Fundamentals of Cryptography — Handout 3.

Yu Yu

Computational security, and pseudorandomness.

1 A Computational Approach to Cryptography

In the previous two lectures we have studied the classical cryptography, and cryptographic schemes that are provably (perfectly or statistically) secure against adversaries with unlimited computational power. However, the strong security requires a long key (of at least the same length as the message) which is not practical for most applications.

Modern cryptography adopts the computational approach to mitigate the problem. Recall one of the six Kerckhoff's principles reads as follows:

"A cipher must be practically, if not mathematically, indecipherable."

which essentially says that it's not necessary to use perfectly-secure encryption, as long as it serves the purposes. For example, it suffices to use an efficient encryption scheme that is insecure against adversaries of unlimited computational power, but cannot be broken with probability better than 10^{-30} in 50 years (as after such a long duration even top secrets may already be disclosed through other channels or declassified due to loss of sensitivity) using the fastest supercomputer.

The computational approach incorporates two relaxations of the notion of perfect security:

1. Security is only preserved against efficient adversaries.
2. Adversaries can potentially succeed with some very small probability (so small that most of the times it's not really happening).

The concrete approach quantifies the security of a given cryptographic scheme by bounding the running time and success probability of the adversary by parameters t and ε respectively. That is, we say that a scheme is (t,ε)-secure if every adversary of running time at most t succeeds in breaking the scheme with probability at most ε. The concrete approach is useful in practice as it directly gives concrete terms of security we care about. However, from the theoretical perspective, its disadvantage is quite obvious. For what range of (t, ε) should we say that a (t, ε)-secure scheme is actually secure? we don't have a good answer for it.

The asymptotic approach , rooted in complexity theory, is the one we will adopt. It views the running time (or circuit size in the non-uniform complexity model) and the success probability as functions of a **security parameter** n[1]. During the initialization of a cryptographic scheme, the challenger chooses an appropriate value for n (one can think of n as the key length), which is also known to the adversary. Then, the running time of the adversary (and the challenger) will be a function of n, i.e., $t(n)$, and so is it with his success probability $\varepsilon(n)$.

[1] We often write n in unary, namely 1^n (a string of n 1's). This is to be in line with the standard convention from complexity theory, where the efficiency is measured as a function of the length of its input.

1. Efficiency. A scheme (e.g., considering the triplet of algorithms Gen, Enc, Dec) is efficient if it can be computed in running time $t = \mathsf{poly}(n)$.

2. Security. A scheme is secure if for every polynomial poly there exists a negligible function negl such that any adversary of running time (no more than) $t = \mathsf{poly}(n)$ succeeds with probability no more than $\mathsf{negl}(n)$ for all large enough n's.
   Equivalently, a scheme is considered as secure if it is $(t(n),\varepsilon(n))$-secure (as defined in the concrete approach) for some super-polynomial t and negligible function ε.

Definition 1. A private-key encryption scheme is a tuple of probabilistic polynomial-time (PPT) algorithms (Gen,Enc,Dec) such that:

1. (Key generation). The key-generation algorithm Gen takes as input the security parameter 1^n and outputs a key k, i.e., $k \leftarrow \mathsf{Gen}(1^n)$.

2. (Encryption). The (possibly probabilistic) encryption algorithm Enc takes as input a key k and a plaintext message $m \in \{0,1\}^{\ell(n)}$, and outputs a ciphertext c, i.e., $c \leftarrow \mathsf{Enc}_k(m)$.

3. (Decryption). The (deterministic) decryption algorithm Dec takes as input a key k and a ciphertext c, and outputs a message m, i.e., $m := \mathsf{Dec}_k(c)$.

Correctness (Soundness) requires that for every n, every key k generated by $\mathsf{Gen}(1^n)$, and every $m \in \{0,1\}^{\ell(n)}$, it holds that $\mathsf{Dec}_k(\mathsf{Enc}_k(m)) = m$.

