COS 533: Advanced Cryptography | Princeton University
Lecture 3 (Sept. 20)
Lecturer: Mark Zhandry | Scribe: Jessica Shi

Notes for Lecture 3

Last time, we built the following the following:

$$\text{PRG} \Rightarrow \text{PRF} \Rightarrow \text{CPA-sk}.$$

This lecture, we will build:

OWP (one-way permutation)[1] $\Rightarrow$ PRG (using the Goldreich-Levin Theorem).

We now have the following relevant definitions:

Definition 1. A *pseudorandom generator* (PRG) is a function $G : \{0,1\}^\lambda \to \{0,1\}^{\lambda+s(\lambda)}$ such that $s(\lambda) \geq 1$, G is deterministic and in polytime, and for all PPT A, there exists negligible ε such that

$$|Pr_{x \xleftarrow{\$} \{0,1\}^\lambda}[1 \leftarrow A(G(x))] - Pr_{y \xleftarrow{\$} \{0,1\}^{\lambda+s(\lambda)}}[1 \leftarrow A(y)]| < \varepsilon(\lambda).$$

Definition 2. A *one-way function* (OWF) is a function $f : \{0,1\}^\lambda \to \{0,1\}^{n(\lambda)}$ that is deterministic and in polytime, such that for all PPT A, there exists negligible ε such that

$$Pr_{x \xleftarrow{\$} \{0,1\}^\lambda}[f(A(f(x))) = f(x)] < \varepsilon(\lambda).$$ [2]

Definition 3. A *one-way permutation* (OWP) is a OWF f such that $n(\lambda) = \lambda$ and f is a bijection.

Note that any PRG is a OWF. We don't prove this claim in full, but this can essentially be done by proving the contrapositive; if f is not a OWF, then there exists an inverter A (correct with non-negligible probability), and we can construct a distinguisher A' for f as a PRG by feeding any query through the inverter and seeing if it returns a valid input.

We need one further construction to discuss OWP $\Rightarrow$ PRG.

[1]This can be improved to be a OWF (one-way function), but this is beyond the scope of this course.

[2]Note that we test for this condition instead of simply $A(f(x)) = x$, because in the cases of many-to-one functions, this makes the notion of security interesting.

Definition 4. Let f be a OWF. $h : \{0,1\}^\lambda \to \{0,1\}$ is a *hardcore bit* (HC bit) for f if f is deterministic and in polytime, and for all PPT A, there exists negligible ε such that

$$Pr_{x \xleftarrow{\$} \{0,1\}^\lambda}[A(f(x)) = h(x)] \leq 1/2 + \varepsilon(\lambda).$$

Equivalently, we have

$$|Pr_{x \xleftarrow{\$} \{0,1\}^\lambda}[1 \leftarrow A(f(x), h(x))] - Pr_{x \xleftarrow{\$} \{0,1\}^\lambda, b \xleftarrow{\$} \{0,1\}}[1 \leftarrow A(f(x), b)]| < \varepsilon'(\lambda)$$

where ε' is negligible.

Now we prove a simpler result, namely that a OWP with a hardcore bit can generate a PRG.

Theorem 5. OWP with a HC bit $\Rightarrow$ PRG.

Proof. Let $f : \{0,1\}^\lambda \to \{0,1\}^\lambda$ be a OWP, and let $h : \{0,1\}^\lambda \to \{0,1\}$ be its HC bit. We claim that $G : \{0,1\}^\lambda \to \{0,1\}^{\lambda+1}$ where $G(x) = (f(x), h(x))$ is a PRG.

Assume that G is not a PRG. Then, there exists a PPT A such that

$$|Pr_{x \xleftarrow{\$} \{0,1\}^\lambda}[1 \leftarrow A(G(x))] - Pr_{y \xleftarrow{\$} \{0,1\}^{\lambda+s(\lambda)}}[1 \leftarrow A(y)]| \geq \varepsilon(\lambda)$$

where ε is non-negligible. Note that

$$Pr_{y \xleftarrow{\$} \{0,1\}^{\lambda+s(\lambda)}}[1 \leftarrow A(y)] = Pr_{x \xleftarrow{\$} \{0,1\}^\lambda, b \xleftarrow{\$} \{0,1\}}[1 \leftarrow A(f(x), b)]$$

since f is a permutation, so we have

$$|Pr_{x \xleftarrow{\$} \{0,1\}^\lambda}[1 \leftarrow A(f(x), h(x))] - Pr_{x \xleftarrow{\$} \{0,1\}^\lambda, b \xleftarrow{\$} \{0,1\}}[1 \leftarrow A(f(x), b)]| \geq \varepsilon(\lambda).$$

This directly contradicts the definition of a HC bit. Thus, G must be a PRG. □

We provide the following example without proof:

Example 6. For a prime p, let g be a generator for $\mathbb{Z}_p^*$ and let $f : \mathbb{Z}_p \to \mathbb{Z}_p$ such that $f(x) = g^x \mod p$[3]. Let

$$h(x) = \begin{cases} 1 \text{ if } x \leq p/2 \\ 0 \text{ if } x > p/2 \end{cases}.$$

We claim that f is a OWP under the discrete logarithm assumption, and we can use this to show that h is a HC bit for f. The concatenation gives us a PRG, as proven above.

