Understanding the definition:

Can we learn the least significant bit of a message given only the ciphertext (assuming a semantically-secure cipher)

No! Suppose we could. Then, adversary can choose two messages m_0, m_1 that differ in their least significant bit and distinguish with probability 1.

This generalizes to any efficiently-computable property of the two messages.

How does semantic security relate to perfect secrecy?

Theorem. If a cipher satisfies perfect secrecy, then it is semantically secure.

Proof. Perfect secrecy means that $\forall m_0, m_1 \in \mathcal{M}$, $c \in C$:

$$\Pr[k \xleftarrow{R} K : \text{Encrypt}(k, m_0) = c] = \Pr[k \xleftarrow{R} K : \text{Encrypt}(k, m_1) = c]$$

Equivalently, the distributions

$$\underbrace{\{k \xleftarrow{R} K : \text{Encrypt}(k, m_0)\}}_{D_0} \quad \text{and} \quad \underbrace{\{k \xleftarrow{R} K : \text{Encrypt}(k, m_1)\}}_{D_1}$$

are identical ($D_0 \equiv D_1$). This means that the adversary's output b' is identically distributed in the two experiments, and so

$$\text{SSAdv}[A, \Pi_{SE}] = |W_0 - W_1| = 0.$$

Corollary. The one-time pad is semantically secure.

$$c \leftarrow G(s) \oplus m$$
$$m \leftarrow G(s) \oplus c$$

encryption key (PRG seed)

seems straightforward, but takes some care to prove

Theorem. Let G be a secure PRG. Then, the resulting stream cipher constructed from G is semantically secure.

Proof. Consider the semantic security experiments:

Experiment 0: Adversary chooses m_0, m_1 and receives $c_0 = G(s) \oplus m_0$

Experiment 1: Adversary chooses m_0, m_1 and receives $c_1 = G(s) \oplus m_1$

} Want to show that adversary's output in these two experiments are indistinguishable

Let $W_0 = \Pr[A \text{ outputs } 1 \text{ in Experiment } 0]$

$W_1 = \Pr[A \text{ outputs } 1 \text{ in Experiment } 1]$

Goal: Show that if G is a secure PRG, then for all efficient adversaries A, $|W_0 - W_1| = \text{negl}(\lambda)$.

Idea: If $G(s)$ is uniform random string (i.e., one-time pad), then $W_0 = W_1$. But $G(s)$ is like a one-time pad!

Define Experiment 0': Adversary chooses m_0, m_1 and receives $c_0 = t \oplus m_0$ where $t \xleftarrow{R} \{0,1\}^n$

Experiment 1': Adversary chooses m_0, m_1 and receives $c_1 = t \oplus m_1$ where $t \xleftarrow{R} \{0,1\}^n$

} called "hybrid experiments"

Define W_0', W_1' accordingly.

Now we can write

$$|W_0 - W_1| = |W_0 - W_0' + W_0' - W_1' + W_1' - W_1|$$
$$\leq |W_0 - W_0'| + \underbrace{|W_0' - W_1'|}_{} + |W_1' - W_1| \quad \text{by triangle inequality}$$

$W_0' = W_1'$ (for all adversaries A)

since OTP satisfies perfect secrecy

Suffices to show that for all efficient adversaries, $|W_0 - W_0'| = \text{negl}(\lambda)$ and $|W_1 - W_1'| = \text{negl}(\lambda)$.

Show. If G is a secure PRG, then for all efficient A, $|W_0 - W_0'| = \text{negl}$.

Common proof technique: prove the contrapositive.

Contrapositive: If A can distinguish Experiments 0 and 0', then G is not a secure PRG.

Suppose there exists efficient A that distinguishes Experiment 0 from 0'

$\Rightarrow$ We use A to construct efficient adversary B that breaks security of G.

