CS 171: Problem Set 6

Due Date: March 14th, 2024 at 8:59pm via Gradescope

1 One-Way Functions

Let $f : \{0,1\}^n \to \{0,1\}^n$ be a one-way function, and let

$$g(x) = f(f(x))$$

Is g necessarily a one-way function? Prove your answer. In your answer, you may use a OWF $h : \{0,1\}^{n/2} \to \{0,1\}^{n/2}$.

Tip: Your answer should have one of the following forms. Only one of them is possible:

- Prove that if f is a OWF, then g is also a OWF.
- (1) Construct a function f. (2) Prove that f is a one-way function. (3) Then prove that when g is constructed from this choice of f, g is not a one-way function.

Also, you may cite without proof any theorems proven in discussion or lecture.

Solution

Theorem 1.1 *g is not necessarily a one-way function.*

Proof

1. We will construct a OWF f such that $g(x) = f(f(x))$ is not a OWF. First, let h be a OWF that maps $\{0,1\}^{n/2} \to \{0,1\}^{n/2}$. Second, let the input to f take the form $x = (x_0, x_1) \in \{0,1\}^{n/2} \times \{0,1\}^{n/2}$. Finally,

$$\text{let } f(x) = 0^{n/2} || h(x_0)$$

2. We proved in discussion 7 that f is a OWF.

3. Next, we will show that for this choice of f, g is not a OWF. Observe that:

$$\begin{aligned} g(x) &= f(f(x)) \\ &= f\big(0^{n/2} || h(x_0)\big) = 0^{n/2} || h(0^{n/2}) \end{aligned}$$

Next, note that $g(x)$ is a constant – $0^{n/2} || h(0^{n/2})$ – that is the same for all x.

Now it is easy to construct an adversary $\mathcal{A}$ that breaks the OWF security of g. $\mathcal{A}$ outputs an arbitrary value of x', such as $x' = 0^n$. Then $\mathcal{A}$ will win the OWF security game with certainty because for any x chosen by the challenger, $g(x) = g(x')$.

■ ■

2 Concatenated Hash Functions

Let $\mathcal{H}_1 = (\mathsf{Gen}_1, H_1)$ and $\mathcal{H}_2 = (\mathsf{Gen}_2, H_2)$ be two fixed-length hash functions that take inputs of length $3n$ bits and produce outputs of length n bits. Only one of $\mathcal{H}_1$ and $\mathcal{H}_2$ is collision resistant; the other one is not collision-resistant, and you don't know which is which.

Next, we define two new hash functions $\mathcal{H}_3 = (\mathsf{Gen}_3, H_3)$ and $\mathcal{H}_4 = (\mathsf{Gen}_4, H_4)$ below:

$\underline{\mathcal{H}_3}$:

1. $\mathsf{Gen}_3(1^n)$: Sample $s_1 \leftarrow \mathsf{Gen}_1(1^n)$ and $s_2 \leftarrow \mathsf{Gen}_2(1^n)$. Output $s = (s_1, s_2)$.
2. $H_3^s(x)$: Output $H_1^{s_1}(x)||H_2^{s_2}(x)$.

Note that $H_3^s : \{0,1\}^{3n} \to \{0,1\}^{2n}$.

$\underline{\mathcal{H}_4}$:

1. $\mathsf{Gen}_4(1^n)$: Sample $s_1 \leftarrow \mathsf{Gen}_1(1^n)$ and $s_2 \leftarrow \mathsf{Gen}_2(1^n)$. Output $s = (s_1, s_2)$.
2. $H_4^s(x)$: Let $x = (x_1, x_2) \in \{0,1\}^{3n} \times \{0,1\}^{3n}$. Output $H_1^{s_1}(x_1)||H_2^{s_2}(x_2)$.

Note that $H_4^s : \{0,1\}^{6n} \to \{0,1\}^{2n}$.

Question: For each of $\mathcal{H}_3$ and $\mathcal{H}_4$, determine whether the hash function is collision-resistant, and prove your answer.

Solution

Theorem 2.1 *$\mathcal{H}_3$ is collision resistant.*

Proof

1. Assume toward contradiction that $\mathcal{H}_3$ is not collision resistant. Then there is an adversary $\mathcal{A}$ that finds a collision in H_3^s with non-negligible probability. Then we will use $\mathcal{A}$ to construct adversaries $\mathcal{B}_1$ and $\mathcal{B}_2$ that find collisions in $H_1^{s_1}$ and $H_2^{s_2}$, respectively, with non-negligible probability.
2. $\underline{\mathcal{B}_1(s_1)}$:
   (a) In the security game for $\mathcal{H}_1$, the challenger samples $s_1 \leftarrow \mathsf{Gen}_1(1^n)$ and sends s_1 to $\mathcal{B}_1$.
   (b) $\mathcal{B}_1$ samples $s_2 \leftarrow \mathsf{Gen}_2(1^n)$, and sets $s = (s_1, s_2)$.
   (c) Then $\mathcal{B}_1$ runs $\mathcal{A}(s)$ until it outputs (x, x'), which will be a collision in H_3^s with non-negligible probability.
   (d) $\mathcal{B}_1$ outputs (x, x') as its guess for a collision in $H_1^{s_1}$.

