPT '09*. Springer, 2009. Preprint received from the authors 1-29-09.

[Ble98] Daniel Bleichenbacher. Chosen ciphertext attacks against protocols based on the rsa encryption standard pkcs #1. In *CRYPTO '98*, volume 1462 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 1998.

[BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *CCS '93*, pages 62–73. ACM Press, 1993.

[CDNO97] Ran Canetti, Cynthia Dwork, Moni Naor, and Rafail Ostrovsky. Deniable encryption. In *CRYPTO '97: Proceedings of the 17th Annual International Cryptology Conference on Advances in Cryptology*, pages 90–104, London, UK, 1997. Springer-Verlag.

[CFGN96] Ran Canetti, Uri Feige, Oded Goldreich, and Moni Naor. Adaptively secure multi-party computation. In *STOC '96: Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 639–648, New York, NY, USA, 1996. ACM.

[CHK05] Ran Canetti, Shai Halevi, and Jon Katz. Adaptively-secure, non-interactive public-key encryption. In *TCC '05*, number 3378 in Lecture Notes in Computer Science, pages 150–168. Springer-Verlag, 2005.

[CIO98] Giovanni Di Crescenzo, Yuval Ishai, and Rafail Ostrovsky. Non-interactive and non-malleable commitment. In *STOC '98*. ACM, 1998.

[CKN03] Ran Canetti, Hugo Krawczyk, and Jesper Buus Nielsen. Relaxing chosen-ciphertext security. In *CRYPTO '03*, volume 2729 of *Lecture Notes in Computer Science*, pages 565–582. Springer, 2003.

[CMO00] Giovanni Di Crescenzo, Tal Malkin, and Rafail Ostrovsky. Single database private information retrieval implies oblivious transfer. In *EUROCRYPT '00*, volume 1807 of *Lecture Notes in Computer Science*, pages 122–138. Springer Berlin / Heidelberg, 2000.

[DDN91] Danny Dolev, Cynthia Dwork, and Moni Naor. Non-malleable cryptography. In *STOC '91*, pages 542–552, 1991.

[DJ01] Ivan Damgård and Mads Jurik. A generalisation, a simplification and some applications of paillier's probabilistic public-key system. In *PKC '01: Proceedings of the 4th International Workshop on Practice and Theory in Public Key Cryptography*, pages 119–136, London, UK, 2001. Springer-Verlag.

[DNRS03] Cynthia Dwork, Moni Naor, Omer Reingold, and Larry Stockmeyer. Magic functions: In memoriam: Bernard m. dwork 1923–1998. *Journal of the ACM*, 50(6):852–921, 2003.

[GOS06] Jens Groth, Rafail Ostrovsky, and Amit Sahai. Perfect non-interactive zero knowledge for np. In *Proceedings of Eurocrypt 2006, volume 4004 of LNCS*, pages 339–358. Springer, 2006.

[Gro04] Jens Groth. Rerandomizable and replayable adaptive chosen ciphertext attack secure cryptosystems. In *TCC '04*, volume 2951 of *Lecture Notes in Computer Science*, pages 152–170. Springer, 2004.

[HO09] Brett Hemenway and Rafail Ostrovsky. Re-randomizable encryption implies selective opening security. Cryptology ePrint Archive, Report 2009/088, 2009. http://eprint.iacr.org/2009/088.

[Hof08] Dennis Hofheinz. Possibility and impossibility results for selective decommitments. Cryptology ePrint Archive, Report 2008/168, 2008. http://eprint.iacr.org/2008/168.

[IKO05] Yuval Ishai, Eyal Kushilevitz, and Rafail Ostrovsky. Sufficient conditions for collision resistant hashing. In *TCC '05*, volume 3378, pages 445–456. Springer Berlin / Heidelberg, 2005.

[JJS04] Markus Jakobsson, Ari Juels, and Paul Syverson. Universal re-encryption for mixnets. In *In Proceedings of the 2004 RSA Conference, Cryptographers track*, pages 163–178. Springer-Verlag, 2004.

