Foundation of Cryptography, Fall 2016 Beny Applebaum and Iftach Haitner

Problem set 2

November 22, 2016 Due: December 8

- Please submit the handout in class, or email me, in case you write in LaTeX.
- Write clearly and shortly using sub claims if needed. The emphasize in most questions is on the proofs (no much point is writing a “solution" w/o proving its correctness).
- For Latex users, a solution example can be found in the course web site.
- It is ok to work in (small) groups, but please write the id list of your partners in the solution file, and each student should write her solution by *herself* (joint effort is only allowed in the “thinking phase").
- The notation we use appears in the introduction part of the first lecture (*Notation* section).

1. Prove that the existence of pseudorandom generators implies the existence of one-way functions.

2. (a) Let $\{X_n, Z_n\}_{n\in\mathbb{N}}$ be distribution ensemble, where $\mathrm{Supp}(X_n) = \{0,1\}$ and $\mathrm{Supp}(Z_n) = \{0,1\}^n$ (i.e., X_n is a bit and Z_n is an n-bit string). Assume there exists a PPT A, function $\varepsilon\colon \mathbb{N} \mapsto [0,1]$ and set $\mathcal{I} \subseteq \mathbb{N}$, such that

$$\Pr[\mathsf{A}(Z_n) = X_n] \geq \frac{1}{2} + \varepsilon(n)$$

for every $n \in \mathcal{I}$. Prove there exists PPT B such that

$$\Pr[\mathsf{B}(Z_n, X_n) = 1] - \Pr[\mathsf{B}(Z_n, U_1) = 1] \geq \varepsilon(n)$$

for every $n \in \mathcal{I}$, where U_1 is uniformly distributed over $\{0,1\}$ (independently, of (X_n, Z_n).

(b) Use (a) to show that if b is *not* an hardcore predicate of $f\colon \{0,1\}^n \mapsto \{0,1\}^n$, then $(f(U_n), b(U_n))$ is computationally *distinguishable* from $(f(U_n), b(U_1))$ — there exists a PPT that distinguishes between

$$\{\{(f(x), b(x))\}_{x\leftarrow\{0,1\}^n}\}_{n\in\mathbb{N}} \quad \text{and} \quad \{\{(f(x), c)\}_{x\leftarrow\{0,1\}^n, c\leftarrow\{0,1\}}\}_{n\in\mathbb{N}}$$

with $1/p(n)$ advantage, for some $p \in \mathrm{poly}$, for infinitely many n's.

3. A function family is a *non-adaptive* PRF, if it is a PRF according to the definition given in class, but its security should only holds against *non-adaptive* distinguishers: distinguishers that choose all queries to the oracle *before* making the first query (alternatively, they make all there queries at once).

   Assuming the existence of OWFs, prove that there exists a non-adaptive PRF that is *not* an (adaptive) PRF.

4. (a) Let X_1 and X_2 be jointly distributed random variables over domain $\mathcal{U}$, and let D_1 and D_2 be their marginal distributions, respectively.

   Prove that $\mathrm{SD}(D_1, D_2) \leq \Pr[X_1 \neq X_2]$.

   (b) For $k \in \mathbb{N}$, let P_k be the distribution (over $\{0,1\}^k$) induced by concatenating k unbiased independent coins, i.e., each coin is taking the value 1 with probability $\frac{1}{2}$ and 0 otherwise. Let Q_k be the distribution induced by concatenating k ε-biassed independent coins, i.e., each coin is taking the value 1 with probability $\frac{1}{2} + \varepsilon$ and 0 otherwise. We would like to bound the statistical distance between P_k and Q_k.

      i. Describe a process S (i.e., an algorithm) that outputs a pair of values (x, y), and let (X, Y) be the random variable describing the output of S in a random execution. Prove that the marginal distribution of X is P_k, and that the marginal distribution of Y is Q_k.

      Hint: consider the following method for sampling an ε-biassed coin: with probability 2ε output 1, and otherwise, output an unbiased coin.

ii. Bound the probability that $X \neq Y$.

iii. Use the first part of the question to bound $\mathrm{SD}(P_k, Q_k)$.

5. In this question we prove several useful probabilistic facts.

   (a) (Markov's inequality) Prove that for any non-negative random variable X and positive real $a > 0$ it holds
   $$\Pr[X \geq a] \leq \frac{\mathsf{E}[X]}{a}.$$

   (b) Prove Chebyshev's inequality (Lemma 22 in the slides). Hint: Use Markov's inequality with respect to the random variable $Y = (\mathsf{E}[X] - X)^2$ and recall that the expected value of Y is the variance of X.

   (c) Let M be an $m \times n$ binary matrix in which every subset of k rows are linearly independent (over the binary field). Let $x = (x_1, \ldots, x_n)$ be a uniformly chosen n-bit string, viewed as a column vector over the binary field, and consider the jointly distributed random variables $y = (y_1, \ldots, y_m)$ defined by $y = Mx$. Prove that y are *k-wise independent*, i.e., for every k distinct indices $i_1 < \cdots < i_k \in [m]$ and every k bits $b_1, \ldots, b_k$ it holds that
   $$\Pr[y_{i_1} = b_1 \wedge \cdots \wedge y_{i_k} = b_k] = 2^{-k}.$$

   Note: This part of the question is unrelated to the first two parts.