MIT 6.875J/18.425J and Berkeley CS276
Foundations of Cryptography (Fall 2020)

Problem Set 1: Released September 2, Due September 15

The problem set is due on **Tuesday September 15 at 10pm ET/7pm PT**. Please make sure to upload **each problem as a separate pdf document** to the Gradescope course webpage by the deadline (all registered students should have access to this webpage on Gradescope). Typed solutions using LaTeX are strongly encouraged (template provided on the course webpage).

Collaboration is permitted; however, you must write up your own solutions individually and acknowledge all of your collaborators.

Problem 1. Secure Communication without Keys

Alice and Bob share a perfect communication channel and Alice wants to send a message to Bob. However an (unbounded) adversary Eve is observing a noisy version of their channel: for each bit b sent through the channel, Eve observes b with probability p and observes $1-b$ with probability $1-p$. Alice wants to send bits through the channel to Bob such that Bob can recover the message, but Eve does not learn anything about the message.

1.1 Say Alice simply sends the message through the channel to Bob. Show that if $p = 1/2$, then this achieves perfect secrecy against Eve: formally, for every two messages $m_0, m_1 \in \{0,1\}^\ell$ and unbounded adversary Eve

$$\Pr[b \leftarrow \{0,1\}, \text{Alice sends } m_b \ ; \ \text{Eve}(c) = b] = \frac{1}{2} \ , \tag{1}$$

where c denotes all the bits that Eve observes in the noisy channel and $\text{Eve}(c)$ guesses which message was sent.

1.2 Show that if $p = 0.9$, then simply sending the message does not achieve perfect secrecy against Eve (i.e. for every $\ell \geq 1$ there exists some $m_0, m_1 \in \{0,1\}^\ell$ and unbounded adversary Eve such that Equation (1) does not hold).

1.3 Say $p = 0.9$ and for simplicity Alice's message is a single bit. Furthermore, Alice relaxes her goal to achieving *almost perfect secrecy*: for $m_0 = 0, m_1 = 1$ and unbounded adversary Eve

$$\Pr[b \leftarrow \{0,1\}, \text{Alice}(1^n, m_b) = e \ ; \ \text{Eve}(c) = b] \leq \frac{1}{2} + \frac{1}{2^n} \ , \tag{2}$$

where $\text{Alice}(1^n, m_b)$ is a p.p.t. algorithm that computes an encoding e that is sent through the channel to Bob, and c denotes all the bits that Eve observes in the noisy channel.

Design $\text{Alice}(1^n, m_b)$ to generate an encoding e such that if e is sent through the channel to Bob:

- Bob can always recover the message m_b
- this encoding achieves almost perfect secrecy against Eve (Equation (2)).

Problem 2. Negligible or Not?

Recall the definition of a negligible function: a function $f : \mathbb{N} \to \mathbb{R}_{\geq 0}$ is negligible if for every constant $c \geq 0$ there exists $n_c \in \mathbb{N}$ such that $f(n) < \frac{1}{n^c}$ for all $n \geq n_c$.

This question asks you to identify properties of negligible functions. Let f and g be negligible functions. Your task is to determine, for each h below, whether h is a negligible function. In each case, you have to either prove that h is negligible, or come up with a counterexample, that is, negligible functions f (and, when relevant, g) such that the resulting h is *not* a negligible function.

2.1 $h = p \circ f$ where p is a polynomial function with constant term 0 and non-negative coefficients (i.e. $p(x) = \sum_{i=1}^{d} a_i \cdot x^i$ for $a_i \in \mathbb{R}_{\geq 0}$). Here, $\circ$ denotes composition of functions, that is, $h(n) = p(f(n))$.

2.2 $h = f + g$. That is, $h(n) = f(n) + g(n)$.

2.3 $h = f/g$. That is, $h(n) = f(n)/g(n)$.

2.4 $h = f^{1/2}$. That is, $h(n) = f(n)^{1/2}$.

2.5 $h = f^{1/n}$. That is, $h(n) = f(n)^{1/n}$.

Problem 3. Subtle Definitions

Recall the security definition of a one-way function.

Definition 1. *A function $f : \{0,1\}^* \to \{0,1\}^*$ is* one-way *if for every p.p.t. algorithm $\mathcal{A}$ there exists a negligible function $\epsilon : \mathbb{N} \to [0,1]$ such that for all $n \in \mathbb{N}$:*

$$\Pr[x \leftarrow \{0,1\}^n \; ; \; x' \leftarrow \mathcal{A}(1^n, f(x)) \; : \; f(x') = f(x)] \leq \epsilon(n) \; .$$

Are the following variants of this definition equivalent to the original definition? Either prove the original definition implies the variant and vice versa, or provide a counterexample that one definition does not imply the other (assuming the existence of length-preserving one-way functions).

3.1 There exists a negligible function $\epsilon(n)$ and a nonempty finite set $S \subset \mathbb{N}$ such that for every p.p.t. algorithm $\mathcal{B}$ and $n \in S$:

$$\Pr[x \leftarrow \{0,1\}^n \; ; \; x' \leftarrow \mathcal{B}(1^n, f(x)) \; : \; f(x') = f(x)] \leq \epsilon(n) \; .$$

3.2 There exists an infinite set $S \subseteq \mathbb{N}$ such that for every p.p.t. algorithm $\mathcal{B}$ there exists a negligible function $\epsilon(n)$ such that for all $n \in S$:

$$\Pr[x \leftarrow \{0,1\}^n \; ; \; x' \leftarrow \mathcal{B}(1^n, f(x)) \; : \; f(x') = f(x)] \leq \epsilon(n) \; .$$

3.3 For every p.p.t. algorithm $\mathcal{B}$, there exists a negligible function $\epsilon(n)$ such that for all $n \in \mathbb{N}$:

$$\Pr[x \leftarrow \{0,1\}^n \; ; \; x' \leftarrow \mathcal{B}(f(x)) \; : \; f(x') = f(x)] \leq \epsilon(n) \; .$$

Problem 4. One Way (Functions) or Another

Assume that f is a one-way function which is length-preserving, i.e., for every $x \in \{0,1\}^*$ it holds that $|f(x)| = |x|$. Note that $\bar{x}$ denotes the bitwise complement of x, $|$ denotes concatenation, e.g., $1011|\overline{1011} = 10110100$, and $x_{[1:a]}$ denotes the first a bits of the string x.

