

Introduction to Cryptography

CSCI-GA.3210-001

Instructor: Nir Bitansky

Lecture 1: Introduction


1 What is Cryptography?

- According to Wikipedia: *Cryptography (from the Greek words kryptós, "hidden, secret," and graphein, "writing") is the practice and study of techniques for secure communication in the presence of third parties called adversaries.*
- Private communication is indeed the oldest and most basic problem in cryptography, it will also be the starting point for this course.
- As we shall see the scope of crypto has greatly expanded beyond this problem, addressing a host of problems (in the digital domain).
The common goal: "Prevent the adversary from executing his evil bidding."

2 The Crypto Way of Thinking: The Case of Encryption

The problem of communicating secretly is as old as humanity itself. Today we'll touch snippets of history to demonstrate how it transformed from an art into a well-founded mathematical concept.

The Basic Setting:



Main question: can we prevent Eve from learning about m .


Solutions from the Old World - Encryption as an Art. We'll now cover a few examples/anecdotes of "artistic encryption".

- *The Caesar cipher* shifts every letter by a fixed amount. Identifying the alphabet with $[26]$, a key $s \in [26]$ encrypts a letter T as

$$\tilde{T} = T + s \bmod 26.$$

Since there are only 26 possible shifts, an attacker can simply try every option and look for the resulting meaningful plaintext.

- *Substitution cipher* used by queen Mary:



In 1587, Mary the queen of Scots, and the heir to the throne of England, wanted to arrange the assassination of her cousin, queen Elisabeth I of England, so that she could ascend to the throne and finally escape the house arrest under which she has been for the last 18 years. She sent an encrypted letter to Sir Anthony Babington. It is a substitution cipher where each letter is transformed into a different symbol according to an arbitrary secret mapping.

Nevertheless, Elisabeth's spies broke the cipher, and Mary and her coconspirators were executed! (An important lesson — bad crypto seldom goes unpunished.) How did the spies break the cipher?

- The *Vigenère cipher* was invented in the 16th century (it was actually invented by Italian cryptographer Bellaso) and lasted for roughly 300 years!!

It (or rather, its English adaptation) works as follows. The secret key consists of an integer n and random numbers $S_1, \dots, S_n \leftarrow [26]$. To encrypt a text $T_1, \dots, T_N$ where each $T_i \in [26]$ is a letter. Let the encrypted letter be $\tilde{T}_i = T_i + S_{i \bmod n} \bmod 26$.

Harder to apply frequency analysis - the eventual text is “smoother”. Can you still break it when $N \gg n$? The idea is to find the key length n , and then the ciphertext can be divided into n sub-texts, where each one is of length N/n and implements a simple substitution cipher, where each letter is shifted by some S (in a cyclic manner).

Kasiski examination is a method for finding the key length n . The method looks for repetitions of (short) strings in the ciphertext. Such repetitions are likely to correspond to repetitions of common words in the underlying plaintext, in which case the *distance* between each two copies is a multiple of the key length n .

- There are many other famous historical ciphers, which we won't cover, and all have been by now broken. Perhaps the most famous one is the German Enigma (watch *The Imitation Game*).

Indeed, throughout most of history, cryptographic systems had limited life spans. They would be invented, broken, fixed, broken, reinvented, and so on.

Is this inherent? Could we have unbreakable crypto?

This is the goal of modern cryptography — construct cryptographic systems whose security we can mathematically prove. Jumping ahead, sometimes, even most times, we won't be able to give unconditional proofs, and we will rely on certain assumptions. We'll defer this discussion for later.

So, we want *provably-secure* encryption. We first need to understand what this means — we need a definition. Indeed, the first step in designing crypto systems is coming up with “the right” definition, which is often challenging on its own.

Definition 2.1 (Secret Key Encryption).

- **Syntax:** An encryption scheme consists of three (possibly probabilistic) polynomial-time (PPT) algorithms (G, E, D) .
 - $G(1^n)$ is a key generation algorithm that takes as input a key-length, and outputs a key sk of length n .¹
 - $E_{sk}(M)$ is an encryption algorithm that takes as input the key sk and a message $m \in \{0, 1\}^*$, and outputs a ciphertext ct .
 - $D_{sk}(ct)$ is a decryption algorithm that takes as input the key sk and a ciphertext ct , and outputs a decrypted m .
- **Correctness:** for any key sk output by G , and any message m , it is always the case that:

$$D_{sk}(E_{sk}(m)) = m .$$

- **Security:** ???

Defining security is usually the more tricky part and requires some more thought. The first thing we need to ask is *what does the adversary (in our case, Eve) know?* Clearly, Eve sees the ciphertext ct , and does not get the secret key sk . What about the algorithms (G, E, D) themselves?

