ETH Zürich, D-INFK
Spring 2018

Prof. Ueli Maurer
Christian Badertscher
Fabio Banfi

Cryptography Foundations

Exercise 11

11.1 Conditional Probability Distributions

Goal*: We repeat the basics on conditional probability distributions and how they are composed to define a random experiment.*

In this task, we define a random experiment using a uniform random permutation $\mathbf{P}_n$ as defined in the lecture notes. We consider an environment $\mathbf{E}$ that issues two queries to $\mathbf{P}_n$ and then stops. More specifically, $\mathbf{E}$ chooses the first input uniformly and sets the second input to be equal to the first output of $\mathbf{P}_n$ and then stops after receiving the reply from $\mathbf{P}_n$. The random experiment defined by $\mathbf{E}$ and $\mathbf{P}_n$ then evolves as follows:

1. $\mathbf{E}$ chooses the first input X_1 uniformly at random from $\{0,1\}^n$.
2. $\mathbf{P}_n$ obtains the input X_1 and responds with the output Y_1.
3. $\mathbf{E}$ obtains the output Y_1 and then sets $X_2 = Y_1$.
4. $\mathbf{P}_n$ obtains the input X_2 and, based on X_1, Y_1, and X_2, responds with the output Y_2.
5. $\mathbf{E}$ outputs $\dashv$.

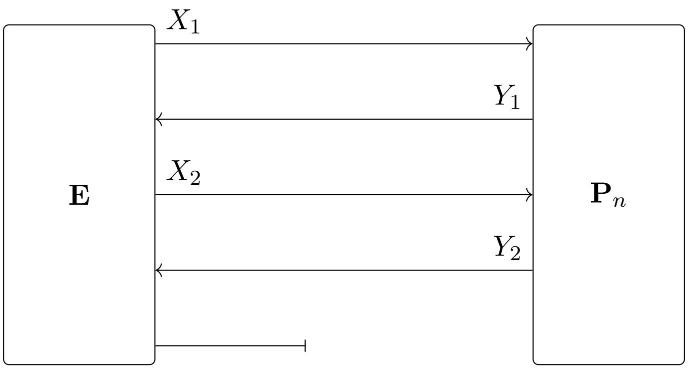

a) Describe the behavior of $\mathbf{E}$ in terms of conditional probability distributions and compute the resulting distribution $\mathrm{Pr}^{\mathbf{EP}_n}_{X_1X_2Y_1Y_2}$ of the transcript.

b) What is the distribution of the outputs Y_1 and Y_2 of $\mathbf{P}_n$, i.e., what is $\mathrm{Pr}^{\mathbf{EP}_n}_{Y_1Y_2}$? What is the conditional distribution $\mathrm{Pr}^{\mathbf{EP}_n}_{Y_1Y_2|X_1X_2}$?

11.2 Distinguishing URFs and URPs

Goal*: In this exercise, we complete the proof of Lemma 4.19.*

a) Prove the statement from Example 4.15, i.e., for the uniform random function $\mathbf{R}_{n,n}$ and the uniform random permutation $\mathbf{P}_n$, formalize an MBO $A_1, A_2, \ldots$ such that $\hat{\mathbf{R}}_{n.n} \mathrel{|\!\!\!\equiv} \mathbf{P}_n$ and prove the conditional equivalence.

b) Prove Lemma 4.18. In more detail, let $X_1, \ldots, X_q$ be uniformly-distributed independent random variables on some set $\mathcal{X}$ with $|\mathcal{X}| = t$. Denote by $\mathsf{p}_{\mathrm{coll}}(q,t)$ the probability that

there exists a *collision*, i.e., there exist indices i, j with $1 \leq i < j \leq q$ and $X_i = X_j$. Show that
$$\mathsf{p}_{\text{coll}}(q,t) \leq \frac{1}{2} q^2/t.$$
Hint: What is the probability (for some $i \neq j$) that $X_i = X_j$? How many such pairs $i \neq j$ are there?

11.3 Distinguishing Systems Adaptively

Goal: *To sharpen the view on (non-)adaptive distinguishers, we examine an easy example to see how adaptivity can help.*

Consider the following two discrete random systems $\mathbf{S}_0$ and $\mathbf{S}_1$. System $\mathbf{S}_0$ accepts n-bit strings as inputs and, upon receiving such an input, ignores its value and returns a uniformly distributed random n-bit string. System $\mathbf{S}_1$ behaves similarly to $\mathbf{S}_0$, but whenever an input is equal to one of the previous outputs, it outputs a special fixed symbol $\perp$.

Assume that you are given either $\mathbf{S}_0$ or $\mathbf{S}_1$. Your task is to devise a distinguisher $\mathbf{D}$ that tries to distinguish $\mathbf{S}_0$ and $\mathbf{S}_1$ by providing inputs to the system and seeing the corresponding output values.

a) Is it possible to distinguish the two systems $\mathbf{S}_0$ and $\mathbf{S}_1$ if you are only allowed to provide a single input to the system?

b) What is the best strategy if you are allowed any number of queries?

c) What happens if you have to fix all your inputs in advance (before seeing any output)?

11.4 Expansion of PRGs

Goal: *We analyze a construction to expand pseudo-randomness and use the Constructive Cryptography framework.*

Let $g\colon \{0,1\}^k \to \{0,1\}^{2k}$ be a function (think of a PRG). We want to use this function to generate $4k$ (pseudo-random) bits from a k-bit seed and formulate it as a construction (cf. Example 5.1 of the lecture notes).

Let $\mathbf{G}$ be the resource that on the first input $s \in \{0,1\}^k$ returns $g(s)$ (and ignores subsequent inputs). Let further denote $\mathbf{U}_n$ the resource that upon the first invocation outputs a uniformly distributed random n-bitstring. Finally, let $\alpha[\mathbf{U}_k, \mathbf{G}]$ be the resource that on the first activation outputs $g(s)$ for a uniformly random k-bitstring s (implemented by a converter α that routes the output of $\mathbf{U}_k$ as input to $\mathbf{G}$).

a) Describe the specification that we aim to construct (using Section 5.3.5 from the lecture notes) as a generic relaxation of $\{\mathbf{U}_{4k}\}$ that contains all systems $\mathbf{S}$ such that the distinction problem $\langle \alpha[\mathbf{U}_k, \mathbf{G}] \mid \mathbf{U}_{2k} \rangle$ reduces to the distinction problem $\langle \mathbf{U}_{4k} \mid \mathbf{S} \rangle$ (for some reduction ρ with performance-translation λ).

b) Describe the assumed specification $\mathcal{R}$ based on the above resources. Then give a converter β and show which specification $\mathcal{S}$ (of the type defined in **a)**) is constructed (cf. Definition 5.6 of the lecture notes) by providing the concrete reduction and performance-translation functions.

Hint: Think of the following construction: compute $s_1|s_2 := g(s)$ ($s_1, s_2 \in \{0,1\}^k$) and output $g(s_1)|g(s_2)$. Note that $\mathcal{R}$ can also be a singleton set.

Discussion of solutions:

Tuesday, 22.5.2018 (Tasks 11.1, 11.2, and 11.3)

28/29.5.2018 (Task 11.4)

The Monday and Tuesday sessions of each week cover the same material.