Lecture 2: Pseudorandom Generators

MIT - 6.5620
Lecture 2 (September 14, 2026)

> ***Warning:*** This document is a rough draft and may contain bugs. Please feel free to email us with corrections.

Recap

Recall that in the last lecture we covered the following:

1. Defined the notion of a perfect secure encryption scheme. We saw two equivalent security definitions:
   - **Shannon security:** For any probability distribution M over the plaintext space $\mathcal{M}$ and every plaintext $m \in \mathcal{M}$ and ciphertext $c \in \mathcal{C}$,
   $$\Pr[M = m] = \Pr_{k \leftarrow \mathcal{K}}[M = m | \mathsf{Enc}(k, M) = c].$$
   - **Perfect indistinguishability:** For every message $m_0, m_1 \in \mathcal{M}$ and every ciphertext $c \in \mathcal{C}$
   $$\Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, m_0) = c] = \Pr_{k \leftarrow \mathcal{K}}[\mathsf{Enc}(k, m_1) = c].$$
   Equivalently, for every $m_0, m_1 \in \mathcal{M}$ and for every adversary $\mathcal{A}$,
   $$\Pr[A(\mathsf{Enc}(k, m_b)) = b] = 1/2.$$
2. Presented the One-Time Pad construction which achieves this definition. In this construction
   $$\mathcal{K} = \mathcal{M} = \mathcal{C} = \{0,1\}^n$$
   and
   $$\mathsf{Enc}(k, m) = k \oplus m \quad \text{and} \quad \mathsf{Dec}(k, c) = k \oplus c.$$
3. We argued that this scheme is only one-time secure, since given $\mathsf{Enc}(k, m_0)$ and $\mathsf{Enc}(k, m_1)$ one can learn $m_0 \oplus m_1$.

   This is a security breach since if the adv happened to know m_0 then he automatically learns m_1.[1]
4. Proved Shannon's theorem: Any scheme that achieves this definition has the drawback that $|\mathcal{K}| \geq |\mathcal{M}|$; namely, the key must grow with the number of bits encrypted!

[1] This is not only a theoretical attack; similar attacks have been executed in reality.

Overcoming Shannon's Conundrum

Recall that Shannon's theorem was proved by showing that if $|\mathcal{K}| < |\mathcal{M}|$ then there is an attack on the perfect security of the scheme, but this attack takes time roughly $|\mathcal{K}|$, which is huge![2]

[2] One possible attack: given a ciphertext ct, we computed all the possible messages it could have encrypted, i.e. $\{\mathsf{Dec}(k, \mathsf{ct})\}_{k \in \mathcal{K}}$, and checked if m_0 or m_1 was in this set. If both were in the set then we guessed randomly, but if only m_b is in the set then we guessed b. 2^{256} ops is way beyond any computing machine of today, and is designed to give a wide margin of error in security.

Key idea: Consider only computationally bounded adversaries! One option is to limit the runtime of the adversary by, say at most 2^{256} bit operations, and construct an encryption scheme where the key consists of much more than 256 bits, meaning that brute forcing the space of all keys is outside the capabilities of the attacker. Such a definition is perfectly reasonable, and in fact, this is how one thinks about security in practice.

But, to provide theoretical guarantees, it is much better to work with an *asymptotic* definition. Indeed, in complexity theory, the notion of efficiency is asymptotic (efficient $\equiv$ polynomial time). So, instead of trying to design schemes that have some concrete level of security, we design schemes whose security level is governed by a parameter known as the *security parameter*, which we will denote by λ.[3] The idea is that the security of the scheme should increase with a larger security parameter.

[3] Sometimes, the security parameter is also denoted by the letter n.

For now, we will think of the secret key as being a random λ-bit value, i.e., $k \leftarrow \{0,1\}^\lambda$. Later we will disentangle the security parameter from the length of the secret key, and also allow the key to be structured and generated via a "key generation algorithm."

