CS 171: Discussion Section 8 (March 11)

1 CPA-Secure Public-Key Encryption From Two-Round Key Exchange

Question: Given a two-round key-exchange protocol with keyspace $\mathcal{K} = \{0,1\}^n$, construct a CPA-secure public-key encryption (PKE) scheme for n-bit messages and prove its security. Do not use any other cryptographic primitive.

1.1 Two-Round Key Exchange

A two-round key-exchange protocol comprises three randomized algorithms (P_1, P_2, P_3) and has the following form:

1. Alice computes $(\mathsf{msg}_1, \mathsf{st}) \leftarrow P_1(1^n)$ and sends msg_1 to Bob.
2. Bob computes $(\mathsf{msg}_2, k) \leftarrow P_2(\mathsf{msg}_1)$. Then he sends msg_2 to Alice and outputs k.
3. Alice computes $k \leftarrow P_3(\mathsf{st}, \mathsf{msg}_2)$ and outputs k.

1.1.1 Definition of Security

We will define security for key exchange below. Our definition of security is equivalent to the one given in lecture 13, slide 26.

Consider the following security game.

$\underline{G_{\mathcal{B},\Pi}(n,b)}$:

1. The challenger executes the key exchange protocol Π to produce $(\mathsf{msg}_1, \mathsf{msg}_2, k)$.
2. If $b = 0$, the challenger sets $\hat{k} = k$. If $b = 1$, they sample $\hat{k} \leftarrow \mathcal{K}$. Then the adversary $\mathcal{B}$ is given $(\mathsf{msg}_1, \mathsf{msg}_2, \hat{k})$.
3. $\mathcal{B}$ outputs a bit b', which is the output of the game as well.

We say that a key-exchange protocol is **secure** if for all PPT adversaries $\mathcal{B}$, there exists a negligible function negl such that:

$$\left| \Pr[G_{\mathcal{B},\Pi}(n,0) \to 1] - \Pr[G_{\mathcal{B},\Pi}(n,1) \to 1] \right| = \mathsf{negl}(n)$$

1.2 Definition of CPA security for PKE

Let's write the definition of CPA security for public-key encryption. It will resemble the definition we've seen previously for secret-key encryption.

Given an adversary $\mathcal{A}$, define the following game:

$\underline{\mathsf{PubK}_{\mathcal{A},\Pi}(n)}$:

1. The challenger samples the keys $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^n)$. Then they give $(1^n, \mathsf{pk})$ to the adversary $\mathcal{A}$.
2. $\mathcal{A}$ outputs a pair of messages (m_0, m_1) such that $|m_0| = |m_1|$.
3. The challenger samples $b \leftarrow \{0,1\}$ and computes the challenge ciphertext:
$$c \leftarrow \mathsf{Enc}(\mathsf{pk}, m_b) \tag{1.1}$$
Then they give c to $\mathcal{A}$.
4. $\mathcal{A}$ outputs a bit b'. The output of the experiment is 1 if $b = b'$ and 0 otherwise.

A public-key encryption scheme is **CPA-secure** if for any probabilistic polynomial-time adversary $\mathcal{A}$, there is a negligible function negl such that:

$$\Pr[\mathsf{PubK}_{\mathcal{A},\Pi}(n) \to 1] = \frac{1}{2} + \mathsf{negl}(n)$$

Solution

1.3 Construction of a PKE Scheme:

1. $\mathsf{Gen}(1^n)$: Compute $(\mathsf{msg}_1, \mathsf{st}) \leftarrow P_1(1^n)$. Output $\mathsf{pk} = \mathsf{msg}_1$ and $\mathsf{sk} = \mathsf{st}$.
2. $\mathsf{Enc}(\mathsf{pk}, m)$: Compute $(\mathsf{msg}_2, k) \leftarrow P_2(\mathsf{msg}_1)$. Output $c = (\mathsf{msg}_2, k \oplus m)$.
3. $\mathsf{Dec}(\mathsf{sk}, c)$: parse c as (msg_2, c'); compute $k \leftarrow P_3(\mathsf{st}, \mathsf{msg}_2)$ and output $k \oplus c'$.