[KN08] Gillat Kol and Moni Naor. Cryptography and game theory: Designing protocols for exchanging information. In *TCC '08 : Proceedings of the fifth annual Theory of Cryptography Conference*, pages 320–339. Springer Berlin / Heidelberg, 2008.

[KO97] Eyal Kushilevitz and Rafail Ostrovsky. Replication is not needed: Single database, computationally-private information retrieva. In *STOC '97*, pages 364–373. ACM, 1997.

[Lin06] Yehuda Lindell. A simpler construction of cca2-secure public-key encryption under general assumptions. *Journal of Cryptology*, 19(3):359–377, 2006.

[Man98] Eran Mann. Private access to distributed information. Master's thesis, Technion - Israel Institute of Technology, 1998.

[NP01] Moni Naor and Benny Pinkas. Efficient oblivious transfer protocols. In *SODA '01*, pages 448–457. ACM/SIAM, 2001.

[NY90] Moni Naor and Moti Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC '90*, pages 427–437, 1990.

[Pai99] Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *Advances in Cryptology – EUROCRYPT '99*, volume 1592 of *Lecture Notes in Computer Science*, pages 223–238. Springer Berlin / Heidelberg, 1999.

[PR07] Manoj Prabhakaran and Mike Rosulek. Rerandomizable RCCA encryption. In *CRYPTO '07*, volume 4622 of *Lecture Notes in Computer Science*, pages 517–534. Springer Berlin / Heidelberg, 2007.

[PVW08] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In David Wagner, editor, *CRYPTO '08*, volume 5157 of *Lecture Notes in Computer Science*, pages 554–571. Springer, 2008.

[PW08] Chris Peikert and Brent Waters. Lossy trapdoor functions and their applications. In *STOC '08: Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 187–196, New York, NY, USA, 2008. ACM.

[RS91] Charles Rackoff and Daniel Simon. Non-interactive zero-knowledge proof of knowledge and chosen ciphertext attack. In *CRYPTO '91*, pages 433–444, 1991.

[Sah99] Amit Sahai. Non-malleable non-interactive zero-knowledge, and adaptive chosen-ciphertext security. In *FOCS '99*, pages 543–553, 1999.

[SCO+01] Alfredo De Santis, Giovanni Di Crescenzo, Rafail Ostrovsky, Giuseppe Persiano, and Amit Sahai. Robust non-interactive zero knowledge. In *CRYPTO '01*, volume 2139 of *Lecture Notes in Computer Science*, pages 566–598. Springer Berlin / Heidelberg, 2001.

Appendix

A Selective Opening Secure Commitments

A.1 Re-randomizable One-Way Functions

A family of functions $\mathcal{F}$, indexed by a security parameter λ is called a *re-randomizable one-way function* family if the following conditions are satisfied

- **Efficiently Computable:** For all $f \in \mathcal{F}$, the function

  $$f : M \times R \to Y$$

  is efficiently computable.

- **One-Way:** For all PPT adversaries $A = (A_1, A_2)$,

  $$\Pr\left[f \leftarrow \mathcal{F}; (m_0, m_1, st) \leftarrow A_1(f); b \leftarrow \{0,1\}; r \leftarrow R; b' \leftarrow A_2(f(m_b, r), st) : b = b'\right] < \frac{1}{2} + \nu$$

  for some negligible function ν (of λ).

- **Injective on the first input:** For all $m \neq m' \in M$, and $r, r' \in R$,

  $$f(m, r) \neq f(m', r').$$

  This is equivalent to the statement

  $$f(m, R) \cap f(m', R) = \emptyset$$

  for all $m \neq m'$.

- **Re-randomizable:** For each f, there exists and efficient function ReRand such that for all $m \in M$ and $r_0 \in R$, we have

  $$\{r \leftarrow R; f(m, r)\} \approx_s \{r \leftarrow \mathsf{coins}(\mathsf{ReRand}); \mathsf{ReRand}(f(m, r_0), r)\}.$$

It is easy to see that the encryption algorithm from a re-randomizable encryption scheme is immediately a re-randomizable one-way function. We note, however, that re-randomizable one-way functions are a significantly weaker primitive since we do not require any kind of trapdoor.

