Preamble

This sheet is split into two parts. Page one should be attempted after the lesson of Monday 8/10/2018 and page two should be attempted after the lesson of Thursday 11/10/2018.

Problem 1

Definition (Key Equivocation measure)

Let X and Y be random variables and let $x_1, \ldots, x_n$ be the possible values of X and $y_1, \ldots, y_n$ the possible values of Y. The equivocation, or conditional entropy, of X on Y is the quantity $H(X|Y)$ defined by

$$H(X|Y) := -\sum_{i=1}^{n}\sum_{j=1}^{m} f_Y(y_j) \cdot f_{X|Y}(x_i|y_j) \cdot \log_2(f_{X|Y}(x_i|y_i))$$

where f_X, f_Y and $f_{X|Y}$ are the density functions for the corresponding distributions.

The **Key Equivocation** *quantity measures the total information about the key revealed by the ciphertext and is formally defined as the quantity $H(K|C)$.*

Suppose that the key equivocation of a cryptosystem vanishes, ie. that $H(K|C) = 0$. Prove that even a single observed ciphertext uniquely determines which key was used.

Problem 2

Definition (Probability ensemble)

If for every natural number $n \in \mathbb{N}$ we have a probability distribution X_n, then $\mathcal{X} = \{X_n\}_{n\in\mathbb{N}}$ is a probability ensemble.

Definition (Distinguisher)

A probabilistic polynomial-time algorithm $\mathcal{D}$ that attempts to distinguish whether a sample from a set S came from one of two probability distributions on S is called a Distinguisher. If $\mathcal{D}$ guesses the correct probability distribution from which the sample was made, we say that $\mathcal{D}$ outputs 1 and outputs 0 otherwise.

Definition (Computational Indistinguishability)

Two probability ensembles $\mathcal{X} = \{X_n\}_{n\in\mathbb{N}}$ and $\mathcal{Y} = \{Y_n\}_{n\in\mathbb{N}}$ are computationally indistinguishable, denoted $\mathcal{X} \stackrel{c}{\equiv} \mathcal{Y}$, if for every probabilistic polynomial-time distinguisher $\mathcal{D}$ there exists a negligible function negl *such that*

$$\left| \Pr_{x\leftarrow X_n}[\mathcal{D}(1^n, x) = 1] - \Pr_{y\leftarrow Y_n}[\mathcal{D}(1^n, y) = 1] \right| \le \mathsf{negl}(n).$$

Let $X = \{X_n\}_{n\in\mathbb{N}}$ and $X = \{Y_n\}_{n\in\mathbb{N}}$ be computationally indistinquishable probability ensembles.

(a) Prove that for any probabilistic polynomial-time algorithm $\mathcal{A}$ it holds that $\{\mathcal{A}(X_n)\}_{n\in\mathbb{N}}$ and $\{\mathcal{A}(Y_n)\}_{n\in\mathbb{N}}$ are computationally indistinguishable, where $\mathcal{A}(X)$ denotes the distribution generated by running $\mathcal{A}(x)$ on all samples $x \leftarrow X$.

(b) Prove that the above may no longer hold if $\mathcal{A}$ does not run in polynomial-time.

Problem 3

If the best algorithm today for finding the prime factors of an n-bit number takes $2^{c \cdot n^{\frac{1}{3}} (\log n)^{\frac{2}{3}}}$ clock cycles, then (assuming that $c = 1$) estimate the size of numbers which cannot be factored in the next 100 years on a 4Ghz[1] computer.

[1] Ghz is shorthand for Giga-hertz and is a measure of how many clock-cycles a computer can perform a second. 1Ghz = 10^9 clock cycles.

Problem 4
Let G be a pseudorandom generator where $|G(s)| \geq 2 \cdot |s|$.

(a) Define $G' := G(s0^{|s|})$. Is G' necessarily a pseudorandom generator?

(b) Define $G'' := G(s_1 \ldots s_{n/2})$ where $s = s_1 \ldots s_n$. Is G'' necessarily a pseudorandom generator?

Problem 5
Let G be a pseudorandom generator and define $G'(s)$ to be the output of G truncated to n bits (where $|s|$ is of length n). Prove that the function $F_k(x) = G'(k) \oplus x$ is not pseudorandom.

Problem 6
Let $\Pi_1 := (\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1)$ and $\Pi_2 := (\mathsf{Gen}_2, \mathsf{Enc}_2, \mathsf{Dec}_2)$ be two encryption schemes for which it is known that at least one is CPA-secure. It is unfortunately unknown whether Π_1 or Π_2 is insecure. Show how to construct an encryption scheme Π that is guaranteed to be CPA-secure as long as at least one of Π_1 or Π_2 is CPA-secure. Try to provide a full proof of your answer.

Problem 7
Let G be a pseudorandom generator. Prove that

$$G'(x_1, \ldots, x_n) := G(x_1)||G(x_2)|| \ldots ||G(x_n)$$

where $|x_1| = \ldots |x_n|$ is a pseudorandom generator.

Problem 8
In the lectures you were given the definition of the $\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n)$:

1. The adversary $\mathcal{A}$ is given input 1^n, and outputs a pair of messages m_0, m_1 with $|m_0| = |m_1|$.
2. A key k is generated by running $\mathsf{Gen}(1^n)$, and a uniform bit $b \in \{0, 1\}$ is chosen. The *Challenge Ciphertext* $c \leftarrow \mathsf{Enc}_k(m_b)$ is computed and given to $\mathcal{A}$.
3. $\mathcal{A}$ outputs a bit b'.
4. The output of the experiment is defined to be 1 if $b' = b$ and 0 otherwise. If $\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n) = 1$ we say that $\mathcal{A}$ succeeds.

Prove that the following definitions are equivalent. This shows, in particular, that the first is an equivalent definition of perfect secrecy.

Definition (Indistinquishability in the presence of an eavesdropper)
A private-key encryption scheme $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *is* EAV-secure*, if for any adversary* $\mathcal{A}$ *it holds that*

$$\Pr\left[\mathsf{PrivK}^{\mathsf{eav}}_{\mathcal{A},\Pi}(n) = 1\right] = 1/2.$$

Definition (Perfect Secrecy)
A private-key encryption scheme $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *is perfectly secret, if for every probability distribution over* $\mathcal{M}$*, every* $m \in \mathcal{M}$*, and every ciphertext* $c \in \mathcal{C}$ *for which* $\Pr[C = c] > 0$*:*

$$\Pr[M = m \mid C = c] = \Pr[M = m].$$