Theorem 1.1. *The construction of PKE given above is CPA-secure.*

Proof. Given a PPT adversary $\mathcal{A}$, let us compare the following hybrids:

- $\mathcal{H}_0$: Is $\mathsf{PubK}_{\mathcal{A},\Pi}(n)$, with the PKE construction given in section 1.3:
  1. The challenger computes $(\mathsf{pk}, \mathsf{sk}) = (\mathsf{msg}_1, \mathsf{st}) \leftarrow P_1(1^n)$.
  2. The adversary $\mathcal{A}$ is given input 1^n and msg_1. Then $\mathcal{A}$ outputs a pair of messages (m_0, m_1) such that $|m_0| = |m_1|$.
  3. The challenger computes the *challenge ciphertext*:
  $$\begin{aligned} b &\leftarrow \{0,1\} \\ (\mathsf{msg}_2, k) &\leftarrow P_2(\mathsf{msg}_1) \\ c &= (\mathsf{msg}_2, k \oplus m_b) \end{aligned}$$
  Then they give c to $\mathcal{A}$.
  4. $\mathcal{A}$ outputs a bit b'. The output of the hybrid is 1 if $b = b'$ and 0 otherwise.
- $\mathcal{H}_1$: Is the same as $\mathcal{H}_0$, except the challenge ciphertext is computed as follows:
$$\begin{aligned} b &\leftarrow \{0,1\} \\ r &\leftarrow \{0,1\}^n \\ (\mathsf{msg}_2, k) &\leftarrow P_2(\mathsf{msg}_1) \\ c &= (\mathsf{msg}_2, r \oplus m_b) \end{aligned}$$

Lemma 1.2. $\Big|\Pr[\mathcal{H}_0 \to 1] - \Pr[\mathcal{H}_1 \to 1]\Big| \leq \mathsf{negl}(n)$

Proof. This follows from the security of the key-exchange protocol.

1. Overview: Assume toward contradiction that there's an adversary $\mathcal{A}$ such that $\Big|\Pr[\mathcal{H}_0 \to 1] - \Pr[\mathcal{H}_1 \to 1]\Big|$ is non-negligible. Then we'll construct an adversary $\mathcal{B}$ that can break the security of the key-exchange protocol with non-negligible advantage.

2. Construction of $\mathcal{B}$:

   (a) $\mathcal{B}$ receives from the key exchange challenger the transcript $(\mathsf{msg}_1, \mathsf{msg}_2)$ and a string $\hat{k}$ that could be k or a random string $r \leftarrow \{0,1\}^n$.

   (b) $\mathcal{B}$ sends $(1^n, \mathsf{msg}_1)$ to $\mathcal{A}$. When $\mathcal{A}$ outputs (m_0, m_1), $\mathcal{B}$ samples $b \leftarrow \{0,1\}$ and sends to $\mathcal{A}$:
   $$(\mathsf{msg}_2, \hat{k} \oplus m_b)$$

   (c) Finally, $\mathcal{A}$ outputs a bit b'. $\mathcal{B}$ checks whether $b = b'$. If so, $\mathcal{B}$ outputs 1, and if not, $\mathcal{B}$ outputs 0.

3. Analysis: When $\hat{k} = k$, $\mathcal{B}$ correctly simulates $\mathcal{H}_0$. When $\hat{k} = r \leftarrow \{0,1\}^n$, $\mathcal{B}$ correctly simulates $\mathcal{H}_1$. Since $\mathcal{A}$ distinguishes between $\mathcal{H}_0$ and $\mathcal{H}_1$ with non-negligible advantage, $\mathcal{B}$ distinguishes whether $\hat{k} = k$ or $\hat{k} = r$ with the same advantage. More formally:

$$\begin{aligned} \Pr[G_{\mathcal{B},\Pi}(n,0) \to 1] &= \Pr[\mathcal{H}_0 \to 1] \\ \Pr[G_{\mathcal{B},\Pi}(n,1) \to 1] &= \Pr[\mathcal{H}_1 \to 1] \\ \Big|\Pr[G_{\mathcal{B},\Pi}(n,0) \to 1] - \Pr[G_{\mathcal{B},\Pi}(n,1) \to 1]\Big| &= \Big|\Pr[\mathcal{H}_0 \to 1] - \Pr[\mathcal{H}_1 \to 1]\Big| \\ &= \mathsf{non\text{-}negl}(n) \end{aligned}$$