We restrict our adversaries to run in probabilistic polynomial time (PPT) in the security parameter λ, following the philosophy in complexity theory that polynomial-time algorithms are "feasible" while any super-polynomial-time algorithm is "infeasible."[4] More generally, we model our adversaries as polynomial-size circuits, which are equivalent to polynomial-time Turing machines with non-uniform advice. This is more general than PPT adversaries since the non-uniformity allows us to hardwire the "best" randomness for the adversary. Often, we refer to poly-size adversaries as "efficient adversaries."

[4] Of course, this is a proxy for what is feasible vs. infeasible in reality. Polynomial-time does not necessarily mean "feasible" in reality; and super-polynomial time does not necessarily mean infeasible in reality, especially for relatively small key sizes for which the asymptotics haven't kicked in.

Our cryptographic schemes will also be restricted to run in time polynomial in the security parameter (as otherwise they will be considered infeasible to implement). We always consider adversaries that are more powerful than the honest players. For example, we may consider honest players that run in time λ^2 but require security to hold against all PPT adversaries, e.g. ones that run in time λ^5 or even λ^{100}.

Definition 1 (Take 1:)**.** A *computationally secure* encryption scheme is associated with a key space $\mathcal{K} = \{\mathcal{K}_\lambda\}_{\lambda \in \mathbb{N}}$, a message space $\mathcal{M} =$

$\{\mathcal{M}_\lambda\}_{\lambda\in\mathbb{N}}$, and a ciphertext space $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda\in\mathbb{N}}$, where there exists a polynomial $\ell : \mathbb{N} \to \mathbb{N}$ such that for every $k \in \mathcal{K}_\lambda$, $m \in \mathcal{M}_\lambda$ and $c \in \mathcal{C}_\lambda$ it holds that $|k|, |m|, |c| \leq \ell(\lambda)$.[5]

[5] Often $\mathcal{K}_\lambda = \{0,1\}^\lambda$. In addition, often $\mathcal{M}_\lambda$ does not depend on λ. For example, as we will see, for some of our schemes $\mathcal{M} = \{0,1\}$, independent of λ.

It is also associated with two polynomial time algorithms $\mathsf{Enc} = \{\mathsf{Enc}_\lambda\}_{\lambda\in\mathbb{N}}$ and $\mathsf{Dec} = \{\mathsf{Dec}_\lambda\}_{\lambda\in\mathbb{N}}$, such that:

1. $\mathsf{Enc}_\lambda : \mathcal{K}_\lambda \times \mathcal{M}_\lambda \to \mathcal{C}_\lambda$.
2. $\mathsf{Dec}_\lambda : \mathcal{K}_\lambda \times \mathcal{C}_\lambda \to \mathcal{M}_\lambda$.

It satisfies the following two properties:

1. **Correctness:** For every $\lambda \in \mathbb{N}$, every $k \in \mathcal{K}_\lambda$, and every $m \in \mathcal{M}_\lambda$
$$\mathsf{Dec}_\lambda(k, \mathsf{Enc}_\lambda(k, m)) = m.$$
2. **Computational security:** For every poly-size adversary $\mathcal{A}$, every $\lambda \in \mathbb{N}$ and every $m_0, m_1 \in \mathcal{M}_\lambda$,
$$\Pr[\mathcal{A}(\mathsf{Enc}_\lambda(k, m_b)) = b] = 1/2.$$

Remark. We often omit the security parameter from Enc_λ and Dec_λ, as we can assume without loss of generality that it can be inferred from the key k.

Remark. Notice that the security definition requires security to hold for every two messages $m_0, m_1 \in \mathcal{M}_\lambda$, even adversarially chosen. We let the adversary have control of as many things as possible to make our security definition as strong as possible.

