MTAT.07.003 Cryptology II
Spring 2010 / Exercise session III / Example solution

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

Solution. Any such problem can be split into three conceptual parts: formalisation of the attack scenario, game manipulation, and final probability computations. One possible formalisation of the attack scenario is given below

$$
\begin{array}{l}
\mathcal{G}_0^{\mathcal{A}} \\
\left[\begin{array}{l}
x_1 \leftarrow \mathcal{X}_0 \\
x_2 \leftarrow \mathcal{X}_0 \\
x_3 \leftarrow \mathcal{X}_1 \\
\pi \xleftarrow[u]{} \mathsf{Perm}(\{1,2,3\}) \\
i \leftarrow \mathcal{A}(x_{\pi(1)}, x_{\pi(2)}, x_{\pi(3)}) \\
\textbf{return } [\pi(i) \stackrel{?}{=} 3]
\end{array}\right.
\end{array}
$$

The fourth line in the game models random shuffling of the values x_1, x_2, x_3. If we choose uniformly a permutation π over $\{1, 2, 3\}$, the elements $x_{\pi(1)}, x_{\pi(2)}, x_{\pi(3)}$ are in a random order. Obviously, the guess of $\mathcal{A}$ is correct if and only if $\pi(i) = 3$. As a second step, we modify the game in the following way

$$
\begin{array}{lcl}
\mathcal{G}_0^{\mathcal{A}} & & \mathcal{G}_1^{\mathcal{A}} \\
\left[\begin{array}{l}
x_1 \leftarrow \mathcal{X}_0 \\
x_2 \leftarrow \mathcal{X}_0 \\
x_3 \leftarrow \mathcal{X}_1 \\
\pi \xleftarrow[u]{} \mathsf{Perm}(\{1,2,3\}) \\
i \leftarrow \mathcal{A}(x_{\pi(1)}, x_{\pi(2)}, x_{\pi(3)}) \\
\textbf{return } [\pi(i) \stackrel{?}{=} 3]
\end{array}\right.
& \xRightarrow{IND} &
\left[\begin{array}{l}
x_1 \leftarrow \mathcal{X}_0 \\
x_2 \leftarrow \mathcal{X}_0 \\
x_3 \leftarrow \mathcal{X}_0 \\
\pi \xleftarrow[u]{} \mathsf{Perm}(\{1,2,3\}) \\
i \leftarrow \mathcal{A}(x_{\pi(1)}, x_{\pi(2)}, x_{\pi(3)}) \\
\textbf{return } [\pi(i) \stackrel{?}{=} 3]
\end{array}\right.
\end{array}
$$

Note that the games differ only in a single line: x_3 is chosen either from $\mathcal{X}_0$ or from $\mathcal{X}_1$ depending on the game. The latter allows us to use the entire game as a distinguisher for $\mathcal{X}_0$ and $\mathcal{X}_1$. Namely, let us define a new adversary

$$\begin{array}{l}\mathcal{B}(x)\\ \left[\begin{array}{l}x_1 \leftarrow \mathcal{X}_0\\ x_2 \leftarrow \mathcal{X}_0\\ x_3 \leftarrow x\\ \pi \xleftarrow{u} \mathsf{Perm}(\{1,2,3\})\\ i \leftarrow \mathcal{A}(x_{\pi(1)}, x_{\pi(2)}, x_{\pi(3)})\\ \textbf{return } [\pi(i) \stackrel{?}{=} 3]\end{array}\right.\end{array}$$

against the indistinguishability games

$$\begin{array}{l}\mathcal{Q}_0^{\mathcal{B}}\\ \left[\begin{array}{l}x \leftarrow \mathcal{X}_0\\ \textbf{return } \mathcal{B}(x)\end{array}\right.\end{array} \qquad \begin{array}{l}\mathcal{Q}_1^{\mathcal{B}}\\ \left[\begin{array}{l}x \leftarrow \mathcal{X}_1\\ \textbf{return } \mathcal{B}(x)\end{array}\right.\end{array}$$

By the (t, ε)-indistinguishability assumptions

$$\mathsf{Adv}^{\mathsf{ind}}_{\mathcal{X}_0,\mathcal{X}_1}(\mathcal{B}) = \left|\Pr\left[\mathcal{Q}_0^{\mathcal{B}} = 1\right] - \Pr\left[\mathcal{Q}_1^{\mathcal{B}} = 1\right]\right| \leq \varepsilon$$

for any t-time adversary $\mathcal{B}$. Let us estimate the behaviour of our concrete adversary by inserting its definition into the games $\mathcal{Q}_0$ and $\mathcal{Q}_1$:

$$\begin{array}{l}\mathcal{Q}_0^{\mathcal{B}}\\ \left[\begin{array}{l}x \leftarrow \mathcal{X}_0\\ x_1 \leftarrow \mathcal{X}_0\\ x_2 \leftarrow \mathcal{X}_0\\ x_3 \leftarrow x\\ \pi \xleftarrow{u} \mathsf{Perm}(\{1,2,3\})\\ i \leftarrow \mathcal{A}(x_{\pi(1)}, x_{\pi(2)}, x_{\pi(3)})\\ \textbf{return } [\pi(i) \stackrel{?}{=} 3]\end{array}\right.\end{array} \qquad \begin{array}{l}\mathcal{Q}_1^{\mathcal{B}}\\ \left[\begin{array}{l}x \leftarrow \mathcal{X}_1\\ x_1 \leftarrow \mathcal{X}_0\\ x_2 \leftarrow \mathcal{X}_0\\ x_3 \leftarrow x\\ \pi \xleftarrow{u} \mathsf{Perm}(\{1,2,3\})\\ i \leftarrow \mathcal{A}(x_{\pi(1)}, x_{\pi(2)}, x_{\pi(3)})\\ \textbf{return } [\pi(i) \stackrel{?}{=} 3]\end{array}\right.\end{array}$$

By doing simple syntactic changes that do not alter the behaviour of games, we can convert $\mathcal{Q}_0^{\mathcal{B}}$ to $\mathcal{G}_1^{\mathcal{A}}$ and $\mathcal{Q}_1^{\mathcal{B}}$ to $\mathcal{G}_0^{\mathcal{A}}$. Hence, we have established that

$$\left|\Pr\left[\mathcal{G}_0^{\mathcal{A}} = 1\right] - \Pr\left[\mathcal{G}_1^{\mathcal{A}} = 1\right]\right| = \left|\Pr\left[\mathcal{Q}_1^{\mathcal{B}} = 1\right] - \Pr\left[\mathcal{G}_0^{\mathcal{B}} = 1\right]\right| \leq \varepsilon$$

provided that the running-time of $\mathcal{B}$ is less than t. Let $t_{\mathcal{A}}$ be the running-time of $\mathcal{A}$ and t_{s} time needed to get a sample from $\mathcal{X}_0$ or $\mathcal{X}_1$. Then the running time of $\mathcal{B}$ is $2t_{\mathrm{s}} + t_{\mathcal{A}} + \mathrm{O}(1)$. Hence, for all $t - 2t_{\mathrm{s}} - \mathrm{O}(1)$ time adversaries

$$\left|\Pr\left[\mathcal{G}_0^{\mathcal{A}} = 1\right] - \Pr\left[\mathcal{G}_1^{\mathcal{A}} = 1\right]\right| \leq \varepsilon\ . \tag{1}$$

By doing syntactic changes that do not alter the behaviour of the game, we can rewrite the game $\mathcal{G}_1$ even further

$$\begin{array}{lcl}
\mathcal{G}_1^{\mathcal{A}} & & \mathcal{G}_2^{\mathcal{A}} \\
\left[\begin{array}{l}
x_1 \leftarrow \mathcal{X}_0 \\
x_2 \leftarrow \mathcal{X}_0 \\
x_3 \leftarrow \mathcal{X}_0 \\
\pi \xleftarrow[u]{} \mathsf{Perm}(\{1,2,3\}) \\
i \leftarrow \mathcal{A}(x_{\pi(1)}, x_{\pi(2)}, x_{\pi(3)}) \\
\textbf{return } [\pi(i) \stackrel{?}{=} 3]
\end{array}\right.
& \xRightarrow{Syntax} &
\left[\begin{array}{l}
x_1 \leftarrow \mathcal{X}_0 \\
x_2 \leftarrow \mathcal{X}_0 \\
x_3 \leftarrow \mathcal{X}_0 \\
i \leftarrow \mathcal{A}(x_1, x_2, x_3) \\
\pi \xleftarrow[u]{} \mathsf{Perm}(\{1,2,3\}) \\
\textbf{return } [\pi(i) \stackrel{?}{=} 3]
\end{array}\right.
\end{array}$$

Note that the behaviour of the game does not change since $\mathcal{A}$ gets the same input distribution $\mathcal{X}_0 \times \mathcal{X}_0 \times \mathcal{X}_0$ in both games. As the output of $\mathcal{A}$ is fixed before the permutation is chosen, we get

$$\Pr\left[\mathcal{G}_2^{\mathcal{A}} = 1\right] = \frac{1}{3} \ . \tag{2}$$

By combing (1) and (2) we obtain

$$\Pr\left[\mathcal{G}_0^{\mathcal{A}} = 1\right] \leq \frac{1}{3} + \varepsilon$$

provided that the running-time of $\mathcal{A}$ is $t - 2t_{\mathrm{s}} - \mathrm{O}(1)$.

Comments. if distributions $\mathcal{X}_0$ and $\mathcal{X}_1$ are (t, ε)-indistinguishable, it is always possible to change the game by replacing a line $x \leftarrow \mathcal{X}_0$ with a line $x \leftarrow \mathcal{X}_1$. The total time-complexity of the game sets limitations on the overall running time of the adversary, as the corresponding distinguisher $\mathcal{B}$ must simulate the game inside its code. By applying such rewriting rules long enough, we can prove computational closeness of many complex games.