Cryptography, winter term 16/17: Sample solution to assignment 4

Cornelius Brand, Marc Roth

Exercise 4.1 (Be nice to your tutors and TAs, 1 Bonus Point) Write the name and matriculation number of every author as well as number, time slot and the name of the tutor of your tutorial group on the first page of your solution. Furthermore write in english and staple *all* sheets of your solution.

Exercise 4.2 (CTR mode, 4 Points) Let $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be the CTR mode encryption scheme and let $\widetilde{\Pi} = (\widetilde{\mathsf{Gen}}, \widetilde{\mathsf{Enc}}, \widetilde{\mathsf{Dec}})$ be the encryption scheme obtained from Π by using a truly random function f instead of a pseudorandom function F_k.[1] Show that there is a negligible function negl, such that for any PPT adversary $\mathcal{A}$, it holds that

$$\left|\Pr[\mathsf{PrivK}^{\mathsf{cpa}}_{\mathcal{A},\Pi}(n) = 1] - \Pr[\mathsf{PrivK}^{\mathsf{cpa}}_{\mathcal{A},\widetilde{\Pi}}(n) = 1]\right| \leq \mathsf{negl}(n)$$

Solution 4.2 (CTR mode, 4 Points) The proof goes by reduction (or in plain terms, by contraposition), along the lines of the argument for the corresponding statement in the proof of Theorem 3.31: Assume the statement from the exercise were false. That is, for some PPT adversary $\mathcal{A}$,

$$\left|\Pr[\mathsf{PrivK}^{\mathsf{cpa}}_{\mathcal{A},\Pi}(n) = 1] - \Pr[\mathsf{PrivK}^{\mathsf{cpa}}_{\mathcal{A},\widetilde{\Pi}}(n) = 1]\right| > t(n) \tag{1}$$

for some non-negligible $t(n)$. We want to show how to construct a PPT distinguisher D contradicting the requirement from the definition of pseudorandom functions, i.e. it should hold that

$$\left|\Pr[D^{F_k(\cdot)}(1^n) - \Pr[D^{f(\cdot)}(1^n)]\right| > t(n)\,. \tag{2}$$

We define the working of the distinguisher D as follows, where D is given access to some oracle $\mathcal{O} : \{0,1\}^n \to \{0,1\}^n$ and receives an input 1^n.

(1) Run $\mathcal{A}(1^n)$, and when $\mathcal{A}(1^n)$ queries its oracle to the encryption function for the i-th time with a message made up from ℓ_i message blocks $m_1, \ldots, m_{\ell_i}$, do the following:
   (a) Choose a uniform initial value $\mathsf{ctr}_i \in \{0,1\}^m$
   (b) Query $\mathcal{O}$ for $j = 1, \ldots, \ell_i$ to obtain $y_j := \mathcal{O}(\mathsf{ctr}_i + j)$
   (c) Return the ciphertext blocks $\langle \mathsf{ctr}_i, c_1, \ldots, c_{\ell_i} \rangle := \langle \mathsf{ctr}_i, y_1 \oplus m_1, \ldots, y_{\ell_i} \oplus m_{\ell_i} \rangle$ to $\mathcal{A}$

(2) Once $\mathcal{A}$ outputs the messages m_0, m_1 consisting of ℓ^* blocks $m_{0,1}, \ldots, m_{0,\ell^*}, m_{1,1}, \ldots, m_{1,\ell^*}$, respectively, choose a uniform bit $b \in \{0,1\}$ and do the following:

[1] That is to say that $\widetilde{\mathsf{Gen}}$ picks uniformly $f \in \mathsf{Func}_n$ where Gen picks uniformly $k \in \{0,1\}^n$, and $\widetilde{\mathsf{Enc}}$ uses f where Enc uses F_k.

(a) Choose a uniform initial value $\mathsf{ctr}^* \in \{0,1\}^m$

(b) Query $\mathcal{O}$ for $j = 1, \ldots, \ell^*$ to obtain $y_j^* := \mathcal{O}(\mathsf{ctr}^* + j)$ Return the challenge ciphertext blocks $\langle \mathsf{ctr}^*, c_1^*, \ldots, c_{\ell^*}^* \rangle := \langle \mathsf{ctr}^*, y_1 \oplus m_{b,1}, \ldots, y_{\ell^*} \oplus m_{b,\ell^*} \rangle$ to $\mathcal{A}$

(3) Answer queries to the encryption oracle as above, until $\mathcal{A}$ produces an output bit b'. Then, output 1 if $b = b'$, and 0 otherwise

We first argue that D is PPT: Each of the above steps clearly only incurs polynomial overhead for each of the oracle calls from $\mathcal{A}$, and as a PPT adversary, $\mathcal{A}$ may only pose a polynomial number of queries to the encryption oracle and may itself only run in polynomial time, hence D also runs in polynomial time.

