0368.4162 Foundations of Cryptography

Solution 2

Eliran Kachlon

Question 1

We are given an ensemble of functions $\{f_n : \{0,1\}^{\ell(n)} \to \{0,1\}^{\ell'(n)}\}_{n\in\mathbb{N}}$ that is one-way, where ℓ, ℓ' are two polynomials. Then we know that there exists a poly-time algorithm $f(1^n, x)$ that given 1^n and $x \in \{0,1\}^{\ell(n)}$ outputs $f_n(x)$. Assume without loss of generality that $\ell(\cdot)$, $\ell'(\cdot)$ are strictly monotone, and note that given some $n \in \mathbb{N}$, we may find in time $poly(n)$ the maximal $k \in \mathbb{N}$ for which $\ell(k) \leq n$, by simply going through all $k \leq poly(n)$ (for some polynomial that depends on ℓ), computing $\ell(k)$, thus finding the maximal k (if we want to be . (fancy– since ℓ is monotone we may even use binary search, thus improving the running time

Now, we construct a function $g : \{0,1\}^* \to \{0,1\}^*$ in the following way: given input x of length n, let $k =$ he prefix oft$k(n)$ be the largest integer for which $\ell(k) \leq n$. Then g outputs $f(1^k, x_{[1,\ldots,\ell(k)]})$ where $x_{[1,\ldots,\ell(k)]}$ is . length $\ell(k)$ of x

First, note that the function is well defined: We take k to be the largest integer for which $\ell(k) \leq n$, so the eprefix of length $\ell(k)$ does exists. Secondly, g may be computed in time $poly(n)$ since we've shown that k may b computed in time $poly(n)$, and we know that f is computed in time $poly(k)$, which is bounded from above by . $poly(n)$, as required. We claim that g is a OWF

Assume towards contradiction that there exists some non-uniform PPT adversary $\mathcal{A} = \{\mathcal{A}_n\}_{n\in\mathbb{N}}$, s.t. for infinitely many n's it holds that

$$\Pr_{x\leftarrow\{0,1\}^n}[\mathcal{A}_n(g(x)) \in g^{-1}(g(x))] > 1/p(n)$$

for some polynomial $p(\cdot)$. Note that since the interval $[\ell(k), \ell(k+1)) = \{\ell(k), \ell(k)+1, \ldots, \ell(k+1)-1\}$ is finite for every k, it means that for infinitely many such intervals, there exists some n in the interval for which $\Pr_{x\leftarrow\{0,1\}^n}[\mathcal{A}_n(g(x)) \in g^{-1}(g(x))] > 1/p(n)$, and we will use this fact to invert f on infinitely many (legal) input lengths.

Fix some k for which there exists some n in the interval $[\ell(k), \ell(k+1))$ s.t. $\Pr_{x\leftarrow\{0,1\}^n}[\mathcal{A}_n(g(x)) \in g^{-1}(g(x))] > 1/p(n)$. We denote this n by n^*. We will show how to invert f on inputs of length $\ell(k)$. We construct adversary $\mathcal{B}$: on input y of length $\ell(k)$, for every $n \in [\ell(k), \ell(k+1))$, $\mathcal{B}$ computes $x_n := \mathcal{A}_n(y)$. If there exists some n for which $y = g(x_n)$ then $\mathcal{B}$ returns the $\ell(k)$-length prefix x_n. Otherwise it returns $0^{\ell(k)}$. Note that $\mathcal{B}$ is PPT.

