Perfectly Secret Encryption

Jacob Benjamin Cholewa

Stud. MSc. Software Development

March 2017

Outline

Outline

Definitions

Common terms

- Gen is a probabilistic algorithm that outputs some $k \in \mathcal{K}$.
- $Enc_k(m)$ takes some message $m \in \mathcal{M}$ and encrypts it with k.
- $Dec_k(c)$ takes some ciphertext $c \in \mathcal{C}$ and decrypts it with k.

Definitions

Probability

- The probability that *Gen* outputs some k is denoted by $Pr[K = k] = p$
- Perfect correctness entails that $Pr\left[Dec_k(Enc_k(m)) = m\right] = 1$

Definitions

Perfect Secrecy

Definition 2.3

An encryption schema (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 $Pr[C = c] > 0$:

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

Alternative definition

For every $m, m' \in \mathcal{M}$ and every $c \in \mathcal{C}$

$$Pr[Enc_k(m) = c] = Pr[Enc_k(m') = c]$$

Outline

Definitions

Perfect indistinguishability

The adversarial indistinguishability experiment $PrivK^{eav}_{\mathcal{A},\Pi}$

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

Definitions

Perfect indistinguishability (continued)

Definition 2.5

Encryption schema $\Pi = (Gen, Enc, Dec)$ with message space $\mathcal{M}$ is perfectly indistinguishable if for every $\mathcal{A}$ it holds that

$$Pr\left[PrivK^{eav}_{\mathcal{A},\Pi} = 1\right] = \frac{1}{2}$$

Definitions

Perfect indistinguishability (example)

Exercise

Let Π denote the Mono-Alphabetic Substitution cipher. How can we construct an adversary $\mathcal{A}$ such that

$$Pr\left[PrivK^{eav}_{\mathcal{A},\Pi} = 1\right] > \frac{1}{2}$$

and thereby proof that Π is not not perfectly indistinguishable?

Definitions

Perfect indistinguishability (example)

Answer

Adversary $\mathcal{A}$ does:

1. Output $m_0 = aa$ and $m_1 = \text{ab}$.
2. Upon receiving the challenge ciphertext $c = c_1 c_2$, do the following: if $c_1 = c_2$ output 0, else output 1.

It is trivial to see that the adversary always succeeds and thus

$$Pr\left[PrivK^{eav}_{\mathcal{A},\Pi} = 1\right] = 1$$

Outline

The One-Time Pad (Vernam Cipher)

Definition

Construction 2.8

Fix an integer $\ell > 0$. The message space $\mathcal{M}$, key space $\mathcal{K}$, and ciphertext space $\mathcal{C}$ are all equal to $\{0,1\}^\ell$.

- *Gen*: the key-generation algorithm chooses a key from $\mathcal{K} = \{0,1\}^\ell$ according to the uniform distribution.
- *Enc*: given a key $k \in \{0,1\}^\ell$ and a message $m \in \{0,1\}^\ell$, then the encryption algorithm outputs the ciphertext $c := k \oplus m$.
- *Dec*: given a key $k \in \{0,1\}^\ell$ and a ciphertext $c \in \{0,1\}^\ell$, the decryption algorithm outputs the message $m := k \oplus c$.

The One-Time Pad (Vernam Cipher)

Proof

Exercise

How can we proof that the One-Time Pad has perfect secrecy?
What assumptions are needed?

Definitions

Perfect Secrecy

Definition 2.3

An encryption schema (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 $Pr[C = c] > 0$:

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

Alternative definition

For every $m, m' \in \mathcal{M}$ and every $c \in \mathcal{C}$

$$Pr[Enc_k(m) = c] = Pr[Enc_k(m') = c]$$

The One-Time Pad (Vernam Cipher)

Proof (continued)

Proof

$$\begin{aligned} Pr[M = m \mid C = c] &= \frac{Pr[C = c \mid M = m] \cdot Pr[M = m]}{Pr[C = c]} \\ &= \frac{2^{-\ell} \cdot Pr[M = m]}{2^{-\ell}} \\ &= Pr[M = m] \end{aligned}$$

Assumptions

What happens if we use the same key more that once?

$$c \oplus c' = (m \oplus k) \oplus (m' \oplus k) = m \oplus m'$$

Outline

The key and message space

Definition

Theorem 2.10

If (Gen, Enc, Dec) is a perfectly secret encryption schema with message space $\mathcal{M}$ and key space $\mathcal{K}$, then $|\mathcal{K}| \geq |\mathcal{M}|$.

The key and message space

Proof

Theorem 2.10

Assume $|\mathcal{K}| < |\mathcal{M}|$ and let $\mathcal{M}(c)$ be the set of all possible messages that are possible decryptions of c.

$$\mathcal{M}(c) \stackrel{def}{=} \{m \mid m = Dec_k(c) \text{ for some } k \in \mathcal{K}\}$$

If Dec is deterministic then clearly $|\mathcal{M}(c)| \leq |\mathcal{K}|$. If $|\mathcal{K}| < |\mathcal{M}|$, then there is some $m' \in \mathcal{M}$ such that $m' \notin \mathcal{M}(c)$. But then

$$Pr[M = m' \mid C = c] = 0 \neq Pr[M = m']$$

Shannon's Theorem

Definition

Theorem 2.11

Let (Gen, Enc, Dec) be an encryption schema with message space $\mathcal{M}$, for which $|\mathcal{M}| = |\mathcal{K}| = |\mathcal{C}|$. The schema is perfectly secret if and only if:

1. Every key $k \in \mathcal{K}$ is chosen with (equal) probability $1/|\mathcal{K}|$ by algorithm Gen
2. For every $m \in \mathcal{M}$ and every $c \in \mathcal{C}$, there exists a unique key $k \in \mathcal{K}$ such that $Enc_k(m)$ outputs c.

Shannon's Theorem

Proof

For any distribution over $\mathcal{M}$, any $m \in \mathcal{M}$ with probability $Pr[M = m] \neq 0$, and any $c \in \mathcal{C}$, we have

$$\begin{aligned} Pr[M = m \mid C = c] &= \frac{Pr[C = c \mid M = m] \cdot Pr[M = m]}{Pr[C = c]} \\ &= \frac{Pr[Enc_k(m) = c] \cdot Pr[M = m]}{Pr[C = c]} \\ &= \frac{|\mathcal{K}|^{-1} \cdot Pr[M = m]}{|\mathcal{K}|^{-1}} = Pr[M = m] \end{aligned}$$