Is this definition achievable for $|\mathcal{M}_\lambda| > |\mathcal{K}_\lambda|$? It turns out that it is not! In fact, it is subject to the same Shannon impossibility result! Namely, the adversary $\mathcal{A}$ given a ciphertext $c \in \mathcal{C}_\lambda$ will choose a random $k \leftarrow \mathcal{K}_\lambda$, and compute $m = \mathsf{Dec}(k, c)$. If there exists $b \in \{0,1\}$ such that $m = m_b$ then output b, and otherwise output a random $b' \leftarrow \{0,1\}$. This adversary will guess b correctly with probability $\geq 1/2 + 2^{-\lambda}$, as long as

$$\{\mathsf{Enc}(k, m_0) : \ k \in \{0,1\}^\lambda\} \neq \{(\mathsf{Enc}(k, m_1) : \ k \in \{0,1\}^\lambda\}$$

and by the correctness of the encryption scheme, and assuming that $|\mathcal{M}_\lambda| > |\mathcal{K}_\lambda|$, there must exist $m_0, m_1 \in \mathcal{M}_\lambda$ that satisfy the equation above.

To get around this impossibility result we just need to allow the adversary to have a tiny advantage.

Definition 2. A function $\mu : \mathbb{N} \to [0,1]$ is said to be negligible if for every polynomial $p : \mathbb{N} \to \mathbb{N}$ there exists $\Lambda \in \mathbb{N}$ such that for every $\lambda \geq \Lambda$,
$$\mu(\lambda) < \frac{1}{p(\lambda)}.$$

Intuitively, an event that occurs with negligible probability looks to a poly-time algorithm like it never occurred.

Examples

1. $\mu(\lambda) = 2^{-\lambda}$ is negligible.
2. $\mu(\lambda) = 1/\lambda^2$ is non-negligible.

Definition 3. A *computationally secure* encryption scheme is defined as above but where computational security is defined as follows:

Computational security (take 2): For every poly-size adversary $\mathcal{A}$ there exists a negligible function $\mu : \mathbb{N} \to [0,1]$ such that for every $\lambda \in \mathbb{N}$, every $m_0, m_1 \in \mathcal{M}_\lambda$,

$$\Pr[\mathcal{A}(\mathsf{Enc}(k, m_b)) = b] \leq 1/2 + \mu(\lambda).$$

Remark. We often refer to μ, which is the difference between the adversary's success probability and $1/2$, as the adversary's advantage.

Notice that the runtime and success probability are measured as a function of the security parameter. Importantly:

- Honest algorithms run in fixed polynomial time in λ.
- Adversaries may run in (arbitrary) polynomial time in λ, and should have only negligible advantage.
- Honest algorithms are uniform, whereas the adversary is allowed to be non-uniform.

Can we construct an encryption scheme that achieves this definition??[6]

Not if $\mathsf{NP} = \mathsf{P}$! Indeed, there is a non-deterministic algorithm that guesses the key and checks if the decryption of the given challenge ciphertext is m_0 or m_1. If $\mathsf{NP} = \mathsf{P}$, then this algorithm has an efficient implementation! But not to worry, it is widely believed that $\mathsf{NP} \neq \mathsf{P}$ (although proving it seems out of reach), so this attack cannot in general be efficiently implemented.

Anyway, if we construct a computationally secure encryption scheme it will immediately imply that $\mathsf{NP} \neq \mathsf{P}$, so this seems out of reach. Instead, we will construct such a scheme and prove that it is secure under some computational assumption. Namely, we will only prove *conditional* security.

One could hope to prove the security of our scheme under the assumption that $\mathsf{NP} \neq \mathsf{P}$. We do not know how to do that. But we do have schemes that are secure assuming the existence of *one-way functions*, which is known to be equivalent to the existence of a computationally secure encryption scheme.

[6] We will construct an encryption scheme that achieves an even stronger definition, where the messages $(m_{b,1}, \ldots, m_{b,\ell})$ can be chosen adaptively depending on the ciphertexts. Namely, we allow the adversary to have black-box access to $\mathsf{Enc}(k, \cdot)$ and guarantee that for every m_0, m_1 that the adversary did not query it still holds that he cannot distinguish between $\mathsf{Enc}(k, m_0)$ and $\mathsf{Enc}(k, m_1)$, even with all the ciphertexts it obtained from the oracle.

