COMP 8920: Cryptography Winter 2025

Lecture 6 — Jan 23, 2025

Prof. Curtis Bright *Scribe: Aidan Bennett*

Overview

In the last lecture, we introduced the notion of perfect security using conditional probability.

In this lecture, we continue to discuss perfect security, and some properties of perfectly secure cryptosystems.

Perfect security (continued)

Recall the following example from the previous lecture; consider the cryptosystem where $\mathcal{P} = \{a, b\}$, $\mathcal{C} = \{1, 2, 3, 4\}$, $\mathcal{K} = \{k_1, k_2, k_3\}$, and encryption is done according to Table 1. Assume that $\Pr[a] = \frac{1}{4}$, $\Pr[b] = \frac{3}{4}$, $\Pr[k_1] = \frac{1}{2}$, $\Pr[k_2] = \frac{1}{4}$, $\Pr[k_3] = \frac{1}{4}$.

What is the probability that a given message sent was a, given that the ciphertext is 3? Note that $\{k : 3 = e_k(a)\} = \{k_3\}$, so

$$\Pr[a \mid 3] = \frac{\Pr[a] \sum_{e_k(a)=3} \Pr[\mathbb{K} = k]}{\Pr[3]} = \frac{\frac{1}{4} \cdot \frac{1}{4}}{\frac{1}{4}} = \frac{1}{4} = \Pr[a]$$

Eve has no new information; it is just as likely as before that a was sent. A cryptosystem has **perfect security** if this is true for all $x \in \mathcal{P}, y \in \mathcal{C}$.

We can show that the shift cipher has perfect security when the key is at least as long as the plaintext. Even with exhaustive key search, an attacker could not deduce any information about what was encrypted; i.e., if Alice were to send 'c' after applying the shift cipher, Eve could not determine what letter was originally sent, even after exhaustively searching the keyspace.

Theorem 1. *Suppose the 26 keys in the shift cipher are used with equal probability. Then for any distribution of plaintexts, the shift cipher has perfect security.*

Proof. Recall that $\mathcal{C} = \mathcal{P} = \mathcal{K} = \mathbb{Z}_{26}$, and for $0 \le k < 26$, $e_k(x) = x + k$, and $d_k(y) = y - k$. Let $y \in \mathcal{C}$. Then

$$\begin{aligned}\Pr[y] &= \sum_{\{k : y \in C(k)\}} \Pr[\mathbb{K} = k] \Pr[\mathbf{X} = d_k(y)] \\ &= \frac{1}{26} \sum_{k \in \mathbb{Z}_{26}} \Pr[\mathbf{X} = y - k]\end{aligned}$$

	a	c
k_1	1	2
k_2	2	3
k_3	3	4

Table 1: Encryption matrix

since $\{k : y \in C(k)\} = \mathcal{K} = \mathbb{Z}_{26}$.
$y - k$ takes all possible values $x \in \mathbb{Z}_{26}$ as k ranges over all $k \in \mathbb{Z}_{26}$, so

$$\sum_{k \in \mathbb{Z}_{26}} \Pr[\mathbf{X} = y - k] = 1$$

and thus $\Pr[y] = \frac{1}{26}$. Next,

$$\Pr[y \mid x] = \sum_{\{k : y = e_k(x)\}} \Pr[\mathbb{K} = k] = \frac{1}{26}$$

Finally by Baye's Theorem,

$$\begin{aligned} \Pr[x \mid y] &= \frac{\Pr[x]\Pr[y \mid x]}{\Pr[y]} \\ &= \frac{\Pr[x] \cdot \frac{1}{26}}{\frac{1}{26}} \\ &= \Pr[x] \end{aligned}$$

Therefore, it is perfectly secure. □

Exhaustive key search doesn't help an attacker, as trying all keys k will generate all possible plaintexts x.

If a system is perfectly secure and we suppose $\Pr[y] > 0$ for all possible ciphertexts y, then we must have $|\mathcal{K}| \geq |\mathcal{C}|$. This is because Baye's theorem says that if $\Pr[x \mid y] = \Pr[x]$, then $\Pr[y \mid x] = \Pr[y]$. Since $\Pr[y] > 0$, we have $\Pr[y \mid x] > 0$. Then given any $x \in \mathcal{P}$, there is a $k \in \mathcal{K}$ such that $e_k(x) = y$. Therefore, $|\mathcal{K}| \geq |\mathcal{C}|$, as otherwise there would be some y with no k sending an x to y.

Also, since e_k is injective in any cryptosystem, $|\mathcal{C}| \geq |\mathcal{P}|$.

Theorem 2. *Let $(\mathcal{P}, \mathcal{C}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ be a cryptosystem. If $|\mathcal{P}| = |\mathcal{C}| = |\mathcal{K}|$, then the cryptosystem is perfectly secure if and only if every key has probability $\frac{1}{|\mathcal{K}|}$, and for all $(x, y) \in \mathcal{P} \times \mathcal{C}$, there is a unique key k such that $e_k(x) = y$.*

Proof. Suppose the cryptosystem is perfectly secure. Then $\forall (x, y) \in \mathcal{P} \times \mathcal{C}$, $\exists k \in \mathcal{K}$ such that $e_k(x) = y$. So $|\mathcal{C}| = |\{e_k(x) : k \in \mathcal{K}\}|$ for all $x \in \mathcal{P}$.
But $|\mathcal{C}| = |\mathcal{K}|$, so $|\mathcal{K}| = \{e_k(x) : k \in \mathcal{K}\}$, and thus cannot have two distinct keys k_1, k_2 such that $e_{k_1}(x) = e_{k_2}(x)$. So there is exactly one k for which $e_k(x) = y$.

We now must show $\Pr[k] = \frac{1}{n}$, where $n = |\mathcal{K}|$. Let $\mathcal{P} = \{x_i : 1 \leq i \leq n\}$. We can write $\mathcal{K} = \{k_1, \ldots, k_n\}$, such that $e_{k_i}(x_i) = y$, for some fixed y. From Baye's theorem,

$$\Pr[x_i \mid y] = \frac{\Pr[y \mid x_i]\Pr[x_i]}{\Pr[y]} = \frac{\Pr[\mathbb{K} = k_i]\Pr[x_i]}{\Pr[y]} = \frac{\Pr[\mathbb{K} = k_i]\Pr[x_i \mid y]}{\Pr[y]}$$

Note that the last equation holds due to the perfect security condition. Cancelling $\Pr[x_i \mid y]$, we find that $\Pr[\mathbb{K} = k_i] = \Pr[y]$. This has to be fixed as $\frac{1}{n}$, since $\sum_{i=1}^{n} \Pr[k_i] = 1$.

Proof continued in next lecture.