

Homework 7

Instructor: Anna Lysyanskaya

Due: March 14, 2007

Problem 1: Pseudo Random Generator and the Hybrid Argument

In class we discussed a pseudorandom generator which outputs only 2 bits, G_{2B} . Let f be a one way permutation, and let H be a hardcore bit function for f . On seed $x \in \text{domain of } f$, the generator will produce: $H(x), H(f(x))$. We showed that even when we also include $f(f(x))$, the output of this generator on a random seed is indistinguishable from a random 2 bit string, i.e.:

Define the following 2 distributions:

$$D_{PR2B} = \{PK \leftarrow GPK(1^k); x \leftarrow \text{dom}(f_{PK}) : (PK, f(f(x)), H(x), H(f(x)))\}$$

and

$$D_{Rand2B} = \{PK \leftarrow GPK(1^k); x \leftarrow \text{dom}(f_{PK}); b_0, b_1 \leftarrow \{0, 1\} : (PK, x, b_0, b_1)\}$$

We showed that $D_{PR2B} \approx D_{Rand2B}$ and thus G_{2B} is a valid 2 bit PRG. We did this by considering a third distribution:

$$D_{Hyb2B} = \{PK \leftarrow GPK(1^k); x \leftarrow \text{dom}(f_{PK}); b_0 \leftarrow \{0, 1\} : (PK, f(x), b_0, H(x))\}$$

We showed that $D_{PR2B} \approx D_{Hyb2B}$, and $D_{Hyb2B} \approx D_{Rand2B}$, and we argued that that implied that $D_{PR2B} \approx D_{Rand2B}$.

For a review of the proof that $D_{PR2B} \approx D_{Hyb2B}$, see page 3.

Now, what if we want a PRG that outputs more than 2 bits? We discussed the solution in class, but lets work out the details.

Our n bit PRG will take a seed $x \in \text{dom}(f)$, and output $H(x), H(f(x)), \dots, H(f^{n-1}(x))$, where f^n represents the operation where f is applied n times (we require that n be polynomial in our security parameter k). To show that this is a valid PRG, it is enough to show that the following 2 distributions are indistinguishable:

$$D_{PRnB} = \{PK \leftarrow GPK(1^k); x \leftarrow \text{dom}(f_{PK}) : (PK, f^n(x), H(x), H(f(x)), \dots, H(f^{n-1}(x)))\}$$

and

$$D_{RandnB} = \{PK \leftarrow GPK(1^k); x \leftarrow \text{dom}(f_{PK}); b_0, b_1, \dots, b_n \leftarrow \{0, 1\} : (PK, x, b_0, b_1, \dots, b_n)\}$$

(Note that this is actually slightly stronger than what we need for pseudorandomness since here we also give the adversary $f^n(x)$).

We will do that by introducing hybrid distributions for $i = 1, \dots, n-1$. Hybrid i will choose the first i bits of its output at random and then will compute the next $n-i$ bits as $H(x), \dots, H(f^{n-i-1}(x))$. More formally:

$$D_{Hybi} = \{PK; x; b_1, \dots, b_i \leftarrow \{0, 1\} : (PK, f^{n-i}(x), b_1, \dots, b_i, H(x), H(f(x)), \dots, H(f^{n-i-1}(x)))\}$$

Note now that $D_{Hyb0} = D_{PRnB}$ and $D_{Hybn} = D_{RandnB}$.

As discussed in class, we know that if we can show that there is some negligible ν such that for all i , no adversary can distinguish D_{Hybi} from D_{Hybi+1} with advantage greater than $\nu(x)$, and if n is polynomial in k , then that means that $D_{PRnB} \approx D_{RandnB}$.

Thus all we need to do is prove that hybrids i and $i+1$ are indistinguishable. Let's prove that now. (Hint: we will generalize the proof given in the Class Review section.)

