Introduction to Cryptology

2.1 - Perfect Secrecy

Federico Pintore

Mathematical Institute, University of Oxford (UK)

Security Definitions

We want to formulate a definition of security for (symmetric-key) encryption schemes.

Security Definitions

We want to formulate a definition of security for (symmetric-key) encryption schemes.

Security guarantee:

Security Definitions

We want to formulate a definition of security for (symmetric-key) encryption schemes.

Security guarantee: when is an encryption secure?

- impossible to recover the key
- impossible to recover the entire plaintext
- impossible to recover any information on the plaintext

Security Definitions

We want to formulate a definition of security for (symmetric-key) encryption schemes.

Threat model: what are the abilities of an adversary $\mathcal{A}$?

Security Definitions

We want to formulate a definition of security for (symmetric-key) encryption schemes.

Threat model: what are the abilities of an adversary $\mathcal{A}$?

- Ciphertext-only attack: $\mathcal{A}$ knows some ciphertexts $c_1 = E_k(m_1), \ldots, c_r = E_k(m_r)$ generated using a key k.

Security Definitions

We want to formulate a definition of security for (symmetric-key) encryption schemes.

Threat model: what are the abilities of an adversary $\mathcal{A}$?

- Ciphertext-only attack: $\mathcal{A}$ knows some ciphertexts $c_1 = E_k(m_1), \ldots, c_r = E_k(m_r)$ generated using a key k.
- Known-plaintext attack: $\mathcal{A}$ learns a number of pairs $(m_i, c_i = E_k(m_i))$ generated using a key k.

Security Definitions

We want to formulate a definition of security for (symmetric-key) encryption schemes.

Threat model: what are the abilities of an adversary $\mathcal{A}$?

- Ciphertext-only attack: $\mathcal{A}$ knows some ciphertexts $c_1 = E_k(m_1), \ldots, c_r = E_k(m_r)$ generated using a key k.
- Known-plaintext attack: $\mathcal{A}$ learns a number of pairs $(m_i, c_i = E_k(m_i))$ generated using a key k.
- Chosen-plaintext attack (CPA): same as above, but $\mathcal{A}$ chooses the plaintexts m_i.

Security Definitions

We want to formulate a definition of security for (symmetric-key) encryption schemes.

Threat model: what are the abilities of an adversary $\mathcal{A}$?

- Ciphertext-only attack: $\mathcal{A}$ knows some ciphertexts $c_1 = E_k(m_1), \ldots, c_r = E_k(m_r)$ generated using a key k.
- Known-plaintext attack: $\mathcal{A}$ learns a number of pairs $(m_i, c_i = E_k(m_i))$ generated using a key k.
- Chosen-plaintext attack (CPA): same as above, but $\mathcal{A}$ chooses the plaintexts m_i.
- Chosen-ciphertext attack (CCA): $\mathcal{A}$ gets the decryption, via D_k, of ciphertexts of their choice.

Perfect Secrecy (informal)

Security guarantee: impossible to recover any information on the plaintext m

Threat model:

- Ciphertext-only attack, where $\mathcal{A}$ knows a single ciphertext $c = E_k(m)$.
- $\mathcal{A}$ knows the probability distribution over $\mathcal{M}$.

Probability Review

Discrete Probability

Let Ω be a finite set of outcomes, known as *sample space*.

The *event space* $\mathcal{E}$ is a subset of $\mathcal{P}(\Omega)$[1] s.t.

- $\Omega \in \mathcal{E}$
- if $X \in \mathcal{E}$, then $\bar{X} \in \mathcal{E}$ ($\bar{X}$ denotes the complementary in Ω)
- if $X_1, X_2 \in \mathcal{E}$, then $X_1 \cup X_2 \in \mathcal{E}$

Usually, $\mathcal{E} = \mathcal{P}(\Omega)$.

[1] Power set of Ω.

Discrete Probability

A probability distribution for $\mathcal{E}$ is a map $\Pr : \mathcal{E} \to [0, 1]$, s.t.

- $\Pr(\Omega) = 1$
- if $X_1, \ldots, X_t$ are pairwise disjoint events, then

$$\Pr(\cup_{i=1}^{t} X_i) = \sum_{i=1}^{t} \Pr(X_i)$$

Discrete Probability

Let Ω be finite and X, Y in $\mathcal{E}$.

