MTAT.07.003 Cryptology II
Spring 2009 / Exercise Session III

1. Recall that a game is a two-party protocol between the challenger $\mathcal{G}$ and an adversary $\mathcal{A}$ and that the output of the game $\mathcal{G}^{\mathcal{A}}$ is always determined by the challenger. Prove the following claims:

   (a) Any hypothesis testing scenario $\mathcal{H}$ can be formalised as a game $\mathcal{G}$ such that $\Pr[\mathcal{A} = b|\mathcal{H}] = \Pr[\mathcal{G}^{\mathcal{A}} = b]$ for all adversaries $\mathcal{A}$.

   (b) For two simple hypotheses $\mathcal{H}_0$ and $\mathcal{H}_1$, there is a game $\mathcal{G}$ such that

   $$\mathsf{cd}_{\star}^{t}(\mathcal{H}_0, \mathcal{H}_1) = 2 \cdot \max_{\mathcal{A} \text{ is } t\text{-time}} \left|\Pr[\mathcal{G}^{\mathcal{A}} = 1] - \tfrac{1}{2}\right| \enspace .$$

   (c) The computational distance between games defined as follows

   $$\mathsf{cd}_{\star}(\mathcal{G}_0, \mathcal{G}_1) = \max_{\mathcal{A} \text{ is } t\text{-time}} \left|\Pr[\mathcal{G}_0^{\mathcal{A}} = 1] - \Pr[\mathcal{G}_1^{\mathcal{A}} = 1]\right| \enspace .$$

   Show that this quantity is indeed a pseudo-metric:

   $$\mathsf{cd}_{\star}^{t}(\mathcal{G}_0, \mathcal{G}_1) = \mathsf{cd}_{\star}^{t}(\mathcal{G}_1, \mathcal{G}_0) \enspace ,$$
   $$\mathsf{cd}_{\star}^{t}(\mathcal{G}_0, \mathcal{G}_2) \leq \mathsf{cd}_{\star}^{t}(\mathcal{G}_0, \mathcal{G}_1) + \mathsf{cd}_{\star}^{t}(\mathcal{G}_1, \mathcal{G}_2) \enspace .$$

   When is the computational distance a proper metric, i.e.,

   $$\mathsf{cd}_{\star}^{t}(\mathcal{G}_0, \mathcal{G}_1) \neq 0 \qquad \Leftrightarrow \qquad \mathsf{sd}_{\star}(\mathcal{G}_0, \mathcal{G}_1) \neq 0 \enspace ?$$

   (?) Usually, the statistical distance $\mathsf{sd}_{\star}(\mathcal{G}_0, \mathcal{G}_1)$ is defined as a limiting value $\mathsf{sd}_{\star}(\mathcal{G}_0, \mathcal{G}_1) = \lim_{t\to\infty} \mathsf{cd}_{\star}^{t}(\mathcal{G}_0, \mathcal{G}_1)$. Give an alternative interpretation in terms of output distributions.

2. Let $\mathcal{A}$ be a t-time distinguisher and let $\alpha(\mathcal{A}) = \Pr[\mathcal{A} = 1|\mathcal{H}_0]$ and $\beta(\mathcal{A}) = \Pr[\mathcal{A} = 0|\mathcal{H}_1]$ be the ratios of false negatives and false positives. Show that for any c there exists a $t + \mathrm{O}(1)$-time adversary $\mathcal{B}$ such that

   $$\alpha(\mathcal{B}) = (1 - c) \cdot \alpha(\mathcal{A}) \qquad \text{and} \qquad \beta(\mathcal{B}) = c + (1 - c) \cdot \beta(\mathcal{A}) \enspace .$$

   Are there any practical settings where such trade-offs are economically justified? Give some real world examples.

   Hint: What happens if you first throw a fair coin and run $\mathcal{A}$ only if you get tail and otherwise output 0?

3. Let $\mathcal{X}_0$ and $\mathcal{X}_1$ efficiently samplable distributions that are (t, ε)-indistinguishable. Show that distributions $\mathcal{X}_0$ and $\mathcal{X}_1$ remain computationally indistinguishable even if the adversary can get n samples.

(a) First estimate computational distances between following games

$$\begin{array}{lll} \mathcal{G}_{00}^{\mathcal{A}} & \mathcal{G}_{01}^{\mathcal{A}} & \mathcal{G}_{11}^{\mathcal{A}} \\ \left[\begin{array}{l} x_0 \leftarrow \mathcal{X}_0 \\ x_1 \leftarrow \mathcal{X}_0 \\ \textbf{return } \mathcal{A}(x_0, x_1) \end{array}\right. & \left[\begin{array}{l} x_0 \leftarrow \mathcal{X}_0 \\ x_1 \leftarrow \mathcal{X}_1 \\ \textbf{return } \mathcal{A}(x_0, x_1) \end{array}\right. & \left[\begin{array}{l} x_0 \leftarrow \mathcal{X}_1 \\ x_1 \leftarrow \mathcal{X}_1 \\ \textbf{return } \mathcal{A}(x_0, x_1) \end{array}\right. \end{array}$$

Hint: What happens if you feed a sample $x_0 \leftarrow \mathcal{X}_0$ together an unknown sample $x_1 \leftarrow \mathcal{X}_i$ to $\mathcal{A}$ and use the reply to guess i.