A.2 Commitments from Re-randomizable One-Way Functions

We begin by describing a construction of a simple bit commitment scheme that arises from any re-randomizable one-way function. Let $\mathcal{F}$ be a re-randomizable one-way function family. Then we define

Parameter Generation:

- $(f, \mathsf{ReRand}) \leftarrow \mathcal{F}(1^\lambda)$,
- $r_0, r_1 \leftarrow R$,
- $c_0 = f(b_0, r_0)$,
  $c_1 = f(b_1, r_1)$.

The public parameters are $(f, \mathsf{ReRand}, c_0, c_1)$.

Commitment:

- $r' \leftarrow \mathsf{coins}(\mathsf{ReRand})$,
- $\mathrm{Com}(b, r') = \mathsf{ReRand}(c_b, r')$.

Decommitment:
To decommit, simply reveal the randomness r'.

This scheme has a number of nice properties. If $b_0 = b_1$ then the scheme is statistically hiding by the properties of ReRand. Alternatively, if $b_0 \neq b_1$ then the scheme is perfectly binding by the injectivity of f on its first input. Now, the two modes are indistinguishable by the one-wayness of the f, combining this with the preceding observations, we obtain the following consequences. If $b_0 = b_1$ then the scheme is computationally binding, and if $b_0 \neq b_1$ the scheme is computationally hiding.

The security analysis is very straightforward, but as this will be the foundation of all our constructions we include it.

Lemma 3. If $b_0 = b_1$, the scheme described in Appendix A.2 is statistically hiding and if $b_0 \neq b_1$, this scheme is perfectly binding.

Proof. If $b_0 = b_1$, the distributions

$$\{r' \leftarrow \mathsf{coins}(\mathrm{Com}) : \mathrm{Com}(0, r')\} \approx_s \{s' \leftarrow \mathsf{coins}(\mathrm{Com}) : \mathrm{Com}(1, s')\},$$

by the definition of ReRand. On the other hand, if $b_0 \neq b_1$, $\mathrm{Com}(0, r) \in f(b_0, R)$, and $\mathrm{Com}(1, s) \in f(b_1, R)$, but by the injectivity on the first input, these sets are disjoint . □

Lemma 4. The schemes when $b_0 = b_1$ and when $b_0 \neq b_1$ are computationally indistinguishable.

Proof. This is exactly the one-way property of f. □

Corollary 5. If $b_0 = b_1$, this scheme is computationally binding, and if $b_0 \neq b_1$, this scheme is computationally hiding.

Proof. Since the scheme is perfectly binding when $b_0 \neq b_1$, breaking the binding property amounts to a proof that $b_0 = b_1$. Since the two modes are computationally indistinguishable, no computationally bounded adversary can create such a "proof." Similarly, since the scheme is perfectly hiding when $b_0 = b_1$, breaking the hiding property amounts to showing that $b_0 \neq b_1$, since the two modes are computationally indistinguishable, no probabilistic polynomial-time adversary can break the hiding property. □

The ability to choose whether the commitment scheme is statistically hiding or perfectly binding is a valuable property, but it is the fact that this choice can be hidden *from the committer* that makes this construction truly useful.

A.3 Selective Opening Secure Commitments

A.4 Definitions

Definition 8. (Indistinguishability under selective openings/IND-SO-COM).