↳ this step is a reduction

[we show how adversary (i.e., algorithm) for distinguishing Exp. 0 and 0' $\Rightarrow$ adversary for PRG]

Algorithm B (PRG adversary):

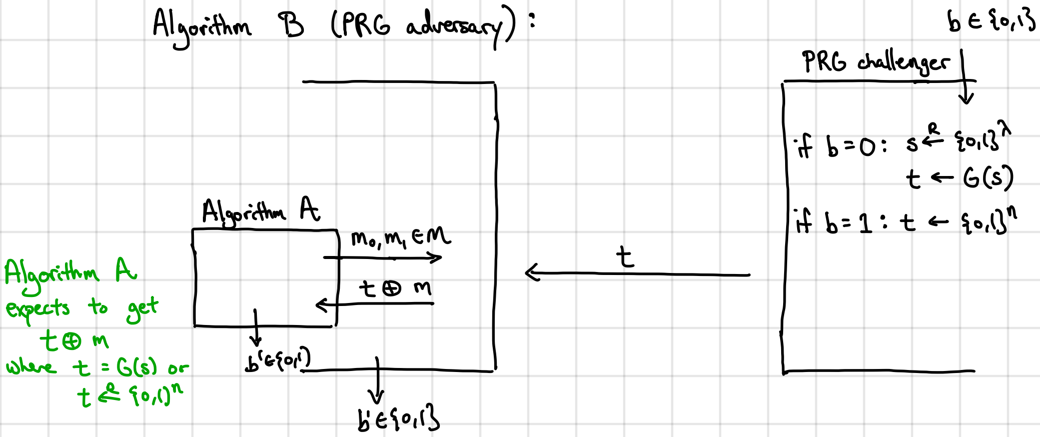

Running time of B = running time of A = efficient

Compute $\text{PRGAdv}[B, G]$.

$\Pr[B \text{ outputs } 1 \text{ if } b=0] = W_0$ ← if $b=0$, then A gets $G(s) \oplus m$ which is precisely the behavior in Exp. 0

$\Pr[B \text{ outputs } 1 \text{ if } b=1] = W_0'$ ← if $b=1$, then A gets $t \oplus m$ which is precisely the behavior in Exp. 0'

$\Rightarrow \text{PRGAdv}[B, G] = |W_0 - W_0'|$, which is non-negligible by assumption. This proves the contrapositive.

Important note: Security of above schemes shown assuming message space is $\{0,1\}^n$ (i.e., all messages are n-bits long)

In practice: We have variable-length messages. In this case, security guarantees indistinguishability from other messages of the same length, but length itself is leaked [inevitable if we want short ciphertexts]

↳ can be problematic — see traffic analysis attacks!

So far, we have shown that if we have a PRG, then we can encrypt messages efficiently (stream cipher).

Do PRGs exist? We don't know! More difficult problem than resolving P vs. NP!

However, it is not hard to see that if PRGs exist, then $P \neq NP$. [Try proving this yourself]

↳ What we can say is that if "one-way functions" (OWF) exist, then there exists a PRG that stretches the seed by 1 bit (e.g., λ-bit seed $\rightarrow$ $(\lambda+1)$-bit string)

function that is "easy" to compute but "hard" to invert

↳ will define more formally later in the course

a PRG is an example of such a function

given $s \in \{0,1\}^\lambda$, evaluating $G(s) \in \{0,1\}^n$ is easy

given $G(s) \in \{0,1\}^n$ for random $s \in \{0,1\}^\lambda$, computing s is hard (why?)

But what if we want PRGs with longer stretch? For example, can we build PRGs with stretch $\ell(\lambda) = \text{poly}(\lambda)$ for arbitrary polynomials?