(b) Generalise the argumentation to the case, where the adversary $\mathcal{A}$ gets n samples from a distribution $\mathcal{X}_i$. That is, define the corresponding sequence of games $\mathcal{G}_{00\ldots0}, \ldots, \mathcal{G}_{11\ldots1}$.

(c) Why do we need to assume that distributions $\mathcal{X}_0$ and $\mathcal{X}_1$ are efficiently samplable?

4. Consider the following game, where an adversary $\mathcal{A}$ gets three values x_1, x_2 and x_3. Two of them are sampled from the efficiently samplable distribution $\mathcal{X}_0$ and one of them is sampled from the efficiently samplable distribution $\mathcal{X}_1$. The adversary wins the game if it correctly determines which sample is taken from $\mathcal{X}_1$.

   (a) Find an upper bound to the success probability if distributions $\mathcal{X}_0$ and $\mathcal{X}_1$ are (t, ε)-indistinguishable.

   (b) How does the bound on the success change if we modify the game in the following manner. First, the adversary can first make its initial guess i_0. Then the challenger reveals $j \neq i_0$ such that x_j was sampled from $\mathcal{X}_0$ and then the adversary can output its final guess i_1.
   Hint: How well the adversary can perform if the challenger gives no samples to the adversary? How can you still simulate the game to the adversary who expects these samples?

5. A predicate $\pi : \{0,1\}^n \to \{0,1\}$ is said to be a ε-*regular* if the output distribution for uniform input distribution is nearly uniform:

$$|\Pr[s \xleftarrow{u} \{0,1\}^n : \pi(s) = 0] - \Pr[s \xleftarrow{u} \{0,1\}^n : \pi(s) = 1]| \leq \varepsilon \ .$$

   A predicate π is a (t, ε)-*unpredictable* also known as (t, ε)-*hardcore predicate* for a function $f : \{0,1\}^n \to \{0,1\}^{n+\ell}$ if for any t-time adversary

$$\mathsf{Adv}_f^{\text{hc-pred}}(\mathcal{A}) = 2 \cdot \left|\Pr[s \xleftarrow{u} \{0,1\}^n : \mathcal{A}(f(s)) = \pi(s)] - \tfrac{1}{2}\right| \leq \varepsilon \ .$$

   Prove the following statements.

   (a) Any (t, ε)-hardcore predicate is 2ε-regular.

   (b) For a function $f : \{0,1\}^n \to \{0,1\}^{n+\ell}$, let $\pi_k(s)$ denote the kth bit of $f(s)$ and $f_k(s)$ denote the output of $f(s)$ without the kth bit. Show that if f is a (t, ε)-secure pseudorandom generator, then π_k is (t, ε)-hadcore predicate for f_k.

($\star$) If a function $f : \{0,1\}^n \to \{0,1\}^{n+\ell}$ is (t, ε_1)-pseudorandom generator and $\pi : \{0,1\}^n \to \{0,1\}$ is efficiently computable predicate (t, ε_1)-hadcore , then a concatenation $f_*(s) = f(s)||\pi(s)$ is $(t, \varepsilon_1+\varepsilon_2)$-pseudorandom generator.

6. Let $\mathcal{F}$ be a (t, q, ε)-pseudorandom function family that maps a domain $\mathcal{M}$ to the range $\mathcal{C}$. Let $g : \mathcal{M} \to \{0,1\}$ be an arbitrary predicate. What is the success probability of a t-time adversary $\mathcal{A}$ in the following games?

$$\mathcal{G}_0^{\mathcal{A}} \quad \left[\begin{array}{l} m \xleftarrow[u]{} \mathcal{M} \\ f \xleftarrow[u]{} \mathcal{F} \\ c \leftarrow f(m) \\ \textbf{return } [\mathcal{A}(c) \stackrel{?}{=} m] \end{array}\right. \qquad \mathcal{G}_1^{\mathcal{A}} \quad \left[\begin{array}{l} m \xleftarrow[u]{} \mathcal{M} \\ f \xleftarrow[u]{} \mathcal{F} \\ c \leftarrow f(m) \\ \textbf{return } [\mathcal{A}(c) \stackrel{?}{=} g(m)] \end{array}\right.$$

Establish the same result by using the IND$\Longrightarrow$SEM theorem. More precisely, show that the hypothesis testing games

$$\mathcal{G}_{m_0}^{\mathcal{A}} \quad \left[\begin{array}{l} f \xleftarrow[u]{} \mathcal{F} \\ c \leftarrow f(m_0) \\ \textbf{return } \mathcal{A}(c) \end{array}\right. \qquad \mathcal{G}_{m_1}^{\mathcal{A}} \quad \left[\begin{array}{l} f \xleftarrow[u]{} \mathcal{F} \\ c \leftarrow f(m_1) \\ \textbf{return } \mathcal{A}(c) \end{array}\right.$$

are $(t, 2\varepsilon)$-indistinguishable for all $m_0, m_1 \in \mathcal{M}$.