[3]Note that this isn't quite a permutation because 0 is not a valid output, but it is close enough to one that this is irrelevant.

Now, we discuss the Goldreich-Levin Theorem. Informally, this theorem claims that all OWFs have a HC bit. More formally, we have the following.

Theorem 7 (Goldreich-Levin)**.** Let f be a OWF. Define $f'(x,r) = (f(x), r)$ where $x, r \overset{\{}{\leftarrow} 0,1\}^\lambda$ (note that f' is also a OWF). Let $h(x,r) = \langle x, r\rangle := \sum_i x_i r_i \mod 2$ (where $x = (x_1, \ldots, x_\lambda)$ and $r = (r_1, \ldots, r_\lambda)$). Then, h is a HC bit for f'.[4]

Proof. Assume that there exists PPT A and non-negligible ε such that

$$Pr_{r,x \overset{\$}{\leftarrow} \{0,1\}^\lambda}[A(f(x), r) = \langle x, r\rangle] \geq 1/2 + \varepsilon(\lambda),$$

that is to say, h is not a HC bit for f').

We prove this theorem in several stages.

Step 1. (Easy) Suppose $\varepsilon(\lambda) = 1/2$, so A guesses h with probability 1.

Let $e_i = 0^{i-1}10^{\lambda-i}$. Then, note that $A(f(x), e_i) = \langle x, e_i\rangle = x_i$. Thus, we can construct an attacker A' that takes as input $f(x)$ and simply applies $A(f(x), e_i)$ for all i, which gives us the value of x. A' then outputs this x, so we have $Pr_{x \overset{\$}{\leftarrow} \{0,1\}^\lambda}[f(A'(f(x))) = f(x)] = 1$, which contradicts the fact that f is a OWF.

We have shown the result for if $\varepsilon(\lambda) = 1/2$.

Step 2. (Medium) Suppose $\varepsilon(\lambda) = 1/4 + \gamma(\lambda)$ where γ is non-negligible.

Note that we cannot simply query with e_i as in the previous step, because these e_i are not chosen at random, so we have no guarantees on A returning a correct output with e_i. We claim the following (without proof):

$$Pr_{x \overset{\$}{\leftarrow} \{0,1\}^\lambda}[Pr_{r \overset{\$}{\leftarrow} \{0,1\}^\lambda}[A(f(x), r) = \langle x, r\rangle] \geq \frac{3}{4} + \frac{\gamma}{2}] \geq \frac{\gamma}{2}.$$

The proof of this statement follows from the Markov inequality.

We also call any given $x \leftarrow \{0,1\}^\lambda$ *good* if the inner condition is true, that is to say, if

$$Pr_{r \overset{\$}{\leftarrow} \{0,1\}^\lambda}[A(f(x), r) = \langle x, r\rangle] \geq \frac{3}{4} + \frac{\gamma}{2}.$$

By our earlier claim, x is good with non-negligible probability, so we can from here fix some good x.

Define $H(r) = A(f(x), r)$, so we have $Pr_{r \overset{\$}{\leftarrow} \{0,1\}^\lambda}[H(r) = \langle x, r\rangle] \geq 3/4 + \gamma/2$.

For each $i \in [\lambda]$, we choose a random $r \overset{\$}{\leftarrow} \{0,1\}^\lambda$ and apply $H(r) \oplus H(r \oplus e_i)$.

[4] Note that this implies that from any OWP, we can build the OWP f', use this theorem to receive a HC bit h, and following previous constructions, build a PRG, as desired.

Note that if H succeeds on both inputs, this gives us $H(r) \oplus H(r \oplus e_i) = \langle x, r \rangle \oplus \langle x, r \oplus e_i \rangle = \langle x, e_i \rangle$, which is exactly what we want.

Now, the probability that H succeeds on both inputs is given as follows:

$$\begin{aligned} Pr[H(r) \oplus H(r \oplus e_i) = \langle x, e_i \rangle] &\geq Pr[H(r) = \langle x, r \rangle \wedge H(r \oplus e_i) = \langle x, r \oplus e_i \rangle] \\ &= 1 - Pr[H(r) \neq \langle x, r \rangle \vee H(r \oplus e_i) \neq \langle x, r \oplus e_i \rangle] \\ &\geq 1 - Pr[H(r) \neq \langle x, r \rangle] - Pr[H(r \oplus e_i) \neq \langle x, r \oplus e_i \rangle] \\ &\geq 1 - (\frac{1}{4} - \frac{\gamma}{2}) - (\frac{1}{4} - \frac{\gamma}{2}) \\ &\geq \frac{1}{2} + \gamma. \end{aligned}$$

Thus, we can repeatedly choose r for each i, and take the value that $H(r) \oplus H(r \oplus e_i)$ outputs most often to be x_i; the number of repetitions can be such that we have a large amount of confidence in each x_i, and the probability that an attacker A' will guess x given $f(x)$ will be relatively high/non-negligible, giving us the desired contradiction.