- a. Fill in the blanks in the following proof outline: We want to show that: $\underline{\hspace{2cm}} \implies \underline{\hspace{2cm}}$
 We will do that by showing the contrapositive, i.e. that: $\neg \underline{\hspace{2cm}} \implies \neg \underline{\hspace{2cm}}$.

This means we assume we have an adversary \mathcal{A} , which takes as input $\underline{\hspace{2cm}}$ and outputs 0 or 1.

Let $p_{\underline{\hspace{2cm}},0}$ be the probability that \mathcal{A} outputs 0 when given a random element from distribution $\underline{\hspace{2cm}}$, and let $p_{\underline{\hspace{2cm}},1}$ be the probability that \mathcal{A} outputs 0 when given a random element from distribution $\underline{\hspace{2cm}}$. Then we are assuming that there exists nonnegligible ϵ such that $|p_{\underline{\hspace{2cm}},0} - p_{\underline{\hspace{2cm}},1}| = \epsilon(k)$. Assume WLOG that $p_{\underline{\hspace{2cm}},0} > p_{\underline{\hspace{2cm}},1}$, i.e. that \mathcal{A} is more likely to output a 0 when given a random element of distribution $\underline{\hspace{2cm}}$ than when it is given a random element of distribution $\underline{\hspace{2cm}}$.

Now we want to show that we can build an algorithm \mathcal{B} which takes as input $\underline{\hspace{2cm}}$ and tries to output 0 if it was given an instance of distribution $\underline{\hspace{2cm}}$ and 1 if it was given an instance of distribution $\underline{\hspace{2cm}}$, and succeeds in distinguishing with nonnegligible advantage, ϵ' . Basically, we want to show that

$$|Pr[e \leftarrow D_{\underline{\hspace{2cm}}}; b \leftarrow \mathcal{B}(e) : b = 0] - Pr[e \leftarrow D_{\underline{\hspace{2cm}}}; b \leftarrow \mathcal{B}(e) : b = 1]| = \epsilon'(k)$$

for nonnegligible ϵ' .

- b. What should algorithm \mathcal{B} do? What arguments should it send to \mathcal{A} ? What should it output if \mathcal{A} returns 0? What should it output if \mathcal{A} returns 1?
- c. What is \mathcal{B} 's advantage, i.e. show that ϵ' as defined above is nonnegligible.

Problem 2: Public Key Encryption

Suppose we use an f which is a trapdoor permutation to build a PRG using the construction above. Then we can use this to construct a public key encryption scheme:

- Bob chooses $(PK, s) \leftarrow GPK(1^k)$ and publishes PK
- Alice wants to encrypt an n bit message m . She chooses a random $x \leftarrow \text{domain of } f_{PK}$. Let m_i be the i th bit of the message m . Alice computes $C = (C_0 = f^n(x), c_1 = H(x) \oplus m_1, c_2 = H(f(x)) \oplus m_2, \dots, c_{n-1} = H(f^{n-2}(x)) \oplus m_{n-1}, c_n = H(f^{n-1}(x)) \oplus m_n)$, and sends that to Bob.
- Bob receives $C = (C_0, c_1, \dots, c_n)$. Bob uses his secret key to compute $f^{-1}(C_0) = f^{n-1}(x), f^{-1}(f^{-1}(C_0)) = f^{n-2}(x), \dots, (f^{-1})^n(C_0) = x$. Then Bob uses these values to compute $H(x), H(f(x)), \dots, H(f^{n-1}(x))$. Finally, Bob computes $H(x) \oplus c_1 = m_1, H(f(x)) \oplus c_2 = m_2, \dots, H(f^{n-1}(x)) \oplus c_n = m_n$ and thus recovers m .

Now we will prove that this encryption scheme is secure. Recall our definition of security for encryption. We said an encryption scheme (GPK, Enc, Dec) is secure if there exists an algorithm FakeCiphertext such that for all m , the following two distributions are equivalent:

$$D_{Enc}(m) = \{(PK, SK) \leftarrow GPK(1^k); c \leftarrow Enc(PK, m); (c, m, PK)\}$$

and

$$D_{Fake}(m) \{(PK, SK) \leftarrow GPK(1^k); \hat{c} \leftarrow FakeCiphertext(PK, |m|); (\hat{c}, m, PK)\}.$$

Note that FakeCiphertext is only given the length of the message, and yet it produces something indistinguishable from the encryption of the message. This captures our idea that encryption should hide all information about the message (besides its length).