Blum-Micali PRG: Suppose $G: \{0,1\}^\lambda \rightarrow \{0,1\}^{\lambda+1}$ is a secure PRG. We build a PRG with stretch $\ell(\lambda) = \text{poly}(\lambda)$ as follows:

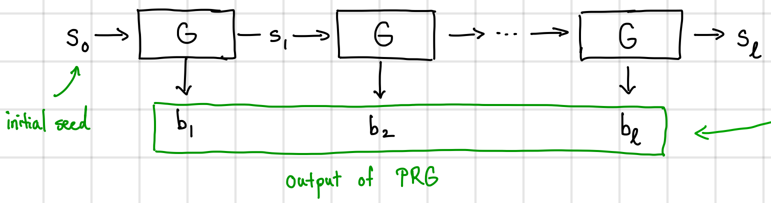

denote the Blum-Micali construction $G^{(\ell)}: \{0,1\}^\lambda \rightarrow \{0,1\}^{\ell(\lambda)}$

Why is this constructing a secure PRG?

↳ Intuitively, if s_0 is uniformly random, then $G(s_0) = (b_1, s_1)$ is uniformly random so we can feed s_1 into the PRG and take b_1 as the first output bit of the PRG $\Rightarrow$ iterate until we have ℓ output bits

Theorem. If $G: \{0,1\}^\lambda \rightarrow \{0,1\}^{\lambda+1}$ is a secure PRG, then the Blum-Micali generator $G^{(\ell)}: \{0,1\}^\lambda \rightarrow \{0,1\}^{\ell(\lambda)}$ is also a secure PRG for all $\ell = \text{poly}(\lambda)$.

Proof. Consider the following experiments:

Experiment H_0: Sample $s_0 \xleftarrow{R} \{0,1\}^\lambda$ and adversary is given $G^{(\ell)}(s_0)$

Experiment H_1: Sample $t \xleftarrow{R} \{0,1\}^{\ell(\lambda)}$ and adversary is given t

For an adversary A, define

$W_0 := \Pr[A \text{ outputs } 1 \text{ in } H_0]$

$W_1 := \Pr[A \text{ outputs } 1 \text{ in } H_1]$

Goal: Show that if G is secure, then for all efficient adversaries A, $|W_0 - W_1| = \text{negl}(\lambda)$.

We will use a "hybrid" argument. Specifically, we first define a sequence of intermediate experiments, where each adjacent pair of experiments is easy to reason about (i.e., directly reduces to security of G)

