Tutorial 1: Perfect Secrecy

Submission Guidelines All problems must be solved in class today. Searching on the internet for solutions is strictly discouraged.

1. The shift cipher (also called Caesar cipher) works as follows. The English alphabet is represented by numbers from 0 to 25 i.e., $\{A, B, \dots, Z\}$ are mapped to $\{0, 1, \dots, 25\}$ in the same order.

   Define $\mathscr{K} = \{0, 1, 2, \dots, 25\}$, $\mathscr{M} = \mathscr{C} = \{0, 1, 2, \dots, 25\}^*$
   $\mathsf{Gen}()$: $k \xleftarrow{\text{U}} \mathscr{K}$
   $\mathsf{Enc}(k, \mathbf{m} = m_1 m_2 \cdots m_n)$: Set $c_i \leftarrow m_i + k \bmod 26$. Ciphertext is given by $\mathbf{c} = c_1 c_2 \cdots c_n$.
   $\mathsf{Dec}(k, \mathbf{c} = c_1 c_2 \cdots c_n)$: Recover message components as $m_i \leftarrow c_i - k \bmod 26$.

   (a) Is it perfectly secret?

   A: It is not perfectly secret. In fact, it is not secure at all. Trying all 26 possibilities for the key will help in deciphering the message.

   (b) Can you modify the description to make it perfectly secret?

   A: The scheme becomes perfectly secure when only one character is encrypted i.e., $\mathscr{M} = \mathscr{C} = \{0, 1, \dots, 25\}$ (in other words, a different key is chosen for each letter). We have, for arbitrary $m \in \mathscr{M}$ and $c \in \mathscr{C}$,

   $$\Pr[\mathrm{C} = c | \mathrm{M} = m] = \Pr[\mathrm{K} = (c - m) \bmod 26] = \frac{1}{26}$$

   and

   $$\begin{aligned} \Pr[\mathrm{C} = c] &= \sum_{k \in \mathscr{K}} Pr[\mathrm{K} = k] \cdot \Pr[\mathrm{M} = \mathsf{Dec}(k, c)] \\ &= \sum_{k \in \mathscr{K}} \frac{1}{26} \cdot \Pr[\mathrm{M} = c - k \bmod 26] \\ &= \frac{1}{26} \sum_{m \in \mathscr{M}} \Pr[\mathrm{M} = m \bmod 26] \\ &= \frac{1}{26} \end{aligned}$$

   In other words, we have $\Pr[\mathrm{C} = c | \mathrm{M} = m] = \Pr[\mathrm{C} = c]$ for any $m \in \mathscr{M}$ and $c \in \mathscr{C}$. Therefore the modified scheme is perfectly secret.

2. Prove or refute: Every encryption scheme for which the size of the key space $\mathscr{K}$ equals the size of the message space $\mathscr{M}$ and for which the key is chosen uniformly from $\mathscr{K}$, is perfectly secret.

   A: The statement is false. We show this by providing a counter example. Define $\mathscr{M} = \{a, b\}$, $\mathscr{K} = \{k_1, k_2\}$, $\mathscr{C} = \{0, 1\}$. Let $\mathsf{Enc}(k, a) = 0$ and $\mathsf{Enc}(k, b) = 1$ for $k = k_1, k_2$. Dec algorithm will return a on input ciphertext 0 and b on input ciphertext 1. Clearly, the scheme is correct.

   $$\Pr[\mathrm{C} = 1 | \mathrm{M} = a] = 1 \neq 0 = \Pr[\mathrm{C} = 1 | \mathrm{M} = b],$$

   thus showing that the scheme is not perfectly secret.

3. Let $\varepsilon < 1$ be a constant. An encryption scheme $\mathcal{E} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ over $(\mathscr{K}, \mathscr{M}, \mathscr{C})$ is called ε-almost perfectly secret if for any distribution over $\mathscr{M}$, any $m \in \mathscr{M}$ and any $c \in \mathscr{C}$,

   $$|\Pr[\mathrm{M} = m | \mathrm{C} = c] - \Pr[\mathrm{M} = m]| < \varepsilon.$$

Show that if $\mathcal{E}$ is an ε-almost perfectly secret encryption scheme, then $|\mathcal{K}| \geq (1-\varepsilon)|\mathcal{M}|$.

A: Suppose that $|\mathcal{K}| < (1-\varepsilon)|\mathcal{M}|$. Consider the uniform distribution on $\mathcal{M}$ and let $c \in \mathcal{C}$ be a ciphertext that occurs with non-zero probability. Let $\mathcal{M}(c)$ denote the set of all possible decryptions of c. That is,

$$\mathcal{M}(c) = \{m : m = \mathsf{Dec}(k, c) \text{ for some } k \in \mathcal{K}\}.$$

Clearly, $|\mathcal{M}(c)| \leq \mathcal{K}$ and from our assumption, $|\mathcal{M}(c)| < (1-\varepsilon)|\mathcal{M}(c)|$. So there exists $m' \in \mathcal{M} \setminus \mathcal{M}(c)$. We have

$$\Pr[\mathrm{M} = m' | \mathrm{C} = c] = 0.$$

Also,

$$\Pr[\mathrm{M} = m'] \geq \frac{\varepsilon|\mathcal{M}|}{|\mathcal{M}|},$$

since the messages are uniformly distributed and m' could be any message out of a fraction of at least ε of the messages. Combining the above we have

$$|\Pr[\mathrm{M} = m' | \mathrm{C} = c] - \Pr[\mathrm{M} = m']| \geq \varepsilon,$$

contradicting the fact that $\mathcal{E}$ is ε-almost perfectly secret.

4. Modify the one-time pad as follows. Let $\mathcal{K} \subseteq \{0,1\}^\ell$ and the key generation be such that first a key $\tilde{k}$ is uniformly chosen from $\{0,1\}^{\ell/2}$ and then the key is defined by $k = \tilde{k} \| \tilde{k}$ where $\|$ denotes concatenation. (In addition, define $\mathcal{M} = \mathcal{C} = \{0,1\}^\ell$.) Is this scheme perfectly secure? Justify your answer.

A: No, this scheme is not perfectly secure. Let $m_0 = 0^{\ell/2} \| 0^{\ell/2}$, $m_1 = 1^{\ell/2} \| 0^{\ell/2}$ and $c = 0^{\ell/2} \| 1^{\ell/2}$. We have

$$\Pr[\mathrm{C} = c | \mathrm{M} = m_0] = \Pr[\mathrm{K} = c \oplus m_0] = \Pr[\mathrm{K} = 0^{\ell/2} \| 1^{\ell/2}] = 0,$$

given the way keys are generated. On the other hand,

$$\Pr[\mathrm{C} = c | \mathrm{M} = m_1] = \Pr[\mathrm{K} = c \oplus m_1] = \Pr[\mathrm{K} = 1^{\ell/2} \| 1^{\ell/2}] = \frac{1}{2^{\ell/2}}.$$

Hence, there exist $m_0, m_1 \in \mathcal{M}$ and $c \in \mathcal{C}$ such that

$$\Pr[\mathrm{C} = c | \mathrm{M} = m_0] \neq \Pr[\mathrm{C} = c | \mathrm{M} = m_1],$$

implying that the scheme is not perfectly secret.