3. Analysis of $\mathcal{B}_1$: $\mathcal{B}_1$ correctly simulates the collision-resistance security game for $\mathcal{H}_3$ because s is sampled from the same distribution as the one used by $\mathsf{Gen}_3(1^n)$. That ensures that when $\mathcal{B}_1$ runs $\mathcal{A}$, $\mathcal{A}$ will output a collision in H_3^s with non-negligible probability. In this case, $x \neq x'$, and:

$$H_3^s(x) = H_3^s(x')$$
$$H_1^{s_1}(x)||H_2^{s_2}(x) = H_1^{s_1}(x')||H_2^{s_2}(x')$$

This implies that (x, x') is also a collision in $H_1^{s_1}$ because $H_1^{s_1}(x) = H_1^{s_1}(x')$.

4. We can find collisions in $H_2^{s_2}$ using a similar procedure. We will describe $\mathcal{B}_2$, the algorithm that does so, but it is almost identical to $\mathcal{B}_1$.[1]

   $\mathcal{B}_2(s_2)$:

   (a) In the security game for $\mathcal{H}_2$, the challenger samples $s_2 \leftarrow \mathsf{Gen}_2(1^n)$ and sends s_2 to $\mathcal{B}_2$.
   (b) $\mathcal{B}_2$ samples $s_1 \leftarrow \mathsf{Gen}_1(1^n)$, and sets $s = (s_1, s_2)$.
   (c) Then $\mathcal{B}_2$ runs $\mathcal{A}(s)$ until it outputs (x, x'), which will be a collision in H_3^s with non-negligible probability.
   (d) $\mathcal{B}_2$ outputs (x, x') as its guess for a collision in $H_2^{s_2}$.

   By a similar argument to the one above, we can show that $\mathcal{B}_2$ finds a collision in $H_2^{s_2}$ with non-negligible probability.

5. Now we can finish the proof. We have constructed adversaries that break the collision-resistance security of $\mathcal{H}_1$ and $\mathcal{H}_2$. However, we know that one of these two hash functions is collision-resistant, so we've reached a contradiction. That means our initial assumption was false, and in fact, $\mathcal{H}_3$ is collision-resistant.

■

Theorem 2.2 $\mathcal{H}_4$ *is not collision-resistant.*

Proof

1. We know that one of $\mathcal{H}_1$ or $\mathcal{H}_2$ is not collision resistant. First, we will prove that if $\mathcal{H}_2$ is not collision-resistant, then neither is $\mathcal{H}_4$.

2. If $\mathcal{H}_2$ is not collision-resistant, then there is an algorithm $\mathcal{A}$ that finds collisions in $H_2^{s_2}$ with non-negligible probability. Then we can construct an adversary $\mathcal{B}$ that finds collisions in H_4^s:

   $\mathcal{B}(s)$:

   (a) In the security game for $\mathcal{H}_4$, the challenger samples $s = (s_1, s_2) \leftarrow \mathsf{Gen}_4(1^n)$ and sends s to $\mathcal{B}$.

[1] Students do not have to describe $\mathcal{B}_2$ in so much detail. They can just say that $\mathcal{B}_2$ works analogously to $\mathcal{B}_1$.

(b) $\mathcal{B}$ runs $\mathcal{A}(s_2)$ until it outputs (x, x'), which will be a collision in $H_2^{s_2}$ with non-negligible probability.

(c) $\mathcal{B}$ outputs $0^{3n}||x$ and $0^{3n}||x'$ as its guess for a collision in H_4^s.

3. Analysis of $\mathcal{B}$: $\mathcal{B}$ correctly simulates the security game for $\mathcal{H}_2$ with $\mathcal{A}$ as the adversary because s_2 is sampled from the same distribution as in $\mathsf{Gen}_2(1^n)$. That ensures that with non-negligible probability, (x, x') are a collision in $H_2^{s_2}$.

   In this case, $x \neq x'$, and $H_2^{s_2}(x) = H_2^{s_2}(x')$. Then this means that $0^{3n}||x$ and $0^{3n}||x'$ are a collision in H_4^s because:

$$0^{3n}||x \neq 0^{3n}||x', \text{ and}$$
$$H_4^s(x) = H_1^{s_1}(0^{3n})||H_2^{s_2}(x) = H_1^{s_1}(0^{3n})||H_2^{s_2}(x') = H_4^s(x')$$

   In conclusion, $\mathcal{B}$ finds a collision in H_4^s with non-negligible probability.

