MAD 4471: Cryptography and Coding Theory **Fall 2020**

Lecture 2: The security of encryption schemes

Lecturer: Jean-François Biasse *TA: William Youmans*

2.1 Perfect secrecy

Informally: The cipher text reveals no information about the plain text.

Definition 2.1 (Perfect secrecy) *$\forall m_0, m_1 \in \mathcal{M}$ with $|m_0| = |m_1|, \forall c \in \mathcal{C}$,*

$$P_{k \xleftarrow{R} \mathcal{K}}(E(k, m_0) = c) = P_{k \xleftarrow{R} \mathcal{K}}(E(k, m_1) = c),$$

where $k \xleftarrow{R} \mathcal{K}$ means that k is sampled uniformly at random in $\mathcal{K}$.

Question : Is the Caesar cipher perfectly secret?
Answer : No!

For example, choose $c = AAA, m_0 = BBB, m_1 = ABA$. The only k such that $ENC(m_0, k) = c$ is $k = 1$ (shift of 1). There is no k such that $Enc(m_1, k) = c$. Therefore: $P(Enc(k, m_0) = c) = 1/26$ and $P(Enc(k, m_1) = c) = 0$.

There is a perfectly secret encryption scheme: the "One-Time-Pad".

Definition 2.2 (One Time Pad) *In the One-Time-Pad, $\mathcal{M} = \{0,1\}^n, \mathcal{K} = \{0,1\}^n, \mathcal{C} = \{0,1\}^n$. Let $m = m_1, \cdots, m_n \in \mathcal{M}$, $Enc(m, k)$ is defined by*

$$m \bigoplus k = m_1 \bigoplus k_1, m_2 \bigoplus k_2, \cdots, m_n \bigoplus k_n,$$

where $0 \bigoplus 0 = 0, 0 \bigoplus 1 = 1, 1 \bigoplus 0 = 1, 1 \bigoplus 1 = 0$. Then $Dec(c, k)$ is defined by

$$c \bigoplus k = c_1 \bigoplus k_1, c_2 \bigoplus k_2, \cdots, c_n \bigoplus k_n.$$

Proposition 2.3 *The One Time Pad has perfect secrecy.*

Proof: Let $m \in \mathcal{M}, k \in \mathcal{K}, c \in \mathcal{C}$, $Enc(m, k) = c \Leftrightarrow k \bigoplus m = c \Leftrightarrow k = c \bigoplus m$. So for each m, c there is one and only one k such that $Enc(m, k) = c$. Therefore $\forall m, c$, $P(Enc(k, m) = c) = 1/|\mathcal{K}|$ and thus $\forall m_0, m_1 \in \mathcal{M}$ with $|m_0| = |m_1|, \forall c \in \mathcal{C}$,

$$P_{k \xleftarrow{R} \mathcal{K}}(E(k, m_0) = c) = P_{k \xleftarrow{R} \mathcal{K}}(E(k, m_1) = c).$$

■

Remark 1 *The One Time Pad is not very practical due to the length of keys: they are as long as the messages.*

Remark 2 *Just because the One-Time-Pad is perfectly secure does not mean there are no other attacks against it (in particular active attacks).*

Theorem 2.4 (Shannon) *If a cipher has perfect secrecy, then* $|\mathcal{K}| \geq |\mathcal{M}|$.

Therefore, the one-time-pad is optimal in that regard. No "efficient" cipher can have perfect secrecy.

2.2 Stream ciphers

We want to use smaller keys, which requires to relax the security requirements. More specifically, we want to use a function $G : \{0,1\}^s \to \{o,1\}^n$ with $n \gg s$ where $\{0,1\}^s$ is the seed space such that $G(x)$ "looks random". The *keystream* $G(s)$ is XORed to the message m to produce an encryption c of m. It is similar to the One Time Pad, the only difference is that we use $G(s)$ instead of the key.

Definition 2.5 (Pseudo Random Generator) *Let* s, n *with* $s \ll n$, *a Pseudo Random Generator is a function* $G : \{0,1\}^s \to \{0,1\}^n$.

We will see later how to quantify how random $G(x)$ looks. Now, to use PRG to encrypt data, we proceed just like for the One-Time-Pad. $\mathcal{K} = \{0,1\}^s, \mathcal{M} = \{0,1\}^n = \mathcal{C}$ with $n \gg s$.

$$Enc(m,k) = m \bigoplus G(k) = m_1 \bigoplus G(k)|_1, m_2 \bigoplus G(k)|_2, \cdots, m_n \bigoplus G(k)|_n,$$

$Dec(c,k) = c \bigoplus G(k)$, which is correct because

$$\begin{aligned} Dec(Enc(m,k),k) &= Enc(m,k) \bigoplus G(k) \\ &= (m \bigoplus G(k)) \bigoplus G(k) \\ &= m \bigoplus (G(k) \bigoplus G(k)) \\ &= m, \end{aligned}$$

as we have $(G(k) \bigoplus G(k)) = 0$.

