Supplementary Proofs for Lecture 4

CY 4770, Spring 2021

Instructor: Ran Cohen.
TA: Eysa Lee.

1 CPA-Secure Private-Key Encryption from a PRF

Let $F : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^{\ell(n)}$ be a pseudorandom function. In class we defined the following private-key encryption scheme, denoted by Π_F.

- $\mathsf{Gen}(1^n)$ outputs a key $k \leftarrow \{0,1\}^n$.
- Given $m \in \{0,1\}^{\ell(n)}$, $\mathsf{Enc}_k(m)$ samples $r \leftarrow \{0,1\}^n$ and outputs $(r, F_k(r) \oplus m)$.
- Given $c = (r,s) \in \{0,1\}^n \times \{0,1\}^{\ell(n)}$, $\mathsf{Dec}_k(c)$ outputs $F_k(r) \oplus s$.

Theorem 1.1. *Let $F : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^{\ell(n)}$ be a pseudorandom function. Then Π_F is CPA-secure.*

Proof. We will show that for every PPT adversary $\mathcal{A}$ there exists a negligible function $\mathsf{negl}(n)$ such that

$$\Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_F,\mathcal{A}}(n) = 1\right] \leq 1/2 + \mathsf{negl}(n).$$

We start by considering the scheme Π_h which is defined just like Π_F with the exception that instead of the pseudorandom function F, a truly random function $h : \{0,1\}^n \to \{0,1\}^{\ell(n)}$ is used. Note that for every PPT adversary $\mathcal{A}$ it holds that

$$\begin{aligned}\Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_F,\mathcal{A}}(n) = 1\right] &= \Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_F,\mathcal{A}}(n) = 1\right] - \Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_h,\mathcal{A}}(n) = 1\right] + \Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_h,\mathcal{A}}(n) = 1\right] \\ &\leq \left|\Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_F,\mathcal{A}}(n) = 1\right] - \Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_h,\mathcal{A}}(n) = 1\right]\right| + \Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_h,\mathcal{A}}(n) = 1\right],\end{aligned}$$

where the inequality follows from the triangular inequality. Indeed, note that for every $x, y \in \mathbb{R}$ it holds that

$$|x| - |y| = |x - y + y| - |y| \leq |x - y| + |y| - |y| = |x - y|.$$

We proceed to prove the following two claims.

Claim 1.2. *For every PPT adversary $\mathcal{A}$ there exists a negligible function $\mathsf{negl}(n)$ such that*

$$\left|\Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_F,\mathcal{A}}(n) = 1\right] - \Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_h,\mathcal{A}}(n) = 1\right]\right| \leq \mathsf{negl}(n).$$

Proof. Let $\mathcal{A}$ be a PPT adversary and denote by $\varepsilon(n)$ the advantage of $\mathcal{A}$ in winning the experiment $\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_F,\mathcal{A}}(n)$, i.e., for every n

$$\Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_F,\mathcal{A}}(n)=1\right]\leq\varepsilon(n).$$

We will construct a distinguisher D that succeeds in distinguishing between F and a random function h with probability $\epsilon(n)$. The distinguisher D has an oracle access to a function $\mathcal{O}$ that is either the PRF $F_k:\{0,1\}^n\to\{0,1\}^{\ell(n)}$ for a random $k\leftarrow\{0,1\}^n$, or a truly random function $h:\{0,1\}^n\to\{0,1\}^{\ell(n)}$.

1. On input 1^n, the distinguisher D invokes $\mathcal{A}$ on 1^n.
2. Whenever $\mathcal{A}$ queries its encryption orale for the i^{th} time with a message $\mathsf{msg}_i\in\{0,1\}^{\ell(n)}$, D samples a random $r_i\leftarrow\{0,1\}^n$, queries its oracle on r_i, and returns to $\mathcal{A}$ the ciphertext $(r_i,\mathcal{O}(r_i)\oplus\mathsf{msg}_i)$.
3. Once $\mathcal{A}$ returns a pair of messages (m_0,m_1), D samples a random $b\leftarrow\{0,1\}$ and a random $r^*\leftarrow\{0,1\}^n$, queries its oracle on r^* and returns the ciphertext $(r^*,\mathcal{O}(r^*)\oplus m_b)$.
4. Once $\mathcal{A}$ outputs a bit b', D outputs 1 if $b=b'$ and outputs 0 otherwise.

