TD 8: Public Key Encryption (corrected version)

Exercise 1.

Let $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ be a correct public-key encryption scheme. Let us assume moreover that Enc is deterministic.

1. Show that this scheme is not CPA-secure.

   ☞ Let us consider the following adversary $\mathcal{A}$. On input the public key pk, it chooses two messages $m_0 \neq m_1$ uniformly chosen over the message space $\mathcal{M}$ as its challenge query. It gets c. It computes $c_b = \mathsf{Enc}(\mathsf{pk}, m_b)$ and returns b such that $c_b = c$. Since the scheme is deterministic, there exists such a $b \in \{0, 1\}$. Moreover, by correctness, the probability that $c_0 = c_1$ is negligible. Its advantage is close to one.

Exercise 2.

Let $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a public-key encryption scheme. The One-Wayness against Chosen Plaintext Attack (OW-CPA) security notion is the following. The challenger samples $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^\lambda)$ and $\mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{pk}, m)$, where $m \leftarrow U(\mathcal{M})$ and $\mathcal{M}$ is the message space. The adversary wins if it outputs a message m' such that $m = m'$.

A scheme is said OW-CPA secure if no ppt adversary wins with non-negligible probability.

1. Write a formal definition of the OW-CPA security. Can a scheme be OW-CPA secure if the message space is $\mathcal{M} = \{0, 1\}$?

   ☞ Define the following game:

$$
\begin{array}{ccc}
\mathcal{C} & & \mathcal{A} \\
\hline
(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^\lambda) & & \\
m \leftarrow U(\mathcal{M}) & & \\
& \xrightarrow{(\mathsf{pk}, c := \mathsf{Enc}(\mathsf{pk}, m))} & \\
& \xleftarrow{\quad m' \quad} &
\end{array}
$$

   The advantage of the adversary is defined as $\mathsf{Adv}(\mathcal{A}) := \Pr(m' = m)$. A PKE scheme is OW-CPA secure if no ppt adversary has non-negligible advantage.

   An adversary guessing uniformly which message was encrypted has advantage $1/|\mathcal{M}|$. As such, if the message space does not have cardinality such that $1/|\mathcal{M}| = \mathsf{negl}(\lambda)$, the scheme is trivially not OW-CPA secure.

2. Show that if $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is IND-CPA secure and has exponential message space, then it is OW-CPA secure.

   ☞ Assume that the scheme is IND-CPA secure, but not OW-CPA secure. Let $\mathcal{A}$ be an adversary with non-negligible advantage in the OW-CPA game. We build the following adversary against the IND-CPA security:

$$
\begin{array}{ccc}
\mathcal{C} & \mathcal{B} & \mathcal{A} \\
\hline
(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^\lambda) & & \\
\text{Send } \mathsf{pk} \text{ to } \mathcal{B} & & \\
& \text{Sample and send } m_0, m_1 \leftarrow U(\mathcal{M}) \text{ as challenge queries} & \\
\text{Sample } b \leftarrow U(\{0, 1\}) & & \\
\text{Return } c^* := \mathsf{Enc}(\mathsf{pk}, m_b) & & \\
& \text{Send } (\mathsf{pk}, c^*) \text{ to } \mathcal{A} & \\
& & \text{Output } m' \\
& \text{If } m' \neq m_0 \text{ and } m_1, \text{ return a random bit.} & \\
& \text{Output } b' \text{ s.t. } m' = m_{b'}. & \\
& (\text{If } m_0 = m_1, \text{ output } 1) &
\end{array}
$$

   Since m_b is uniformly sampled, adversary $\mathcal{A}$'s view is exactly the same as in the OW-CPA game. As such, it outputs m_b with probability $\mathsf{Adv}(\mathcal{A})$.

   It holds that $\Pr(b' = 1 | b = 1) = \frac{1}{2} \Pr(m' \neq m_0 \wedge m' \neq m_1) + \Pr(m' = m_1 | b = 1)$: the second term is the probability of the random bit being right, if we output a random bit, and the second term is actually the advantage of $\mathcal{A}$.

