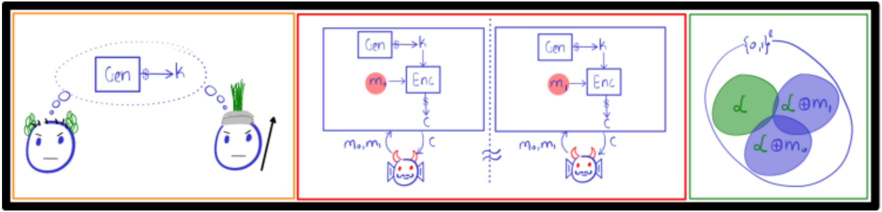

CS409m: Introduction to Cryptography

Lecture 03 (06/Aug/25)

Instructor: Chethan Kamath

Annonuncement

- Hands-on Exercise 1 will be out this Friday (08/Aug)
- Please register on https://cs409m.ctfd.io/ by Thursday (07/Aug)

Recall from Lecture 01...

- Classical vs modern cryptography
- Guiding principles for modern cryptography:
  1. Identify the task and specify syntax
  2. Come up with precise threat model M (a.k.a security model)
     - Attack model: What are the adversary's capabilities?
     - Break model: What does it mean to be secure?
  3. Construct a scheme Π
  4. Formally prove that Π in secure in threat model M
- Classical ciphers: shift, substitution, polyalphabetic shift
- Saw informally why these are insecure by modern standards
  - Ciphertext leaks some information about the message

Plan for This Lecture...

secret communication with shared keys

- Guiding principles for modern cryptography:
  1. Identify the task and specify syntax
  2. Come up with precise threat model M (a.k.a security model) 
     - Attack model: What are the adversary's capabilities? ← Eavesdropper
     - Break model: What does it mean to be secure? ↺ Perfect secrecy
  3. Construct a scheme Π ← One-time pad
  4. Formally prove that Π in secure in threat model M

Plan for This Lecture...

1. Syntax of Shared/Symmetric-Key Encryption (SKE)

2. Perfect Secrecy and One-Time Pad (OTP)

   +First proof

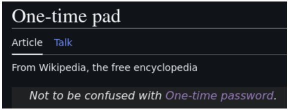

3. Limitations of Perfect Secrecy: Shannon's Impossibility

   -First impossibility

Plan for This Lecture

1. Syntax of Shared/Symmetric-Key Encryption (SKE)

2. Perfect Secrecy and One-Time Pad (OTP)

   +First proof

3. Limitations of Perfect Secrecy: Shannon's Impossibility

   −First impossibility

Some Notation and Conventions...

- Sets:
  - Denoted using calligraphic font: e.g., $\mathcal{M}$, $\mathcal{C}$
  - Sampling *uniformly at random* from a set denoted by '$\leftarrow$'
    - E.g., $k \leftarrow \{0,1\}^{\ell}$ and $m \leftarrow \mathcal{M}$
- Probability notation:
  - For a distribution/random variable M over a set $\mathcal{M}$ and element $m \in \mathcal{M}$, $m = \mathrm{M}$ denotes the *event*: 'a random sample from M equals m''
  - Following denotes probability that $\mathsf{A}(x) = 1$ when $x \leftarrow \{0,1\}^n$:

$$\Pr_{x \leftarrow \{0,1\}^n}[\mathsf{A}(x) = 1]$$

Syntax of Shared/Symmetric-Key Encryption

Definition 1 (Shared/Symmetric-Key Encryption (SKE))

An SKE Π for message space $\mathcal{M}$ is a triple of efficient algorithms $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ with the following syntax:

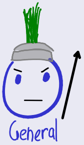

Syntax of Shared/Symmetric-Key Encryption

Definition 1 (Shared/Symmetric-Key Encryption (SKE))

An SKE Π for message space $\mathcal{M}$ is a triple of efficient algorithms $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ with the following syntax:

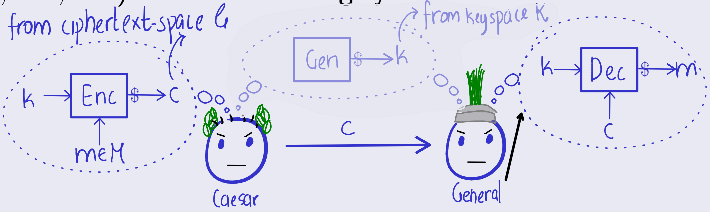

- Correctness of decryption: for all message $m \in \mathcal{M}$,

$$\Pr_{k \leftarrow \mathsf{Gen}, c \leftarrow \mathsf{Enc}(k,m)}[\mathsf{Dec}(k,c) = m] = 1$$

? Why can we assume that Dec is *deterministic* w.l.o.g.?

Example: Shift Cipher (Caesar Cipher)

Construction 1 (for message space $\{a, \cdots, z\}^\ell$)

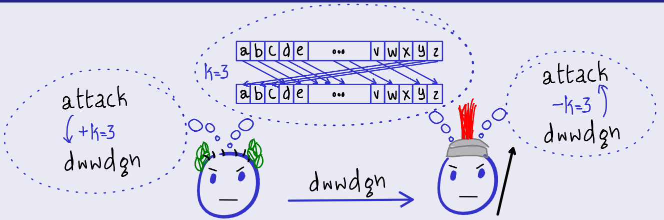

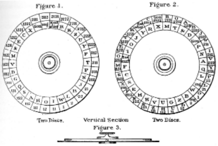

Pseudocode 1 (Message space $\{0, \cdots, 25\}^\ell \leftrightarrow \{a, \cdots, z\}^\ell$)

- Key generation, Gen: output $k \leftarrow \{0, \cdots, 25\}$
- Encryption, $\mathsf{Enc}(k, m = m_1 \| \cdots \| m_\ell)$:
  - Output $c := c_1 \| \cdots \| c_\ell$, where $c_i := m_i + k \bmod 26$
- Decryption, $\mathsf{Dec}(k, c = c_1 \| \cdots \| c_\ell)$:
  - Output $m := m_1 \| \cdots \| m_\ell$, where $m_i := c_i - k \bmod 26$

Why does correctness of decryption hold?

Plan for This Lecture

1 Syntax of Shared/Symmetric-Key Encryption (SKE)

2 Perfect Secrecy and One-Time Pad (OTP)

+First proof

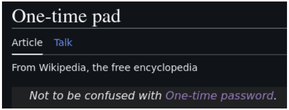
One-time pad

Article Talk

From Wikipedia, the free encyclopedia

Not to be confused with One-time password.

3 Limitations of Perfect Secrecy: Shannon's Impossibility

−First impossibility

General *template*: secret communication with shared keys

1. Identify the task and specify syntax
2. Come up with precise threat model M (a.k.a security model)

    - Attack model: What are the adversary's capabilities? Eavesdropper
    - Break model: What does it mean to be secure? Perfect secrecy
3. Construct a scheme Π One-time pad
4. Formally prove that Π in secure in threat model M

Attack Model and Break Model

Attack Model: Eavesdropping

1. How powerful is Eve?
   - Computationally unbounded
2. What attack can Eve do?
   - Only eavesdrop and obtain ciphertext (ciphertext-only attack)
3. Is Eve randomised?
   - ?

Break Model:

- Attempt 1: Eve must find key
  - $\mathrm{Enc}(k, m) := m$ secure!
- Attempt 2: Eve must recover m
  - What if ciphertext leaks first few bits of the message?
- Shannon's take
  - Ciphertext must reveal *no information* about the message

How to Model 'No Information Learnt'?