Note that by the definition of g, for every $n \in [\ell(k), \ell(k+1))$, the distributions $g(U_n)$ and $f(1^k, U_{\ell(k)})$ are identically distributed, since $g(U_n) = f(1^k, (U_n)_{[1,\ldots,\ell(k)]}) = f(1^k, U_{\ell(k)})$. Note that any x_n for which $y = g(x_n)$ can be transformed to an x of length $\ell(k)$ s.t. $f(1^k, x) = y$ by simply taking the prefix of length $\ell(k)$ of x_n (this follows because $y = g(x_n) = f(1^k, (x_n)_{[1,\ldots,\ell(k)]})$). Moreover, we know that for n^* we have $\Pr_{x\leftarrow\{0,1\}^{n^*}}[\mathcal{A}_{n^*}(g(x)) \in g^{-1}(g(x))] > 1/p(n^*)$, and since $\mathcal{B}$ tries every $n \in [\ell(k), \ell(k+1))$, it also tries n^*, so the success probability of $\mathcal{B}$ is at least the success probability of $\mathcal{A}_{n^*}$. From all the above, we conclude that

$$\Pr_{x\leftarrow\{0,1\}^{\ell(k)}}[\mathcal{B}(f(1^k, x)) \in f^{-1}(f(1^k, x))] \geq \Pr_{x\leftarrow\{0,1\}^{n^*}}[\mathcal{A}_{n^*}(g(x)) \in g^{-1}(g(x))] > 1/p(n^*) > 1/p(\ell(k+1)) = 1/poly(k)$$

in contradiction to the security of f. Thus g must be a OWF, as required.

A few comments about the proof above:

- We may assume that $\ell(n), \ell'(n)$ are strictly monotone since for every polynomial $p(\cdot)$ there exists some n_0 s.t. for every $n > n_0$ we have $p(n+1) > p(n)$.
- In the definition of adversary $\mathcal{B}$, the proof above uses the (implicit) simplifying assumption that $\mathcal{A}_n$ returns x_n of length $\geq \ell(k(n))$. In general, $\mathcal{B}$ should return the $\ell(k(|x|))$-length prefix of x, where the rest of the analysis follows similarly.

Question 2

1. Assume towards contradiction that there exists a non-uniform PPT $\mathcal{A} = \{\mathcal{A}_n\}_{n\in\mathbb{N}}$ such that for infinitely many n's it holds that

$$\Pr_{x\leftarrow\{0,1\}^n}[\mathcal{A}(G(x)) \in G^{-1}(G(x))] > 1/p(n)$$

for some polynomial $p(\cdot)$. Fix such n. We define the following adversary $\mathcal{B}$: given input y, $\mathcal{B}$ computes $z := \mathcal{A}(y)$. If $G(z) = y$ it returns 1, and 0 otherwise. Note that $\mathcal{B}$ is a non-uniform PPT. We have:

$$\begin{aligned}\Pr[\mathcal{B}(G(U_n)) = 1] &= \Pr[G(\mathcal{A}(G(U_n))) = G(U_n)] \\ &= \Pr[\mathcal{A}(G(U_n)) \in G^{-1}(G(U_n))] \\ &> 1/p(n) \ .\end{aligned}$$

On the other hand we have:

$$\begin{aligned}\Pr[\mathcal{B}(U_{n+\ell}) = 1] &= \Pr[\mathcal{A}(U_{n+\ell}) \in G^{-1}(U_{n+\ell})] \\ &\leq \Pr[\exists x \in \{0,1\}^{n+\ell} \ \ s.t. \ G(x) = U_{n+\ell}] \\ &\leq 2^n/2^{n+\ell} \\ &= neg(n) \ ,\end{aligned}$$

where the first inequality follows since that claim "$\exists x \in \{0,1\}^{n+\ell}$ $s.t.$ $G(x) = U_{n+\ell}$" follows from the claim "$\mathcal{A}(U_{n+\ell}) \in G^{-1}(U_{n+\ell})$", the second inequality follows since $U_{n+\ell}$ is uniformly distributed over $\{0,1\}^{n+\ell}$ and the size of the image of G is at most 2^n , and the last equality follows since $\ell = \omega(\log n)$. In conclusion, we have

$$\Pr[\mathcal{B}(G(U_n)) = 1] - \Pr[\mathcal{B}(U_{n+\ell}) = 1] > 1/p(n) - neg(n) > 1/q(n)$$

for some polynomial $q(\cdot)$ (since inverse-polynomial minus negligible is still lower-bounded by inverse polynomial), in contradiction to the security of G as a PRG. Thus G is also a OWF.