We'll follow the principle saying that the scheme's algorithms themselves are always assumed to be public. This principle was first suggested by Prussian general Kerckhoffs (1883) who rightfully argued that the system/method of encryption could quite easily be stolen by the enemy. So what is secret? and how can we keep using a system that is known to the adversary? The imminent answer is *randomness*. In particular, as long as the secret key is chosen at random, we can hope for security. If the key is stolen, we can sample a fresh random key, and we don't have to throw away the entire system!!

With this principle in mind, let's try to define security. We're going to consider several attempts, some will look quite stupid, but are useful for our understanding.

- **Security, 1st attempt:** (G, E, D) is secure if no adversary A can recover the key sk better than by guessing: for all messages $m \in \{0, 1\}^*$,

$$\Pr \left[A(ct) = sk \mid \begin{array}{l} sk \leftarrow G(1^n) \\ ct \leftarrow E_{sk}(m) \end{array} \right] \leq 2^{-n} .$$

Is this a good definition?

Claim 2.2. The scheme where $sk \leftarrow \{0, 1\}^n$ is chosen at random, $E_{sk}(m) = m$, and D is the identity is secure according to this definition.

- **Security, 2nd attempt:** (G, E, D) is secure if no adversary A can recover the message m better than by guessing the key: for all messages $m \in \{0, 1\}^*$,

$$\Pr \left[A(ct) = m \mid \begin{array}{l} sk \leftarrow G(1^n) \\ ct \leftarrow E_{sk}(m) \end{array} \right] \leq 2^{-n} .$$

Is this a good definition?

Claim 2.3. No encryption scheme satisfies this definition.

Clearly, an adversary that outputs some fixed message, for instance 0^n , guesses $m = 0^n$ with probability one. But should this be considered an attack? Does this adversary learn anything about the message that she did not know before?

¹For secret key encryption, we can assume w.l.o.g that sk is drawn uniformly at random from $\{0, 1\}^n$, denoted $sk \leftarrow \{0, 1\}^n$. This is because we can always treat the random coins used by G as the secret key, and apply G as part of the encryption algorithm.

Shannon's Perfect Secrecy. In 1949, Shannon formalized the notion of *perfectly secret encryption* (this is often viewed as the dawn of modern cryptography). The leading principle was indeed that *a ciphertext shouldn't add information about the underlying plaintext messages beyond what is already known*. It turns out that there are several equivalent ways to formulate this, and we can gain different valuable intuitions from them.

- **Perfect Secrecy 1:** (G, E, D) is perfectly secure for messages of size $\ell = \ell(n)$ if for any two messages $m_0, m_1 \in \{0, 1\}^\ell$, no adversary A can distinguish an encryption of one from the other:

$$\Pr \left[A(ct) = b \mid \begin{array}{l} sk \leftarrow G(1^n) \\ b \leftarrow \{0, 1\} \\ ct \leftarrow E_{sk}(m_b) \end{array} \right] \leq \frac{1}{2} .$$

- **Perfect Secrecy 2:** (G, E, D) is perfectly secure for messages of size $\ell = \ell(n)$ if there exists a fixed distribution S_ℓ such that for all $m \in \{0, 1\}^\ell$:

$$\left\{ ct \mid \begin{array}{l} sk \leftarrow G(1^n) \\ ct \leftarrow E_{sk}(m) \end{array} \right\} \equiv \{ \tilde{ct} \mid \tilde{ct} \leftarrow S_\ell \} .$$

The above means that, without the secret key, the ciphertext distribution that the adversary sees is completely independent of the encrypted message up to its length. (Indeed, the definition does allow leaking the length of messages. In fact, some leakage on the length is inherent, think why...)

Claim 2.4. *The above two definitions are equivalent.*

Proof sketch. Assume that (1) holds. Define

$$S_\ell := \left\{ \tilde{ct} \mid \begin{array}{l} sk \leftarrow G(1^n) \\ \tilde{ct} \leftarrow E_{sk}(0^\ell) \end{array} \right\} .$$

We need to show that for all $m \in \{0, 1\}^\ell$, the above is distributed like

$$S_\ell(m) := \left\{ \tilde{ct} \mid \begin{array}{l} sk \leftarrow G(1^n) \\ \tilde{ct} \leftarrow E_{sk}(m) \end{array} \right\} .$$

Assume that for some m^* , $S_\ell \not\equiv S_\ell(m^*)$. We will show an adversary A that distinguishes encryptions of $m_0 = 0^\ell$ from ones of $m_1 = m^*$. Let ct^* be such that

$$\Pr[S_\ell = ct^*] - \Pr[S_\ell(m^*) = ct^*] > 0 \quad (\text{why does it exist?}).$$

And let us denote this quantity by ε .