Thus, we have shown the result in the case when $\varepsilon(\lambda) = 1/4 + \gamma(\lambda)$.

Step 3. (Hard) Suppose ε is any non-negligible function.

Note that the previous step fails in this case if $\varepsilon(\lambda) < 1/4$, because the last probability we calculated will be $< 1/2$. We make the following claim instead (again without proof):

$$Pr_{x \xleftarrow{\$} \{0,1\}^\lambda}[Pr_{r \xleftarrow{\$} \{0,1\}^\lambda}[A(f(x), r) = \langle x, r \rangle] \geq \frac{1}{2} + \frac{\varepsilon}{2}] \geq \frac{\varepsilon}{2}.$$

We again call any given $x \leftarrow \{0,1\}^\lambda$ *good* if the inner condition is true, that is to say, if

$$Pr_{r \xleftarrow{\$} \{0,1\}^\lambda}[A(f(x), r) = \langle x, r \rangle] \geq \frac{1}{2} + \frac{\varepsilon}{2}.$$

By our earlier claim, x is good with non-negligible probability, so we can from here fix some good x.

Again, define $H(r) = A(f(x), r)$, so we have $Pr_{r \xleftarrow{\$} \{0,1\}^\lambda}[H(r) = \langle x, r \rangle] \geq 1/2 + \varepsilon/2$. Note that we have a problem: given such a H, x is not uniquely defined, since such an H can correspond to multiple values of x. The solution is to find a list L of x that satisfies this H, and then output a random element of L; this list will be polynomial in λ, so it is acceptable to do this.

Our strategy here is to use H to implement H', where $Pr_{r \xleftarrow{\$} \{0,1\}^\lambda}[H'(r) = \langle x, r \rangle] \geq 7/8$. We choose some random $r_1, \ldots, r_k \xleftarrow{\$} \{0,1\}^\lambda$ for some k, and pretend that we can magically obtain $b_i = \langle x, r_i \rangle$ (for now). We then define

$$H'(r) = \text{maj}_{i \in [k]}(H(r \oplus r_i) \oplus b_i),$$

that is to say, the majority of all values $H(r \oplus r_i) \oplus b_i$ over all i.

Note that

$$Pr_{r_1,\ldots,r_k,r \xleftarrow{\$} \{0,1\}^\lambda}[H(r \oplus r_i) \oplus b_i = \langle x, r\rangle] = Pr_{r_1,\ldots,r_k,r \xleftarrow{\$} \{0,1\}^\lambda}[H(r \oplus r_i) = \langle x, r \oplus r_i\rangle]$$
$$\geq \frac{1}{2} + \frac{\varepsilon}{2},$$

by definition. Thus, if k is high enough, then H' will give us the right answer with high probability. More formally, for high enough k, we can obtain

$$Pr_{r_1,\ldots,r_k,r \xleftarrow{\$} \{0,1\}^\lambda}[H'(r) = \langle x, r\rangle] \geq \frac{31}{32},$$

so by the Markov inequality, we have

$$Pr_{r_1,\ldots,r_k \xleftarrow{\$} \{0,1\}^\lambda}[Pr_{r \xleftarrow{\$} \{0,1\}^\lambda}[H'(r) = \langle x, r\rangle] \geq \frac{7}{8}] \geq \frac{3}{4}.$$

Now, we can use H' as the H defined in Step 2, and proceed with the rest of the proof as in Step 2.

The only unresolved issue here is how to obtain the magic b_i. We can do this as follows. Let $k = 2^\ell$, and choose some random $r'_1, \ldots, r'_\ell \xleftarrow{\$} \{0,1\}^\lambda$. Let $S \subseteq \{1, \ldots, \ell\}$, and let $r_S = \bigoplus_{i \in S} r'_i$. Assume that we can magically obtain $b'_i = \langle x, r'_i \langle$ for each $i \in [\ell]$, and let $b_S = \bigoplus_{i \in S} b'_i$. Thus, we have $b_S = \langle r_S, x\rangle$, and since we have $k = 2^\ell$ such S, this gives us the values b_S as desired. We have essentially reduced the initial k magic steps into $\log k$ magic steps, and we can repeat this problem to solve the progressively smaller magic steps. Note that the values b_S are correlated; this is not actually an issue in the proof.

$\square$