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

Mid-semester Exam: 24 points — Date: September 15, 2025

1. [3 points] If $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is a perfectly secret encryption scheme with message space $\mathcal{M}$ and key space $\mathcal{K}$, then prove that $|\mathcal{K}| \geq |\mathcal{M}|$.

   Note: 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. The shift cipher has message space $\mathcal{M} = \{0, 1, \ldots, 25\}^n$, ciphertext space $\mathcal{C} = \{0, 1, \ldots, 25\}^n$, and keyspace $\mathcal{K} = \{0, 1, \ldots, 25\}$ where $n \geq 1$. For $m_1 m_2 \ldots m_n \in \mathcal{M}$, the ciphertext is given by $c_1 c_2 \ldots c_n$ where $c_i = m_i + k \bmod 26$ and k is chosen uniformly from $\mathcal{K}$.

   (a) [2 points] For $n = 1$, show that the shift cipher is perfectly secret.

   (b) [1 point] For $n > 1$, by question 1 the shift cipher cannot be perfectly secret. Give an example of a probability distribution over $\mathcal{M}$, a message $m \in \mathcal{M}$, and a ciphertext $c \in \mathcal{C}$ for which $\Pr[C = c] > 0$ such that $\Pr[M = m \mid C = c] \neq \Pr[M = m]$.

3. [6 points] If $F : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$ is a length-preserving keyed pseudorandom function, then prove that the below construction is a CPA-secure private-key encryption scheme for messages of length n.

   - `Gen`: On input 1^n, choose k uniformly from $\{0,1\}^n$.
   - `Enc`: Given $k \in \{0,1\}^n$ and message $m \in \{0,1\}^n$, choose uniform $r \in \{0,1\}^n$ and output the ciphertext
   $$c := \langle r, F_k(r) \oplus m \rangle.$$
   - `Dec`: Given $k \in \{0,1\}^n$ and ciphertext $c = \langle r, s \rangle$, output the plaintext message
   $$m := F_k(r) \oplus s.$$

   Note 1: F is a pseudorandom function if for any PPT distinguisher D, there is a negligible function $\mathtt{negl}$ such that:
   $$\left| \Pr\left[D^{F_k(\cdot)}(1^n) = 1\right] - \Pr\left[D^{f(\cdot)}(1^n) = 1\right] \right| \leq \mathtt{negl}(n),$$
   where the first probability is taken over uniform choice of $k \in \{0,1\}^n$ and the randomness of D, and the second probability is taken over uniform choice of $f \in \mathtt{Func}_n$ and the randomness of D. The set $\mathtt{Func}_n$ is the set of all functions with domain and range equal to $\{0,1\}^n$. By $D^{F_k(\cdot)}(1^n)$ and $D^{f(\cdot)}(1^n)$, we mean distinguishers D who have oracle access to F_k and f respectively.

   Note 2: 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.

4. Recall that the PKCS #7 padding scheme is used to pad a message $\vec{x}$ having length some integral number of bytes into a *encoded data* $\vec{m}$ having length jL bytes where L is the block length in bytes. The number of bytes which are appended to $\vec{x}$ to get $\vec{m}$ is b where $1 \leq b \leq L$. Each of these padding bytes is equal to the byte representation of the integer b. Assume that $L < 256$ so b can fit in a single byte.

   Suppose the encoded data $\vec{m}$ has length $2L$ bytes, i.e. $\vec{m} = (m_1, m_2)$ where $|m_i| = L$ bytes for $i = 1, 2$. Now suppose the encoded data is encrypted using CTR mode where F is a length-preserving pseudorandom function as shown below. The input and output lengths of F_k are both equal to $n = 8L$ bits. Here the value IV is uniformly chosen from $\{0,1\}^{\frac{3n}{4}}$.

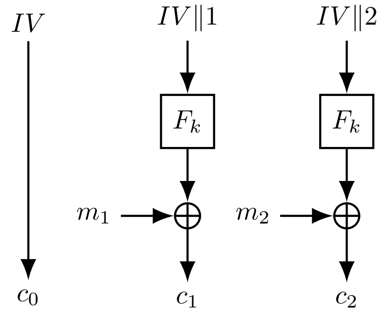

Suppose an adversary can observe valid ciphertexts (c_0, c_1, c_2). Additionally, the adversary has access to a padding oracle. On input some ciphertext block $\vec{c} = (c'_0, c'_1, c'_2)$, the padding oracle only returns a message from the set $\{\mathsf{ok}, \mathsf{padding\_error}\}$. The ok is returned when there is no padding error in the encoded data $\vec{m}'$ obtained from $\vec{c}$. Note that the adversary can also send $\vec{c} = (c'_0, c'_1)$ to the padding oracle.

(a) [2 points] Describe a procedure by which the adversary can recover the **length** b of the padding in the encoded data $\vec{m}$.

(b) [2 points] Suppose $b \leq L - 2$. Describe a procedure by which the adversary can recover the **last two message bytes** in m_2. By last two bytes, we mean the rightmost two non-padding bytes in m_2.

(c) [2 points] Describe a procedure by which the adversary can recover the **last message byte** in m_1. By last byte, we mean the rightmost byte in m_1.

5. (a) [1 point] Show that the encryption scheme in question 3 is not CCA-secure.

   (b) [1 point] Show that the CBC block cipher mode encryption scheme is not CCA-secure.
   Note: Cipher Block Chaining (CBC) mode works as follows:
   - Let $m = m_1, m_2, \ldots, m_l$ where $m_i \in \{0,1\}^n$.
   - Let F be a length-preserving pseudorandom permutation with block length n.
   - A uniform *initialization vector (IV)* of length n is first chosen.
   - Set $c_0 = IV$. For $i = 1, \ldots, l$, set $c_i := F_k(c_{i-1} \oplus m_i)$. The ciphertext is $(c_0, c_1, \ldots, c_l)$.
   - For $i = 1, 2, \ldots, l$, $m_i := F_k^{-1}(c_i) \oplus c_{i-1}$.

   (c) [1 point] Show that the CTR block cipher mode encryption scheme is not CCA-secure.
   Note: Counter (CTR) mode works as follows:
   - Let $m = m_1, m_2, \ldots, m_l$ where $m_i \in \{0,1\}^n$.
   - Let F be a length-preserving pseudorandom function with block length n.
   - A uniform value IV of length $\frac{3n}{4}$ is first chosen.
   - Set $c_0 = IV$. For $i = 1, \ldots, l$, set $c_i := F_k(IV\|i) \oplus m_i$. The ciphertext is $(c_0, c_1, \ldots, c_l)$.
   - For $i = 1, 2, \ldots, l$, $m_i := F_k(IV\|i) \oplus c_i$.

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

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

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

1. A key k is generated by running $\mathsf{Gen}(1^n)$.
2. The adversary $\mathcal{A}$ is given 1^n and oracle access to $\mathsf{Enc}_k(\cdot)$ and $\mathsf{Dec}_k(\cdot)$, 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. Challenge ciphertext $c^* \leftarrow \mathsf{Enc}_k(m_b)$ is computed and given to $\mathcal{A}$.
4. The adversary $\mathcal{A}$ continues to have oracle access to $\mathsf{Enc}_k(\cdot)$ and $\mathsf{Dec}_k(\cdot)$, but is not allowed to query Dec_k on the challenge ciphertext itself. Eventually, $\mathcal{A}$ 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.

6. [3 points] Using a length-preserving pseudorandom function $F : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$, construct a CCA-secure encryption scheme for messages of length n. Describe the algorithms $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$. You can use results discussed in class without proof.