4.1 Let $g(x) = f\left(\overline{f(x)}\right)$ where f is a one-way permutation. Is g a one-way permutation?

4.2 Let $g(x) = f(x)|x_{[1:\log(n)]}$. Is g a one-way function?

4.3 Let $g(x) = f(x)_{[1:n-1]}$. Is g a one-way function?

4.4 Let f be a length-preserving one-way function. Construct a one-way function g such that for every $n \in \mathbb{N}$, $g(0^n) = 0^n$.

4.5 Let f be a one-way permutation. Construct a one-way permutation g such that for every $n \in \mathbb{N}$, $g(0^n) = 0^n$.

Hint for 4.1-4.3:

- In the case that g is one-way: you should show a reduction from inverting f to inverting g; and
- In the case that g is not always one-way: you should do the following. Assuming only that length-preserving one-way functions exist, you should construct a length-preserving one-way function f such that when instantiated with f, g is not a one-way function. That is, you should show a polynomial-time inverting algorithm for g.

  For example, to prove that the construction of $g(x) = f(x_{[1:n/2]}|0^{n/2})$ is not always one-way, the argument goes as follows. Assume that h is an arbitrary length-preserving one-way function. Construct the function

$$f(x) = \begin{cases} 0^n & \text{if } x_{[n/2+1,n]} = 0^{n/2} \\ h(x) & \text{otherwise} \end{cases}$$

  f is one-way assuming that h is one-way (you have to show this by a reduction). Furthermore, g constructed using this f is not one-way (you have to show this by exhibiting a polynomial-time attack). In general, your solutions should contain complete proofs that the counterexamples are valid.

Problem 5. Does one imply the other?

5.1 Let $G : \{0,1\}^n \to \{0,1\}^{n+1}$ be a pseudorandom generator. Is G necessarily a one-way function? If yes, prove it; and if not, construct a counterexample, namely, a function G that is a pseudorandom generator but not a one-way function.

5.2 Assume that g is a one-way function which is length-preserving, i.e., for every $x \in \{0,1\}^*$ it holds that $|g(x)| = |x|$. Show there exists a length-expanding one-way function $F : \{0,1\}^n \to \{0,1\}^{n+1}$ that is not a pseudorandom generator.

Although not all one-way functions are pseudorandom generators, some are. Next we will prove that the subset-sum one-way function from class is a pseudorandom generator.

For random $\vec{a} = (a_1, \ldots, a_n) \in \{0, \ldots, p-1\}^n$ the subset-sum one-way function $f_{\vec{a}}$ maps subsets $S \subseteq [n]$ to $\{0, \ldots, p-1\}$ as follows

$$f_{\vec{a}}(S) = \sum_{i \in S} a_i \pmod{p} \ .$$

The subset $S \subseteq [n]$ is represented by a string $(S_1, \ldots, S_n) \in \{0,1\}^n$ such that for every $i \in [n]$, $S_i = 1$ if and only if $i \in S$.

The one-wayness of this function is based on the hardness of the subset-sum problem (for random $\vec{a}$). For this problem, assume the prime $p \in [2^{2n-1}, 2^{2n} - 1]$ (so its bitwise representation is in $\{0,1\}^{2n}$).

Assume $f_{\vec{a}}$ is not pseudorandom: i.e. there exists an adversary $\mathcal{A}$ that distinguishes between outputs of $f_{\vec{a}}$ and the uniform distribution over $\{0, \ldots, p-1\}$. For this problem, we will make the stronger assumption

that $\mathcal{A}$, given $\vec{a}$ and y, perfectly distinguishes whether y is in the image of $f_{\vec{a}}$ or not (i.e. for every $\vec{a} \in \{0,\ldots,p-1\}^n, y \in \{0,\ldots,p-1\}$, if there exists $S \subseteq [n]$ such that $y = f_{\vec{a}}(S)$ then $\mathcal{A}(\vec{a}, y) = 1$ and otherwise $\mathcal{A}(\vec{a}, y) = 0$).

Next we will construct a reduction $\mathcal{B}$ that given as input $y = f_{\vec{a}}(S)$, outputs an inverse S' with non-negligible probability. This contradicts the one-wayness of $f_{\vec{a}}$. Hence there does not exist such an $\mathcal{A}$.

5.3 Design a reduction $\mathcal{B}$ that given as input $y = f_{\vec{a}}(S_1 \ldots S_n)$ outputs an inverse S'.
Hint: First design a reduction $\mathcal{B}'$ that given as input $y = f_{\vec{a}}(S_1 \ldots S_n)$, uses $\mathcal{A}$ to produce a guess for S_1.

5.4 Show your reduction $\mathcal{B}$ outputs an inverse S' with non-negligible probability, i.e.

$$\Pr[\vec{a} \leftarrow \{0,\ldots,p-1\}^n \; ; \; S \leftarrow \{0,1\}^n \; ; \; S' \leftarrow \mathcal{B}(1^n, f_{\vec{a}}(S)) \; : \; f_{\vec{a}}(S') = f_{\vec{a}}(S)] = \epsilon(n)$$

where $\epsilon(n)$ is non-negligible.