- If $\{w\} \in \mathcal{E}$ for all $w \in X$, we have $\Pr(X) = \sum_{w \in X} \Pr(w)$.
- Union Formula: $\Pr(X \cup Y) = \Pr(X) + \Pr(Y) - \Pr(X \cap Y)$.
- Union Bound: $\Pr(X \cup Y) \leq \Pr(X) + \Pr(Y)$.

Discrete Probability

- Conditional Probability: $\Pr(X|Y) := \Pr(X \cap Y)/\Pr(Y)$ (provided $\Pr(Y) > 0$).
- X and Y are said independent if $\Pr(X \cap Y) = \Pr(X) \cdot \Pr(Y)$.
- Bayes' Theorem: $\Pr(X|Y) = \dfrac{\Pr(X) \cdot \Pr(Y|X)}{\Pr(Y)}$ (provided $\Pr(X) \cdot \Pr(Y) > 0$)

Perfect secrecy

Syntax of Symmetric-key Encryption Schemes

A symmetric-key encryption scheme consists of three algorithms $(\text{KeyGen}, \text{Enc}, \text{Dec})$:

- $k \leftarrow \text{KeyGen}(n)$: a randomised algorithm that takes the security parameter n and outputs a key k[2].
- $c \leftarrow \text{Enc}(k, m)$: an algorithm (often randomised) that, on input a key k and a message $m \in \{0,1\}^*$, outputs a ciphertext c[3].
- $m \leftarrow \text{Dec}(k, c)$: a deterministic algorithm that, given a key k and a ciphertext c, returns a message m.

such that $\text{Dec}(k, \text{Enc}(k, m)) = m \quad \forall\ m \in \{0,1\}^*, k \in \mathcal{K}$.

[2] KeyGen determines the key space $\mathcal{K}$.

[3] The message space $\mathcal{M}$ is $\{0,1\}^*$. Enc determins the ciphertext space $\mathcal{C}$.

Notation

Given an encryption scheme E, a probability distribution $\Pr_{\mathcal{K}\times\mathcal{M}}$ for $\mathcal{P}(\mathcal{K}\times\mathcal{M})$ can be deduced from

- $(\mathcal{K}, \mathcal{P}(\mathcal{K}), \Pr_{\mathcal{K}})$
- $(\mathcal{M}, \mathcal{P}(\mathcal{M}), \Pr_{\mathcal{M}})$.

It is defined as:

$$\Pr_{\mathcal{K}\times\mathcal{M}}(X) = \sum_{(k,m)\in X} \Pr_{\mathcal{K}}(k)\Pr_{\mathcal{M}}(m)$$

- $\Pr(K = k) := \Pr_{\mathcal{K}\times\mathcal{M}}(\{k\}\times\mathcal{M}) = \Pr_{\mathcal{K}}(k)$
- $\Pr(M = m) := \Pr_{\mathcal{K}\times\mathcal{M}}(\mathcal{K}\times\{m\}) = \Pr_{\mathcal{M}}(m)$
- $\Pr(C = c) := \Pr_{\mathcal{K}\times\mathcal{M}}(X)$ with $X = \{(k,m) \mid \mathrm{Enc}(k,m) = c\}$

Perfect Secrecy (Shannon 1949)

It formalises the idea that the ciphertext should reveal no information about the plaintext.

Perfect Secrecy (Shannon 1949)

It formalises the idea that the ciphertext should reveal no information about the plaintext.

Perfect Secrecy (or Information-theoretic Security)

An encryption scheme E is perfectly secret if, for every probability distribution $\Pr_{\mathcal{M}}$ over $\mathcal{M}$, we have

$$\Pr(M = m | C = c) = \Pr(M = m) \tag{1}$$

for every $m \in \mathcal{M}$ and every $c \in \mathcal{C}$ s.t. $\Pr(C = c) > 0$.

Perfect Secrecy (Shannon 1949)

It formalises the idea that the ciphertext should reveal no information about the plaintext.

Perfect Secrecy (or Information-theoretic Security)

An encryption scheme E is perfectly secret if, for every probability distribution $\Pr_{\mathcal{M}}$ over $\mathcal{M}$, we have

$$\Pr(M = m | C = c) = \Pr(M = m) \tag{1}$$

for every $m \in \mathcal{M}$ and every $c \in \mathcal{C}$ s.t. $\Pr(C = c) > 0$.