Let Com be a commitment scheme, we say that Com is indistinguishable under selective openings (IND-SO-COM secure) if for every PPT message distribution M and every PPT adversary A, we have that

$$\left|\Pr\left[A^{\mathsf{ind\text{-}so\text{-}real}} = 1\right] - \Pr\left[A^{\mathsf{ind\text{-}so\text{-}ideal}} = 1\right]\right| < \nu$$

for some negligible function ν, and where the games ind-so-real and ind-so-ideal are defined as follows

IND-SO-COM (Real)

- $(m_1, \ldots, m_n) \leftarrow M$
- $r_1, \ldots, r_n \leftarrow \mathsf{coins}(\mathrm{Com})$
- $I \leftarrow A((\mathrm{Com}(m_1, r_1), \ldots, \mathrm{Com}(m_n, r_n))$
- $b \leftarrow A(\mathrm{Dec}(\mathrm{Com}(m_i, r_1))_{i \in I}, (m_1, \ldots, m_n))$

IND-SO-COM (Ideal)

- $(m_1, \ldots, m_n) \leftarrow M$
- $r_1, \ldots, r_n \leftarrow \mathsf{coins}(\mathrm{Com})$
- $I \leftarrow A((\mathrm{Com}(m_1, r_1), \ldots, \mathrm{Com}(m_n, r_n))$
- $(m'_1, \ldots, m'_n) \leftarrow M|M_I$
- $b \leftarrow A(\mathrm{Dec}(\mathrm{Com}(m_i, r_i))_{i \in I}, (m'_1, \ldots, m'_n))$

More explicitly, in the real game,

- The challenger samples messages $(m_1, \ldots, m_n) \leftarrow M$, from the joint message distribution.
- The challenger generates randomness $r_1, \ldots, r_n \leftarrow \mathsf{coins}(\mathrm{Com})$.
- The challenger sends $(\mathrm{Com}(m_1, r_1), \ldots, \mathrm{Com}(m_n, r_n))$ to A.
- The adversary A responds with a subset $I \subset \{1, \ldots, n\}$, with $|I| = n/2$.
- The challenger decommits $(\mathrm{Com}(m_i, r_i))_{i \in I}$.
- The challenger sends $(m_1, \ldots, m_n)$ to the adversary.
- The adversary outputs a bit b.

In the ideal game,

- The challenger samples messages $(m_1, \ldots, m_n) \leftarrow M$, from the joint message distribution.
- The challenger generates randomness $r_1, \ldots, r_n \leftarrow \mathsf{coins}(\mathrm{Com})$.
- The challenger sends $(\mathrm{Com}(m_1, r_1), \ldots, \mathrm{Com}(m_n, r_n))$ to A.
- The adversary A responds with a subset $I \subset \{1, \ldots, n\}$, with $|I| = n/2$.
- The challenger decommits $(\mathrm{Com}(m_i, r_i))_{i \in I}$.
- The challenger samples a new vector $m' \leftarrow M|M_I$, from M conditioned on the fact that $m_i = m'_i$ for $i \in I$, and sends M' to A.
- The adversary outputs a bit b.

A.5 IND-SO-COM Constructions from Re-randomizable One-Way Functions

To construct an IND-SO-COM secure commitment scheme, it is enough to create a statistically hiding commitment scheme, since Bellare, Hofheinz and Yilek showed

Theorem 5. (Theorem 6 From [BHY09]).
Statistically-hiding commitment schemes are IND-SO-COM secure.

Proof. We follow the general form of the proof from [BHY09], but by restricting ourselves to non-interactive commitments we can slightly simplify the exposition. We begin by defining an (inefficient) algorithm called opener, which tries to open a commitment c to a specified message m. In particular

$$\mathsf{opener}(c, m) = \begin{cases} r & \text{s.t. } \mathrm{Com}(m, r) = c, \\ \perp & \text{if no such } r \text{ exists .} \end{cases}$$

Now, we proceed in a sequence of games. Let Game_{-1} be the real IND-SO-COM game. Let Game_0 be the game, where the challenger uses opener to decommit the commitments $(\mathrm{Com}(m_i, r_i))_{i \in I}$. Notice that the views of the adversary in Game_{-1} and Game_0 are identical (but Game_0 is no longer efficiently implementable). In particular

$$\Pr[A^{\mathrm{Game}_{-1}} = 1] = \Pr[A^{\mathrm{Game}_0} = 1].$$

Next we describe Game j for $j \in [n]$. The only difference between Game_j and Game_0 is that in Game_j for $i \leq j$, the challenger sends the vector

