Private-Key Encryption

Ali El Kaafarani[1,2]

[1]Mathematical Institute
[2] PQShield Ltd.

Outline

Course Main Reference

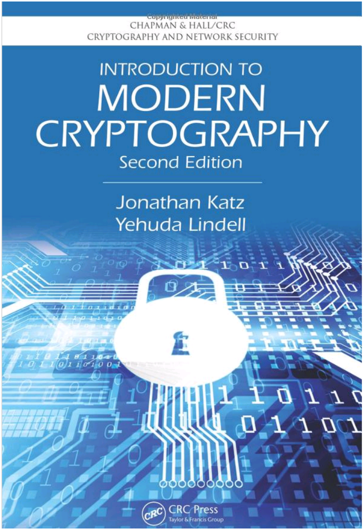

Caesar Cipher (100-44BC)

Example

- Plaintext: ABCD $\cdots$ WXYZ.
- Shift:$+3 \mod 26$
- Ciphertext: DEFG $\cdots$ ZABC.

Cryptanalysis:

Caesar Cipher (100-44BC)

Example

- Plaintext: ABCD $\cdots$ WXYZ.
- Shift:$+3 \mod 26$
- Ciphertext: DEFG $\cdots$ ZABC.

Cryptanalysis:

- Brute Force (trying every possible key): key space size is $|\mathcal{K}| = 26$.

Caesar Cipher (100-44BC)

Example

- Plaintext: ABCD $\cdots$ WXYZ.
- Shift:$+3 \mod 26$
- Ciphertext: DEFG $\cdots$ ZABC.

Cryptanalysis:

- Brute Force (trying every possible key): key space size is $|\mathcal{K}| = 26$.
- *Sufficient key-space principle: Any secure symmetric key encryption scheme must have a key space that is sufficiently large to make an exhaustive-search attack infeasible* (e.g. $|\mathcal{K}| \geq 2^{70}$).

Caesar Cipher (100-44BC)

Example

- Plaintext: ABCD $\cdots$ WXYZ.
- Shift:$+3 \mod 26$
- Ciphertext: DEFG $\cdots$ ZABC.

Cryptanalysis:

- Brute Force (trying every possible key): key space size is $|\mathcal{K}| = 26$.
- *Sufficient key-space principle: Any secure symmetric key encryption scheme must have a key space that is sufficiently large to make an exhaustive-search attack infeasible* (e.g. $|\mathcal{K}| \geq 2^{70}$).
- Is it a sufficient condition?

Substitution Cipher (mono-alphabetic)

Example

- Plaintext: ABCZ
- Substitution: (A$\rightarrow$ T, B$\rightarrow$ N, C$\rightarrow$ L, . . . , Z$\rightarrow$ O)
- Ciphertext:TNLO

Cryptanalysis:

Substitution Cipher (mono-alphabetic)

Example

- Plaintext: ABCZ
- Substitution: (A$\rightarrow$ T, B$\rightarrow$ N, C$\rightarrow$ L, . . . , Z$\rightarrow$ O)
- Ciphertext:TNLO

Cryptanalysis:

- Brute Force: Key space size is $|\mathcal{K}| = 26! \approx 2^{88}$.

Substitution Cipher (mono-alphabetic)

Example

- Plaintext: ABCZ
- Substitution: (A$\to$ T, B$\to$ N, C$\to$ L, . . . , Z$\to$ O)
- Ciphertext:TNLO

Cryptanalysis:

- Brute Force: Key space size is $|\mathcal{K}| = 26! \approx 2^{88}$.
- Frequency analysis:
  - Frequency of English letters

Substitution Cipher (mono-alphabetic)

Example

- Plaintext: ABCZ
- Substitution: (A$\rightarrow$ T, B$\rightarrow$ N, C$\rightarrow$ L, . . . , Z$\rightarrow$ O)
- Ciphertext:TNLO

Cryptanalysis:

- Brute Force: Key space size is $|\mathcal{K}| = 26! \approx 2^{88}$.
- Frequency analysis:
  - Frequency of English letters
  - Frequency of pairs (or more) of letters, e.g. digrams, trigrams, etc.

