ETH Zürich, D-INFK
Spring 2018

Prof. Ueli Maurer
Christian Badertscher
Fabio Banfi

Cryptography Foundations

Exercise 13

13.1 Key-Agreement Using a Trapdoor One-Way Permutation in the ROM

Goal*: We prove the security of the key-agreement protocol based on a trapdoor one-way permutation that is described in the lecture notes in the random oracle model.*

Let the resource $\mathbf{R} := [\bullet\!\!\longrightarrow, \longleftarrow\!\!\bullet, \mathbf{PO}_k]$ and let further $\mathbf{S} := \overset{\{0,1\}^k}{\bullet\!\!=\!\!=\!\!\bullet}$ be a shared secret key as described in the lecture notes. Let $\boldsymbol{\pi} = (\pi_1, \pi_2)$ be the protocol described in Section 7.3.2 of the lecture notes. Further let $\mathbf{G}$ be the inversion game for the underlying trapdoor one-way permutation. Provide a simulator σ and an explicit reduction ρ such that $\overline{\langle \pi_1^A \pi_2^B \mathbf{R} \mid \sigma^E \mathbf{S} \rangle} \leq \overline{\mathbf{G}}\rho$.

13.2 Entropies and Information-Theoretical Key Agreement

Goal*: This task exemplifies the use of entropy diagrams for information-theoretical statements. We then use this to show (basic) statements about the security of key agreement protocols.*

a) Let A_1, A_2, and A_3 be independent and uniformly distributed random variables on $\{0, 1\}$. Draw the entropy diagram for the random variables

$$X := [A_1, A_2], \quad Y := [A_2, A_3], \quad \text{and} \quad Z := [A_1, A_3].$$

b) To finish the proof of Theorem 7.3, show that deletion of information by A does not increase the conditional mutual information of A and B under the knowledge of E.

Hint: Use entropy diagrams.

c) Let $B_1, \ldots, B_6$ be independent and uniformly distributed random bits. We define the random variables X, Y, and Z as follows:

$$X := [B_1 \oplus B_6, B_2, B_3, B_4, B_5, B_6],\ Y := [B_3, B_4, B_5 \oplus B_6], \text{ and } Z := [B_2 \oplus B_4, B_3, B_6].$$

We assume that Alice, Bob, and Eve obtain the values X, Y, and Z, respectively. Furthermore, Alice and Bob are connected via authenticated channels $\bullet\!\!\longrightarrow$ and $\longleftarrow\!\!\bullet$.

What is the maximal length of a shared key that can be generated by Alice and Bob in this setting if Eve is required to have no information about the key? How can Alice and Bob generate a key of this length?

13.3 Privacy Amplification

Goal*: Explore an application of privacy amplification in constructive cryptography.*

For a random variable X over an m-ary alphabet $\mathcal{X}$ let the distance from uniform $d(X)$, the maximal probability $p_{\max}(X)$, and the collision probability $p_{\text{coll}}(X)$ be defined as in Section 7.2.3 of the lecture notes.

Hint: Recall Jensen's inequality, especially the following variant of it: for non-negative $\lambda_1, \dots, \lambda_n$ with $\sum_{i=1}^{n} \lambda_i = 1$ and for all $x_1, \dots, x_n \in \mathbb{R}_{\geq 0}$ we have

$$f\left(\sum_{i=1}^{n} \lambda_i x_i\right) \leq \sum_{i=1}^{n} \lambda_i f(x_i) \qquad \text{for convex functions } f\colon \mathbb{R}_{\geq 0} \to \mathbb{R} \quad \text{and}$$

$$f\left(\sum_{i=1}^{n} \lambda_i x_i\right) \geq \sum_{i=1}^{n} \lambda_i f(x_i) \qquad \text{for concave functions } f\colon \mathbb{R}_{\geq 0} \to \mathbb{R}.$$

a) Prove Lemma 7.6, i.e., show that $\frac{1}{|\mathcal{X}|} \leq p_{\text{coll}}(X) \leq p_{\max}(X)$.

b) Prove Lemma 7.7, i.e., show that $d(X) \leq \frac{1}{2}\sqrt{|\mathcal{X}| \cdot p_{\text{coll}}(X) - 1}$.

In the following, we investigate the construction that corresponds to Example 7.1 from the lecture notes. That is, we exemplify how privacy amplification can be used to construct a shared secret r-bit key from an authenticated channel and a shared uniformly random n-bit key of which the adversary knows at most m bits. Now consider the authenticated channel $\bullet\!\!\longrightarrow$ from the lecture and the following two resources for a nonempty set $\mathcal{K}$ and a function f with domain $\mathcal{K}$:

The key $\overset{\mathcal{K}}{\bullet\!\!=\!\!=\!\!\bullet}$

- chooses $k \in \mathcal{K}$ uniformly at random
- On input getKey at interface A or B, output k at the same interface.

The key $\overset{\mathcal{K},f}{\bullet\!\!=\!\!=\!\!\bullet}$ with leakage

- chooses $k \in \mathcal{K}$ uniformly at random
- On input getKey at interface A or B, output k at the same interface
- On input getLeakage at interface E, output $f(k)$ at E.

c) Let $\mathbf{R} := \left[\overset{\{0,1\}^n,f}{\bullet\!\!=\!\!=\!\!\bullet}, \bullet\!\!\longrightarrow\right]$ and $\mathbf{S} := \overset{\{0,1\}^r}{\bullet\!\!=\!\!=\!\!\bullet}$. Let $\mathcal{G}$ denote a known universal class of functions $\{0,1\}^n \to \{0,1\}^r$. Specify a protocol $\boldsymbol{\pi} = (\pi_1, \pi_2)$ and for all $f\colon \{0,1\}^n \to \{0,1\}^m$ a simulator σ_f in order to prove that $\Delta(\pi_1^A \pi_2^B \mathbf{R}, \sigma_f^E \mathbf{S}) \leq \frac{1}{2}\sqrt{2^{r+m-n}}$.