Definition 2.6 (Stream Cipher) *Let* G *be a Pseudo Random Generator, and* $\mathcal{M} = \mathcal{C} = \{0,1\}^n, \mathcal{K} = \{0,1\}^s$. *The stream cipher defined by* G *is the encryption scheme using* $Enc(m) := m \bigoplus G(k)$ *and* $Dec(c) := c \bigoplus G(k)$.

2.3 The security of a pseudo random generator

Before talking about the encryption of m and how secure it is, let us quantify what it means to say that $G(n)$ "looks random". We will characterize it in two (equivalent) ways. First, we intuitively require that a random sequence be "unpredictable". Indeed, if $G(k)$ can be guessed from its first $i < n$ bits, then the knowledge

of the beginning of the message suffices to recover the entire transmission. It is not entirely unrealistic to imagine that an adversary knows $m_{1,1,\ldots,1}$ as a lot of protocols use a predictable formats.

Definition 2.7 (Predictability) *We say that a PRG G is predictable if: $\exists$ an efficient algorithm $\mathcal{A}$ and $1 \leq i \leq n-1$ such that $Pr_{k \xleftarrow{R} \mathcal{K}}[\mathcal{A}(G(k)|_{1,\ldots,i} = G(k)|_{i+1}] \geq \frac{1}{2} + \epsilon$ for a non-negligible $\epsilon > 0$.*

This means, that the algorithm $\mathcal{A}$ guesses the $i+1$ bit of $G(k)$ from $G(h)|_{1,\ldots,i}$ with probability significantly better than a coin toss. In class, we defined

$$Adv_{PRG}[\mathcal{A}, G] = \left| Pr_{k \xleftarrow{R} \mathcal{K}}[\mathcal{A}(G(k)|_{1,\ldots,i} = G(k)|_{i+1}] - \frac{1}{2} \right|.$$

G is unpredictable if and only if $Adv_{PRG}[\mathcal{A}, G]$ is negligible.

Remark 3 *Whether an event has negligible depends on the context.*

- *An event of probability $\epsilon = \frac{1}{2^{30}}$ is likely to happen every 1 GB of data.*
- *An event of probability $\epsilon = \frac{1}{2^{80}}$ is unlikely to happen in the lifetime of the key.*

Example 1 *Assume that G satisfies: $G(x)|_1 \bigoplus G(x)|_2 \bigoplus \cdots G(x)|_{i+1} = 0$. Then for all x, $G(x)|_1 \bigoplus \cdots G(x)|_i = G(x)|_{i+1}$. We can define $\mathcal{A}(G(x)|_{1,\ldots,i} := G(x)|_1 \bigoplus \cdots \bigoplus G(x)|_i$. With this $\mathcal{A}$,*

$$Pr_{k \xleftarrow{R} \mathcal{K}}[\mathcal{A}(G(k)|_{1,\ldots,i} = G(k)|_{i+1}] = 1.$$

More generally, we need to prevent any kind of statistical test to distinguish $G(k)$ from "true" randomness.

Definition 2.8 (statistical test) *A statistical test is an efficient algorithm $\mathcal{A}$ on $\{0,1\}^n$ such that $\mathcal{A}(n) = 0$ or 1. By convention, we can think of 0 as 'random', 1 as 'not random'.*

The advantage of a given statistical test $\mathcal{A}$ against the PRG G is defined as

$$Adv_{Random}[\mathcal{A}, G] = |Pr_{k \xleftarrow{R} \mathcal{K}}[\mathcal{A}(G(k)) = 1] - Pr_{r \xleftarrow{R} \{0,1\}^n}[\mathcal{A}(r) = 1]|.$$

- An advantage close to 0 means that the test $\mathcal{A}$ does not distinguish $G(k)$ from random.
- An advantage close to 1 means that the test $\mathcal{A}$ allows to distinguish $G(k)$ from random.

Example 2 *If a PRG G is such that $G(k)|_n = 1$ for $2/3$ of the keys k.*
Then we define $\mathcal{A}(x) = 1$ if $x|_n = 1$, and the advantage of $\mathcal{A}$ is

$$\begin{aligned} Adv_{Random}[\mathcal{A}, G] &= |Pr_{k \xleftarrow{R} \mathcal{K}}[\mathcal{A}(G(k)) = 1] - Pr_{r \xleftarrow{R} \{0,1\}^n}[\mathcal{A}(r) = 1]| \\ &= |\frac{2}{3} - \frac{1}{2}| \\ &= \frac{1}{6} \end{aligned}$$

Definition 2.9 (Secure PRG) *A PRG G is secure if for all efficient statistical test $\mathcal{A}$, $Adv_{PRG}[\mathcal{A}, G]$ is negligible.*

Question: How do the notions of unpredictability and security of a PRG relate to each other?
Answer: They are equivalent.