Substitution Cipher (mono-alphabetic)

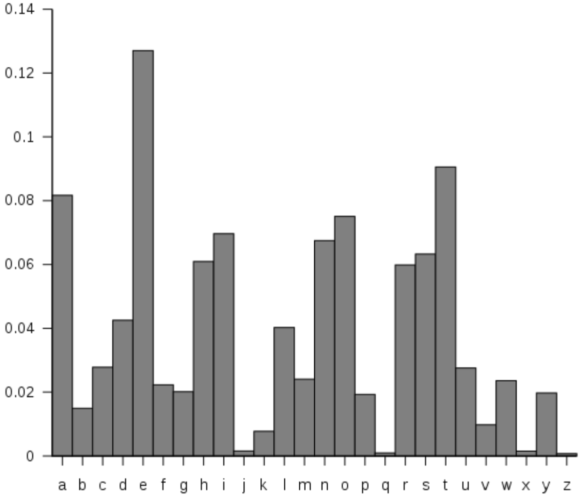

Vigenere Cipher (1553)

Example

- Poly-alphabetic shift:

Plaintext m:	TOBEORNOTTOBE
key k:(+ mod 26)	CRYPTOCRYPTOC
Ciphertext c:	VFZTHFPFRIHPG

- Cryptanalysis:

Vigenere Cipher (1553)

Example

- Poly-alphabetic shift:

Plaintext m:	TOBEORNOTTOBE
key k:(+ mod 26)	CRYPTOCRYPTOC
Ciphertext c:	VFZTHFPFRIHPG

- Cryptanalysis:
  - If the length of the key, say n, is known, then break ciphertext into blocks of size n, and solve each block similar to Caesar cipher and using letter-frequency analysis.
  - If n is not known, use Kasiski method (Kasiski 1863) or *index of coincidence method* to find n, and do the rest as in the first case. (What if $n = |c| = |k|$?)

Kerckhoff's Principle (1883):

Definition

The cipher must NOT be required to be secret and it must be able to fall into the hands of the enemy without inconvenience.

Modern Cryptography:

Kerckhoff's Principle (1883):

Definition

The cipher must NOT be required to be secret and it must be able to fall into the hands of the enemy without inconvenience.

Modern Cryptography:

- The encryption scheme's algorithms should be public. (Standardized, etc.)

Outline

Discrete Probability

Let Ω be the set of outcomes (sample space), define $\Pr : \Omega \to [0, 1]$ such that $\Pr(\omega) =$"probability that outcome ω occurs". Note that $0 \leq \Pr(\omega) \leq 1$, $\forall \omega \in \Omega$.

- Let $A \subseteq \Omega$, $\Pr(A) = \sum_{\omega \in A} \Pr(w)$.
- Union Formula: $\Pr(A \cup B) = \Pr(A) + \Pr(B) - \Pr(A \cap B)$.
- Union Bound: $\Pr(A \cup B) \leq \Pr(A) + \Pr(B)$.
- Conditional Probability: $\Pr(A|B) = \Pr(A \cap B)/\Pr(B)$.
- A and B are independent $\Leftrightarrow \Pr(A \cap B) = \Pr(A) \cdot \Pr(B)$.
- Bayes' Theorem: $\Pr(A|B) = \dfrac{\Pr(A) \cdot \Pr(B|A)}{\Pr(B)}$

Random Variables

- A coin is tossed 100 times. The variable X is the number of tails that are noted. X can *only* take the values $0, 1, ..., 100$. The variable X is called a *discrete random variable*.
- A random variable is a function $X : \Omega \to S$ that associates a unique numerical value with every outcome of an experiment.
- The probability distribution of a discrete random variable X is a list of probabilities associated with each of its possible values.
- If these probabilities are equal, the distribution is called a *Uniform distribution over S.*
- $\Pr(X = x) = \sum_{X(\omega)=x} \Pr(\omega)$.

Expected Value and Variance

- The expected value $E(X)$ of a random variable X indicates its average or central value; $E(X) = \sum_{\omega \in \Omega} X(\omega)\,\mathrm{Pr}(\omega)$,
- Property: $E(X+Y) = E(X) + E(Y)$.
- The Variance $V(X)$ is a measure of the “spread” of a distribution about its average value $E(X)$; $V(X) = E((X - E(X))^2) = E(X^2) - E(X)^2$.

