University of Maryland
CMSC858K — Introduction to Cryptography
Professor Jonathan Katz

Problem Set 3 — Solutions

1. (a) This f' is not (in general) a one-way function. To see this, take $f = g$ (i.e., set them to be the same function). Then f' maps all points to the all-0 string, and is certainly not one-way.

   (b) This f' is not (in general) a one-way function. For example, let g be a one-way function and define f as follows:
   $$f(x_1 \| x_2) = g(x_2) \| 0^n,$$
   where $|x_1| = |x_2| = n$. It is not hard to see that f is one-way (a proof is left as an exercise). On the other hand, f' as defined in the problem maps all inputs to the constant value $g(0^n)\|0^n$, and so is not one-way.

   Interestingly, if f is a one-way *permutation* then f' must be one-way. A proof of this is also left as an exercise.

   (c) This f' **is** one-way. In fact, this holds even if only f is one-way (regardless of g, as long as g is efficiently-computable). To see this, fix a PPT adversary $\mathcal{A}'$ and let
   $$\epsilon(n) \stackrel{\text{def}}{=} \Pr[\mathcal{A}'(f'(x)) \text{ outputs an inverse of } f'(x)],$$
   where the probability is taken over uniform choice of x and the random coins of $\mathcal{A}'$. Consider the following PPT adversary $\mathcal{A}$: given input y_1 (which is equal to $f(x_1)$ for randomly-chosen x_1), choose random x_2, compute $y_2 := g(x_2)$, and run $\mathcal{A}'(y_1\|y_2)$. Then output the first half of the string output by $\mathcal{A}'$. It is not hard to see that (1) the input $y_1\|y_2$ given to $\mathcal{A}'$ is distributed identically to $f'(x_1\|x_2)$ for randomly-chosen x_1, x_2. This implies that $\mathcal{A}'$ inverts its input with probability $\epsilon(n)$. Furthermore, (2) whenever $\mathcal{A}'$ successfully inverts its input, $\mathcal{A}$ successfully inverts its own input. We conclude that $\mathcal{A}$ outputs an inverse of y_1 with probability at least $\epsilon(n)$, showing that ϵ must be negligible.

2. (a) It is immediate that $\mathsf{bit}(1, \cdot)$ is not hard-core for the given function f', so we just prove that f' is one-way. Fix some PPT adversary $\mathcal{A}'$ and let
   $$\epsilon(n) \stackrel{\text{def}}{=} \Pr[\mathcal{A}'(f'(x)) \text{ outputs an inverse of } f'(x)].$$
   Construct the following adversary $\mathcal{A}$:

   > Given input y (which is equal to $f(x)$ for random x), choose a random bit b and run $\mathcal{A}'(y\|b\|1)$ to get x. Output x.

   To analyze the behavior of $\mathcal{A}$, note that $b = \mathsf{bit}(1, x)$ with probability at least 1/2. (It can occur with higher probability if f is not one-to-one.) Furthermore, if $\mathcal{A}'$ outputs an inverse of $y\|b\|1$ then $\mathcal{A}$ correctly inverts its given input y. We conclude that $\mathcal{A}$ outputs an inverse of y with probability at least $\epsilon(n)/2$, and so ϵ must be negligible.

   (b) One possibility is to define $f'(x, i) = f(x) \,\|\, \mathsf{bit}(i, x) \,\|\, i$. Any bit of the input can be guessed with probability at least $1/2 + O(1/n)$ (where $|x| = n$), but it is possible to prove (as in part (a)) that f' is still one-way.

3. We want to prove that for all ppt $\mathcal{A}$, the following is negligible:

$$\epsilon(n) \stackrel{\text{def}}{=} \left| \Pr[\mathcal{A}(G(x_1) \,\|\, G(x_2)) = 1] - \Pr[\mathcal{A}(r_1 \,\|\, r_2) = 1] \right|,$$

where x_1, x_2 are chosen uniformly from $\{0,1\}^n$ and r_1, r_2 are chosen uniformly from $\{0,1\}^{n+1}$. We prove it in two steps.