2 Computationally indistinguishable encryptions.

Similar to the statistical security introduced in Handout 2, computationally indistinguishability for any encryption scheme Π=(Gen,Enc,Dec) can based on the following experiment involving an adversary A (decoupled into a pair of algorithms A_1 and D) and a challenger C.

Challenger C		Adversary $\mathsf{A} = (\mathsf{A}_1, \mathsf{D})$
$k \leftarrow \mathsf{Gen}(1^n)$;	$\xleftarrow{(m_0, m_1)}$	$(m_0, m_1, state) \leftarrow \mathsf{A}_1(1^n)$
$b \leftarrow U_1$; $c := \mathsf{Enc}_k(m_b)$	$\xrightarrow{c}$	$b' \leftarrow \mathsf{D}(c, state)$
$\mathsf{PrivK}^{eav}_{\mathsf{A},\Pi}$ outputs 1 iff $b = b'$.		

Figure 1: The adversarial indistinguishability experiment $\mathsf{PrivK}^{eav}_{\mathsf{A},\Pi}$ between A and C.

Note that A_1 is possibly probabilistic so she may passes some state information (such as the choices m_0, m_1 and intermediate results) to the distinguisher D.

Definition 2 (computationally indistinguishable encryption)**.** An encryption scheme $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ has indistinguishable encryptions in the presence of an eavesdropper if for all PPT adversaries A there exists a negligible function negl such that

$$\Pr[\mathsf{PrivK}^{eav}_{\mathsf{A},\Pi} = 1] \leq \frac{1}{2} + \mathsf{negl}(n) \ .$$

where the probability is taken over the random coins of A as well as those used by the experiment (key generation, random bit b, and any randomness used by the encryption algorithm).

Definition 3 (computationally indistinguishable encryption – alternative definition)**.** An encryption scheme $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ has indistinguishable encryptions in the presence of an eavesdropper if for all PPT adversaries $\mathsf{A} = (\mathsf{A}_1, \mathsf{D})$ there exists a negligible function negl such that if for every $m_0 \neq m_1 \in \mathcal{M}$ it holds that

$$\left| \Pr[\ \mathsf{D}(\mathsf{Enc}_k(m_0), state) = 1\ ] \ - \ \Pr[\ \mathsf{D}(\mathsf{Enc}_k(m_1), state) = 1\ ] \right| \ \leq \ \mathsf{negl}(n) \quad . \tag{1}$$

where the probability is taken over $k \leftarrow \mathsf{Gen}(1^n)$, $(m_0, m_1, state) \leftarrow \mathsf{A}_1(1^n)$ and the random coins of D and Enc (if probabilistic).

Lemma 1. *Definition 2 and Definition 3 are equivalent.*

Proof. Similar to Lemma 1 from Handout 2, it suffices to prove that

$$\Pr[b = b'] \ = \ \frac{1}{2} + \frac{1}{2} \cdot \left(\Pr[\mathsf{D}(\mathsf{Enc}_k(m_1), state) = 1] \ - \ \Pr[\mathsf{D}(\mathsf{Enc}_k(m_0), state) = 1] \right) \ .$$

□

Claim 1.1 (Claim 3.11 from KL book)**.** *Let (*Gen,Enc,Dec*) be a private-key encryption scheme that has indistinguishable encryptions in the presence of an eavesdropper. Then for all PPT adversaries* $\tilde{\mathsf{A}}$ *and all* i [2]*, there exists a negligible function* negl *such that:*

$$\Pr[\ \tilde{\mathsf{A}}(1^n, \mathsf{Enc}_k(m)) \ = \ m^i\ ] \ \leq \ \frac{1}{2} + \mathsf{negl}(n)$$

where m *is chosen uniformly at random over* $\mathcal{M}$*,* m^i *denotes the* i^{th} *bit of* m*, and the probability is taken over the random coins of* $\tilde{\mathsf{A}}$*, the choice of* m *and the key* k*, and any random coins used by* Enc*.*