Statistical Distance/Indistinguishability

Definition

Statistical distance
Let X and Y be two random variables distributed according to the distributions D_1 and D_2 respectively. The statistical distance between X and Y can be defined as:

$$\Delta(X, Y) = \frac{1}{2} \sum_{v \in X \cup Y} |\Pr(X = v) - \Pr(Y = v)|$$

Definition

Statistical Indistinguishability
Let X and Y be two random variables distributed according to distributions D_1 and D_2. We say that D_1 and D_2 are statistically indistinguishable if $\Delta(X, Y)$ is negligible.

Entropy

Figure: Unsurprised women watching the ticker tape in 1918.

`https://plus.maths.org/content/information-surprise`

Entropy

Definition

Let $(\Pr, \Omega)$ *be a discrete probability on a sample space* Ω *where* $A \subseteq \Omega$*. We define the* ***information of A*** *as*

$$I(A) = -\log_2 \Pr(A).$$

Definition

The entropy $H(X)$ *of a discrete random variable* X *on a sample space* Ω *is the average amount of information conveyed by it.*

$$H(X) = E(I(X = x)) = -\sum_x \Pr(X = x) \cdot \log_2 \Pr(X = x).$$

- Entropy Demo: `http://www.math.ucsd.edu/~crypto/java/ENTROPY/`

Entropy

Theorem

If X is a random variable, $X : \Omega \to S$, then $H(X) \leq \log(|S|)$.

Theorem

Minimum entropy
$H(X) \geq k \Leftrightarrow \forall x, \Pr(X = x) \leq 2^{-k}$

Definition

Negligible function
A function ϵ is negligible iff $\forall c \in \mathbb{N}\ \exists n_0 \in \mathbb{N}$ such that $\forall n \geq n_0, \epsilon(n) \leq n^{-c}$.

Examples

- Maximum entropy is achieved when all events are equally likely, in this case $H = log(|S|)$.
- Minimum entropy happens when one event is certain and the others are impossible, in this case $H = 0$.
- In theory: $2^{-n}, 2^{-\sqrt{n}}$ and $n^{-\log n}$ are negligible functions.
- In practice: $\epsilon \geq 1/2^{30}$ is non-negligible, whereas $\epsilon \leq 1/2^{80}$ is negligible.

Outline

Syntax of Private Key Encryption Schemes

Any encryption scheme consists of three algorithms:

- $k \leftarrow \mathsf{KeyGen}(n)$: It takes the security parameters n and outputs the key k. We assume that $|k| \geq n$.
- $c \leftarrow \mathsf{Enc}(k, m \in \mathcal{M})$: An algorithm (often randomized) that takes the encryption key k and the message and outputs the ciphertext c.
- $m \leftarrow \mathsf{Dec}(k, c)$: An algorithm (always deterministic) that takes the key and ciphertext and gives back the message.

Definition

Correctness: An encryption scheme is correct iff

$$\forall k \in \mathcal{K}, \forall m \in \mathcal{M}, \mathsf{Dec}(k, \mathsf{Enc}(k, m)) = m.$$

Security Definitions

What is a secure encryption scheme (security goals)?

Security Definitions

What is a secure encryption scheme (security goals)?

- Adversaries cannot compute the plaintext.

Security Definitions

What is a secure encryption scheme (security goals)?

- Adversaries cannot compute the plaintext.
- Adversaries cannot compute the secret key.

Security Definitions

What is a secure encryption scheme (security goals)?

- Adversaries cannot compute the plaintext.
- Adversaries cannot compute the secret key.
- Adversaries cannot compute information about the plaintext.

Security Definitions

What is a secure encryption scheme (security goals)?

- Adversaries cannot compute the plaintext.
- Adversaries cannot compute the secret key.
- Adversaries cannot compute information about the plaintext.
- Adversaries cannot compute any function about the plaintext.

On the other hand: what are the adversaries’ abilities (or threat models)?

Security Definitions

What is a secure encryption scheme (security goals)?

