Warning: This document is a rough draft, so it may contain bugs. Please feel free to email me with corrections.

Course Staff

- Instructor: Yael Tauman Kalai
- TAs: Aparna Gupte and Andrew Huang

Course Website

All the information about the class can be found in the course website: `mit6875.github.io`.

Grading

- 40% of the grade: Midterm.
- 60% of the grade: Psets. We will have 5 Psets and will count your best 4 (so 15% for each Pset).
- You can collaborate on the Psets in groups of up to three, but must write your solutions separately, in your own words.
- If you need an extension (for any reason!) just email the staff, and you will be granted 72 additional hours, no questions asked! If you need more than that, please contact S^3 if you are an undergraduate and your academic advisor if you are a graduate student.

What is this class about?

This class is a foundations class where we will learn fundamental concepts in cryptography. We will see three themes:

- **Definitions:** We will learn how to think adversarially: How to model the adversary, its goals and its capabilities. We will focus on coming up with the "correct definitions" that capture the real world. We will see that often when trying to achieve a cryptographic goal, be it secrecy, integrity, zero-knowledge proofs etc., we often hit an impossibility result. Cryptography is the art of overcoming such barriers. This is often achieved by carefully choosing the definitions and models.

- **Hardness assumptions:** Most of cryptography relies on hardness assumptions, since information theoretic security is often impossible to achieve. These hardness assumptions come from various branches of mathematics: number theory, group theory, elliptic curves, lattices, and coding theory.
- **Reductions:** We prove the security of our schemes via reductions: We prove that if there exists an adversary that breaks our scheme then we can reduce this adversary to a break of the underlying hardness assumption. So "science wins either way!" (quote of Silvio Micali).

We will use these concepts to do magic! We will see how we can communicate in a secret and authenticated manner without ever meeting to share a secret! We will show how to compute on encrypted data, how to prove statements without revealing any information about why the statement is true, how to shrink proofs, and much more!

Today: Perfect Security and the One-Time Pad

Claude Shannon was the first to give a rigorous definition of a secure encryption scheme [1], and his definition is now commonly referred to as *perfect security*. He also constructed a very simple encryption scheme that satisfies this definition.

Defining Encryption Schemes

In what follows, we present Shannon's definition of a perfectly secure encryption scheme.

Definition 1. An encryption scheme is associated with a message space $\mathcal{M}$ (also referred to as the plaintext space), a ciphertext space $\mathcal{C}$ and a key space $\mathcal{K}$, and two polynomial time algorithms $(\mathsf{Enc}, \mathsf{Dec})$ with the following syntax:

- $\mathsf{Enc} : \mathcal{K} \times \mathcal{M} \to \mathcal{C}$
- $\mathsf{Dec} : \mathcal{K} \times \mathcal{C} \to \mathcal{M}$

The encryption scheme is required to satisfy the following properties:

- **Correctness:** For every $m \in \mathcal{M}$ and $k \in \mathcal{K}$,
$$\mathsf{Dec}(k, \mathsf{Enc}(k, m)) = m.$$
- **Shannon Security:** For any probability distribution M over the plaintext space $\mathcal{M}$ and every plaintext $m \in \mathcal{M}$ and ciphertext $c \in \mathcal{C}$,
$$\Pr[M = m] = \Pr_{k \leftarrow \mathcal{K}}[M = m | \mathsf{Enc}(k, M) = c].$$

Notice that there is no reference to an adversary in the security definition. However, this definition intuitively captures that the adversary (Eve) knows exactly as much about the plaintext after seeing the ciphertext as she did before. In other words, Eve does not gain any information about the plaintext m from the ciphertext c.

It turns out that Shannon security is equivalent to the following (IMO, more intuitive) definition.

Definition 2 (Perfect Indistinguishability)**.** An encryption scheme $(\mathsf{Enc}, \mathsf{Dec})$ is said to have perfect indistinguishability if for every message $m_0, m_1 \in \mathcal{M}$ and every ciphertext $c \in \mathcal{C}$

$$\Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, m_0) = c] = \Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, m_1) = c]$$

Claim 1. An encryption scheme is Shannon secure if and only if it is perfectly indistinguishable.

The proof is straightforward, and is essentially just a single application of Bayes' theorem. But since this is the first lecture we will do it carefully in class.

