EE 720: Introduction to Number Theory and Cryptography (Autumn 2025)
Instructor: Saravanan Vijayakumaran
Indian Institute of Technology Bombay

Quiz 1: 20 points Date: September 4, 2025

---

1. [5 points] Prove that the one-time pad is perfectly secret.

   Note 1: The one-time pad has $\mathcal{M} = \mathcal{K} = \mathcal{C} = \{0,1\}^n$. Gen chooses key k uniformly from $\mathcal{K}$. $\mathsf{Enc}_k(m) = k \oplus m$ and $\mathsf{Dec}_k(c) = k \oplus c$.

   Note 2: An encryption scheme (Gen, Enc, Dec) with message space $\mathcal{M}$ is **perfectly secret** if for every probability distribution over $\mathcal{M}$, every message $m \in \mathcal{M}$, and every ciphertext $c \in \mathcal{C}$ for which $\Pr[C = c] > 0$:

$$\Pr[M = m \mid C = c] = \Pr[M = m].$$

2. [5 points] Let $\Pi = (\mathtt{Gen}, \mathtt{Enc}, \mathtt{Dec})$ be an encryption scheme with message space $\mathcal{M}$ and key space $\mathcal{K}$. Suppose that $\mathtt{Enc}$ is a deterministic function of the key $k \in \mathcal{K}$ and message $m \in \mathcal{M}$. Prove that Π is **not CPA-secure**.

   Note: A private-key encryption scheme $\Pi = (\mathtt{Gen}, \mathtt{Enc}, \mathtt{Dec})$ has indistinguishable encryptions under a chosen-plaintext attack, or is **CPA-secure**, if for all probabilistic polynomial-time adversaries $\mathcal{A}$ there is a negligible function $\mathtt{negl}$ such that, for all n,

$$\Pr\left[\mathtt{PrivK}^{\mathtt{cpa}}_{\mathcal{A},\Pi}(n) = 1\right] \leq \frac{1}{2} + \mathtt{negl}(n).$$

   Specification of the experiment $\mathtt{PrivK}^{\mathtt{cpa}}_{\mathcal{A},\Pi}(n)$:

   1. A key k is generated by running $\mathtt{Gen}(1^n)$.
   2. The adversary $\mathcal{A}$ is given 1^n and oracle access to $\mathtt{Enc}_k(\cdot)$, and outputs a pair of messages $m_0, m_1 \in \mathcal{M}$ with $|m_0| = |m_1|$.
   3. A uniform bit $b \in \{0,1\}$ is chosen. Ciphertext $c \leftarrow \mathtt{Enc}_k(m_b)$ is computed and given to $\mathcal{A}$.
   4. The adversary $\mathcal{A}$ continues to have oracle access to $\mathtt{Enc}_k(\cdot)$, and outputs a bit b'.
   5. The output of the experiment is defined to be 1 if $b' = b$, and 0 otherwise. If output is 1, we say that $\mathcal{A}$ succeeds.

3. [5 points] Alice has a length-preserving pseudorandom function $F : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$. She wants to encrypt messages of length $2n$. Let $m \in \{0,1\}^{2n}$ denote the message. Let $m_1 \in \{0,1\}^n$ denote the first n bits of m and let $m_2 \in \{0,1\}^n$ denote the last n bits of m. Alice uses the encryption scheme $\Pi = (\mathtt{Gen}, \mathtt{Enc}, \mathtt{Dec})$ where:

   - $\mathtt{Gen}$: Key k is chosen uniformly from $\{0,1\}^n$.
   - $\mathtt{Enc}$: The message space $\mathcal{M} = \{0,1\}^{2n}$. A string r is chosen uniformly from $\{0,1\}^n$ and the ciphertext $c \in \{0,1\}^{3n}$ corresponding to $m = (m_1, m_2) \in \{0,1\}^{2n}$ is given by

$$c := \langle r, m_1 \oplus F_k(r), m_2 \oplus F_k(r) \rangle.$$

   - $\mathtt{Dec}$: Given key k and ciphertext $c = \langle r, c_1, c_2 \rangle \in \{0,1\}^{3n}$, the message $m = (m_1, m_2)$ is decrypted using $m_1 = c_1 \oplus F_k(r)$ and $m_2 = c_2 \oplus F_k(r)$.