- a. If we want to show that the encryption scheme described above is secure according to this definition, we need to show that there exists an algorithm *FakeCiphertext* as described above. What should *FakeCiphertext*(PK, l) do?
- b. Now let's show that this is indistinguishable from a real encryption. First let's consider an alternate encryption algorithm *RandEnc* which on input (PK, m) chooses a random $x \leftarrow$ domain of f_{PK} as above, and then chooses random string b_1, \dots, b_n . It outputs $C = (C_0 = x, c_1 = m_1 \oplus b_1, \dots, c_n = m_n \oplus b_n)$. Show that this distribution is identical to that generated by *FakeCiphertext* (i.e. that for all m , $D_{RandEnc}(m) \approx D_{fake}(m)$).
- c. Show that the above distribution is indistinguishable from that generated by the true encryption algorithm. Recall that we proved in problem 1 that $D_{Random} \approx D_{PR}$. Let's use this to show that for all m , $D_{Enc}(m) \approx D_{RandEnc}(m)$. Fill in the blanks:

We want to show that: _____ \implies _____

We will do that by showing the contrapositive, i.e. that: \neg _____ \implies \neg _____.

This means we assume we have an adversary \mathcal{A} , which takes as input _____ and outputs 0 or 1.

Let $p_{_,0}$ be the probability that \mathcal{A} outputs 0 when given a random element from distribution _____, and let $p_{_,1}$ be the probability that \mathcal{A} outputs 0 when given a random element from distribution _____. Then we are assuming that there exists nonnegligible ϵ such that $|p_{_,0} - p_{_,1}| = \epsilon(k)$. Assume WLOG that $p_{_,0} > p_{_,1}$, i.e. that \mathcal{A} is more likely to output a 0 when given a random element of distribution _____ than when it is given a random element of distribution _____.

Now we want to show that we can build an algorithm \mathcal{B} which takes as input _____ and tries to output 0 if it was given an instance of distribution _____ and 1 if it was given an instance of distribution _____, and succeeds in distinguishing with nonnegligible advantage, ϵ' . I.e. we want to show that $|Pr[e \leftarrow D_{_}; b \leftarrow \mathcal{B}(e) : b = 0] - Pr[e \leftarrow D_{_}; b \leftarrow \mathcal{B}(e) : b = 1]| = \epsilon'(k)$ for nonnegligible ϵ' .

- d. What should algorithm \mathcal{B} do? What arguments should it send to \mathcal{A} ? What should it output if \mathcal{A} returns 0? What should it output if \mathcal{A} returns 1?
- e. What is \mathcal{B} 's advantage, i.e. show that ϵ' as defined above is nonnegligible.

Class Review

Recall the definition we used in class for a hardcore bit: H is a hardcore bit function for f if the following two distributions are indistinguishable:

$$D_0 = \{PK \leftarrow GPK(1^k); x \leftarrow \text{domain of } f_{PK} : (PK, f(x), H(x))\}$$

and

$$D_{\frac{1}{2}} = \{PK \leftarrow GPK(1^k); x \leftarrow \text{domain of } f_{PK}; b \leftarrow \{0, 1\} : (PK, x, b)\}.$$

We want to show that $H(x)$ is hardcore for $f \implies D_{PR2B} \approx D_{Hyb2B}$
 i.e. that $D_0 \approx D_{\frac{1}{2}} \implies D_{PR2B} \approx D_{Hyb2B}$

We will do that by showing the contrapositive, i.e. that: $\neg(D_{PR2B} \approx D_{Hyb2B}) \implies \neg(D_0 \approx D_{\frac{1}{2}})$.