- We will look at two ways:

"Information theoretic"

$$\Pr[\mathrm{M} = m^* | \mathrm{C} = c^*] = \Pr[\mathrm{M} = m^*]$$

"Two worlds"

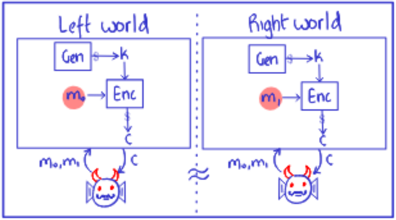

Modelling 'No Information Learnt': Shannon's Take...

Definition 2 (Shannon'49)

Let $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be an SKE with message space $\mathcal{M}$.
Π is perfectly-secret if *for every* message distribution M over $\mathcal{M}$, message $m^* \in \mathcal{M}$ and ciphertext $c^* \in \mathcal{C}$ (in support):

Ciphertext distribution induced by M, Gen & Enc

$$\Pr[\mathrm{M} = m^* | \underbrace{\mathrm{C} = c^*}] = \Pr[\mathrm{M} = m^*]$$

- Intuition: '*observing a ciphertext must have no effect on* Eve*'s knowledge about the message being sent*'
- Definition *does not* refer to Eve at all!

Definition 2 (Shannon'49) $\Pr[\text{attack}] = \frac{1}{2} = \Pr[\text{defend}]$

Let $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be an SKE with message space $\mathcal{M}$.
Π is not perfectly-secret if ~~for every~~ there exists message distribution M over $\mathcal{M}$, message $m^* \in \mathcal{M}$ and ciphertext $c^* \in \mathcal{C}$ (in support):

defend — dwwdgn

$$\Pr[\mathrm{M} = m^* | \mathrm{C} = c^*] \neq \Pr[\mathrm{M} = m^*]$$

0! — 1/2

- Let's see why shift cipher is not perfectly secret.

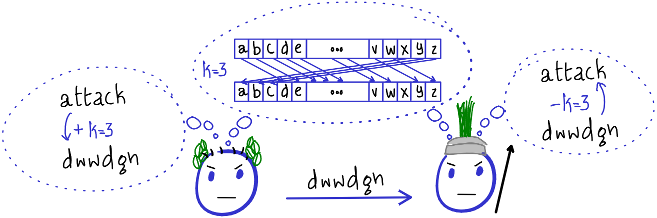

Modelling ‘No Information Learnt’: Shannon’s Take...

Exercise 1

- Formally define substitution cipher using a pseudocode (clearly state key-space etc)
- Show that it is not perfectly secret according to Definition 2

Exercise 2

- Formally define polyalphabetic shift cipher using a pseudocode
- Show that it is not perfectly secret according to Definition 2

- Turing's Imitation Game (Turing Test)

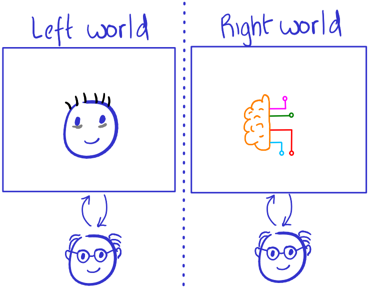

- Turing, on artificial intelligence: *"Are there imaginable digital computers which would do well in the imitation game?"*

- Turing’s Imitation Game (Turing Test)

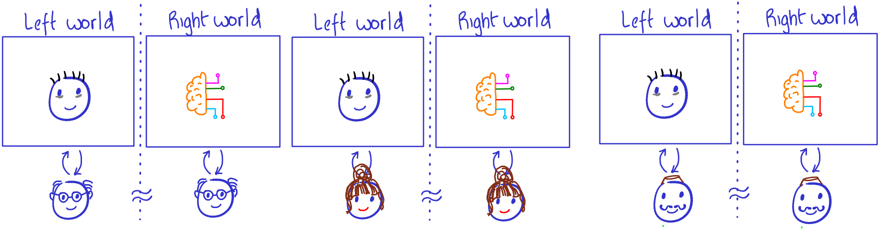

- Turing, on artificial intelligence: *"Are there imaginable digital computers which would do well in the imitation game?"*
- To paraphrase: sign of artificial (human) intelligence if no human can tell the two worlds apart $\approx$

? What are our two worlds?

