Lecture 03: Perfect Security Definition

First Attempt

- Intuitively, we might want to define perfect security of an encryption scheme as follows: Given a ciphertext all messages are equally likely.
- This can be formulated as: For all $m^{(0)}, m^{(1)} \in \mathcal{M}$ and $c \in \mathcal{C}$ we have:

$$\Pr[M = m^{(0)} | C = c] = \Pr[M = m^{(1)} | C = c]$$

- The probability here is over the randomness used in the Gen and Enc algorithms and the probability distribution over the message space
- But this definition has a problem. It might be a priori known that the message $m^{(0)}$ is more likely than $m^{(1)}$. We do not want "seeing the ciphertext" to change this information

Perfect Security of Encryption

- We want the ciphertext to provide no *additional* information about the message

Definition (One: Perfect Security)

For all $m \in \mathcal{M}$ and $c \in \mathcal{C}$, we have:

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

- Here we are assuming that $c \in \mathcal{C}$ has $\Pr[C = c] > 0$. Everywhere this assumption will be implicit

Another Definition

- We want to say that the probability to generate a ciphertext given a message is independent of the message

Definition (Two: Perfect Security)

For all $m \in \mathcal{M}$ and $c \in \mathcal{C}$ we have:

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

- How to show equivalence of Definition 1 and Definition 2? (Hint: Use Bayes' Rule)

How to Show Equivalence of Definitions

- To show Definition 1 implies Definition 2: Assume $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is an encryption scheme that satisfies Definition 1. Then show that it also satisfies Definition 2
- To show Definition 2 implies Definition 1: Assume $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is an encryption scheme that satisfies Definition 2. Then show that it also satisfies Definition 1
- Do this exercise yourself

Another Definition of Perfect Secrecy

- We want to say that the probability of generating a ciphertext given as message $m^{(0)}$, is same as the probability of generating that ciphertext given any other different message $m^{(1)}$

Definition (Three: Perfect Security)

For any messages $m^{(0)}, m^{(1)} \in \mathcal{M}$ and $c \in \mathcal{C}$ we have:

$$\Pr[C = c | M = m^{(0)}] = \Pr[C = c | M = m^{(1)}]$$

- Show the equivalence of Definition 2 and Definition 3 (Hint: Use Bayes' rule and the intuition that $\Pr[C = c]$ is that expectation (or, average) of $\Pr[C = c | M = m]$ over all $m \in \mathcal{M}$)

Game-based Security Definition

- This security is defined by a game between two parties: An honest challenger $\mathcal{H}$ and an adversary $\mathcal{A}$
- The game is defined as follows:
  - The adversary provides two message $m^{(0)}$ and $m^{(1)}$ of its choice to the honest challenger $\mathcal{H}$
  - The honest challenger $\mathcal{H}$ picks $\mathsf{sk} \sim \mathsf{Gen}(1^\lambda)$, $b \xleftarrow{\$} \{0,1\}$ and computes $c \sim \mathsf{Enc}_{\mathsf{sk}}(m^{(b)})$. The honest challenger $\mathcal{H}$ sends c to the adversary $\mathcal{A}$
  - The adversary $\mathcal{A}$ returns back a bit $\widetilde{b} \in \{0,1\}$ to the honest challenger $\mathcal{H}$ that is its guess of the bit b
  - The honest honest challenger $\mathcal{H}$ computes a bit z that is 1 if and only if $b = \widetilde{b}$
- The adversary $\mathcal{A}$ wins the game if $z = 1$ (i.e., its guess of b is correct)

Game-based Security Definition

- Note that it is trivial to obtain $\Pr[Z = 1] = 1/2$
- An adversary is able to distinguish the encryptions of $m^{(0)}$ from the encryptions of $m^{(1)}$ if she is able to ensure $\Pr[Z = 1] > 1/2$
- The *advantage* of an adversary $\mathcal{A}$ is defined to be: $|\Pr[Z = 1] - 1/2|$ (Think: Why do we consider it to be an advantage if an adversary can predict b with probability $< 1/2$?)

Definition (Four: Perfect Security)

For all adversary $\mathcal{A}$, its advantage in the security game define above is 0.

- Exhibit the equivalence of Definition 4 with one of the previous definitions of perfect security

One-time Pad: Perfectly-secure Encryption Scheme

- Consider the scheme defined below:
  - $\mathsf{Gen}(1^\lambda)$: Output $\mathsf{sk} \xleftarrow{\$} \{0,1\}^\lambda$
  - $\mathsf{Enc}_{\mathsf{sk}}(m)$: Output $m + \mathsf{sk}$
  - $\mathsf{Dec}_{\mathsf{sk}}(c)$: Output $c + \mathsf{sk}$
- Prove that this scheme is perfectly secure using all four definitions of perfect security
- Think: What information is leaked if two messages are encrypted using the same one-time pad sk

Looking Ahead: (Im)perfect Security(?)

- Alter the definition of Definition 4 to define some meaningful notion of "imperfect" security

Looking Ahead: More Stringent Security Notions

Current Definition Intuition (Ciphertext-only Attack):

- Given a ciphertext the adversary is not able to distinguish an encryption of $m^{(0)}$ from $m^{(1)}$

Looking Ahead: More Stringent Security Notions

Current Definition Intuition (Ciphertext-only Attack):

- Given a ciphertext the adversary is not able to distinguish an encryption of $m^{(0)}$ from $m^{(1)}$

Known-plaintext Attack:

- Given a ciphertext and a few $(m^{(i)}, c^{(i)})$ pairs, where $c^{(i)}$ is encryption of the message $m^{(i)}$ and $i > 1$, the adversary is not able to distinguish an encryption of $m^{(0)}$ from $m^{(1)}$

Looking Ahead: More Stringent Security Notions

Chosen-plaintext Attack:

- Given a ciphertext, the adversary can ask encryptions of a few other messages $m^{(i)}$, $i > 1$, and obtain their ciphertexts $c^{(i)}$. Even with this additional information, it is not able to distinguish an encryption of $m^{(0)}$ from $m^{(1)}$

Looking Ahead: More Stringent Security Notions

Chosen-ciphertext Attack:

- Given a ciphertext, the adversary can ask decryptions of a few other ciphertexts $c^{(i)}$, $i > 1$, and obtain their plaintexts $m^{(i)}$. Even with this additional information, it is not able to distinguish an encryption of $m^{(0)}$ from $m^{(1)}$

Reductions

- The security notions sorted by their requirement strengths: Ciphertext-only, Known-plaintext, chosen-plaintext, chosen-ciphertext (Prove this statement)
- Stronger securities are more difficult to achieve
- Think: Do any historical encryption schemes discussed earlier satisfy even known-plaintext attacks?
- Think: Attacks on One-time Pad using known-plaintext attacks