Proof. First, we show that any Shannon secure encryption scheme $(\mathsf{Enc}, \mathsf{Dec})$ is perfectly indistinguishable. To this end, fix any two plaintexts $m_0, m_1 \in \mathcal{M}$ and any ciphertext $c \in \mathcal{C}$. We need to prove that

$$\Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, m_0) = c] = \Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, m_1) = c].$$

Let M be the distribution defined by

$$\Pr[M = m_0] = \Pr[M = m_1] = 1/2.$$

Namely, M is the uniform distribution on $\{m_0, m_1\}$. By Shannon security, for any $b \in \{0, 1\}$,

$$\Pr[M = m_b] = \Pr_{k \leftarrow \mathcal{K}}[M = m_b | \mathsf{Enc}(k, M) = c].$$

By Bayes' theorem,

$$\Pr_{k \leftarrow \mathcal{K}}[M = m_b | \mathsf{Enc}(k, M) = c] = \frac{\Pr[M = m_b] \cdot \Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, m_b) = c]}{\Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, M) = c]}$$

Putting these two equalities together, we conclude that

$$\Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, m_b) = c] = \Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, M) = c].$$

By the definition of the distribution M

$$\Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, M) = c] = \frac{1}{2} \cdot \Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, m_0) = c] + \frac{1}{2} \cdot \Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, m_1) = c],$$

which together with the equation above, implies that

$$\Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, m_0) = c] = \Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, m_1) = c],$$

as desired.

Next, suppose that our encryption scheme is perfectly indistinguishable, and we will prove that it is Shannon secure. To this end, let M be any distribution over the plaintext space $\mathcal{M}$, and fix any $m_0 \in \mathcal{M}$ and $c \in \mathcal{C}$. By Bayes' rule,

$$\Pr_{k \leftarrow \mathcal{K}}[M = m_0 | \mathsf{Enc}(k, M) = c] = \frac{\Pr[M = m_0] \cdot \Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, m_0) = c]}{\Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, M) = c]}$$

By perfect indistinguishability,

$$\begin{aligned}
&\Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, M) = c] = \\
&\sum_{m_1 \in \mathcal{M}} \Pr[M = m_1] \cdot \Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, m_1) = c] = \\
&\sum_{m_1 \in \mathcal{M}} \Pr[M = m_1] \cdot \Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, m_0) = c] = \\
&\Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, m_0) = c]
\end{aligned}$$

Plugging this in to the above, we see that

$$\Pr_{k \leftarrow \mathcal{K}}[M = m_0 | \mathsf{Enc}(k, M) = c] = \Pr[M = m_0],$$

as desired. □

Here is yet another definition that is equivalent to the previous two, and where the adversary Eve is considered explicitly. It is also more similar to most of the other definitions that we will see in this course.

Definition 3 (Perfect security against an adversary)**.** An encryption scheme $(\mathsf{Enc}, \mathsf{Dec})$ is perfectly secure against an adversary if for any adversary $\mathcal{E} : \mathcal{C} \to \{0, 1\}$ and any pair of messages $m_0, m_1 \in \mathcal{M}$,

$$\Pr_{b \leftarrow \{0,1\}, k \leftarrow \mathcal{K}}[\mathcal{E}(\mathsf{Enc}(k, m_b)) = b] = 1/2.$$

It is a good exercise to convince yourself that this definition is equivalent to perfect indistinguishability (and thus to Shannon security).

The One-Time Pad

Shannon not only gave the first rigorous definition of a secure encryption scheme. He also constructed a scheme that satisfies this definition. The construction is known as the *One-Time Pad*.

In this scheme, the message space $\mathcal{M}$, the ciphertext $\mathcal{C}$ and the key space $\mathcal{K}$ are all equal to $\{0, 1\}^n$, for any integer n of our choice.

The one-time pad was discovered and used previously to Shannon, dating back to Frank Miller in 1882. However, Shannon was the first to provide formal guarantees.

We simply choose n large enough to accomodate the plaintexts that we'd like to send.

- For every $k, m \in \{0,1\}^n$, $\mathsf{Enc}(k, m) = k \oplus m$.
- For every $k, c \in \{0,1\}^n$, $\mathsf{Dec}(k, c) = k \oplus c$.