D clearly runs in polynomial time since $\mathcal{A}$ is PPT.

We proceed to analyze the behaviour of D.

- **Case 1: $\mathcal{O}=F_k$ for $k\leftarrow\{0,1\}^n$.** In this case, the view of $\mathcal{A}$ when invoked by D is identically distributed as in the experiment $\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_F,\mathcal{A}}(n)$. Therefore

$$\Pr\left[D^{F_k(\cdot)}(1^n)=1\right]=\Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_F,\mathcal{A}}(n)=1\right]. \tag{1}$$

- **Case 2: $\mathcal{O}=h$ for a truly random function.** In this case, the view of $\mathcal{A}$ when invoked by D is identically distributed as in the experiment $\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_h,\mathcal{A}}(n)$, Therefore,

$$\Pr\left[D^{h(\cdot)}(1^n)=1\right]=\Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_h,\mathcal{A}}(n)=1\right]. \tag{2}$$

Combining Equations 1, and 2 we get that

$$\begin{aligned}\left|\Pr\left[D^{F_k(\cdot)}(1^n)=1\right]-\Pr\left[D^{h(\cdot)}(1^n)=1\right]\right|&=\left|\Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_F,\mathcal{A}}(n)=1\right]-\Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_h,\mathcal{A}}(n)=1\right]\right|\\&=\varepsilon(n).\end{aligned}$$

By the assumption that F is a PRF, and since D runs in probabilistic polynomial time, it holds that $\varepsilon(n)$ is negligible. □

Claim 1.3. *Let $q(n)$ denote an upper bound on the number of queries made by $\mathcal{A}(1^n)$ to the encryption oracle. Then,*

$$\Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_h,\mathcal{A}}(n)=1\right]\leq\frac{1}{2}+\frac{q(n)}{2^n}.$$

Proof. Recall that we denote by r^* the random string sampled by D to generate the challenge ciphertext, and by r_i the random string sampled by D to answer the i^{th} encryption oracle query made by $\mathcal{A}$. We will denote by REPEAT the event that there exists i for which $r_i = r^*$. Note that for every i, the probability that $r_i = r^*$ is 2^{-n}, and therefore, by union bound,

$$\Pr[\text{REPEAT}] = \Pr[\exists i \text{ such that } r_i = r^*] \leq \sum_{i=1}^{q(n)} \Pr[r_i = r^*] = \frac{q(n)}{2^n}.$$

Next, note that if the event REPEAT does not occur, i.e., the randomness r^* used for the challenge ciphertext is not used during the encryption queries made by $\mathcal{A}$, then the challenge ciphertext is essentially encrypted using a one-time pad, hence

$$\Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_h,\mathcal{A}}(n) = 1 \mid \overline{\text{REPEAT}}\right] = \frac{1}{2}.$$

By the law of total probability it holds that

$$\begin{aligned}
\Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_h,\mathcal{A}}(n) = 1\right] &= \Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_h,\mathcal{A}}(n) = 1 \mid \overline{\text{REPEAT}}\right] \cdot \Pr[\overline{\text{REPEAT}}] \\
&\quad + \Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_h,\mathcal{A}}(n) = 1 \mid \text{REPEAT}\right] \cdot \Pr[\text{REPEAT}] \\
&\leq \Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_h,\mathcal{A}}(n) = 1 \mid \overline{\text{REPEAT}}\right] + \Pr[\text{REPEAT}] \\
&\leq \frac{1}{2} + \frac{q(n)}{2^n}.
\end{aligned}$$

□

By the two claims it holds that

$$\begin{aligned}
\Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_F,\mathcal{A}}(n) = 1\right] &\leq \left|\Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_F,\mathcal{A}}(n) = 1\right] - \Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_h,\mathcal{A}}(n) = 1\right]\right| + \Pr\left[\mathsf{PrivK}^{\mathsf{CPA}}_{\Pi_h,\mathcal{A}}(n) = 1\right] \\
&\leq \frac{1}{2} + \left(\mathsf{negl}(n) + \frac{q(n)}{2^n}\right).
\end{aligned}$$

Since this holds for an arbitrary PPT adversary $\mathcal{A}$, we conclude that the scheme Π_F is CPA-secure. □