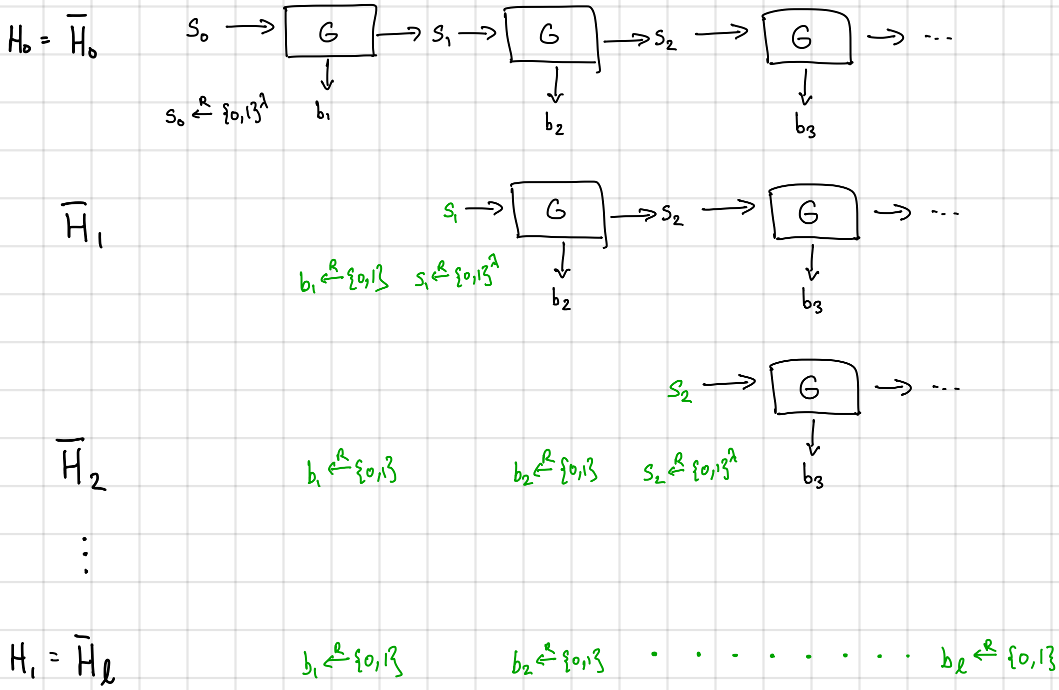

<u>Basic idea</u>: in experiment $\bar{H}_i$ the first i bits of output are generated <u>uniformly at random</u> while the remaining bits are generated using the Blum-Micali generator

In each experiment, adversary is given the sequence of bits $b_1 b_2 \cdots b_\ell$

Let A be an efficient distinguisher. Define $\bar{W}_i := \Pr[\text{A outputs 1 in experiment } \bar{H}_i]$

Then, $\text{PRGAdv}[A,G] = |W_0 - W_1|$

$$= |\bar{W}_0 - \bar{W}_\ell| \quad \text{(by definition)}$$
$$= |\bar{W}_0 - \bar{W}_1 + \bar{W}_1 - \bar{W}_2 + \cdots + \bar{W}_{\ell-1} - \bar{W}_\ell|$$
$$\leq |\bar{W}_0 - \bar{W}_1| + |\bar{W}_1 - \bar{W}_2| + \cdots + |\bar{W}_{\ell-1} - \bar{W}_\ell| \quad \text{(by triangle equality)}$$

<u>Claim</u>. If G is a secure PRG, then for all efficient adversaries A, $|\bar{W}_i - \bar{W}_{i+1}| = \text{negl}(\lambda)$.

<u>Proof</u>. We will show the <u>contrapositive</u>: if A can distinguish experiments $\bar{H}_i$ and $\bar{H}_{i+1}$, then A can break pseudorandomness of G.

Suppose $|\bar{W}_i - \bar{W}_{i+1}| = \varepsilon$. We use A to build a distinguisher B for G. Algorithm B works as follows:

1. On input a string $z \in \{0,1\}^{\lambda+1}$, algorithm B parses z as (b_{i+1}, s_{i+1}) where $b_{i+1} \in \{0,1\}$ and $s_{i+1} \in \{0,1\}^\lambda$
2. Sample $b_1, \ldots, b_i \xleftarrow{R} \{0,1\}$.
3. Compute $b_{i+2}, \ldots, b_\ell$ using Blum-Micali with seed s_{i+1}. Give $b_1 \cdots b_\ell$ to A and output whatever A outputs.

In pictures:

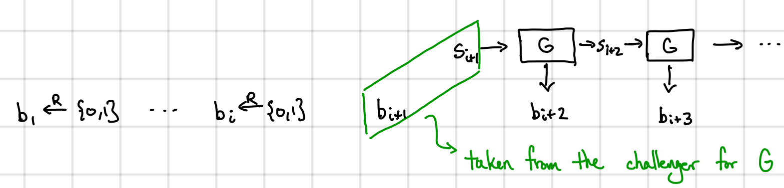

Two possibilities: 1. Suppose $z = G(s_i)$ for some $s_i \xleftarrow{R} \{0,1\}^\lambda$. Then, above picture looks like this:

$$
\begin{tikzcd}
 & s_i \xleftarrow{R} \{0,1\}^\lambda & & & \\
b_1 \xleftarrow{R} \{0,1\} \quad \cdots \quad b_i \xleftarrow{R} \{0,1\} & s_i \arrow[r] & \boxed{G} \arrow[r] \arrow[d] & s_{i+1} \arrow[r] & \boxed{G} \arrow[r, "s_{i+2}"] \arrow[d] & \boxed{G} \arrow[r] \arrow[d] & \cdots \\
 & & b_{i+1} & & b_{i+2} & b_{i+3} &
\end{tikzcd}
$$