4. By a nearly identical argument, we can show that if $\mathcal{H}_1$ is not collision-resistant, then neither is $\mathcal{H}_4$. Since we know that one of $\mathcal{H}_1$ or $\mathcal{H}_2$ is not collision-resistant, we can conclude that $\mathcal{H}_4$ is also not collision-resistant.

■

■

3 Hard-Concentrate Predicates

Let $f : \{0,1\}^n \to \{0,1\}^n$ be an efficiently computable one-to-one function. Prove that if f has a hard-concentrate predicate[2], then f is one-way.

Solution

1. The following definition comes from Katz & Lindell, 3rd edition, definition 8.4.

   Definition 3.1 (Hard-Concentrate Predicate) *A function* $\mathsf{hc} : \{0,1\}^* \to \{0,1\}$ *is a hard-concentrate predicate of a function* f *if* hc *can be computed in polynomial time, and for every probabilistic polynomial-time adversary* $\mathcal{A}$*, there is a negligible function* negl *such that*

   $$\Pr_{x \leftarrow \{0,1\}^n}[\mathcal{A}(1^n, f(x)) = \mathsf{hc}(x)] \leq \frac{1}{2} + \mathsf{negl}(n)$$

   where the probability is taken over the uniform choice of $x \leftarrow \{0,1\}^n$ *and the randomness of* $\mathcal{A}$*.*

2. We will prove the contrapositive of our goal: that if f is not one-way, then f does not have a hard-concentrate predicate.

3. If f is not one-way, then there is an adversary $\mathcal{A}$ that maps $f(x)$ to a pre-image of $f(x)$ with non-negligible probability.

   $$\Pr_{x \leftarrow \{0,1\}^n}[\mathcal{A}(1^n, f(x)) = x' \text{ such that } f(x) = f(x')] \text{ is } \mathsf{non\text{-}negl}(n)$$

   Since f is one-to-one, there is only one preimage for every output value.

   $$\text{If } f(x) = f(x'), \text{ then } x = x'$$

   So with non-negligible probability, $\mathcal{A}$ maps $f(x)$ to x itself.

   $$\Pr_{x \leftarrow \{0,1\}^n}[\mathcal{A}(1^n, f(x)) = x] \text{ is } \mathsf{non\text{-}negl}(n)$$

4. Now we will show that any function $\mathsf{hc} : \{0,1\}^n \to \{0,1\}$ is not a hard-concentrate predicate for f. To do so, we will construct an adversary $\mathcal{B}$ that correctly guesses $\mathsf{hc}(x)$ with non-negligible advantage.

   <u>$\mathcal{B}$:</u>

   (a) In the hard-concentrate predicate security game, the challenger samples $x \leftarrow \{0,1\}^n$ and gives $(1^n, f(x))$ to $\mathcal{B}$.
   (b) $\mathcal{B}$ runs $\mathcal{A}(1^n, f(x))$, to obtain x'.
   (c) $\mathcal{B}$ checks whether $f(x') = f(x)$.
      - If so, $\mathcal{B}$ computes and outputs $\mathsf{hc}(x')$.[3]

[2] Hard-concentrate predicates are defined in Katz & Lindell, 3rd edition, definition 8.4 under the name *hard-core predicate*.

[3] Note that $\mathcal{B}$ is given a description of hc.

- If not, $\mathcal{B}$ samples $b \leftarrow \{0,1\}$ and outputs it.

5. Analysis: With non-negligible probability, $\mathcal{A}$ outputs $x' = x$. In this case, $\mathcal{B}$'s output is $\mathsf{hc}(x)$, as we desired.

   If $\mathcal{A}$ fails to output x, then $\mathcal{B}$ will find that $f(x') \neq f(x)$, so $\mathcal{B}$ will output a random bit b. Then $\Pr[b = \mathsf{hc}(x)] = \frac{1}{2}$.

   In total, $\mathcal{B}$'s success probability is:

$$\begin{aligned}
\Pr_{x \leftarrow \{0,1\}^n}[\mathcal{B}(1^n, f(x)) = \mathsf{hc}(x)] &= 1 \cdot \Pr_{x \leftarrow \{0,1\}^n}[\mathcal{A}(1^n, f(x)) = x] + \frac{1}{2} \cdot \big(1 - \Pr_{x \leftarrow \{0,1\}^n}[\mathcal{A}(1^n, f(x)) = x]\big) \\
&= \frac{1}{2} + \frac{1}{2} \cdot \Pr_{x \leftarrow \{0,1\}^n}[\mathcal{A}(1^n, f(x)) = x] \\
&= \frac{1}{2} + \mathsf{non\text{-}negl}(n)
\end{aligned}$$

6. In summary, we've shown that if f is not one-way, then f does not have a hard-concentrate predicate. Then the contrapositive is also true: if f has a hard-concentrate predicate, then f is one-way.

■