EE 720: Introduction to Number Theory and Cryptography (Autumn 2024)
Instructor: Saravanan Vijayakumaran
Indian Institute of Technology Bombay
Midsemester Exam: 25 points Date: September 20, 2024

1. [5 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 1: 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. [10 points] Let F be a length-preserving **strong** pseudorandom permutation having key length, input length, and output length all equal to n bits. Suppose a fixed-length private key encryption scheme $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is defined as follows:

   - Gen: Key k is chosen uniformly from $\{0,1\}^n$.
   - Enc: The message space $\mathcal{M} = \{0,1\}^{n/2}$. A string r is chosen uniformly from $\{0,1\}^{n/2}$ and the ciphertext $c \in \{0,1\}^n$ corresponding to $m \in \{0,1\}^{n/2}$ is given by
     $$c := F_k(r \| m).$$
     Here $\|$ is the string concatenation operator.
   - Dec: Given key k and ciphertext $c \in \{0,1\}^n$, the message m is obtained by taking the last $n/2$ bits of $F_k^{-1}(c)$.

   Prove that Π is CCA-secure for messages of length $n/2$.

   Note 1: 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.

   Note 2: F is a **strong** pseudorandom permutation if for any PPT distinguisher D, there is a negligible function negl such that:
   $$\left|\Pr\left[D^{F_k(\cdot), F_k^{-1}(\cdot)}(1^n) = 1\right] - \Pr\left[D^{f(\cdot), f^{-1}(\cdot)}(1^n) = 1\right]\right| \leq \mathsf{negl}(n),$$
   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 $f \in \mathsf{Perm}_n$ and the randomness of D. The set Perm_n is the set of all bijections with domain and range equal to $\{0,1\}^n$. By $D^{F_k(\cdot), F_k^{-1}(\cdot)}(1^n)$ and $D^{f(\cdot), f^{-1}(\cdot)}(1^n)$, we mean distinguishers D who have oracle access to F_k, F_k^{-1} and f, f^{-1} respectively.

3. Recall that the PKCS #7 padding scheme is used to pad a message $\vec{x}$ having length some integral number of bytes into an *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 output feedback (OFB) 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 $\{0,1\}^n$.

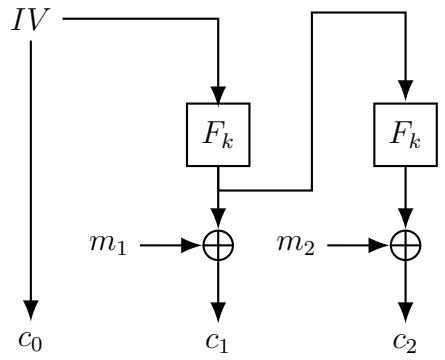

   Suppose an adversary has access to a padding oracle. On input some ciphertext block $\vec{c} = (c'_0, c'_1, c'_2)$, the padding oracle only returns a message from the set $\{\mathsf{ok}, \mathsf{padding\_error}\}$. The ok is returned when there is no padding error in the encoded data $\vec{m}'$ obtained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

   (a) [1 point] 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.

4. Suppose $F : \{0,1\}^n \times \{0,1\}^n \mapsto \{0,1\}^n$ is a length-preserving pseudorandom function. **Using only F, give constructions** of schemes which satisfy the following properties. You can use results discussed in class without proof.

   (a) [1 point] A CPA-secure encryption scheme for messages of length n which is not CCA-secure. Describe the algorithms $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$.

   (b) [1 point] A CPA-secure encryption scheme for messages of length $3n$. Describe the algorithms $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$.

   (c) [1 point] A secure MAC for messages of length $3n$. Describe the algorithms $(\mathsf{Gen}, \mathsf{Mac}, \mathsf{Vrfy})$.

   (d) [2 points] A CCA-secure encryption scheme for messages of length $3n$. Describe the algorithms $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$.