Proof. Suppose towards contradiction that there exists an efficient $\tilde{\mathsf{A}}$ and i and non-negligible function $\varepsilon(\cdot)$ such that

$$\varepsilon(n) \ \stackrel{\text{def}}{=} \ \Pr[\ \tilde{\mathsf{A}}(\ \mathsf{Enc}_k(m)\ ) \ = \ m^i] - \frac{1}{2}$$

where we omit 1^n from. We define two subsets $\mathcal{M}_0$, $\mathcal{M}_1$ of message space $\mathcal{M}$ whose i^{th} bits are 0 and 1 respectively.

$$\mathcal{M}_0 \stackrel{\text{def}}{=} \{m \in \mathcal{M} : m^i = 0\}, \qquad \mathcal{M}_1 \stackrel{\text{def}}{=} \{m \in \mathcal{M} : m^i = 1\} \ .$$

Then define adversary A' that on input 1^n, picks $m_0 \stackrel{\$}{\leftarrow} \mathcal{M}_0$ and $m_1 \stackrel{\$}{\leftarrow} \mathcal{M}_1$ and send them to the challenger. Upon receiving $c = \mathsf{Enc}_k(m_b)$ from the challenger, A' invokes $\tilde{\mathsf{A}}$ on c and let $b' = \tilde{\mathsf{A}}(c)$. Then we have

$$\Pr[\tilde{\mathsf{A}}(\mathsf{Enc}_k(m_b)) = b'] = \frac{\Pr[\tilde{\mathsf{A}}(\mathsf{Enc}_k(m_0)) = 0]}{2} + \frac{\Pr[\tilde{\mathsf{A}}(\mathsf{Enc}_k(m_1)) = 1]}{2} = \Pr[\tilde{\mathsf{A}}(\mathsf{Enc}_k(m)) = m^i] = \frac{1}{2} + \varepsilon(n)$$

which contradicts the assumption that the encryption scheme has indistinguishable encryptions against all PPT adversaries (including the A' as defined above) . □

[2] If the message space is $\mathcal{M} = \{0,1\}^{\ell(n)}$, then $i \in [\ell(n)]$.

Claim 1.2 (Claim 3.12 from the KL book)**.** *Let (*Gen,Enc,Dec*) be a private-key encryption scheme (with message space $\mathcal{M} = \{0,1\}^n$) that has indistinguishable encryptions in the presence of an eavesdropper. Then for every PPT adversary* A *there exists a PPT algorithm* A′ *such that for every polynomial-time computable function f and every efficiently-sampleable set $\mathcal{S}$, there exists a negligible function* negl *such that:*

$$| \Pr[\mathsf{A}(1^n, \mathsf{Enc}_k(m)\ ) = f(m)\ ] \ - \ \Pr[\mathsf{A}'(1^n) = f(m)] \ | \ \leq \ \mathsf{negl}(n) \quad , \tag{2}$$

where m is chosen uniformly at random from $\mathcal{S}_n \stackrel{\text{def}}{=} \mathcal{S} \cap \{0,1\}^n$, and the probability is taken over the choices of m and the k, and any random coins used by A*,* A′*, and* Enc*.*

Proof. For every PPT A on input 1^n (which we omit hereafter) and ciphertext c, we claim that it holds that

$$|\Pr[\mathsf{A}(\mathsf{Enc}_k(m)) = f(m)] \ - \ |\Pr[\mathsf{A}(\mathsf{Enc}_k(0^n))] = f(m)]| \ \leq \ \mathsf{negl}(n) \tag{3}$$

where probability is taken over the coins of k, m, Enc_k, f and A. Then we define $\mathsf{A}'(1^n) \stackrel{\text{def}}{=} \mathsf{A}(1^n, \mathsf{Enc}_k(0^n))$ where k is uniformly sampled, and it is easy to see that A′ satisfies Equation (2). It remains for us to prove Equation (3) as we claimed. It follows from the assumption of indistinguishable encryptions by defining $\tilde{\mathsf{A}} = (\tilde{\mathsf{A}}_1, \mathsf{D})$ as follows:

1. $\tilde{\mathsf{A}}_1$ samples $m_0 = m \leftarrow \mathcal{S}_n$ and $m_1 := 0^n$.
2. Upon receiving $c = \mathsf{Enc}_k(m_b)$ from the challenger, D invokes A and outputs 1 iff $\mathsf{A}(c) = f(m)$.

which completes the proof by Definition 3 whose Equation (1) is essentially Equation (3) by substituting variables. □

We give the following definition about "semantic security" which is more general the than claim proven above. It is known to the equivalent to "indistinguishable encryptions" stated as Theorem 2 (whose proofs follow a similar route of the above claim and are omitted here).

Definition 4 (semantic security (DEF 3.13 from KL))**.** A private-key encryption scheme (Gen,Enc,Dec) is semantically secure in the presence of an eavesdropper if for every PPT algorithm A there exists a probabilistic polynomial-time algorithm A′ such that for every efficiently-sampleable distribution M and all polynomial-time computable functions f and h, there exists a negligible function negl such that

$$\Pr[\mathsf{A}(1^n, \mathsf{Enc}_k(m), h(m)) = f(m)] \ - \ \Pr[\mathsf{A}'(1^n, h(m)) = f(m)] \leq negl(n),$$

where m is sampled from distribution M, and the probabilities are taken over the choice of m and the key k, and any random coins used by A, A′, and the encryption process.

Theorem 2 (The equivalence of semantic security and indistinguishable encryptions)**.** *A private-key encryption scheme has indistinguishable encryptions in the presence of an eavesdropper if and only if it is semantically secure in the presence of an eavesdropper.*

3 Pseudorandom generators and fixed-length indistinguishable encryption schemes

Definition 5 (pseudorandom generator)**.** Let $\ell(\cdot)$ be a polynomial and let g be a **deterministic** polynomial-time algorithm such that upon any input $s \in \{0,1\}^n$, algorithm g outputs a string of

length $\ell(n) > n$. We say that g is a pseudorandom generator (PRG) if for all PPT distinguishers D, there exists a negligible function negl such that:

$$|\Pr[\mathsf{D}(\ U_{\ell(n)}\ ) = 1]\ -\ \Pr[\mathsf{D}(\ g(U_n)\ ) = 1]|\ \leq\ \mathsf{negl}(n),$$

where the probabilities are taken over the random coins used by D and U_n (or $U_{\ell(n)}$). The difference between output and input lengths $\ell(n) - n$ is called the stretch factor of g.

Note that PRGs are not secure against adversaries who have unlimited computing power. For example, an unbounded distinguisher D on input y (which might be sampled from either $U_{\ell(n)}$ or $g(U_n)$), can output 1 iff there exists some s' such that $g(s') = y$. Therefore, D outputs 1 on $g(U_n)$ with probability 1, and in contrast the probability that D outputs 1 on $U_{\ell(n)}$ is at most $2^{-(l(n)-n)} \leq 1/2$, and thus win the above distinguishing game with advantage more than $1/2$, which is far more than negligible.

Definition 6 ($(t(n), \varepsilon(n))$-secure pseudorandom generator)**.** Let $\ell(\cdot)$ and g be as assumed in Definition 5, we say that g is a $(t(n), \varepsilon(n))$-secure PRG if for all probabilistic distinguisher of running time no more than t, it holds that

$$|\Pr[\mathsf{D}(\ U_{\ell(n)}\ ) = 1]\ -\ \Pr[\mathsf{D}(\ g(U_n)\ ) = 1]|\ \leq\ \varepsilon(n),$$

The definition is equivalent to Definition 5 for super-polynomial $t(\cdot)$ and negligible function $\varepsilon(\cdot)$. We often omit n, and write t and ε (instead of $t(n)$ and $\varepsilon(n)$) for brevity.