Consider the adversary A that given ct outputs 0 if $ct = ct^*$, and otherwise outputs 1.² We now analyze this adversary. In what follows, all the probabilities are over $sk \leftarrow G(1^n), b \leftarrow \{0, 1\}, ct \leftarrow E_{sk}(m_b)$. Then

$$\Pr[A(ct) = b] = \tag{1}$$

$$\Pr[b = 0] \Pr[A(ct) = 0 \mid b = 0] + \Pr[b = 1] \Pr[A(ct) = 1 \mid b = 1] ,$$

$$\Pr[b = 0] = \Pr[b = 1] = \frac{1}{2} , \tag{2}$$

$$\Pr[A(ct) = 0 \mid b = 0] = \Pr[ct = ct^* \mid b = 0] = \Pr[S_\ell = ct^*] , \tag{3}$$

$$\Pr[A(ct) = 1 \mid b = 1] = \Pr[ct \neq ct^* \mid b = 1] = 1 - \Pr[S_\ell(m^*) = ct^*] . \tag{4}$$

²Other natural attempts, such as outputting a random bit $b' \leftarrow \{0, 1\}$ when $ct \neq ct^*$, or outputting the bit corresponding to the distribution under which ct is more likely (and outputting 0 in case of a tie), would also work.

From these four equations, we get that

$$\begin{aligned}\Pr[A(ct) = b] &= \frac{1}{2} \Pr[S_\ell = ct^*] + \frac{1}{2} (1 - \Pr[S_\ell(m^*) = ct^*]) \\ &= \frac{1}{2} (1 + \Pr[S_\ell = ct^*] - \Pr[S_\ell(m^*) = ct^*]) \\ &= \frac{1}{2} (1 + \varepsilon) > \frac{1}{2} .\end{aligned}$$

We got that $\Pr[A(ct) = b] > \frac{1}{2}$, which is a contradiction to our assumption that (1) holds.

For the other direction, assume that (2) holds, and we will show that (1) holds. Let $m_0, m_1 \in \{0, 1\}^\ell$ be two arbitrary messages of length ℓ . Note that

$$\Pr[A(E_{sk}(m_b)) = b] = \Pr[A(\tilde{ct}) = b \mid \tilde{ct} \leftarrow S_\ell] = 1/2 .$$

Where the first equality follows from the fact that $E_{sk}(m_b) \equiv S_\ell$, and the latter by the fact that A 's input, and thus also its output, are independent of b . $\square$

Another equivalent definition worth baring in mind:

- **Perfect Secrecy 3:** (G, E, D) is perfectly secure for messages of size $\ell = \ell(n)$ if for any distribution M on messages in $\{0, 1\}^\ell$ any ciphertext ct in the support of E and any message $m' \in \{0, 1\}^\ell$:³

$$\Pr_{m \leftarrow M}[m = m'] = \Pr_{m \leftarrow M, sk \leftarrow G(1^n)}[m = m' \mid E_{sk}(m) = ct]$$

Here the distribution M should be thought of as partial information that the adversary has regarding encrypted messages. For instance, all German messages in WW2 ended with "Heil Hitler". Of course, we cannot expect to hide this fact, again all that we ask is not to add information — namely conditioning on an encryption of a message does not change its distribution.

Another useful definition, which may seem less general, but is equivalent:

- **Perfect Secrecy 4:** (G, E, D) is perfectly secure for messages of size $\ell = \ell(n)$ if for any set $S \subseteq \{0, 1\}^\ell$:

$$\Pr \left[A(ct) = m \mid \begin{array}{l} sk \leftarrow G(1^n) \\ m \leftarrow S \\ ct \leftarrow E_{sk}(m) \end{array} \right] \leq \frac{1}{|S|} .$$

So now that we've defined perfect ciphers, do they exist?

Theorem 2.5 (Vernam 1917). *There exists a perfect cipher for messages of size n with keys and ciphertexts of size n .*

Proof. The cipher is the *one-time pad* (OTP):

- $G(1^n)$ outputs a random "pad" $p \leftarrow \{0, 1\}^n$.
- $E_p(m)$ is the bitwise Xor $p \oplus m$.
- $D_p = E_p$.

This is correct as $p \oplus (p \oplus m) = m$.

It is perfectly secure since for any m , $p \oplus m$ is uniformly random in $\{0, 1\}^n$, when p is chosen at random. $\square$

³Here by the support of E , we mean the support of the distribution over ciphertexts $E_{sk}(m)$ obtained by sampling $sk \leftarrow G(1^n)$ and $m \leftarrow M$.

Why only one time? Can we use OTP two times? It's not hard to see that encrypting two messages using the same key, already leaks information about them (their Xor). Indeed, throughout history, misuse of OTP has led to devastating leakage. This makes OTP very hard to use. Imagine that every time you give your credit card online, you first have to go somewhere physically to take your one-time pad.