   Therefore, $\mathcal{B}$ breaks the security of the key-exchange protocol. This is a contradiction because we know the key-exchange protocol is secure. Therefore, our initial assumption was false, and in fact $\Big|\Pr[\mathcal{H}_0 \to 1] - \Pr[\mathcal{H}_1 \to 1]\Big|$ is negligible.

□

Lemma 1.3. $\Pr[\mathcal{H}_1 \to 1] = \frac{1}{2}$

Proof. This follows from the security of the one-time pad.

Fix any values of $(\mathsf{msg}_1, \mathsf{st}, m_0, m_1, \mathsf{msg}_2, k)$. Then over the randomness of r, the variable $r \oplus m_0$ is uniformly random. So is $r \oplus m_1$. Therefore, $r \oplus m_b$ is independent of b and the other variables $(\mathsf{msg}_1, \mathsf{st}, m_0, m_1, \mathsf{msg}_2, k)$.

The output distribution of $\mathcal{A}$ depends only on the variables $(\mathsf{msg}_1, m_0, m_1, \mathsf{msg}_2)$ and the distribution of $r \oplus m_b$. Therefore, the output distribution of $\mathcal{A}$ is independent of b, so:

$$\Pr[\mathcal{H}_1 \to 1] = \frac{1}{2}$$

□

The previous lemmas imply that

$$\begin{aligned}\Pr[\mathcal{H}_0 \to 1] &\leq \Big|\Pr[\mathcal{H}_0 \to 1] - \Pr[\mathcal{H}_1 \to 1]\Big| + \Pr[\mathcal{H}_1 \to 1] \\ &\leq \frac{1}{2} + \mathsf{negl}(n)\end{aligned}$$

Therefore, the construction in section 1.3 satisfies CPA security.

□ □

2 One-way functions from Pseudorandom Permutations

One-way functions can be constructed from many other cryptographic primitives, including from pseudorandom permutations.

Let $F : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$ be a pseudorandom permutation. This can be written as $F(k, x)$ or equivalently $F_k(x)$, where k is the key. Note that an adversary can compute $F_k^{-1}(\cdot)$ in addition to $F_k(\cdot)$ if they are given the key k.

1. Let $x \in \{0,1\}^n$, and

$$\text{let } f_1(x) = F_{0^n}(x)$$

Show that f_1 is not a one-way function.

Solution We will construct a PPT adversary $\mathcal{A}$ that breaks the one-wayness of f_1.

(a) The OWF challenger samples $x \leftarrow \{0,1\}^n$, computes $y = f_1(x) = F_{0^n}(x)$, and gives $\mathcal{A}$ the input $(1^n, y)$.

(b) $\mathcal{A}$ computes $x' = F_{0^n}^{-1}(y)$, and outputs x'.

It holds that $f_1(x') = F_{0^n}(F_{0^n}^{-1}(y)) = y$, so $\mathcal{A}$ wins the OWF security game with probability 1. This breaks the OWF security of f_1. □

2. Let $x = (x_0, x_1) \in \{0,1\}^n \times \{0,1\}^n$, and

$$\text{let } f_2(x) = F_{x_0}(x_1)$$

Show that f_2 is not a one-way function.

Solution We will construct a PPT adversary $\mathcal{A}$ that breaks the one-wayness of f_2.

(a) The OWF challenger samples $(x_0, x_1) \leftarrow \{0,1\}^n \times \{0,1\}^n$, computes $y = f_2(x_0, x_1) = F_{x_0}(x_1)$, and gives $\mathcal{A}$ the input $(1^n, y)$.