   Prove that Alice's scheme is **not** EAV-secure.

   Note: A private-key encryption scheme $\Pi = (\mathtt{Gen}, \mathtt{Enc}, \mathtt{Dec})$ has indistinguishable encryptions in the presence of an eavesdropper, or is **EAV-secure**, if for all probabilistic polynomial-time adversaries $\mathcal{A}$ there is a negligible function $\mathtt{negl}$ such that, for all n,

$$\Pr\left[\mathtt{PrivK}^{\mathtt{eav}}_{\mathcal{A},\Pi}(n) = 1\right] \leq \frac{1}{2} + \mathtt{negl}(n).$$

   Specification of the experiment $\mathtt{PrivK}^{\mathtt{eav}}_{\mathcal{A},\Pi}(n)$:

1. The adversary $\mathcal{A}$ is given 1^n and outputs a pair of arbitrary messages $m_0, m_1 \in \mathcal{M}$ with $|m_0| = |m_1|$.
2. A key k is generated using $\texttt{Gen}$, and a uniform bit $b \in \{0,1\}$ is chosen. Ciphertext $c \leftarrow \texttt{Enc}_k(m_b)$ is computed and given to $\mathcal{A}$. This ciphertext c is called the *challenge ciphertext*.
3. $\mathcal{A}$ outputs a bit b'.
4. The output of the experiment is defined to be 1 if $b' = b$, and 0 otherwise. We write $\texttt{PrivK}^{\texttt{eav}}_{\mathcal{A},\Pi}(n) = 1$ if the output of the experiment is 1 and in this case we say that $\mathcal{A}$ succeeds.

4. [5 points] Alice has a length-preserving pseudorandom function $F : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$. She wants to authenticate messages of length $2n$. Let $m \in \{0,1\}^{2n}$ denote the message. Let $m_1 \in \{0,1\}^n$ denote the first n bits of m and let $m_2 \in \{0,1\}^n$ denote the last n bits of m. Alice uses the encryption scheme $\Pi = (\texttt{Gen}, \texttt{Mac}, \texttt{Vrfy})$ where:

   - $\texttt{Gen}$: Key k is chosen uniformly from $\{0,1\}^n$.
   - $\texttt{Mac}$: The message space $\mathcal{M} = \{0,1\}^{2n}$. Given key k, the tag corresponding to $m = (m_1, m_2) \in \{0,1\}^{2n}$ is given by a $2n$-bit string t as follows.

   $$\mathsf{Mac}_k(m) = t := \langle t_1, t_2 \rangle = \langle F_k(m_1), F_k(m_1 \oplus m_2) \rangle.$$

   - $\texttt{Vrfy}$: Given key k, message $m = (m_1, m_2)$ and tag $t = \langle t_1, t_2 \rangle$, the verifier recomputes the tag on the message

   $$t' := \langle t'_1, t'_2 \rangle = \langle F_k(m_1), F_k(m_1 \oplus m_2) \rangle.$$

   If $t = t'$, then the verifier outputs 1 and 0 otherwise.

   Prove that Alice's MAC scheme is **not** secure.

   Note: A message authentication code $\Pi = (\texttt{Gen}, \texttt{Mac}, \texttt{Vrfy})$ is existentially unforgeable under an adaptive chosen-message attack, or just **secure**, if for all PPT adversaries $\mathcal{A}$, there is a negligible function $\texttt{negl}$ such that:

   $$\Pr\left[\texttt{Mac-forge}_{\mathcal{A},\Pi}(n) = 1\right] \leq \texttt{negl}(n).$$

   The message authentication experiment $\texttt{Mac-forge}_{\mathcal{A},\Pi}(n)$ is defined as follows:

   1. A key k is generated by running $\texttt{Gen}(1^n)$.
   2. The adversary $\mathcal{A}$ is given input 1^n and oracle access to $\texttt{Mac}_k(\cdot)$. The adversary eventually outputs (m, t). Let $\mathcal{Q}$ denote the set of all queries that $\mathcal{A}$ asked its oracle.
   3. $\mathcal{A}$ succeeds if and only if (1) $\texttt{Vrfy}_k(m, t) = 1$ and (2) $m \notin \mathcal{Q}$. If $\mathcal{A}$ succeeds, the output of the experiment is 1. Otherwise, the output is 0.