


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

Encrypt (k, m) : $c \leftarrow m \oplus G(k)$ 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
 Decrypt (k, c) : $m \leftarrow c \oplus G(k)$

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
- 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 $t \leftarrow \{0,1\}^n$)

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

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

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

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 ,

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

probabilistic polynomial time

\hookrightarrow negligible function (in the input length)

smaller than any inverse polynomial

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

- 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 $\leq 2^{-128}$

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.

- $W_0 = 1$

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

} $\text{PRGAdv}[A, G] = 1 - \frac{1}{2^{n-\lambda}} \approx 1$ 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.

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

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

} $\text{PRGAdv}[A, G] = \frac{1}{6}$ 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)

Exercising the definition: we will now consider an example of proving security of a PRG

Theorem. Suppose $G: \{0,1\}^\lambda \rightarrow \{0,1\}^n$ is a secure PRG. Then, the function $G'(s) := G(s) \oplus 1^n$ is also a secure PRG.


\leftarrow This is a conditional statement! We will discuss this more in the coming lectures.

Proof. To prove this directly seems difficult: must show non-existence of an adversary.

Instead, we consider the contrapositive:

"If G' is not a secure PRG, then G is not a secure PRG"

Suppose G' is not secure. Namely, there exists an efficient adversary A that breaks security of G' with non-negligible advantage ϵ . We use A to construct a new adversary B that breaks security of G :



In Exp_0 , algorithm B invokes algorithm A on the string $G(s) \oplus 1^n$ where $s \xleftarrow{R} \{0,1\}^n$ is random. This is precisely the distribution of Exp_0 for A . Thus,

$$W_0 = \Pr [B \text{ outputs } 1 \text{ in } \text{Exp}_0] = \Pr [A \text{ outputs } 1 \text{ in } \text{Exp}_0]$$

In Exp_1 , algorithm B invokes algorithm A on the string $t \oplus 1^n$ where $t \xleftarrow{R} \{0,1\}^n$ is uniformly random. The distribution of $t \oplus 1^n$ is still uniform:

$$\begin{aligned} \forall u \in \{0,1\}^n: \Pr [t \xleftarrow{R} \{0,1\}^n : t \oplus 1^n = u] \\ = \Pr [t \xleftarrow{R} \{0,1\}^n : t = u \oplus 1^n] = \frac{1}{2^n} \end{aligned}$$

This means

$$W_1 = \Pr [B \text{ outputs } 1 \text{ in } \text{Exp}_1] = \Pr [A \text{ outputs } 1 \text{ in } \text{Exp}_1]$$

We conclude then that

$$\begin{aligned} \text{PRGAdv}[B, G] &= |W_0 - W_1| \\ &= |\Pr [A \text{ outputs } 1 \text{ in } \text{Exp}_0] - \Pr [A \text{ outputs } 1 \text{ in } \text{Exp}_1]| \\ &= \epsilon, \end{aligned}$$

which is non-negligible by assumption. This proves the contrapositive.

The above proof is an example of a security reduction. We show how to reduce the task of breaking G to that of breaking G' . This means an attack on G' implies an attack on G . Correspondingly, if G is secure (i.e., no efficient attacks succeed with non-negligible probability), then the same holds for G' .

Refer to the posted notes on the course website as well as the textbook for more examples. We will see more reductions throughout the course as well.