2. Assume towards contradiction that there exists a non-uniform PPT $\mathcal{A} = \{\mathcal{A}_n\}_{n\in\mathbb{N}}$ such that for infinitely many n's it holds that

$$\Pr_{x\leftarrow\{0,1\}^n}[\mathcal{A}(G(x)) \in G^{-1}(G(x))] = \epsilon(n)$$

where $\epsilon(\cdot)$ is some non-negligible function (i.e. there exists some polynomial $p(\cdot)$ such that $\epsilon(n) > p(n)$ for every n) and assume *wlog* that $\mathcal{A}$ is deterministic (we've seen the justification for this assumption in the last HW). Fix such n. We define the following adversary $\mathcal{B}$: given input y, $\mathcal{B}$ computes $z := \mathcal{A}(y)$. If $G(z) = y$ it returns 1, and 0 otherwise. Note that $\mathcal{B}$ is a non-uniform PPT. We have:

$$\begin{aligned}\Pr[\mathcal{B}(G(U_n)) = 1] &= \Pr[G(\mathcal{A}(G(U_n))) = G(U_n)] \\ &= \Pr[\mathcal{A}(G(U_n)) \in G^{-1}(G(U_n))] \\ &= \epsilon(n) \ .\end{aligned}$$

It remains to bound $\Pr[\mathcal{B}(U_{n+1}) = 1]$. Note that the assumption $\Pr_{x \leftarrow \{0,1\}^n}[\mathcal{A}(G(x)) \in G^{-1}(G(x))] = \epsilon(n)$ means that the fraction of $x \in \{0,1\}^n$ such that $\mathcal{A}(G(x)) \in G^{-1}(G(x))$ is exactly $\epsilon(n)$ (remember that $\mathcal{A}$ is deterministic). Thus the largest number of elements in $\{0,1\}^{n+1}$ that $\mathcal{A}$ successfully inverts is $\epsilon(n) \cdot 2^n$, since this is the largest number of elements which are images of the set of x's on which $\mathcal{A}$ can invert G and there is exactly one element in the image of G for each x for which $\mathcal{A}$ can invert $G(x)$. Moreover, for every element in $\{0,1\}^{n+1}$ that is not in the image of G, $\mathcal{A}$ necessarily fails. We conclude that

$$\begin{aligned}
\Pr[\mathcal{B}(U_{n+1}) = 1] &= \Pr[G(\mathcal{A}(U_{n+1})) = U_{n+1}] \\
&= \Pr[\mathcal{A}(U_{n+1}) \in G^{-1}(U_{n+1})] \\
&\leq \frac{\epsilon(n) \cdot 2^n}{2^{n+1}} \\
&= \epsilon(n)/2 \ .
\end{aligned}$$

Thus we have

$$\Pr[\mathcal{B}(G(U_n)) = 1] - \Pr[\mathcal{B}(U_{n+1})) = 1] \geq \epsilon(n) - \epsilon(n)/2 = \epsilon(n)/2 > \frac{1}{2 \cdot p(n)} \ ,$$

in contradiction.

Question 3

1. The general idea is that given some OWF f, we want to construct a function g, such that on about $(1-1/n)$ fraction of the inputs, g will behave like the identity function, and for all other inputs, g will be hard to invert, using f hardness. We may assume *wlog* that f is length preserving (i.e $|f(x)| = |x|$).[1]

   Since g cannot use randomization, the behaviour of g must depend only on the input. We construct g, so that on input (x, z), where $|x| = |z|$, the function g outputs $(x, f(z))$ if the $\lfloor \log(|x|) \rfloor$-long prefix is the zero string, and otherwise it returns (x, z). Note that g is not OWF, because for inputs of length n, on fraction $1 - 2^{-\lfloor \log(n) \rfloor} \geq 1 - 2^{-\log(n)+1} = 1 - 2/n$, g simply functions as the identity function, hence any adversary which given y outputs y inverts it with probability at least $1 - 2/n$. We claim that g is weakly-OWF.

