

Are we done? We now have a perfectly-secure cipher!

No! Keys are very long! In fact, as long as the message... [if we can share keys of this length, can use same mechanism to share the message itself]

"One-time" restriction

Malleable

Issues with the one-time pad:

- One-time: Very important. Never reuse the one-time pad to encrypt two messages. Completely broken!

Suppose $c_1 = k \oplus m_1$ and $c_2 = k \oplus m_2$

Then, $c_1 \oplus c_2 = (k \oplus m_1) \oplus (k \oplus m_2)$
 $= m_1 \oplus m_2$

← can leverage this to recover messages
← learn the xor of two messages!

One-time pad reuse:

- Project Verona (U.S. counter-intelligence operation against U.S.S.R during Cold War)

↳ Soviets reused some pages in codebook ~ led to decryption of ~ 3000 messages sent by Soviet intelligence over 37-year period [notably exposed espionage by Julius and Ethel Rosenberg]

- Microsoft Point-to-Point Tunneling (MS-PTP) in Windows 98/NT (used for VPN)

↳ Same key (in stream cipher) used for both server → client communication AND for client → server communication ↳ (RC4)

- 802.11 WEP: both client and server use same key to encrypt traffic

many problems just beyond one-time pad reuse (can even recover key after observing small number of frames!)

- Malleable: one-time pad provides no integrity; anyone can modify the ciphertext:

$$m \leftarrow k \oplus c$$

← replace c with $c \oplus m'$

$$\Rightarrow k \oplus (c \oplus m') = m \oplus m' \leftarrow \text{adversary's change now xored into original message}$$

Theorem (Shannon). If a cipher satisfies perfect secrecy, then $|K| \geq |M|$.

Intuition: Every ciphertext can decrypt to at most $|K| < |M|$ messages. This means that ciphertext leaks information about the message (not all messages equally likely). Cannot be perfectly secret.

Proof. We will use a "counting" argument. Suppose $|K| < |M|$. Take any ciphertext $c \leftarrow \text{Encrypt}(k, m)$ for some $k \in K, m \in M$. This ciphertext can only decrypt to at most $|K|$ possible messages (one for each choice of key). Since $|K| < |M|$, there is some message $m' \in M$ such that

$$\forall k \in K : \text{Decrypt}(k, c) \neq m'$$

By correctness of the cipher,

$$\forall k \in K : \text{Encrypt}(k, m') \neq c$$

This means that

$$\left. \begin{array}{l} \Pr[k \xleftarrow{R} K : \text{Encrypt}(k, m') = c] = 0 \\ \Pr[k \xleftarrow{R} K : \text{Encrypt}(k, m) = c] > 0 \end{array} \right\} \text{ Cannot be perfectly secret!}$$

Take-away: Perfect secrecy requires long keys. Very impractical (except in the most critical scenarios - exchanging daily codebooks)

If we want something efficient/usable, we need to compromise somewhere.


- Observe: Perfect secrecy is an information-theoretic (i.e., a mathematical) property

Even an infinitely-powerful (computationally-unbounded) adversary cannot break security

We will relax this property and only require

security against computationally-bounded (efficient) adversaries

Idea: "compress" the one-time pad: we will generate a long random-looking string from a short seed (e.g., $s \in \{0,1\}^{128}$).



Stream cipher: $K = \{0,1\}^\lambda$

$$M = C = \{0,1\}^n$$

$$\text{Encrypt}(k, m): c \leftarrow m \oplus G(k)$$

$$\text{Decrypt}(k, c): m \leftarrow c \oplus G(k)$$

Instead of XORing with the key, we use the key to derive a "stream" of random-looking bits and use that in place of the one-time pad

If $\lambda < n$, then this scheme cannot be perfectly secure! So we need a different notion of security

Intuitively: Want a stream cipher to function "like" a one-time pad to any "reasonable" adversary.



$\Rightarrow$ Equivalently: output of a PRG should "look" like uniformly-random string

What is a "reasonable" adversary?

- Theoretical answer: algorithm runs in (probabilistic) polynomial time (denote $\text{poly}(\lambda)$ where λ is a security parameter)
- Practical answer: runs in time $< 2^{80}$ and space $< 2^{64}$ (can use larger numbers as well)

Goal: Construct a PRG so no efficient adversary can distinguish output from random.