This means we assume we have an adversary \mathcal{A} , who takes as input (PK, y, b_1, b_2) which is either of the form $(PK, f(x), H(x), H(f(x)))$ for randomly chosen x , or of the form $(PK, f(x), b, H(x))$ for randomly chosen bit b and randomly chosen x . \mathcal{A} outputs 0 or 1.

Let $p_{PR2B,0}$ be the probability that \mathcal{A} outputs 0 when given a random element from distribution D_{PR2B} , and let $p_{Hyb2B,0}$ be the probability that \mathcal{A} outputs 0 when given a random element from distribution D_{Hyb2B} . Then we are assuming that there exists nonnegligible ϵ such that $|p_{PR2B,0} - p_{Hyb2B,0}| = \epsilon(k)$. Assume WLOG that $p_{PR2B,0} > p_{Hyb2B,0}$, i.e. that \mathcal{A} is more likely to output a 0 when given a random element of distribution D_{PR2B} than when it is given a random element of distribution D_{Hyb2B} .

Now we want to show that we can build an algorithm \mathcal{B} which takes as input (PK, y, b') which is either of the form $(PK, f(x), H(x))$ for randomly chosen x , or of the form (PK, x, b) for randomly chosen x and random bit b . \mathcal{B} tries to output 0 if it was given an instance of distribution D_0 and 1 if it was given an instance of distribution $D_{\frac{1}{2}}$, and we want to show that it succeeds in distinguishing with nonnegligible advantage, ϵ' . I.e. we want to show that $|Pr[e \leftarrow D_0; \hat{b} \leftarrow \mathcal{B}(e) : \hat{b} = 0] - Pr[e \leftarrow D_{\frac{1}{2}}; \hat{b} \leftarrow \mathcal{B}(e) : \hat{b} = 0]| = \epsilon'(k)$ for nonnegligible ϵ' .

Let's define algorithm \mathcal{B} as follows: on input (PK, y, b') , \mathcal{B} computes $f(y)$ and $H(y)$. It then runs $\mathcal{A}(PK, f(y), b', H(y))$. If \mathcal{A} outputs 0, \mathcal{B} outputs 0, and if \mathcal{A} outputs 1, \mathcal{B} outputs 1.

Note that, if the input \mathcal{B} received was chosen from D_0 , then it was of the form $(PK, f(x), H(x))$ for randomly chosen x . That means that what \mathcal{A} receives was $(PK, f(y) = f(f(x)), b' = H(x), H(y) = H(f(x)))$, which is clearly distributed like D_{PR2B} . Thus, \mathcal{A} will output 0 with probability $p_{PR2B,0}$.

On the other hand, if the input \mathcal{B} received was chosen from D_1 , then it was of the form (PK, x, b) for randomly chosen x and random bit b . That means that what \mathcal{A} received was $(PK, f(y) = f(x), b' = b, H(y) = H(x))$, which is clearly distributed like D_{PR2B} . Thus, \mathcal{A} will output 0 with probability $p_{Hyb2B,0}$.

Thus \mathcal{B} 's advantage will be:

$$\begin{aligned}
& |Pr[e \leftarrow D_0; b \leftarrow \mathcal{B}(e) : b = 0] - Pr[e \leftarrow D_{\frac{1}{2}}; b \leftarrow \mathcal{B}(e) : b = 0]| \\
= & |Pr[(PK, y, b) \leftarrow D_0; b \leftarrow \mathcal{A}(PK, f(y), b, H(y)) : b = 0] \\
& - Pr[(PK, y, b) \leftarrow D_{\frac{1}{2}}; b \leftarrow \mathcal{A}(PK, f(y), b, H(y)) : b = 0]| \\
= & |p_{PR2B,0} - p_{Hyb2B,0}| \text{ by our arguments above} \\
= & \epsilon(k)
\end{aligned}$$

Thus, \mathcal{B} can distinguish D_0 from $D_{\frac{1}{2}}$ with nonnegligible advantage.