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.

- $W_0 = 1$
- $W_1 = \Pr[t \xleftarrow{R} \{0,1\}^n : \exists s \in \{0,1\}^\lambda : G(s) = t] = \frac{1}{2^{n-\lambda}}$

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

- $W_0 = \frac{2}{3}$
- $W_1 = \frac{1}{2}$

$\Big\}$ $\mathrm{PRGAdv}[A,G] = \frac{1}{6}$ 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 11^{th} bit)?

No! If the bits are predictable w.p. $\frac{1}{2} + \varepsilon$, 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 $G: \{0,1\}^\lambda \to \{0,1\}^n$ is a secure PRG. Then, the function $G'(s) := G(s) \oplus 1^n$ is also a secure PRG.

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:

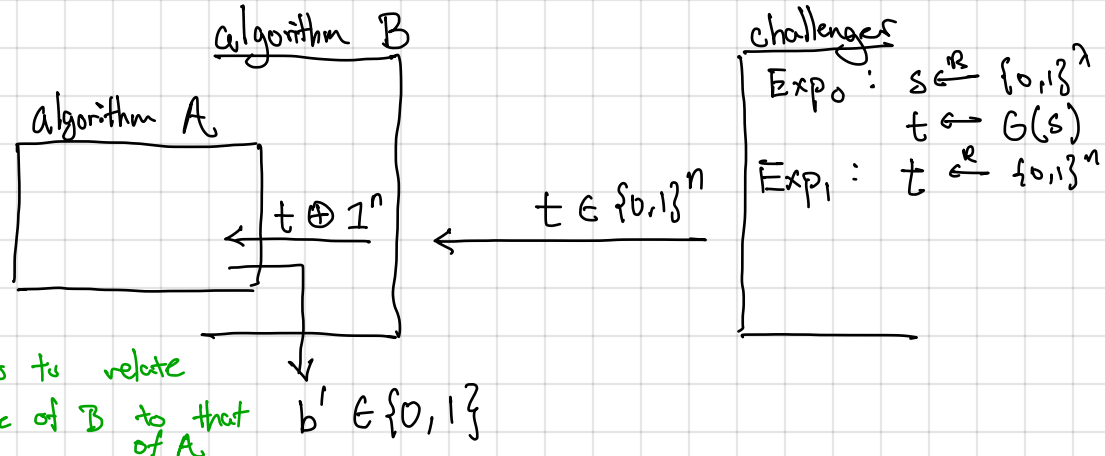

algorithm B basically runs A under the hood

→ our goal is to relate the advantage of B to that of A

In Exp_0, algorithm B invokes algorithm A on the string $G(s) \oplus 1^n$ where $s \xleftarrow{R} \{0,1\}^\lambda$ 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:

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

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

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

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

The above proof is an example of a security <u>reduction</u>. 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.

Now we will return to the notion of a secure 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 (Encrypt, Decrypt) be a cipher. We define two experiments (parameterized by $b \in \{0,1\}$):

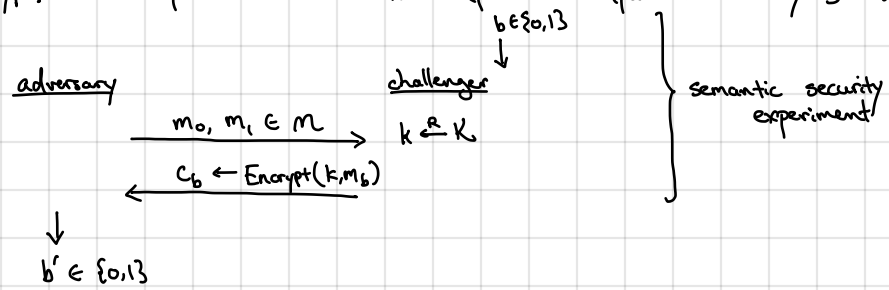

Adversary chooses two messages and receives encryption of one of them. Needs to guess which one (i.e., distinguish encryption of m_0 from encryption of m_1)

Let $W_0 := \Pr[b' = 1 \mid b = 0]$
$W_1 := \Pr[b' = 1 \mid b = 1]$
} probability that adversary guesses 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,

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

↑ λ is a security parameter (here, models the bit-length of the key)