Claim 1 $\epsilon_1(n) \stackrel{\text{def}}{=} \left| \Pr[\mathcal{A}(G(x_1) \,\|\, G(x_2)) = 1] - \Pr[\mathcal{A}(G(x_1) \,\|\, r_2) = 1] \right|$ *is negligible.*

Construct a PPT adversary $\mathcal{A}'$ as follows: on input y_2, choose random $x_1 \in \{0,1\}^n$ and output whatever $\mathcal{A}(G(x_1) \,\|\, y_2)$ outputs. We have

$$\begin{aligned} \epsilon'(n) &\stackrel{\text{def}}{=} \left| \Pr[\mathcal{A}'(G(x)) = 1] - \Pr[\mathcal{A}'(r) = 1] \right| \\ &= \left| \Pr[\mathcal{A}(G(x_1) \,\|\, G(x)) = 1] - \Pr[\mathcal{A}(G(x_1) \,\|\, r) = 1] \right| \\ &= \epsilon_1(n). \end{aligned}$$

The claim follows by security of G.

Claim 2 $\epsilon_2(n) \stackrel{\text{def}}{=} \left| \Pr[\mathcal{A}(G(x_1) \,\|\, r_2) = 1] - \Pr[\mathcal{A}(r_1 \,\|\, r_2) = 1] \right|$ *is negligible.*

Construct a PPT adversary $\mathcal{A}'$ as follows: on input y_1, choose random $r_2 \in \{0,1\}^{n+1}$ and output whatever $\mathcal{A}(y_1 \,\|\, r_2)$ outputs. We have

$$\begin{aligned} \epsilon'(n) &\stackrel{\text{def}}{=} \left| \Pr[\mathcal{A}'(G(x)) = 1] - \Pr[\mathcal{A}'(r) = 1] \right| \\ &= \left| \Pr[\mathcal{A}(G(x) \,\|\, r_2) = 1] - \Pr[\mathcal{A}(r \,\|\, r_2) = 1] \right| \\ &= \epsilon_2(n). \end{aligned}$$

The claim follows by security of G.

Finally, we have

$$\begin{aligned} \epsilon(n) &= \Big| \Pr[\mathcal{A}(G(x_1) \,\|\, G(x_2)) = 1] - \Pr[\mathcal{A}(G(x_1) \,\|\, r_2) = 1] \\ &\quad + \Pr[\mathcal{A}(G(x_1) \,\|\, r_2) = 1] - \Pr[\mathcal{A}(r_1 \,\|\, r_2) = 1] \Big| \\ &\le \left| \Pr[\mathcal{A}(G(x_1) \,\|\, G(x_2)) = 1] - \Pr[\mathcal{A}(G(x_1) \,\|\, r_2) = 1] \right| \\ &\quad + \left| \Pr[\mathcal{A}(G(x_1) \,\|\, r_2) = 1] - \Pr[\mathcal{A}(r_1 \,\|\, r_2) = 1] \right| \\ &= \epsilon_1(n) + \epsilon_2(n). \end{aligned}$$

Since ϵ_1, ϵ_2 are negligible, this completes the proof.

4. Fix a PPT algorithm $\mathcal{A}$, and let

$$\epsilon(n) \stackrel{\text{def}}{=} \Pr[\mathcal{A}(G(x)) \text{ inverts } G(x)].$$

Consider the following PPT distinguisher $\mathcal{A}'$: given input $y \in \{0,1\}^{2n}$, run $\mathcal{A}(y)$ to obtain output x. If $G(x) = y$ output 1; otherwise, output 0.

Almost by definition, we have $\Pr[\mathcal{A}'(G(x)) = 1] = \epsilon(n)$. On the other hand

$$\Pr[\mathcal{A}'(r) = 1] \leq \Pr[\exists x \text{ such that } G(x) = r].$$

Since $G(x)$ takes on at most 2^n values, the latter probability is at most $2^n/2^{2n} = 2^{-n}$. Taken together, we have

$$\left|\Pr[\mathcal{A}'(G(x)) = 1] - \Pr[\mathcal{A}'(r) = 1]\right| \geq \epsilon(n) - 2^{-n}\,;$$

since G is a pseudorandom generator, we conclude that ϵ must be negligible.

Interestingly, it is possible to prove that a PRG G is a one-way function even if it only expands by a single bit, though the proof is a bit more challenging.