   Assume that there exists some non-uniform PPT adversary $\mathcal{A}$ s.t.

$$\Pr_{x \leftarrow \{0,1\}^n}[\mathcal{A}(g(x)) \in g^{-1}(g(x))] = \epsilon(n).$$

   We construct an adversary $\mathcal{B}$ which inverts f. $\mathcal{B}$ on input y of length n samples $u \quad \{0,1\}^{n-\lfloor \log(n) \rfloor}$, samples $(x, z) \quad \mathcal{A}(0^{\lfloor \log(|x|) \rfloor}, u, y)$ and outputs z. Note that if $\mathcal{A}$ succeeds in inverting $(0^{\lfloor \log(|x|) \rfloor}, u, y)$ for g, then $\mathcal{B}$ succeeds in inverting y. Hence we have:

$$\Pr_{z \leftarrow \{0,1\}^n}[\mathcal{B}(f(z)) \in f^{-1}(f(z))] \geq \Pr_{\substack{z \leftarrow \{0,1\}^n \\ u \leftarrow \{0,1\}^{n-\lfloor \log(n) \rfloor}}}[\mathcal{A}(0^{\lfloor \log(n) \rfloor}, u, z) \in g^{-1}(g(0^{\lfloor \log(n) \rfloor}, u, z))]$$

   Denote by $Pre(x)$ the prefix of length $\lfloor \log(|x|) \rfloor$ of x. We get

$$\begin{aligned}
\epsilon(n) &= \Pr_{\substack{z \leftarrow \{0,1\}^n \\ x \leftarrow \{0,1\}^n}}[\mathcal{A}(g(x,z)) \in g^{-1}(g(x,z))] \\
&= \Pr_{\substack{z \leftarrow \{0,1\}^n \\ x \leftarrow \{0,1\}^n}}[\mathcal{A}(g(x,z)) \in g^{-1}(g(x,z)) | Pre(x) = 0^{\lfloor \log(n) \rfloor}] \Pr[Pre(x) = 0^{\lfloor \log(n) \rfloor}] \\
&\quad + \Pr_{\substack{z \leftarrow \{0,1\}^n \\ x \leftarrow \{0,1\}^n}}[\mathcal{A}(g(x,z)) \in g^{-1}(g(x,z)) | Pre(x) \neq 0^{\lfloor \log(n) \rfloor}] \Pr[Pre(x) \neq 0^{\lfloor \log(n) \rfloor}] \\
&= \Pr_{\substack{z \leftarrow \{0,1\}^n \\ u \leftarrow \{0,1\}^{n-\lfloor \log(n) \rfloor}}}[\mathcal{A}(0^{\lfloor \log(n) \rfloor}, u, z) \in g^{-1}(g(0^{\lfloor \log(n) \rfloor}, u, z))] \cdot 2^{-\lfloor \log(n) \rfloor} \\
&\quad + \Pr_{\substack{z \leftarrow \{0,1\}^n \\ x \leftarrow \{0,1\}^n}}[\mathcal{A}(g(x,z)) \in g^{-1}(g(x,z)) | Pre(x) \neq 0^{\lfloor \log(n) \rfloor}] \cdot (1 - 2^{-\lfloor \log(n) \rfloor}) \\
&\leq \Pr_{z \leftarrow \{0,1\}^n}[\mathcal{B}(f(z)) \in f^{-1}(f(z))] \cdot \frac{2}{n} + 1 \cdot (1 - \frac{1}{n})
\end{aligned}$$

[1] From any OWF f we may construct a OWF h which is length-preserving in the following way: since f is poly-time computable, there exist some polynomial $p(\cdot)$ that bounds the length expansion of f. Hence the function $f'(x) = x \parallel 0^{p(|x|)-|f(x)|}$ has the property that $|f'(x)| = |f'(y)|$ for every $|x| = |y|$. By defining $h(x, z) := f'(x)$ for every x, z s.t. $|x, z| = p(|x|)$, we get a OWF. Security follows by a similar argument to that given in Question 1.