- 'Left" world: always encrypt m_0
  "Right" world: always encrypt m_1

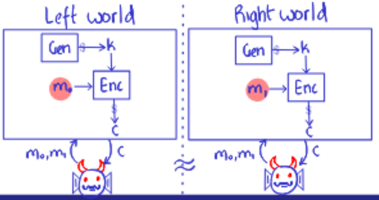

Definition 3 (Two-Worlds Definition)

An SKE $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is perfectly-secret if *for every* eavesdropper *Eve* and every message-pair $(m_0, m_1) \in \mathcal{M}$:

$$\Pr_{\substack{k \leftarrow \mathsf{Gen} \\ c \leftarrow \mathsf{Enc}(k, m_0)}}[\mathsf{Eve}(c) \underbrace{\text{ outputs 'left'}}_{0}] = \Pr_{\substack{k \leftarrow \mathsf{Gen} \\ c \leftarrow \mathsf{Enc}(k, m_1)}}[\mathsf{Eve}(c) = \underbrace{\text{ outputs 'left'}}_{0}]$$

Exercise 3

Show that shift and substitution ciphers are not perfectly secret w.r.to Definition 3

How to Model 'No Information Learnt'?...

- We saw two definitions. There are two more.

"Information theoretic"

$$\Pr[\mathrm{M} = m^* | \mathrm{C} = c^*] = \Pr[\mathrm{M} = m^*]$$

"Two worlds"

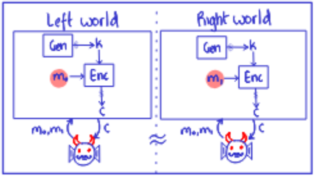

$$\Pr_{k \leftarrow \mathrm{Gen}}[\mathrm{Enc}(k, m_0) = c^*] = \Pr_{k \leftarrow \mathrm{Gen}}[\mathrm{Enc}(k, m_1) = c^*]$$

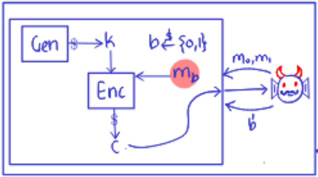

- 'Semantic-security': ciphertext contains no info. about plaintext
- Ciphertext indistinguishability: variant of imitation game

Exercise 4

Show equivalence of all these definitions.

One-Time Pad (Vernam's Cipher)...

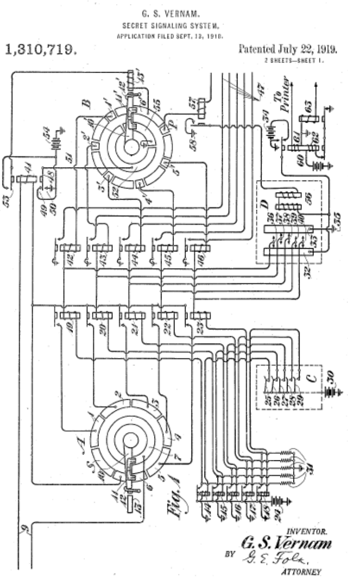

One-Time Pad (Vernam's Cipher)...

Construction 2 (Message space $\{0,1\}^{\ell}$)$^{=6}$

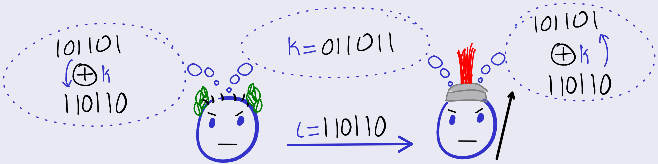

Pseudocode 2 (Message space $\{0,1\}^{\ell}$)

- Key generation Gen: output $k \leftarrow \{0,1\}^{\ell}$
- Encryption Enc(k, m): output $c := k \oplus m$
- Decryption Dec(k, c): output $m := k \oplus c$

Exercise 5

1. Design OTP for message space $\{\mathrm{a}, \cdots, \mathrm{z}\}^{\ell}$
2. How is this different from *polyalphabetic* shift cipher?