$$(\mathrm{Com}(\delta, r_1), \ldots, \mathrm{Com}(\delta, r_j), \mathrm{Com}(m_{j+1}, r_{j+1}), \ldots, \mathrm{Com}(m_n, r_n)),$$

for some fixed dummy message δ. For concreteness, we may set $\delta = 0^\lambda$. Now, the only difference between Game_j and Game_{j-1} is whether the jth commitment is a commitment to δ or m_j. Since Com is *statistically*-hiding, even an *unbounded* adversary has only a negligible probability of distinguishing the two cases. Thus by the triangle inequality

$$|\Pr[A^{\mathrm{Game}_0} = 1] - \Pr[A^{\mathrm{Game}_n} = 1]| < n \cdot \nu = \text{negligible},$$

where ν is the probability that an *unbounded* adversary breaks the hiding property of Com. Thus we obtain

$$|\Pr[A^{IND-SO-COM}_{real} = 1] - \Pr[A^{\mathrm{Game}_n} = 1] = \text{negligible}.$$

Finally, we notice that in Game_n, all the commitments are independent of the message $(m_1, \ldots, m_n)$, so we can repeat the above argument, starting with the ideal IND-SO-COM game, instead of the real game. Thus we obtain

$$|\Pr[A^{IND-SO-COM}_{ideal} = 1] - \Pr[A^{\mathrm{Game}_n} = 1] = \text{negligible}.$$

Thus

$$|\Pr[A^{IND-SO-COM}_{ideal} = 1] - \Pr[A^{IND-SO-COM}_{real} = 1]| = \text{negligible}.$$

□

The commitment scheme constructed in Appendix A.2 is statistically hiding when $b_0 = b_1$, so we obtain the following corollary

Corollary 6. Re-randomizable one-way functions imply non-interactive IND-SO-COM secure commitments.

Since Re-randomizable encryptions imply re-randomizable one-way functions, we have

Corollary 7. Re-randomizable encryption implies non-interactive IND-SO-COM secure commitments.

Perhaps more interesting is the case when $b_0 \neq b_1$. The commitment scheme constructed in Appendix A.2 is no longer perfectly hiding, so Theorem 5 doesn't apply. In this case, we can still achieve IND-SO-COM security, by using the indistinguishability of the two modes. Roughly, this follows because an IND-SO-COM adversary must have similar probabilities of success against both modes, otherwise it could be used to distinguish the modes. Thus we arrive at the following Corollary.

Corollary 8. Re-randomizable one-way functions imply perfectly-binding IND-SO-COM secure commitments.

Since Re-randomizable encryptions imply re-randomizable one-way functions, we have

Corollary 9. Re-randomizable encryption implies perfectly binding non-interactive IND-SO-COM secure commitments.

Proof. We proceed via contradiction. Suppose there exists an IND-SO-COM adversary A that succeeds against the protocol with probability $\frac{1}{2}+\epsilon$ when $b_0 = b_1$. We will use A to construct a distinguisher D for the one-way game against the underlying re-randomizable one-way function f. In the one-wayness game against f, the challenger samples a function f and sends it to D. D will respond by sending $\{0,1\}$ to the one-wayness challenger, and the one-wayness challenger samples $r \leftarrow R$ and sends $e = f(b,r)$ to D. Now, D samples $r' \leftarrow R$, and generates $e' = f(0,r')$. Now, D creates an instantiation of the commitment protocol setting $c_0 = e, c_1 = e'$, and plays the IND-SO-COM game with the adversary A. If A wins, D guesses $b = 1$, and if A loses, D guesses $b = 0$. From Theorem 5 we know that if $b = 0$ then A succeeds with probability ν for some negligible function ν. On the other hand, by hypothesis, if $b = 1$, then A wins the IND-SO-COM game with probability ϵ. Now

$$\begin{aligned}
\Pr[D \text{ wins }] &= \Pr[b = 1 \cap \ A \text{ wins }] + \Pr[b = 0 \cap \ A \text{ loses }] \\
&= \Pr[A \text{ wins} | b = 1] \Pr[b = 1] + \Pr[A \text{ loses} | b = 0] \Pr[b = 0] \\
&= \frac{1}{2}\left(\frac{1}{2} + \epsilon + \frac{1}{2} - \nu\right) \\
&= \frac{1}{2} + \frac{\epsilon - \nu}{2}.
\end{aligned}$$

Since ϵ is non-negligible, and ν is negligible, D breaks the one-way property of f. □

This result is perhaps surprising, since [BHY09] showed a black-box separation between most known cryptographic primitives and Perfectly Binding IND-SO-COM secure commitments.