So we get

$$\Pr_{z\leftarrow\{0,1\}^n}[\mathcal{B}(f(z)) \in f^{-1}(f(z))] \geq \frac{1}{2} - \frac{n}{2}(1-\epsilon).$$

Since f is OWF, the RHS must be negligible, i.e. there exists some negligible function $\mu(n)$ s.t.

$$\frac{1}{2} - \frac{n}{2}(1-\epsilon) \leq \mu(n).$$

So we get $\epsilon \leq 1 - (\frac{1}{n} - \frac{2\mu(n)}{n}) = 1 - \Omega(1/n)$, as required.

2. (a) Fix some n. Denote by X the random variable $(x_1, \ldots, x_t)$, where each x_i is chosen uniformly from $\{0,1\}^n$. We know that $\Pr[\mathcal{A} \text{ inverts } f^{\otimes t}(x_1, \ldots, x_t)] = \epsilon(n)$. Assume towards contradiction that $\forall i \in [t]$, $\Pr_{x\leftarrow\{0,1\}^n}[x \in G_i] < 1 - \log(2/\epsilon)/t$. We have:

$$\begin{aligned}\Pr_X[\mathcal{A} \text{ inverts } f^{\otimes t}(x_1, \ldots, x_t)] = &\Pr_X[\mathcal{A} \text{ inverts } f^{\otimes t}(x_1, \ldots, x_t) \wedge \exists i \in [t] \text{ s.t. } x_i \notin G_i] \\ &+ \Pr_X[\mathcal{A} \text{ inverts } f^{\otimes t}(x_1, \ldots, x_t) \wedge \forall i \in [t], x_i \in G_i].\end{aligned}$$

We shall bound this two terms in order to get a contradiction. For the first term we have:

$$\begin{aligned}\Pr_X[\mathcal{A} \text{ inverts } f^{\otimes t}(x_1, \ldots, x_t) \wedge \exists i \in [t] \text{ s.t. } x_i \notin G_i] &\leq \sum_{1\leq i\leq t} \Pr_X[\mathcal{A} \text{ inverts } f^{\otimes t}(x_1, \ldots, x_t) \wedge x_i \notin G_i] \\ &\leq \sum_{1\leq i\leq t} \sum_{x\notin G_i} \Pr_X[\mathcal{A} \text{ inverts } f^{\otimes t}(x_1, \ldots, x_t) \wedge x_i = x] \\ &= \sum_{1\leq i\leq t} \sum_{x\notin G_i} \Pr_X[\mathcal{A} \text{ inverts } f^{\otimes t}(x_1, \ldots, x_t) | x_i = x] \Pr[x_i = x] \\ &< \sum_{1\leq i\leq t} (\frac{\epsilon}{2t}) \\ &= \epsilon/2,\end{aligned}$$

where the first inequality is due to union bound. For the second term we have

$$\begin{aligned}\Pr_X[\mathcal{A} \text{ inverts } f^{\otimes t}(x_1, \ldots, x_t) \wedge \forall i \in [t], x_i \in G_i] &\leq \Pr[\forall i \in [t], x_i \in G_i] \\ &< (1 - \frac{\log(2/\epsilon)}{t})^t \\ &\leq e^{-\log(2/\epsilon)} \\ &< 2^{-\log(2/\epsilon)} \\ &= \epsilon/2,\end{aligned}$$

where the third inequality is due to the known inequality $(1 - 1/x)^x \leq e^{-1}$. In conclusion, we have