The one-time pad is very elegant, simple, and efficient. It is also very easy to prove that it's perfectly indistinguishable, which immediately implies that it is also Shannon secret (since we proved that the two definitions are equivalent).

Theorem 4. *The one-time pad is perfectly indistinguishable.*

Proof. For any two plaintexts $m_0, m_1 \in \{0,1\}^n$ and any ciphertext $c \in \{0,1\}^n$ there are unique keys $k_0 := m_0 \oplus c$ and $k_1 := m_1 \oplus c$ satisfying $\mathsf{Enc}(k_b, m_b) = c$. Therefore,

$$\Pr_{k \leftarrow \{0,1\}^n}[\mathsf{Enc}(k, m_b) = c] = \Pr_{k \leftarrow \{0,1\}^n}[k = k_b] = 2^{-n},$$

and thus

$$\Pr_{k \leftarrow \{0,1\}^n}[\mathsf{Enc}(k, m_0) = c] = \Pr_{k \leftarrow \{0,1\}^n}[\mathsf{Enc}(k, m_1) = c],$$

as desired. □

The one-time pad can be used only once!

The one-time pad was used significantly in practice (especially by diplomats to transmit classified information for example during World War 2). However, it is important to note that the one-time pad can be used to encrypt only n bits. If Alice and Bob want to exchange a message of length 1 GB then they need to exchange a key of size 1 GB.

Notice that if we use the same key $k \leftarrow \{0,1\}^n$ to encrypt two messages $m_0, m_1 \in \{0,1\}^n$ then security is broken, since

$$\mathsf{Enc}(k, m_0) \oplus \mathsf{Enc}(k, m_1) = m_0 \oplus m_1$$

The sad fact is that this is inherent!

Theorem 5. *If* $(\mathsf{Enc}, \mathsf{Dec})$ *is perfectly indistinguishable then* $|\mathcal{K}| \geq |\mathcal{M}|$.

Proof. For every plaintext $m \in \mathcal{M}$ and every ciphertext c in the image of Enc there must be at least one key that maps m to c, since otherwise the scheme is not perfectly indistinguishable. Since two distinct plaintexts cannot map to the same ciphertext under the same key (because then we could not possibly have correctness), we must have at least as many keys as ciphertexts. Since there must be at least one distinct ciphertext for each plaintext, this implies that we must have as many keys as plaintexts. □

Remark. We will later consider randomized encryption algorithms. We mention that the same impossibility result holds for randomized encryptions, but the proof is slightly more delicate.

The above impossibility result in unacceptable! We would like to agree on a single key $k \leftarrow \{0,1\}^n$ and then encrypt arbitrarily many messages! What can we do (given the impossibility result above)? Clearly, we should somehow weaken the security definition! To see how, let's examine the attack:

suppose that we have some encryption scheme for which $|\mathcal{K}| < |\mathcal{M}|$, and let's try to understand what the above proof tells us about Eve's ability to break this scheme. Recall that Eve's goal is to take as input a ciphertext c and guess whether it is an encryption of m_0 or m_1, with success probability better than $1/2$. Given a ciphertext $c = \mathsf{Enc}(k, m_b)$, Eve will compute the set $M_c := \{\mathsf{Dec}(k', c) : k' \in \mathcal{K}\}$. If $m_0 \in M_c$ and $m_1 \notin M_c$, then Eve will output output 0. Similarly, if $m_1 \in M_c$ and $m_0 \notin M_c$, then Eve will output 1. If $m_0, m_1 \in M_c$ then Eve will output a random bit b'.

The above proof shows that, for at least one pair of messages $m_0, m_1 \in \mathcal{M}$, there is a non-zero probability $p > 0$ that one of the plaintexts will not lie in M_c, in which case Eve will succeed with probability at least $(1 + p)/2 > 1/2$.

But, computing whether $m_b \in M_c$ is extremely challenging, and takes time roughly 2^n, which even for $n = 256$ is more than the number of molecules on earth! Recall that Eve is meant to represent some entity in the real world, so let's model her as such, and restrict her running time to be significantly less than 2^n. This turns out to be a very good idea!

References

[1] Claude E. Shannon. Communication theory of secrecy systems. *Bell System Technical Journal*, 28(4):656–715, 1949.