One-way functions

Intuitively, a one-way function is a function that is easy to compute, but hard to invert. Namely, you can compute it in one direction, but not the other.

Definition 4. A function $f : \{0,1\}^* \to \{0,1\}^*$ is one-way if it satisfies the following two conditions:

- **Easy to compute:** There exists a poly-time algorithm $\mathcal{B}$ such that for every $x \in \{0,1\}^*$, $\mathcal{B}(x) = f(x)$.
- **Hard to invert:** For every poly-size adversary $\mathcal{A}$ there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$

$$\Pr_{x \leftarrow \{0,1\}^\lambda}[\mathcal{A}(f(x)) = x' \text{ s.t. } f(x') = f(x)] \leq \mu(\lambda)$$

We will see one-way functions in a later lecture, but will use a related object called a pseudorandom generator to construct computationally secure encryption schemes (also called just secure encryption schemes from now on).

Our first cryptographic tool: Pseudo Random Generator (PRG)

Our goal is to encrypt messages of length n with a key of length $\lambda \ll n$. Suppose that we lived in a magical world where we could take one secret key $k \leftarrow \{0,1\}^\lambda$ and stretch it into a random key of length n! Then we could use this stretched key as a one-time pad, and hence securely encrypt a message of length n!

You may think that it is impossible to generate randomness out of thin air! But cryptography is indistinguishable from magic and makes it possible. This is one of the most beautiful gifts that cryptography bestowed upon us.

Definition of a PRG

A pseudorandom generator (PRG) is a deterministic function $G : \{0,1\}^* \to \{0,1\}^*$ that takes as input a (short) random seed $s \leftarrow \{0,1\}^\lambda$ and stretches it into a longer string $G(s) \in \{0,1\}^n$ that "looks random."

Definition 5. An efficient (poly-time computable) deterministic function $G : \{0,1\}^* \to \{0,1\}^*$ is said to be a pseudorandom generator if the following two conditions hold:

- It is expanding in the sense that there exists an expansion function $n : \mathbb{N} \to \mathbb{N}$ such that for every $\lambda \in \mathbb{N}$ it holds that $n(\lambda) > \lambda$ and G takes inputs in $\{0,1\}^\lambda$ to outputs in $\{0,1\}^{n(\lambda)}$.

We often abuse notation and denote $G : \{0,1\}^* \to \{0,1\}^*$, although G takes as input strings of arbitrary length.

- It is pseudorandom, i.e.

$$\{G(U_\lambda)\}_{\lambda\in\mathbb{N}} \approx \{U_{n(\lambda)}\}_{\lambda\in\mathbb{N}},$$

where U_ℓ is the uniform distribution over $\{0,1\}^\ell$.

An equivalent definition of the pseudorandom property is following next-bit unpredictability definition, which states that no poly-size adversary can predict the $i+1$'st bit of the output of a pseudorandom generator with probability better than $1/2 + \text{negl}(\lambda)$, where negl denotes a negligible function.

Definition 6 (Next-bit unpredictability)**.** An expanding function $G : \{0,1\}^* \to \{0,1\}^*$, with expansion $n = n(\lambda)$, is next-bit unpredictable if for every poly-size adversary $\mathcal{A}$ there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$ and every $i = i(\lambda) \in [n(\lambda)]$

$$\Pr_{r\leftarrow\{0,1\}^\lambda}[\mathcal{A}(G(r)_{[i]}) = G(r)_{i+1}] = 1/2 + \mu(\lambda)$$

where $G(r)_{[i]}$ denotes the first i bits of $G(r)$ and $G(r)_{i+1}$ denotes the $i+1$'st bit of $G(r)$.