   Moreover, $\Pr(b' = 1 | b = 0) = \frac{1}{2} \Pr(m' \neq m_0 \wedge m' \neq m_1) + \Pr(m' = m_1 | b = 0)$. All that is left is to estimate the quantity $\Pr(m' = m_1 | b = 0)$.

Note that in the case $b = 0$, m' and m_1 are two independent random variables. As such, whatever the value of m', the message m_1 has probability $\frac{1}{|\mathcal{M}|}$ of being equal to it, and thus gives $\Pr(m' = m_1 | b = 0) = \frac{1}{|\mathcal{M}|}$ (this can be seen by using the law of total probabilities on the possible values of m').

Then the advantage of $\mathcal{B}$ is:

$$\mathsf{Adv}(\mathcal{B}) = \left| \mathsf{Adv}(\mathcal{A}) - \frac{1}{|\mathcal{M}|} \right|.$$

However, recall that in the first question we have proven that if the size of the message space is polynomial in λ then the notion of $OW - CPA$ does not exist. As such, we restrict our reduction to the case where $\frac{1}{|\mathcal{M}|} = \mathsf{negl}(\lambda)$, and this concludes the reduction.

3. Let $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be an IND-CPA secure encryption scheme with message space $\mathcal{M}$ such that it has cardinality $|\mathcal{M}| = 2^\lambda$, where λ is the security parameter. Show that a small modification of the scheme leads to an encryption scheme $(\mathsf{Gen}, \mathsf{Enc}', \mathsf{Dec}')$ that is OW-CPA secure but not IND-CPA secure anymore.

☞ **First method:** determinizing the scheme.

As we have already done a few times, let us assume that Enc uses ℓ bits of uniform randomness, and we make them explicit: there exists some deterministic algorithm Enc' such that $\mathsf{Enc}'(pk, m, S)$ behaves exactly as $\mathsf{Enc}(pk, m)$ when S is uniformly sampled over $\{0, 1\}^\ell$. We simply choose to put the seed S as a part of the message: the new message space is $\mathcal{M}' = \mathcal{M} \times \{0, 1\}^\ell$.

No adversary in the OW-CPA setting sees any difference, as the seed is uniformly sampled (as the rest of the message): $(\mathsf{Gen}, \mathsf{Enc}', \mathsf{Dec})$ is still OW-CPA secure.

However, no scheme with deterministic encryption is IND-CPA secure, proving that this scheme is not IND-CPA secure.

Second method: Let m be a fixed message and choose any ciphertext c. Then set $\mathsf{Enc}'(pk, m) = c$ for any pk, and do the following (to keep perfect correctness):
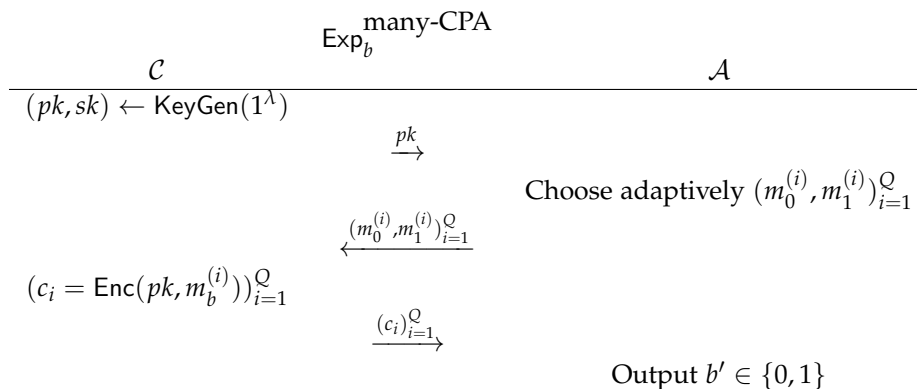
$$\forall pk, m' \, \mathsf{Enc}'(pk, m') = \begin{cases} \mathsf{Enc}(pk, m') & \text{if } \mathsf{Enc}(pk, m') \neq c \\ \mathsf{Enc}(pk, m) & \text{if } \mathsf{Enc}(pk, m') = c \end{cases}$$