Thanks to Bayes' theorem, (1) is equivalent to

$$\Pr(C = c | M = m) = \Pr(C = c) \tag{2}$$

Perfect Indistinguishability

Perfect Indistinguishability Experiment $\mathrm{PrivK}^{\mathrm{perfect-ind}}_{\mathcal{A},E}$

Challenger Ch | Adversary $\mathcal{A}$

$\xleftarrow{m_0, m_1, |m_0| = |m_1|}$

$b \leftarrow_{\$} \{0,1\}$

$\xrightarrow{c = \mathrm{Enc}(k, m_b)}$

Outputs their guess b'

Perfect Indistinguishability

Perfect Indistinguishability Experiment $\mathrm{PrivK}^{\mathrm{perfect-ind}}_{\mathcal{A},E}$

Challenger Ch		Adversary $\mathcal{A}$
	$\xleftarrow{m_0, m_1, \lvert m_0\rvert = \lvert m_1\rvert}$	
$b \leftarrow_{\$} \{0,1\}$		
	$\xrightarrow{c = \mathrm{Enc}(k, m_b)}$	
		Outputs their guess b'

Definition

An encryption scheme E is perfectly indistinguishable if, for every adversary $\mathcal{A}$, the following holds:

$$\Pr(\mathrm{PrivK}^{\mathrm{perfect-ind}}_{\mathcal{A},E} = 1) = 1/2,$$

where $\mathrm{PrivK}^{\mathrm{perfect-ind}}_{\mathcal{A},E} = 1$ if $b' = b$, and 0 otherwise.

Perfect Secrecy

Theorem (Perfect Secrecy)

An encryption scheme $(\mathrm{KeyGen}, \mathrm{Enc}, \mathrm{Dec})$ *has perfect secrecy iff, for every probability distribution* $\Pr_{\mathcal{M}}$ *over* $\mathcal{M}$*, we have*

$$\Pr(C = c | M = m_0) = \Pr(C = c | M = m_1),$$

for every $m_0, m_1 \in \mathcal{M}$ *s.t.* $|m_0| = |m_1|$ *and every* $c \in \mathcal{C}$ *s.t.* $\Pr(C = c) > 0$.

Proof.

$(\Rightarrow)\ \Pr(C = c|M = m_0) = \Pr(C = c) = \Pr(C = c|M = m_1)$

$(\Leftarrow)$ For any message m_0, it holds

$$\begin{aligned}
\Pr(C = c) &= \sum_m \Pr(C = c|M = m) \cdot \Pr(M = m) \\
&= \sum_m \Pr(C = c|M = m_0) \cdot \Pr(M = m) \\
&= \Pr(C = c|M = m_0) \cdot \sum_m \Pr(M = m) \\
&= \Pr(C = c|M = m_0)
\end{aligned}$$

□

Further Reading I

N. J. AlFardan, D. J. Bernstein, K. G. Paterson, B. Poettering, and J. C. Schuldt.
On the security of RC4 in TLS.
In 22nd USENIX Security Symposium (USENIX Security 13), pages 305–320, 2013.

B. Barak and S. Halevi.
A model and architecture for pseudo-random generation with applications to/dev/random.
In Proceedings of the 12th ACM conference on Computer and communications security, pages 203–212. ACM, 2005.

D. J. Bernstein.
The Salsa20 Family of Stream Ciphers.
In New stream cipher designs, pages 84–97. Springer, 2008.

Further Reading II


L. Blum, M. Blum, and M. Shub.
A simple unpredictable pseudo-random number generator.
SIAM Journal on computing, 15(2):364–383, 1986.


C. Cachin.
Entropy measures and unconditional security in cryptography.
PhD thesis, ETH Zurich, 1997.


S. Fluhrer, I. Mantin, and A. Shamir.
Weaknesses in the key scheduling algorithm of RC4.
In Selected areas in cryptography, pages 1–24. Springer, 2001.

Further Reading III


C. Garman, K. G. Paterson, and T. Van der Merwe.
Attacks only get better: Password recovery attacks against RC4 in TLS.
In 24th USENIX Security Symposium (USENIX Security 15), pages 113–128, 2015.


I. Mantin and A. Shamir.
A practical attack on broadcast RC4.
In Fast Software Encryption, pages 152–164. Springer, 2002.