Perhaps there is a better encryption system where we *can* use the same key multiple times? This turns out to be the heart of the matter. Indeed, a main challenge in cryptography is to *recycle keys* or more generally *recycle randomness*. The first hurdle to cross is that *this goal is impossible!* Well, not really... but it is if we insist on perfect secrecy.

Theorem 2.6. *A perfect cipher for messages of length ℓ must have keys of size $n \geq \ell$.*

Proof. Assume $n < \ell$. We will show an adversary A that given a ciphertext of a random message $m \leftarrow \{0, 1\}^\ell$ guesses the message with probability better than $2^{-\ell}$. A simply guesses the key and attempts to decrypt. We know that it succeeds with probability $2^{-n} > 2^{-\ell}$. $\square$

So perfect secrecy is too much to ask for. Still guessing the key does not seem like a realistic attack, assuming its moderately long (e.g., a reasonable choice for a secret key is 256 bits. Look for "how big is 2^{256} ?" on youtube...). However, there are in fact much more devastating attacks.

Claim 2.7. *For any cipher for messages of length ℓ and keys of length $n < \ell - 6$, there exist two messages $m_0 \neq m_1$ and an attacker A such that*

$$\Pr \left[A(ct) = b \mid \begin{array}{l} sk \leftarrow G(1^n) \\ b \leftarrow \{0, 1\} \\ ct \leftarrow E_{sk}(m_b) \end{array} \right] > 0.99 .$$

Proof. For simplicity, we will assume that E does not use any additional randomness beyond sk . Fix a message m_0 and consider the set of all of its encryptions $C_0 = \{E_{sk}(m_0)\}$. We will show that there exists a message m_1 , such that an encryption of m_1 (under a random key sk) falls in C_0 with tiny probability.

Indeed, note that for any key sk , if we choose $m_1 \leftarrow \{0, 1\}^\ell$ at random, then the corresponding ciphertext ct falls in C_0 with probability at most $|C_0|/2^\ell = 2^n/2^\ell < 2^{-6}$. This is true because $E_{sk}(\cdot)$ is an injective function (otherwise decryption would be impossible) and the input is drawn uniformly, so the output also distributes uniformly over a set of size 2^ℓ .

In particular the above holds for a random key. Formally, because:

$$\Pr_{sk, m_1} [ct \in C_0 \mid ct \leftarrow E_{sk}(m_1)] = \mathbb{E}_{sk} \Pr_{m_1} [ct \in C_0 \mid ct \leftarrow E_{sk}(m_1)] \leq 2^{-6}.$$

Now we can fix such specific m_1 . Formally, because:

$$\mathbb{E}_{m_1} \Pr_{sk} [ct \in C_0 \mid ct \leftarrow E_{sk}(m_1)] = \Pr_{sk, m_1} [ct \in C_0 \mid ct \leftarrow E_{sk}(m_1)] \leq 2^{-6}.$$

So there must be some m_1 such that $\Pr_{sk} [ct \in C_0 \mid ct \leftarrow E_{sk}(m_1)] \leq 2^{-6}$, otherwise the expectation would be greater than 2^{-6} .

Now, consider the adversary A that given ct outputs 0 if $ct \in C_0$, and otherwise outputs 1. By the definition of C_0 ,

$$\Pr [A(ct) = 0 \mid b = 0] = \Pr_{sk} [ct \in C_0 \mid ct \leftarrow E_{sk}(m_0)] = 1 .$$

By the definition of m_1 , we have that

$$\Pr [A(ct) = 0 \mid b = 1] = \Pr_{sk} [ct \in C_0 \mid ct \leftarrow E_{sk}(m_1)] \leq 2^{-6} .$$

Overall, it holds that

$$\begin{aligned} \Pr [A(ct) = b] &= \Pr [b = 0] \Pr [A(ct) = 0 \mid b = 0] + \Pr [b = 1] \Pr [A(ct) = 1 \mid b = 1] \\ &= \frac{1}{2} (\Pr [A(ct) = 0 \mid b = 0] + 1 - \Pr [A(ct) = 0 \mid b = 1]) \\ &\geq \frac{1}{2} (1 + 1 - 2^{-6}) > 0.99 . \end{aligned}$$

□

Does this mean that there's no hope to reuse keys? How realistic is the described attack? The point is that this attack takes a very long time, roughly as long as $|C_0| \approx 2^n$ steps. Our salvation comes from the fact that

actual adversaries are computationally bounded.

Most of modern cryptography, and accordingly this course, will rely on this premise. Indeed, as we shall see, there's (almost) no cryptography without computational hardness. This also means that cryptographic theory strongly relies on complexity theory.

One major implication of this is that our lack of understanding of major open questions, P vs NP above them all, translate to cryptography — we cannot prove security unconditionally, but only under computational assumptions. In fact, as we shall see, we will need to assume much beyond $P \neq NP$, and minimizing the assumptions made will of course be a major goal.