In this case, $b_1, \dots, b_\ell$ is distributed exactly as in experiment H_i and so A outputs 1 with prob. $\overline{W}_i$

2. Suppose $z \xleftarrow{R} \{0,1\}^{\lambda+1}$. Then above picture looks like this:

$$
\begin{tikzcd}
 & & s_{i+1} \xleftarrow{R} \{0,1\}^\lambda & & \\
b_1 \xleftarrow{R} \{0,1\} \quad \cdots \quad b_i \xleftarrow{R} \{0,1\} & b_{i+1} \xleftarrow{R} \{0,1\} & s_{i+1} \arrow[r] & \boxed{G} \arrow[r, "s_{i+2}"] \arrow[d] & \boxed{G} \arrow[r] \arrow[d] & \cdots \\
 & & & b_{i+2} & b_{i+3} &
\end{tikzcd}
$$

In this case, $b_1, \dots, b_\ell$ is distributed exactly as in experiment H_{i+1} and so A outputs 1 with prob. $\overline{W}_{i+1}$

Very important to argue that B "simulates" the correct view for A. Otherwise, behavior of A is unknown!

Thus, $\text{PRGAdv}[B,G] = |\overline{W}_i - \overline{W}_{i+1}|$
$= \varepsilon$ since B outputs whatever A outputs

Since B is efficient (assuming A is efficient), by security of G, $\text{PRGAdv}[B,G] = \text{negl}(\lambda)$. Thus, $\varepsilon = |p_i - p_{i+1}| = \text{negl}(\lambda)$, and the claim follows. ∎

To complete the proof of the main theorem, we have that

$$
\begin{aligned}
|\overline{W}_0 - \overline{W}_\ell| &\le |\overline{W}_0 - \overline{W}_1| + \cdots + |\overline{W}_{\ell-1} - \overline{W}_\ell| \\
&\le \ell \cdot \text{negl}(\lambda) \\
&= \text{negl}(\lambda) \quad \text{since } \ell = \text{poly}(\lambda).
\end{aligned}
$$

<u>Proof strategy recap</u>: 1. Hybrid arguments: to argue indistinguishability of a pair of distributions, begin by identifying a <u>simple</u> set of intermediate distributions, and argue that each pair of adjacent distributions is indistinguishable

2. Security reduction (proof by contrapositive): To show a statement of the form "If X is secure, then Y is secure," show instead the the statement "If Y is not secure, then X is not secure." In the proof, show that if there exists an adversary for Y (i.e. Y is not secure), then there exists an adversary for X.

↳ When constructing this adversary, it is important to show that it simulates the <u>correct distribution</u> of inputs to the underlying adversary (i.e., this is essentially showing correctness of the <u>reduction algorithm</u>)

Stream ciphers in practice:

- RC4 (1987) stream cipher (widely used – SSL/TLS protocol, 802.11b)

$$
\begin{tikzcd}
\boxed{\text{128-bits}} \arrow[d] & \text{initial PRG seed} \\
\boxed{\text{2048-bit internal state}} \arrow[d, "G"] & \\
\text{1-byte per round} &
\end{tikzcd}
$$

Numerous problems:
- Bias in initial output: $\Pr[\text{second byte} = 0] = \frac{2}{256} > \frac{1}{256}$
  - ↳ When using RC4, recommendation is to ignore first 256 bytes due to potential bias
  - ↳ Correlations in output: probability of seeing $(0,0)$ in output is $\frac{1}{256^2} + \frac{1}{256^3} > \frac{1}{256^2}$
  - ↳ Given outputs of RC4 with related keys (e.g., keys sharing common suffix), possible to recover keys after seeing few blocks of output
    - ↳ Can be very problematic on weak devices (who may not have good sources of entropy)