(b) $\mathcal{A}$ picks $x'_0 = 0^n$ and computes

$$x'_1 = F_{x'_0}^{-1}(y)$$

Then $\mathcal{A}$ outputs $x' := (x'_0, x'_1)$.

It holds that $f_2(x') = F_{x'_0}(F_{x'_0}^{-1}(y)) = y$, so $\mathcal{A}$ wins the OWF security game with probability 1. This breaks the OWF security of f_2. □

3. Extra problem: Let $x \in \{0,1\}^n$, and

$$\text{let } f_3(x) = F_x(0^n) || F_x(1^n)$$

Show that f_3 is a one-way function.

Solution

(a) We claim that any PRG that maps $\{0,1\}^n \to \{0,1\}^{2n}$ is also a OWF. We will prove a claim essentially the same as this one on HW 7, Q1. Next, to prove that f_3 is a OWF, we just need to prove that f_3 is a PRG.

(b) Assume toward contradiction that f_3 is not a PRG. Then there is an adversary $\mathcal{A}$ that can distinguish $f_3(x)$ (where $x \leftarrow \{0,1\}^n$) from $y \leftarrow \{0,1\}^{2n}$ with non-negligible advantage. Then we will use $\mathcal{A}$ to construct an adversary $\mathcal{B}$ that breaks the PRP security of F.
<u>Construction of $\mathcal{B}$:</u>

i. The PRP challenger gives $\mathcal{B}$ query access to a function, either $F_x(\cdot)$, where $x \leftarrow \{0,1\}^n$, or $R(\cdot)$, where R is a truly random permutation.

ii. $\mathcal{B}$ queries the function on 0^n and 1^n to get outputs y_0 and y_1 respectively. $\mathcal{B}$ runs $\mathcal{A}$ on inputs $(1^n, y_0||y_1)$. $\mathcal{A}$ will output a bit b', which $\mathcal{B}$ outputs as well.

(c) <u>Pseudorandom Case:</u> If $\mathcal{B}$ gets query access to $F_x(\cdot)$, then $y_0||y_1 = f_3(x)$. Then:

$$\Pr_{x \leftarrow \{0,1\}^n}[B^{F_x(\cdot)} \to 1] = \Pr_{x \leftarrow \{0,1\}^n}[\mathcal{A}(f_3(x)) \to 1]$$

(d) <u>Truly Random Case:</u> If $\mathcal{B}$ gets query access to a truly random permutation $R(\cdot)$, then $(y_0||y_1)$ is sampled uniformly at random from all $2n$-bit strings such that the first n bits do not equal the second n bits. The distribution of $(y_0||y_1)$ has negligible statistical distance from the uniform distribution over $\{0,1\}^{2n}$. This is because

$$\Pr_{(y_0||y_1) \leftarrow \{0,1\}^{2n}}[y_0 = y_1] = 2^{-n} = \mathsf{negl}(n)$$

Therefore,

$$\left| \Pr_R[\mathcal{B}^{R(\cdot)} \to 1] - \Pr_{y_0||y_1 \leftarrow \{0,1\}^n}[\mathcal{A}(y_0||y_1) \to 1] \right| \leq 2^{-n}$$

(e) In summary:

$$\begin{aligned}\left| \Pr_x[B^{F_x(\cdot)} \to 1] - \Pr_R[\mathcal{B}^{R(\cdot)} \to 1] \right| &\geq \left| \Pr_{x \leftarrow \{0,1\}^n}[\mathcal{A}(f_3(x)) \to 1] - \Pr_{y_0||y_1 \leftarrow \{0,1\}^n}[\mathcal{A}(y_0||y_1) \to 1] \right| - 2^{-n} \\ &= \mathsf{non\text{-}negl}(n) - \mathsf{negl}(n) = \mathsf{non\text{-}negl}(n)\end{aligned}$$

This means that $\mathcal{B}$ breaks the PRP security of F. But that's a contradiction because we know that F is a secure PRP. Therefore, the initial assumption was false, and in fact f_3 is a PRG.

$\square$