One-Time Pad is Perfectly Secret...

Theorem 1 (Shannon'49)

One-time pad is a perfectly secret SKE according to Definition 3.

Proof.

Goal is to show: $\forall \text{Eve}, \forall m_0, m_1 \in \mathcal{M}$

$$\Pr_{r \leftarrow \{0,1\}^\ell}\left[\text{Eve}(m_0 \oplus r) = \text{"left"}\right] = \Pr_{r \leftarrow \{0,1\}^\ell}\left[\text{Eve}(m_1 \oplus r) = \text{"left"}\right]$$

$$\Leftrightarrow \cancel{\frac{1}{2^\ell}} \sum_{r \in \{0,1\}^\ell} \Pr\left[\text{Eve}(m_0 \oplus r) = \text{"left"}\right] = \cancel{\frac{1}{2^\ell}} \sum_{r \in \{0,1\}^\ell} \Pr\left[\text{Eve}(m_1 \oplus r) = \text{"left"}\right]$$

$$\Leftrightarrow \left|\left\{r : \text{Eve}(m_0 \oplus r) = \text{"left"}\right\}\right| = \left|\left\{r : \text{Eve}(m_1 \oplus r) = \text{"left"}\right\}\right|$$

$$|\mathcal{L} \oplus m_0| = |\mathcal{L}| = |\mathcal{L} \oplus m_1|$$

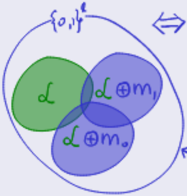

Now consider the set $\mathcal{L} \subseteq \{0,1\}^\ell := \{c : \text{Eve}(c) = \text{"left"}\}$

One-Time Pad is Perfectly Secret...

Exercise 6 (Hint: use Bayes' theorem.)

Show that one-time pad is a perfectly secret SKE according to Definition 2.

OTP IRL

'Red telephone'

Radio Netherlands Archives

THE NETHERLANDS / HISTORY / AFRICA

Operation Vula: A secret Dutch network against apartheid

Published 9th September 1999

Moscow–Washington hotline

Article Talk

From Wikipedia, the free encyclopedia

(Redirected from Moscow-Washington hotline)

Why not use OTP for all purposes?

- Keys are as large as messages $|\mathcal{K}| = |\mathcal{M}|$
- Why not re-use keys? Then it becomes insecure! See Hands-on Exercise 1

The Register®

Declassified files reveal how pre-WW2 Brits smashed Russian crypto

Moscow's agents used one-time pads, er, two times – ой!

Venona project

Article Talk

From Wikipedia, the free encyclopedia

Plan for This Lecture

1 Syntax of Shared/Symmetric-Key Encryption (SKE)

2 Perfect Secrecy and One-Time Pad (OTP)

+ First proof

3 Limitations of Perfect Secrecy: Shannon's Impossibility

- First impossibility

Theorem 2 (Shannon'49)

Let $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *be any perfectly-secret encryption scheme with message space* $\mathcal{M}$ *and key-space* $\mathcal{K}$. *Then* $|\mathcal{K}| \geq |\mathcal{M}|$.

Proof Sketch. Idea: proof by contradiction.

Assume for contradiction that $|\mathcal{K}| < |\mathcal{M}|$

· Goal: show that Π not perfectly secure

Fix any message $m^* \in \mathcal{M}$ and c^* in m^*'s ciphertext-space

Consider set $\mathcal{M}_{c^*} \subseteq \mathcal{M}$ defined as

Why? $\{m \in \mathcal{M} : \exists k \in \mathcal{K} \text{ s.t. } \mathrm{Dec}(k, c^*) = m\}$

Since $|\mathcal{M}_{c^*}| \leq |\mathcal{K}| < |\mathcal{M}|$,

$\exists m' \in \mathcal{M} \setminus \mathcal{M}_{c^*}$: c^* never decrypts to m'

⁝ (1/2)

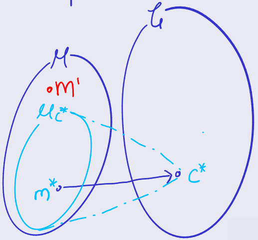

Theorem 2 (Shannon'49)

Let $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *be any perfectly-secret encryption scheme with message space* $\mathcal{M}$ *and key-space* $\mathcal{K}$. *Then* $|\mathcal{K}| \geq |\mathcal{M}|$.