Captured by defining two experiments or games:



the input to the adversary (t) is often called the challenge

Adversary's goal is to distinguish between Experiment 0 (pseudorandom string) and Experiment 1 (truly random string)

$\hookrightarrow$ It is given as input a string t of length n (either $t \leftarrow G(s)$ or $t \xleftarrow{R} \{0,1\}^n$)

$\hookrightarrow$ It outputs a guess (a single bit $b \in \{0,1\}$)

Let $W_0 := \Pr[\text{adversary outputs 1 in Experiment 0}]$
 Let $W_1 := \Pr[\text{adversary outputs 1 in Experiment 1}]$
 define the distinguishing advantage of A as $\text{PRGAdv}[A, G] := |W_0 - W_1|$

Remember: adversary knows the algorithm G ; only seed is hidden!

Do NOT RELY ON SECURITY BY OBSCURITY!

Definition. A PRG $G: \{0,1\}^\lambda \rightarrow \{0,1\}^n$ is secure if for all efficient adversaries A ,

$$\text{PRGAdv}[A, G] = \text{negl}(\lambda)$$

$\hookrightarrow$ negligible function (in the input length)

- Theoretical definition: $f(\lambda)$ is negligible if $f \in o(\lambda^{-c})$ for all $c \in \mathbb{N}$

- Practical "definition": quantity $\leq 2^{-80}$ or $\leq 2^{-128}$

smaller than any inverse polynomial

e.g., $\frac{1}{2^\lambda}$, $\lambda^{-\log \lambda}$

Understanding the definition:

1. Can we ask for security against all adversaries (when $n \gg \lambda$)?

No! Consider inefficient adversary that outputs 1 if t is the image of G and 0 otherwise.

$$- W_0 = 1$$

$$- W_1 = \Pr[t \in \{0,1\}^n : \exists s \in \{0,1\}^\lambda : G(s) = t] = \frac{1}{2^{n-\lambda}}$$

$$\left. \begin{array}{l} - W_0 = 1 \\ - W_1 = \frac{1}{2^{n-\lambda}} \end{array} \right\} \text{PRGAdv}[A, G] = 1 - \frac{1}{2^{n-\lambda}} \approx 1 \text{ if } n \gg \lambda$$

2. Can the output of a PRG be biased (e.g., first bit of PRG output is 1 w.p. $\frac{2}{3}$)?

No! Consider efficient adversary that outputs 1 if first bit of challenge is 1.

$$- W_0 = \frac{2}{3}$$

$$- W_1 = \frac{1}{2}$$

$$\left. \begin{array}{l} - W_0 = \frac{2}{3} \\ - W_1 = \frac{1}{2} \end{array} \right\} \text{PRGAdv}[A, G] = \frac{1}{6} \text{ Not NEGLIGIBLE!}$$

More generally, no efficient statistical test can distinguish output of a secure PRG from random.

3. Can the output of a PRG be predictable (e.g., given first 10 bits, predict the 11th bit)?

No! If the bits are predictable w.p. $\frac{1}{2} + \epsilon$, can distinguish with advantage ϵ (since random string is unpredictable)

In fact: unpredictable $\Rightarrow$ pseudorandom


Take-away: A secure PRG has the same statistical properties as the one-time pad to any efficient adversary.

$\Rightarrow$ Should be able to use it in place of one-time pad to obtain a secure encryption scheme (against efficient adversaries)

Need to define security of an encryption scheme.

Goal is to capture property that no efficient adversary can learn any information about the message given only the ciphertext. Suffices to argue that no efficient adversary can distinguish encryption of message m_0 from m_1 , even if m_0, m_1 are adversarially-chosen.

Let $(\text{Encrypt}, \text{Decrypt})$ be a cipher. We define two experiments (parameterized by $b \in \{0,1\}$):



Adversary chooses two messages and receives encryption of one of them. Needs to guess which one (i.e., distinguish encryption of m_0 from encryption of m_1)

Let $W_0 := \Pr[b' = 1 \mid b = 0]$ } probability that adversary guesses 1
 $W_1 := \Pr[b' = 1 \mid b = 1]$ } (if adversary is good distinguisher, these two should be very different)

Define semantic security advantage of adversary A for cipher $\Pi_{SE} = (\text{Encrypt}, \text{Decrypt})$

$$\text{SSAdv}[A, \Pi_{SE}] = |W_0 - W_1|$$

Definition. A cipher $\Pi_{SE} = (\text{Encrypt}, \text{Decrypt})$ is semantically secure if for all efficient adversaries A ,

$$\text{SSAdv}[A, \Pi_{SE}] = \text{negl}(\lambda)$$

$\leftarrow \lambda$ is a security parameter (here, models the bit-length of the key)