B Homomorphic Encryption

A Public Key Cryptosystem given by algorithms (G, E, D) is called *homomorphic* if

- The plaintext space forms a group X, with group operation $+$.
- The ciphertexts are members of a group Y.
- For all $x_0, x_1 \in X$, and for all $r_0, r_1 \in \mathsf{coins}(E)$, there exists an $r^* \in \mathsf{coins}(E)$ such that

$$E(pk, x_0 + x_1, r^*) = E(pk, x_0, r_0)E(pk, x_1, r_1).$$

Notice that we do not assume that the encryption is also homomorphic over the randomness, as is the case in most homomorphic encryption schemes, e.g. El-Gamal, Paillier, and Goldwasser-Micali. We also do not assume that the image $E(pk, X, R)$ is all of the group Y, only that $E(pk, X, R) \subset Y$. Since the homomorphic property implies closure, we have that $E(pk, X, R)$ is a semi-group. Notice also, that while it is common to use the word "homomorphic" to describe the cryptosystem, encryption is *not* a homomorphism in the mathematical sense.

We now show some basic properties from all homomorphic encryption schemes, these facts are commonly used, but since our definition is weaker than the (implicit) definitions of homomorphic encryption that appear in the literature, it is important to note that they hold under this definition as well.

- $E(pk, X, R)$ is a group.
- $E(pk, 0, R)$ is a subgroup of $E(pk, X, R)$.

- For all $x \in X$, $E(pk, x, R)$ is the coset $E(pk, x, r)E(pk, 0, R)$.
- For all $x_0, x_1 \in X$, $|E(pk, x_0, R)| = |E(pk, x_1, R)|$.
- If y is chosen uniformly from $E(pk, 0, R)$, then $yE(pk, x, r)$ is uniform in $E(pk, x, R)$.
- The group $E(pk, X, R) \simeq X \times E(pk, 0, R)$, and decryption is simply the homomorphism

$$E(pk, X, R) \to E(pk, X, R)/E(pk, 0, R) \simeq X.$$

We call a cryptosystem a *Homomorphic Public Key Cryptosystem* (HPKC), if the cryptosystem is IND-CPA secure, and homomorphic.

If we make the additional assumption that we can sample in a manner statistically close to uniform on the subgroup $E(pk, 0, R)$, then the cryptosystem (G, E, D) will be re-randomizable.

Definition 9. We call a Homomorphic Public Key Cryptosystem *Uniformly Sampleable* if there exists a PPT algorithm sample such that the output of $\mathsf{sample}(pk)$ is statistically close to uniform on the group $E(pk, 0, R)$.

We note, that for all known homomorphic cryptosystems we may define

$$\mathsf{sample}(pk) = \{r \leftarrow \mathsf{coins}(E) : E(pk, 0, r)\}.$$

It is not too hard to see that this property *does not* follow from the definition of Homomorphic Encryption.

B.1 Efficient Re-randomizable Encryption from Uniformly Sampleable Homomorphic Encryption

The scheme described above only allows commitment to single bits. If the underlying cryptosystem (G, E, D), can encrypt more than one bit at a time, we can increase the efficiency of this system, by simply putting $c_0, c_1, \ldots, c_n$ into the public key, and a commitment to i will be $\mathsf{ReRand}(pk, c_i, r)$. In most cases, however, we can increase the size of the committed messages without increasing the public-key.