Remark. An equivalent formulation of the above next-bit unpredictability property is that for every poly-size adversary $\mathcal{A}$ there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$ and every $i = i(\lambda) \in [n(\lambda)]$

$$|\Pr_{r\leftarrow\{0,1\}^\lambda}[\mathcal{A}(G(r)_{[i+1]}) = 1] - \Pr_{r\leftarrow\{0,1\}^\lambda, u_1\leftarrow\{0,1\}}[\mathcal{A}(G(r)_{[i]}, u_1) = 1]| \leq \mu(\lambda).$$

Theorem 7. *An expanding function $G : \{0,1\}^* \to \{0,1\}^*$, with expansion $n = n(\lambda)$, is pseudorandom if and only if it is next-bit unpredictable.*

Proof. On the next lecture. □

What is a PRG good for? It is the bread-and-butter of cryptography! It is precisely what enables the encryption of long messages using a short key!

Overcoming Shannon's Conundrum using a PRG

Consider the following encryption scheme with $\mathcal{K}_\lambda = \{0,1\}^\lambda$ and $\mathcal{M}_\lambda = \mathcal{C}_\lambda = \{0,1\}^{n(\lambda)}$ where $n(\lambda) > \lambda$. The encryption scheme uses a PRG $G : \{0,1\}^\lambda \to \{0,1\}^n$ as a building block, and is defined by

$$\mathsf{Enc}(k, m) = G(k) \oplus m$$

and

$$\mathsf{Dec}(k, c) = G(k) \oplus c.$$

Theorem 8. $(\mathsf{Enc}, \mathsf{Dec})$ *is a (computationally) secure encryption scheme.*

Proof. We prove that the two desired properties, correctness and security are satisfied.

Correctness: For every $m \in \{0,1\}^n$ and every $k \in \{0,1\}^\lambda$

$$\mathsf{Dec}(k, \mathsf{Enc}(k, m)) = G(k) \oplus (G(k) \oplus m) = m.$$

Computational security: This is our first reduction! Suppose for contradiction that there exists a poly-size adversary $\mathcal{A}$ and a non-negligible $\epsilon = \epsilon(\lambda)$ such that for every $\lambda \in \mathbb{N}$ there exist $m_0, m_1 \in \{0,1\}^n$ such that

$$\Pr[\mathcal{A}(\mathsf{Enc}(k, m_b)) = b] \geq 1/2 + \epsilon$$

Namely,

$$\Pr[\mathcal{A}(G(k) \oplus m_b) = b] \geq 1/2 + \epsilon$$

If we replace $G(k)$ with a random string then $\mathcal{A}$ would succeed in guessing only with probability $1/2$ (by the security of the one-time pad). We use this to break the security of the PRG, by constructing a poly-size adversary $\mathcal{B}$ that on input $r \in \{0,1\}^n$ distinguishes between the case that r is random or pseudorandom as follows:

1. Choose at random $b \leftarrow \{0,1\}$.
2. Compute $b' = \mathcal{A}(r \oplus m_b)$.
3. If $b' = b$ output 1 indicating that r is pseudorandom, and if $b' \neq b$ then output 0, indicating that r is random.

Notice that

$$\Pr[\mathcal{B}(G(U_\lambda)) = 1] = \Pr[\mathcal{A}(G(k) \oplus m_b) = b] \geq 1/2 + \epsilon,$$

while

$$\Pr[\mathcal{B}(U_n) = 1] = \Pr[A(r \oplus m_b) = b] = 1/2.$$

Thus,

$$|\Pr[\mathcal{B}(G(U_\lambda)) = 1] - \Pr[\mathcal{B}(U_n) = 1]| \geq \epsilon,$$

contradicting the pseudorandomness property of G. □

Do PRG*s exist?*

Theorem 9. PRG*s exist assuming the existence of one-way functions.*

We will prove an easier theorem:

Theorem 10. PRG*s exist assuming the existence of one-way permutations.*

On the next lecture we will construct a PRG assuming the existence of one-way permutations (as already hinted at the end of the class).