Proof Sketch. Idea: proof by contradiction.

(2/2):

Consider (m^*, m') and $\mathrm{Eve}_{c^*}(c) := \begin{cases} \text{'left'} & \text{if } c = c^* \\ \text{'right'} & \text{otherwise} \end{cases}$

We have:

1) for m^*: $\Pr_{\substack{k \leftarrow \mathsf{Gen} \\ c \leftarrow \mathsf{Enc}(m^*)}} [\mathrm{Eve}_{c^*}(c) = \text{'left'}] > 0$

2) for m': $\Pr_{\substack{k \leftarrow \mathsf{Gen} \\ c \leftarrow \mathsf{Enc}(m')}} [\mathrm{Eve}_{c^*}(c) = \text{'left'}] = 0$

$\Rightarrow$ Π is not perfectly secure ⚡

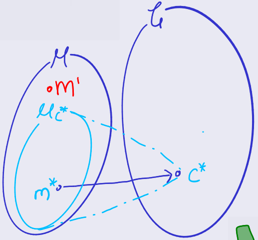

What Do We Do in Face of Shannon's Impossibility?

- You compromise.
  - Kerckhoffs' principle: *"The system should be, if not theoretically unbreakable, unbreakable in practice."*

Definition 3 (Two-Worlds Definition)

An SKE $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is perfectly-secret if *for every eavesdropper Eve* and every message-pair $(m_0, m_1) \in \mathcal{M}$:

$$\Pr_{\substack{k \leftarrow \mathsf{Gen} \\ c \leftarrow \mathsf{Enc}(k, m_0)}}[\mathsf{Eve}(c) \text{ outputs 'left'}] \approx \Pr_{\substack{k \leftarrow \mathsf{Gen} \\ c \leftarrow \mathsf{Enc}(k, m_1)}}[\mathsf{Eve}(c) = \text{ outputs 'left'}]$$

- Compromise two aspects of Definition 3:
  1. Restrict to *computationally*-bounded Eve
  2. Allow "slack": Eve may distinguish, but with "very small" prob.
- Turns out both compromises are necessary!

Next Two Lectures

- How to model computationally-bounded adversaries?
  - Probabilitic polynomial-time (PPT) algorithms

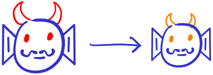

- How to capture "very small" probability?
  - Negligible functions

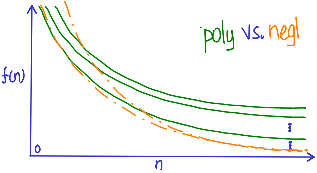

- Pseudo-random generators (PRG)
- Computational OTP

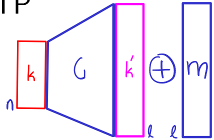

More Questions?

References

1. [KL14, Chapters 1 and 2] for details about this lecture
2. Shannon's paper on perfect secrecy and proof of perfect secrecy one-time pad: [Sha49]
3. Turing's paper on artificial intelligence: [Tur50]
4. David Kahn's *The Codebreakers* for historical aspects of cryptography

Jonathan Katz and Yehuda Lindell.
Introduction to Modern Cryptography (3rd ed.).
Chapman and Hall/CRC, 2014.

C. E. Shannon.
Communication theory of secrecy systems.
The Bell System Technical Journal, 28(4):656–715, 1949.

A. M. Turing.
Computing Machinery and Intelligence.
Mind, LIX(236):433–460, 10 1950.