In particular, if $(G, E, D, \mathsf{sample})$ is a Uniformly Sampleable Homomorphic Encryption scheme and $\mathbb{Z}/N\mathbb{Z} \hookrightarrow X$. Then, we can commit to elements in $\{0, 1, \ldots, N-1\}$ instead of $\{0, 1\}$ by simply taking

Parameter Generation:

- $(pk, sk) \leftarrow G(1^\lambda)$,
- $r \leftarrow \mathsf{coins}(E)$,
- $c = E(pk, b, r)$,

The public parameters are (pk, c).

Encryption:

- $r' \leftarrow \mathsf{coins}(\mathsf{sample})$.
- $c' \leftarrow \mathsf{sample}(pk, r')$.
- return $c^a \cdot c'$.

Decryption:
To decrypt a ciphertext c, simply return $D(c)$.

Now, if $c = E(pk, 0, r)$ the scheme is lossy, since all encryptions will be uniformly distributed in the subgroup $E(pk, 0, R)$, while if $c = E(pk, 1, r)$ the scheme is injective by the correctness of the decryption algorithm. This is the natural construction when working with the Paillier or Damgård-Jurik cryptosystems. We must use caution when applying this to El-Gamal, since the inverse map $\mathbb{Z}/N\mathbb{Z} \hookrightarrow X$ is not

efficiently computable (it's the discrete log). In this context, it will not be a problem, since we never have to compute the inverse to decommit. If, on the other hand, we wanted to view this as an encryption scheme instead of a commitment scheme, then this lack of inverse would be an issue. Fortunately, there is a well known scheme to create a re-randomizable encryption from the DDH assumption that is only a slight modification of the original El-Gamal scheme. See [NP01], [PVW08] or [BHY09] for a description of this scheme. We stress, however, that "plain" El-Gamal is re-randomizable, however, it is slightly less efficient than this modification.

C The Paillier Cryptosystem

We briefly review the Paillier Cryptosystem proposed by Pascal Paillier in [Pai99], and extended by Damgård and Jurik in [DJ01].

The cryptosystem works in $(\mathbb{Z}/N^2\mathbb{Z})^*$. From the Binomial Theorem, we have

$$(1+N)^a = 1 + aN \mod N^2,$$

so $(1+N)$ generates a cyclic subgroup of order N. In this group, we can take Discrete Logs efficiently by $L(x) = \frac{x-1}{N}$, since

$$L((1+N)^a) = L(1+aN) = a.$$

Now, if g generates $\langle 1+N \rangle$, and $c = g^a$, then as with traditional logs

$$a = \frac{L(c)}{L(g)}.$$

Now, we are ready to describe Paillier's Cryptosystem

- **Parameter Generation:**
  - Generates primes p, q of length $\lambda/2$ and sets $N = pq$.
  - Generate $g \in \mathbb{Z}/N^2\mathbb{Z}$ such that N divides the order of g.
    This condition is easy to verify if you have the factorization of N.

  The public parameters are $pk = (N, g)$
  The secret key is $sk = \text{lcm}(p-1, q-1)$.

- **Encryption:**
  - $r \leftarrow \mathbb{Z}/N\mathbb{Z}$,
    (really you want to generate $r \in (\mathbb{Z}/N\mathbb{Z})^*$, but the distributions are statistically close).
  - For $m \in \mathbb{Z}/N\mathbb{Z}$,
    $E(pk, m, r) = g^m r^N \mod N^2$.

- **Decryption:**
  Given a ciphertext $c \in (\mathbb{Z}/N^2\mathbb{Z})^*$,

$$m = \frac{L(c^{sk}) \mod N}{L(g^{sk})} \mod N.$$

This cryptosystem is IND-CPA secure under the Decisional Composite Residuosity Assumption (DCR), which (informally) says

Assumption 1. Decisional Composite Residuosity/(DCR): If N is an λ-bit RSA modulus, (i.e. $N = pq$), then

$$\{g \leftarrow \mathbb{Z}/N^2\mathbb{Z}; g\} \approx_c \{g \leftarrow \mathbb{Z}/N^2\mathbb{Z}; g^N\}.$$