As to why D is in fact a distinguisher, note that D is essentially just the experiment $\mathsf{PrivK}^{\mathsf{cpa}}_{\mathcal{A},\Pi}$ or $\mathsf{PrivK}^{\mathsf{cpa}}_{\mathcal{A},\widetilde{\Pi}}$—depending on which oracle D is given—implemented as an algorithm, where the oracle queries of $\mathcal{A}$ are spelled out step-by-step, with the first step of key generation in the experiment being simulated by uniformly choosing $k \in \{0,1\}^n$ or $f \in \mathsf{Func}_n$, respectively. This is equivalent since this is also how Gen and $\widetilde{\mathsf{Gen}}$ generate the keys by definition of CTR. Therefore, by definition of Π and $\widetilde{\Pi}$, we have that $D^{F_k(\cdot)}(1^n)$ and $\mathsf{PrivK}^{\mathsf{cpa}}_{\mathcal{A},\Pi}(n)$ are identically distributed, and $D^{f(\cdot)}(1^n)$ and $\mathsf{PrivK}^{\mathsf{cpa}}_{\mathcal{A},\widetilde{\Pi}}(n)$ are identically distributed, for uniformly chosen f. In other words,

$$\Pr_{k \leftarrow \{0,1\}^n}\left[D^{F_k(\cdot)}(1^n) = 1\right] = \Pr\left[\mathsf{PrivK}^{\mathsf{cpa}}_{\mathcal{A},\Pi}(n) = 1\right],$$
$$\Pr_{f \leftarrow \mathsf{Func}_n}\left[D^{f(\cdot)}(1^n) = 1\right] = \Pr\left[\mathsf{PrivK}^{\mathsf{cpa}}_{\mathcal{A},\widetilde{Pi}}(n) = 1\right].$$

Thus, (2) follows directly from (1).

Exercise 4.3 (Indistinguishability and CPA-security, 8 Points) Let F be a PRF and G be a PRG with expansion factor $n \mapsto n+1$. For each of the following encryption schemes, decide whether the scheme has indistinguishable encryptions in the presence of an eavesdropper and whether it is CPA-secure. (In each case, the shared key is a uniform $k \in \{0,1\}^n$.) Explain your answer.

(a) To encrypt $m \in \{0,1\}^n$, output the ciphertext $m \oplus F_k(0^n)$.

(b) The one-time pad.

(c) To encrypt $m \in \{0,1\}^{n+1}$, choose uniform $r \in \{0,1\}^n$ and output the ciphertext $\langle r, G(r) \oplus m \rangle$.

(d) To encrypt $m \in \{0,1\}^{2n}$, parse m as $m_1 || m_2$ with $|m_1| = |m_2| = n$, then choose uniform $r \in \{0,1\}^n$ and output the ciphertext $\langle r, m_1 \oplus F_k(r), m_2 \oplus F_k(r+1) \rangle$.

Solution 4.3 (Indistinguishability and CPA-security, 8 Points)

(a) This scheme has indistinguishable encryptions in the presence of an eavesdropper. To see this, we observe that $F_k(0^n)$ is pseudorandom. The formal proof is similar to the proof of Theorem 3.18. The scheme is not CPA-secure as encryption is deterministic.

(b) The one-time pad is perfectly indistinguishable which immediately implies that it has also indistinguishable encryptions in the presence of an eavesdropper. It is not CPA-secure as encryption is deterministic.

(c) This scheme does not even have indistinguishable encryptions in the presence of an eavesdropper, as decryption can be done in polynomial time without knowing the key: On input $(r, G(r) \oplus m)$, we just compute $G(r)$ and then output $G(r) \oplus (G(r) \oplus m) = m$. Therefore, it can not be CPA-secure as well.

(d) This scheme is CPA-secure. The proof is similar to the correctness proof of CTR mode with just two blocks. Therefore it has indistinguishable encryptions in the presence of an eavesdropper as well.

Exercise 4.4 (Birthday paradox, 4 Points) Let $k \leq n$. You are in a room with k people, and everyone is born on one out of n possibly dates, i.e., the typical terrestrial case asks for $n = 365$ (ignoring leap years). What is, for arbitrary n and k, the probability that at least two people have the same birthday, assuming that the dates of birth are uniformly distributed and independent of each other?

Additionally, for $n = 365$, what is the smallest number of people k such that this probability is at least $\frac{1}{2}$?

Solution 4.4 (Birthday paradox, 4 Points) We consider the complementary event, namely that each birthday is unique. Formally, we are dealing with a sample space Ω of size $|n^k|$, where a single outcome $f \in \Omega$ is a mapping $\{1, \ldots, k\} \to \{1, \ldots, n\}$, assigning to each of the k individuals one out of the n birthdays. Let U be the set of such mappings that are injective, so we want to calculate $1 - \Pr[U]$. Since the outcomes are assumed to be uniformly distributed, this is just $1 - \frac{|U|}{|\Omega|} = 1 - \frac{|U|}{n^k}$. To compute $|U|$, notice that each $f \in U$ is uniquely determined by its image $f([k]) \subseteq [n]$ and a permutation of $[k]$ (i.e., a bijection between $f([k])$ and $[k]$), and each such pair of a subset of $[n]$ and a permutation of $[k]$ determines a unique $f \in U$. There are $\binom{n}{k}$ subsets and $k!$ bijections, making $|U| = \binom{n}{k} \cdot k! = \frac{n!}{(n-k)!}$, which yields $1 - \Pr[U] = 1 - \frac{n!}{n^k(n-k)!}$.

As for the case of $n = 365$, a quick calculation shows that $1 - \frac{364}{365} \cdot \frac{363}{365} \cdots \frac{344}{365} < \frac{1}{2} < 1 - \frac{364}{365} \cdot \frac{363}{365} \cdots \frac{343}{365}$, so the first k for which this happens is $k = 23$.