$$\begin{aligned}\epsilon &= \Pr_X[\mathcal{A} \text{ inverts } f^{\otimes t}(x_1, \ldots, x_t)] \\ &= \Pr_X[\mathcal{A} \text{ inverts } f^{\otimes t}(x_1, \ldots, x_t) \wedge \exists i \in [t] \text{ s.t. } x_i \notin G_i] + \Pr_X[\mathcal{A} \text{ inverts } f^{\otimes t}(x_1, \ldots, x_t) \wedge \forall i \in [t], x_i \in G_i] \\ &< \epsilon/2 + \epsilon/2 = \epsilon,\end{aligned}$$

in contradiction. Thus there must exist some $i \in [t]$ s.t. $\Pr_{x\leftarrow\{0,1\}^n}[x \in G_i] \geq 1 - \log(2/\epsilon)/t$.

(b) We define the following adversary A': on input y, repeat for $2tn/\epsilon$ times:

- Sample $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_t$ in uniform from $(\{0,1\}^n)^t$.

- Compute $f(x_1), \ldots, f(x_{i-1}), f(x_{i+1}), \ldots, f(x_t)$.
- Sample $(w_1, \ldots, w_t) \quad \mathcal{A}(f(x_1), \ldots, f(x_{i-1}), y, f(x_{i+1}), \ldots, f(x_t))$.
- If $f(w_i) = y$ return w_i, otherwise repeat the above procedure.
- If failed in all iterations, return 0^n.

Now, given that $y = f(x)$ for some $x \in G_i$, the failure probability is bounded by the probability to fail in all iterations. Hence,

$$\Pr[\mathcal{A}'(f(x)) \notin f^{-1}(f(x))|x \in G_i] < (1 - \frac{\epsilon}{2t})^{2nt/\epsilon} \leq e^{-n} \leq 2^{-n}.$$

Hence we can bound the success probability by:

$$\begin{aligned}
\Pr[\mathcal{A}'(f(x)) \in f^{-1}(f(x))] &= \Pr[\mathcal{A}'(f(x)) \in f^{-1}(f(x))|x \in G_i] \cdot \Pr[x \in G_i] \\
&\quad + \Pr[\mathcal{A}'(f(x)) \in f^{-1}(f(x))|x \notin G_i] \cdot \Pr[x \notin G_i] \\
&\geq \Pr[\mathcal{A}'(f(x)) \in f^{-1}(f(x))|x \in G_i] \cdot \Pr[x \in G_i] \\
&\geq (1 - 2^{-n})(1 - \frac{\log(2/\epsilon)}{t}) \\
&= 1 - \frac{\log(2/\epsilon)}{t} - 2^{-n} + 2^{-n}\frac{\log(2/\epsilon)}{t} \\
&\geq 1 - \frac{\log(2/\epsilon)}{t} - 2^{-n}
\end{aligned}$$

and we got the required success probability. Note that $\mathcal{A}'$ has $2tn/\epsilon$ iterations, in each it invokes $\mathcal{A}$, thus its running time is $O(time(\mathcal{A})2tn/\epsilon)$, where the sampling time is included in the running time of $\mathcal{A}$.

(c) Assume towards contradiction that $\epsilon(n) = n^{-O(1)}$. So the running time of $\mathcal{A}'$ is

$$O(time(\mathcal{A})2tn/\epsilon) = O(time(\mathcal{A})\frac{\log^2(n)p(n)n}{n^{-O(1)}})$$

which is polynomial time since $\mathcal{A}$ is PPT. Moreover, the probability that $\mathcal{A}'$ inverts f is at least

$$1 - \frac{\log(2/\epsilon)}{t} - 2^{-n} = 1 - O(\frac{1}{\log(n)p(n)}) - 2^{-n} = 1 - o(1/p(n))$$

Thus we got a non-uniform PPT adversary which inverts f with probability $> 1 - o(1/p(n))$ for infinitely many n's in contradiction. Thus it must be the case that $\epsilon(n) = n^{-\omega(1)}$.