A good property about PRG is that once we expand the n-bit seed into $(n + s(n))$-bit output (i.e., with a stretch of $s(n)$ bits) that looks random to all efficient algorithms, we can iteratively compose g with itself to get arbitrarily long pseudorandom bits. We state it as the lemma below, whose proof technique (known as the hybrid argument) is attributed to Yao, and independently by Blum and Micali.

Lemma 3 (sequential composition of PRGs)**.** *Let*

$$\begin{aligned} g\ &: \{0,1\}^n \to \{0,1\}^{n+s(n)} \\ &s_i \mapsto (s_{i+1}, r_{i+1})\ \text{, where } s_i, s_{i+1} \in \{0,1\}^n,\ r_{i+1} \in \{0,1\}^{s(n)} \end{aligned}$$

be a $(t(n),\varepsilon(n))$-secure PRG, and for any $q(n) \in \mathbb{N}$ define $g^q : \{0,1\}^n \to \{0,1\}^{n+q(n)\cdot s(n)}$

$$\begin{aligned} g^q\ &: \{0,1\}^n \to \{0,1\}^{n+q(n)\cdot s(n)} \\ &s_0 \mapsto (s_{q(n)}, r_{q(n)}, r_{q(n)-1}, \ldots, r_1)\ , \end{aligned}$$

where for $0 \leq i \leq q(n) - 1$ iteratively compute $(s_{i+1}, r_{i+1}) := g(s_i)$. Then, we have that $g^{q(n)}$ is a $(t(n) - q(n) \cdot \mathsf{poly}(n),\ q(n) \cdot \varepsilon(n))$-secure PRG, where $\mathsf{poly}(n)$ is the running time for computing function g.

Proof. We define the following hybrid distributions:

$$\begin{aligned} H_0 &\stackrel{\mathsf{def}}{=} g^q(U_n) \\ H_1 &\stackrel{\mathsf{def}}{=} (g^{q-1}(U_n), U_s) \\ H_2 &\stackrel{\mathsf{def}}{=} (g^{q-2}(U_n), U_{2s}) \\ &\vdots \\ H_{q-1} &\stackrel{\mathsf{def}}{=} (g^1(U_n), U_{(q-1)s}) \\ H_q &\stackrel{\mathsf{def}}{=} U_{n+qs} \end{aligned}$$

For any probabilistic distinguisher D of running time $t(n)-q\cdot\mathsf{poly}(n)$, we have by triangle inequality

$$|\Pr[\mathsf{D}(H_0)=1] \;-\; \Pr[\mathsf{D}(H_q)=1]| \;\leq\; \sum_{i=0}^{q-1} |\Pr[\mathsf{D}(H_i)=1] \;-\; \Pr[\mathsf{D}(H_{i+1})=1]|$$

Now we claim that for every $0\leq i\leq q-1$ it holds that $|\Pr[\mathsf{D}(H_i)=1] \;-\; \Pr[\mathsf{D}(H_{i+1})=1]| \leq \varepsilon$, and thus the above is bounded by $q\varepsilon$. Thus, it suffices to prove this claim. Suppose towards a contradiction that there exists some D such that

$$\Pr[\mathsf{D}(H_i)=1] \;-\; \Pr[\mathsf{D}(H_{i+1})=1] \;>\; \varepsilon$$

We define another distinguisher $\mathsf{D}' \stackrel{\mathsf{def}}{=} \mathsf{D}(f(s_{i+1},r_{i+1}))$ for $g(U_n)$ from U_{n+s}, where $s_{i+1}\in\{0,1\}^n$, $r_{i+1}\in\{0,1\}^s$, and probabilistic function

$$f(s_{i+1},r_{i+1}) \;\stackrel{\mathrm{def}}{=}\; (\; g^{q-(i+1)}(s_{i+1}), r_{i+1}, U_{i\cdot s}\;)$$