- Adversaries cannot compute the plaintext.
- Adversaries cannot compute the secret key.
- Adversaries cannot compute information about the plaintext.
- Adversaries cannot compute any function about the plaintext.

On the other hand: what are the adversaries' abilities (or threat models)?

- Ciphertext-only attack: one single ciphertext c.

Security Definitions

What is a secure encryption scheme (security goals)?

- Adversaries cannot compute the plaintext.
- Adversaries cannot compute the secret key.
- Adversaries cannot compute information about the plaintext.
- Adversaries cannot compute any function about the plaintext.

On the other hand: what are the adversaries' abilities (or threat models)?

- Ciphertext-only attack: one single ciphertext c.
- Known Plaintext attack: the adversary learns a number of pairs of (c_i, m_i) generated using some key.

Security Definitions

What is a secure encryption scheme (security goals)?

- Adversaries cannot compute the plaintext.
- Adversaries cannot compute the secret key.
- Adversaries cannot compute information about the plaintext.
- Adversaries cannot compute any function about the plaintext.

On the other hand: what are the adversaries' abilities (or threat models)?

- Ciphertext-only attack: one single ciphertext c.
- Known Plaintext attack: the adversary learns a number of pairs of (c_i, m_i) generated using some key.
- Chosen-plaintext attack (CPA): same as above, but the adversary gets to choose the plaintexts this time.

Security Definitions

What is a secure encryption scheme (security goals)?

- Adversaries cannot compute the plaintext.
- Adversaries cannot compute the secret key.
- Adversaries cannot compute information about the plaintext.
- Adversaries cannot compute any function about the plaintext.

On the other hand: what are the adversaries' abilities (or threat models)?

- Ciphertext-only attack: one single ciphertext c.
- Known Plaintext attack: the adversary learns a number of pairs of (c_i, m_i) generated using some key.
- Chosen-plaintext attack (CPA): same as above, but the adversary gets to choose the plaintexts this time.
- Chosen-ciphertext attack (CCA): now, he additionally gets the decryption of ciphertexts of its choice.

Perfect Secrecy (Shannon 1949)

Perfect Secrecy (Shannon 1949)

- *"The ciphertext should reveal no information about the plaintext"*

Perfect Secrecy (Shannon 1949)

- *"The ciphertext should reveal no information about the plaintext"*
- Also called *information theoretic security*.

Definition

Perfect Secrecy
For every probability distribution over the message space $\mathcal{M}$*,*
$\forall m \in \mathcal{M}, \forall c \in \mathcal{C}$ *for which* $\Pr[C = c] > 0$ *we have;*

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

equivalently,

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

Perfect Indistinguishability

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

Challenger Ch **Adversary** A

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

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

$\xrightarrow{c=\mathsf{Enc}(k, m_b)}$ Outputs his guess b'

Definition

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

$$\Pr[\mathsf{PrivK}^{\mathsf{perfect-IND}}_{\mathcal{A},E} = 1] = 1/2$$

Where $\mathsf{PrivK}^{\mathsf{perfect-IND}}_{\mathcal{A},E} = 1$ if $b' = b$, and 0 otherwise.

Perfect Indistinguishability

Theorem

Perfect indistinguishability
An encryption scheme ($\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}$*) has perfect secrecy iff for every probability distribution over* $\mathcal{M}$*,*
$\forall m_0, m_1 \in \mathcal{M}\ s.t.\ |m_0| = |m_1|, \forall c \in \mathcal{C},$
$\Pr[C = c | M = m_0] = \Pr[C = c | M = m_1]$

Proof.

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

$(\Leftarrow)$:

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

which is correct for any m_0 □

Outline

One Time Pad (Vernam 1917 or some 35 years earlier!)

Fix an integer $n > 0$. Let $\mathcal{M} = \mathcal{C} = \mathcal{K} = \{0,1\}^n$.

- **Key Generation:** $\mathsf{KeyGen}(n)$: It produces a random bit string of length n, i.e. $k \in \mathcal{K}$.
- **Encryption:** $\mathsf{Enc} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$, such that $c \leftarrow \mathsf{Enc}(k,m) = k \oplus m$.
- **Decryption:** $\mathsf{Dec} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$, such that $m \leftarrow \mathsf{Dec}(k,c) = k \oplus c$.