Theorem 2.10 *A secure PRG is unpredictable.*

Proof: We prove that if the PRG G is predictable, then it is insecure. Suppose there exists an efficient algorithm $\mathcal{A}$ such that $Pr_{k \xleftarrow{R} \mathcal{K}}[\mathcal{A}(G(k)|_{1,\ldots,i}) = G(k)|_{i+1}] = \frac{1}{2} + \epsilon$ for a non-negligible ϵ.
Then we describe the statistical test $\mathcal{B}$ by:
$\mathcal{B}(x) = \begin{cases} 1, & \text{if } \mathcal{A}(x_{1,\ldots,i}) = x|_{i+1} \\ 0, & \text{otherwise} \end{cases}$
The statistical advantage of this test satisfies

$$\begin{aligned} Adv_{Random}[\mathcal{B}, G] &= \left| Pr_{k \xleftarrow{R} \mathcal{K}}[\mathcal{B}(G(k)) = 1] - Pr_{n \xleftarrow{R} \{0,1\}^n}[\mathcal{B}(n) = 1] \right| \\ &= \left| \frac{1}{2} + \epsilon - \frac{1}{2} \right| \\ &= \epsilon \end{aligned}$$

Therefore G is insecure. ■

Theorem 2.11 (Yao 1982) *An unpredictable PRG is secure.*

Example 3 *Given the above statement, if G is such that one can guess the first $\frac{n}{2}$ bits from the last $\frac{n}{2}$ bits of $G(h)$, then G is predictable. This is counter-intuitive given the definition of predictability, by G is clearly insecure, therefore it has to be predictable.*

Generalization: Two distributions $\mathcal{D}_1, \mathcal{D}_2$ over $\{0,1\}^n$ are indistinguishable if $\forall$ efficient test $\mathcal{A}$:

$$\left| Pr_{x \xleftarrow{R} \mathcal{D}_1}[\mathcal{A}(x) = 1] - Pr_{x \xleftarrow{R} \mathcal{D}_2}[\mathcal{A}(x) = 1] \right| < \epsilon$$

for a negligible ϵ. We denote it $\mathcal{D}_1 \approx_p \mathcal{D}_2$. Clearly, a PRG G is secure if and only if $\{G(h)\}_{h \xleftarrow{R} \mathcal{K}} \approx_p$ Uniform distribution over $\{0,1\}^n$.

2.4 Semantic security (one time key)

Informally: an adversary cannot choose two messages m_0 and m_1, and be able to distinguish there's encryption.

We see now how to use secure PRG to describe secure stream cipher. But first we define a secure cipher (semantically). We test over adversary in two different experiments.

Experiment 0: $\mathcal{A}$ stands in front of a black box $\mathcal{B}_0$that chooses $h \xleftarrow{R} \mathcal{K}$ and returns $Enc(m_0, h)$.

Figure 2.1: Experiment 0

CHALLENGER		ADVERSARY
$k \xleftarrow{\$} \mathcal{K}$.	$\xleftarrow{m_0, m_1}$	Choose $m_0, m_1 \in \mathcal{M}$
Computes $c = \mathrm{Enc}(\mathrm{m}_0, \mathrm{k})$.	$\xrightarrow{c}$	Return 0 if it thinks it is the encryption of m_0 and 1 otherwise

Let W_0 be the event "$\mathcal{A}$ returns 1 in Experiment 0".

Experiment 1: $\mathcal{A}$ stands in front of a black box $\mathcal{B}_1$, that chooses $h \leftarrow \mathcal{K}$ and denote by $Enc(m, h)$.

Figure 2.2: Experiment 1

CHALLENGER		ADVERSARY
$k \xleftarrow{\$} \mathcal{K}$.	$\xleftarrow{m_0, m_1}$	Choose $m_0, m_1 \in \mathcal{M}$
Computes $c = \mathrm{Enc}(\mathrm{m}_1, \mathrm{k})$.	$\xrightarrow{c}$	Return 0 if it thinks it is the encryption of m_0 and 1 otherwise

Let W_1 be the event "$\mathcal{A}$ returns 1 in Experiment 1".

Definition 2.12 (Advantage against an encryption scheme) *We define the advantage of $\mathcal{A}$ against the scheme E by*

$$Adv_{ss}[\mathcal{A}, E] = |Pr[W_0] - Pr[W_1]|\,.$$

- An advantage close to 0 means that $mathcalA$ behaves the same in front of $\mathcal{B}_0$ and $\mathcal{B}_1$.
- An advantage close to 1 means that $\mathcal{A}$ distinguishes $\mathcal{B}_0$ and $\mathcal{B}_1$ $\rightarrow$ the scheme is not schematically secure.

Remark 4 *The answer of $\mathcal{A}$ does not matter !. It is the difference between its in front of $\mathcal{B}_0$ and $\mathcal{B}_1$, so that quantifies the security.*

Proposition 2.13 *$\mathcal{A}$ stream cipher used with a secure PRG is semantically secure.*