CAS CS 538. Lecture Notes on Perfect Secrecy (September 5, 2019 lecture by Leo Reyzin)

1 Equivalent definitions of perfect one-time secrecy

The theorems show definitions of perfect one-time secrecy that are equivalent to [Ros19, Definition 2.8].

Theorem 1. *An encryption scheme Σ satisfies perfect one-time secrecy if and only if for every two messages $m_L, m_R \in \Sigma.\mathcal{M}$ and for every ciphertext c,*

$$\Pr_{k\leftarrow \text{KeyGen}()}[\text{Enc}(k, m_L) = c] = \Pr_{k\leftarrow \text{KeyGen}()}[\text{Enc}(k, m_R) = c]\,.$$

(The subscript notation for Pr *means "sample a random variable from a given distribution" or "generate a random variable according to a randomized algorithm.")*

The above theorem is very to easy to prove.

Definition 1. An encryption scheme Σ is secure with respect to a distribution D_M from $\Sigma.\mathcal{M}$, if for every message $g \in \Sigma.\mathcal{M}$ (think of g as "adversarial guess") and for every ciphertext c,

$$\Pr_{m\leftarrow D_M, k\leftarrow \text{KeyGen}()}[D_m = g \,|\, \text{Enc}(k, m) = c] = \Pr_{m\leftarrow D_M}[m = g]\,.$$

(The subscript notation for Pr means "sample a random variable from a given distribution" or "generate a random variable according to a randomized algorithm".)

Theorem 2. *An encryption scheme satisfies perfect one-time secrecy if for every distribution D_M it is secure respect to D_M.*

The above theorem is a bit harder to prove; proving it is a good exercise to check your own understanding of probability.

2 The price of perfection

The one-time-pad has very long keys. Unfortunately, as Shannon [Sha49] has shown, that's not the fault of the one-time-pad. Long keys are an inherent problem for any perfectly secret encryption scheme — not just the one-time-pad. The following theorem says that you need at least as many keys as messages.

Theorem 3. *For any encryption scheme Σ that is perfectly one-time secret, $|\Sigma.\mathcal{K}| \geq |\Sigma.\mathcal{M}|$.*

Proof. Suppose, for purposes of contradiction, that $|\Sigma.\mathcal{K}| < |\Sigma.\mathcal{M}|$

Fix some message $m_L \in \Sigma.\mathcal{M}$ and let c_0 be such that

$$\Pr[\text{Enc}(k, m_L) = c_0] > 0$$

(where the probability is taken over a random key and random choices made in Enc, if any). Consider the set $\{\mathrm{Dec}(k, c_0) \mid k \in \Sigma.\mathcal{K}\}$. It has at most $|\Sigma.\mathcal{K}|$ elements, and therefore there exists at least one element in $\Sigma.\mathcal{M} - \Sigma.\mathcal{K}$. Let m_R be such an element.

Then, by definition of correctness of encryption scheme m_R, would never get encrypted to c_0 (because otherwise you couldn't decrypt it, because no key decrypts c_0 to m_R). So

$$\Pr[\mathrm{Enc}(k, m_R) = c_0] = 0\,.$$

Thus, we could build a distinguisher $\mathcal{A}$ as follows: query (m_L, m_R) to get c; output 1 if $c = c_0$ and 0 otherwise. Then

$$|\Pr[\mathcal{A} \diamond \mathcal{L}^{\Sigma}_{\mathsf{ots\text{-}L}} \Rightarrow 1] = \Pr[\mathrm{Enc}(k, m_L) = c_0] \neq \Pr[\mathrm{Enc}(k, m_R) = c_0] = \Pr[\mathcal{A} \diamond \mathcal{L}_{\mathsf{ots\text{-}R}} \Rightarrow 1]$$

and thus Σ is not perfectly one-time secret by definition. We have reached a contradiction. □

Since the price of perfection is so high, let us try to give up on perfection and aim for something slightly less. We will do so in the next lecture.

References

[Ros19] Mike Rosulek. *The Joy of Cryptography.* 2019. `http://web.engr.oregonstate.edu/~rosulekm/crypto/`.

[Sha49] Claude E. Shannon. Communication theory of secrecy systems. *Bell System Technical Journal*, 28(4):656–715, October 1949. Available at `http://netlab.cs.ucla.edu/wiki/files/shannon1949.pdf`.