One Time Pad (Vernam 1917 or some 35 years earlier!)

Fix an integer $n > 0$. Let $\mathcal{M} = \mathcal{C} = \mathcal{K} = \{0,1\}^n$.

- **Key Generation:** KeyGen(n) : It produces a random bit string of length n, i.e. $k \in \mathcal{K}$.
- **Encryption:** Enc : $\{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$, such that $c \leftarrow \text{Enc}(k,m) = k \oplus m$.
- **Decryption:** Dec : $\{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$, such that $m \leftarrow \text{Dec}(k,c) = k \oplus c$.

It was used between the White House and the Kremlin during the Cold War!

Security of OTP

Theorem

The one time pad (OTP) encryption scheme is perfectly secret.

Proof.

$$\begin{aligned}\Pr[C = c|M = m] &= \Pr[M \oplus k = c|M = m] \\ &= \Pr[m \oplus k = c] \\ &= \Pr[k = m \oplus c] \\ &= \frac{1}{2^n}\end{aligned}$$

because the key k is a uniform n-bit string. Therefore, For any m_0, m_1, we have $\Pr[C = c|M = m_0] = \frac{1}{2^n} = \Pr[C = c|M = m_1]$ □

OTP has perfect secrecy, but is it practical?

Theorem

If an encryption scheme E is perfectly secret, then $|\mathcal{K}| \geq |\mathcal{M}|$.

OTP has perfect secrecy, but is it practical?

Theorem

If an encryption scheme E is perfectly secret, then $|\mathcal{K}| \geq |\mathcal{M}|$.

Proof.

Assume that $|\mathcal{K}| < |\mathcal{M}|$, we will show that E is not perfectly secure. We first fix a uniform distribution over $\mathcal{M}$, and let

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

but $|\mathcal{M}(c)| \leq |\mathcal{K}|$, then there exists $m' \in \mathcal{M}$ s.t. $m' \notin \mathcal{M}(c)$.
Therefore, $\Pr[M = m' | C = c] = 0 \neq \Pr[M = m']$ □

Is there a way to make OTP practical?

From Perfect to Computational Security

- Perfect secrecy: No leakage of information about an encrypted message even to an eavesdropper with *unlimited computational power*.

From Perfect to Computational Security

- Perfect secrecy: No leakage of information about an encrypted message even to an eavesdropper with *unlimited computational power*.
- Computational secrecy: an encryption scheme is still considered to be secure even if it leaks a very small amount of information to eavesdroppers with *limited power*.

From Perfect to Computational Security

- Perfect secrecy: No leakage of information about an encrypted message even to an eavesdropper with *unlimited computational power*.
- Computational secrecy: an encryption scheme is still considered to be secure even if it leaks a very small amount of information to eavesdroppers with *limited power*.
- Real-world application: happy with a scheme that leaks information with probability at most 2^{-60} over 200 years using fastest supercomputers!

Computational Security

Concrete version:

Definition

An encryption scheme is (t, ϵ)-secure if any adversary running for time at most t succeeds in breaking the scheme with probability at most ϵ.

Asymptotic version:

Definition

An encryption scheme is secure if any *probabilistic polynomial-time algorithm in* n (PPT) succeeds in breaking the scheme with at most negligible probability (in n).

Further Reading (1)

Nadhem J AlFardan, Daniel J Bernstein, Kenneth G Paterson, Bertram Poettering, and Jacob CN Schuldt.
On the security of RC4 in TLS.
In *USENIX Security*, pages 305–320, 2013.

Boaz Barak and Shai 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.

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

Further Reading (2)

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

- Christian Cachin.
  Entropy measures and unconditional security in cryptography.
  PhD thesis, SWISS FEDERAL INSTITUTE OF TECHNOLOGY ZURICH, 1997.

- Scott Fluhrer, Itsik Mantin, and Adi Shamir.
  Weaknesses in the key scheduling algorithm of RC4.
  In *Selected areas in cryptography*, pages 1–24. Springer, 2001.

Further Reading (3)

Christina Garman, Kenneth G Paterson, and Thyla van der Merwe.
Attacks only get better: Password recovery attacks against RC4 in TLS.
2015.

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