It is not hard to see (by the definition of f and g^q)

$$\begin{aligned} &\Pr[\mathsf{D}'(g(U_n))=1] - \Pr[\mathsf{D}'(U_{n+s})=1] \\ =\;& \Pr[\mathsf{D}(H_i)=1] \;-\; \Pr[\mathsf{D}(H_{i+1})=1] \;>\; \varepsilon \end{aligned}$$

where the running time of D' is the sum of D (i.e., $t-q\mathsf{poly}(n)$) and f (i.e., no more than $q\mathsf{poly}(n)$), which is no more than t. We complete the proof by reaching a contradiction to the (t,ε)-security of the PRG g. □

We introduce below the fixed-length encryption scheme which is essentially the computational version of one-time pad (aka. Vernam's Cipher).

Theorem 4 (PRG-based fixed-length encryption scheme)**.** *The encryption scheme* $(\mathsf{Gen},\mathsf{Enc},\mathsf{Dec})$ *defined below is s a private-key encryption scheme that has indistinguishable encryptions in the presence of an eavesdropper.*

1. *For any security parameter n, let $\mathcal{K}=\{0,1\}^n$ and $\mathcal{M}=\mathcal{C}=\{0,1\}^{\ell(n)}$, let $g:\{0,1\}^n\to\{0,1\}^{\ell(n)}$ be a PRG, where $\ell(n)>n$.*
2. *On input 1^n, Gen outputs a string $k\in\mathcal{K}$ uniformly at random, i.e., $k\leftarrow U_n$.*
3. *Encryption Enc works by bitwise XORing (exclusive-or) every message bit with corresponding bit of the PRG output, i.e.,*
$$\mathsf{Enc}_k(m) \;=\; m\oplus g(k)\;.$$
4. *Decryption works by bitwise XORing the ciphertext with the PRG output, i.e.,*
$$\mathsf{Dec}_k(c) \;=\; c\oplus g(k)\;.$$

Proof. Assume that the PRG g is $(t(n),\varepsilon(n))$-secure, and we prove the statement with respect to Definition 3. For any probabilistic $\mathsf{A}=(\mathsf{A}_1,\mathsf{D})$ of running time $t(n)$, upon $(m_0,m_1,state)\leftarrow\mathsf{A}_1(1^n)$

we have

$$
\begin{aligned}
& \left| \Pr[\ \mathsf{D}(m_0 \oplus g(U_n), state) = 1\ ] \ - \ \Pr[\ \mathsf{D}(m_1 \oplus g(U_n), state) = 1\ ] \right| \\
\leq \ & \left| \Pr[\ \mathsf{D}(m_0 \oplus g(U_n), state) = 1\ ] \ - \ \Pr[\ \mathsf{D}(U_{\ell(n)}, state) = 1\ ] \right| \\
& + \left| \Pr[\ \mathsf{D}(m_1 \oplus g(U_n), state) = 1\ ] \ - \ \Pr[\ \mathsf{D}(U_{\ell(n)}, state) = 1\ ] \right| \\
= \ & \left| \Pr[\ \mathsf{D}(m_0 \oplus g(U_n), state) = 1\ ] \ - \ \Pr[\ \mathsf{D}(m_0 \oplus U_{\ell(n)}, state) = 1\ ] \right| \\
& + \left| \Pr[\ \mathsf{D}(m_1 \oplus g(U_n), state) = 1\ ] \ - \ \Pr[\ \mathsf{D}(m_1 \oplus U_{\ell(n)}, state) = 1\ ] \right| \\
\leq \ & 2 \cdot \varepsilon(n)
\end{aligned}
$$

where the first inequality is triangle, and the second is due to the assumption of PRG by defining distinguisher $\mathsf{D}_0(y) \stackrel{\text{def}}{=} \mathsf{D}(m_0 \oplus y)$ and $\mathsf{D}_1(y) \stackrel{\text{def}}{=} \mathsf{D}(m_1 \oplus y)$. This completes the proof as $t(\cdot)$ is super-polynomial and $\varepsilon(\cdot)$ is a negligible function. $\square$