Then since we only changed, for a fixed public key, at most two ciphertexts ($\mathsf{Enc}(pk, \cdot)$ is injective if we have perfect correctness), an OW-CPA adversary only sees a difference with probability $1/2|\mathcal{M}|$, which is negligible: the scheme is still OW-CPA secure.

However, an IND-CPA challenger can specifically give as challenge (m, m') and return 0 iff $c^\star = c$, where $c^\star$ is the challenger ciphertext: it has an advantage of 1. This scheme is not IND-CPA secure.

Exercise 3.

Let $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a Public-Key encryption scheme. Let us define the following experiments for $b \in \{0, 1\}$ and $Q = \mathsf{poly}(\lambda)$.

$$\mathsf{Exp}_b^{\text{many-CPA}}$$

$\mathcal{C}$		$\mathcal{A}$
$(pk, sk) \leftarrow \mathsf{KeyGen}(1^\lambda)$		
	$\xrightarrow{pk}$	
		Choose adaptively $(m_0^{(i)}, m_1^{(i)})_{i=1}^Q$
	$\xleftarrow{(m_0^{(i)}, m_1^{(i)})_{i=1}^Q}$	
$(c_i = \mathsf{Enc}(pk, m_b^{(i)}))_{i=1}^Q$		
	$\xrightarrow{(c_i)_{i=1}^Q}$	
		Output $b' \in \{0, 1\}$

The advantage of $\mathcal{A}$ in the many-time CPA game is defined as

$$\mathsf{Adv}^{\text{many-CPA}}(\mathcal{A}) = |\Pr(\mathcal{A} \xrightarrow{\mathsf{Exp}_1^{\text{many-CPA}}} 1) - \Pr(\mathcal{A} \xrightarrow{\mathsf{Exp}_0^{\text{many-CPA}}} 1)|.$$

1. Recall the definition of CPA-security that was given during the lecture. What is the difference?

☞ In the lecture, we fixed $Q = 1$.

2. Show that these two definitions are equivalent.

☞ With what was previously stated, we already see that many-CPA-secure implies CPA-secure. Let us define the following hybrid games: in game H_i, the challenger returns $\mathsf{Enc}(pk, m_0^{(j)})$ for the first i queries, and $\mathsf{Enc}(pk, m_1^{(j)})$ for the following ones. Experiment 0 is H_Q and experiment 1 is H_0. Let us assume that there exists an encryption scheme which is CPA-secure but not many-CPA-secure.

This means that there exists an adversary able to distinguish between hybrids H_0 and H_Q with non-negligible probability. In particular, there exists some index i where this adversary has non-negligible advantage in distinguishing H_i and H_{i+1} by the hybrid argument:

$$\text{Adv}(\mathcal{A}) \leq \sum_{i=0}^{Q-1} |\Pr(\mathcal{A} \xrightarrow{H_i} 1) - \Pr(\mathcal{A} \xrightarrow{H_{i+1}} 1)|.$$

However, the distringuishing game between H_i and H_{i+1} can be simulated by a (one time) CPA adversary: it encrypts the first i and last $n-1-i$ queries itself, and chooses query $i+1$ as its challenge query. Finally it returns the same bit that the simulated adversary returns, and they both have the same advantage, thus breaking the CPA security.

Then these two notions are equivalent.

3. Do we have a similar equivalence in the secret-key setting? ☞ No, recall the one-time pad, which is insecure as soon as two message challenges are allowed. Notably, what fails in the previous proof is that now, the adversary in the middle cannot encrypt by itself.