NYU COMPUTER SCIENCE FALL 2026

Introduction to Cryptography

CSCI-GA.3210-001 Instructor: Nir Bitansky

Problem Set 1

Due: Wednesday, September 23, 2026, at 11:59 p.m.

Instructions:

- Turn in your solutions in LaTeX and submit a PDF through Gradescope.
(You don't have to know latex, can use AI tools like Prism.)
- Solutions will be graded for both correctness and clarity. If you do not know how to solve a question, you may write "I don't know how to do this" to receive 15% credit for that question.
- Discussion with peers is strongly encouraged, but you must write up your solutions on your own and list your discussion partners. Cite any external sources you use.
- AI use is permitted. If you use AI to solve the homework, list the exact model as a collaborator and include your prompts in an appendix. We strongly recommend solving the questions yourself: the final exam will mostly contain variants of homework questions.
- Use the course forum for clarification questions, but do not post solutions there.

1. (32 pts) Recall the following four definitions of perfect secrecy from Lecture 1 for an encryption scheme (E, D) with keys of length n and messages of length $\ell = \ell(n)$. In this question the adversary A is generally unbounded (think about it as an arbitrary function $A : \{0,1\}^* \to \{0,1\}$).

   1. (8 pts) For any two messages $m_0, m_1 \in \{0,1\}^\ell$ and any adversary A,

   $$\Pr\left[A(ct) = b \;\middle|\; \begin{array}{l} sk \leftarrow \{0,1\}^n \\ b \leftarrow \{0,1\} \\ ct \leftarrow E_{sk}(m_b) \end{array}\right] \leq \frac{1}{2} \ .$$

   2. (8 pts) There exists a fixed distribution S_ℓ such that for every $m \in \{0,1\}^\ell$,

   $$\left\{ct \;\middle|\; \begin{array}{l} sk \leftarrow \{0,1\}^n \\ ct \leftarrow E_{sk}(m) \end{array}\right\} \equiv S_\ell \ .$$

   3. (8 pts) For any distribution M over $\{0,1\}^\ell$, sample $m \leftarrow M$, independently sample $sk \leftarrow \{0,1\}^n$, and let $ct \leftarrow E_{sk}(m)$. For every ciphertext c with $\Pr[ct = c] > 0$ and every message $m' \in \{0,1\}^\ell$,

   $$\Pr[m = m' \mid ct = c] = \Pr[m = m'] \ .$$

4. (8 pts) For every nonempty set $S \subseteq \{0,1\}^\ell$ and every adversary A,

$$\Pr\left[A(ct) = m \;\middle|\; \begin{array}{l} sk \leftarrow \{0,1\}^n \\ m \leftarrow S \\ ct \leftarrow E_{sk}(m) \end{array}\right] \leq \frac{1}{|S|} \ ,$$

where $m \leftarrow S$ denotes a uniformly random message from S, chosen independently of sk.

We showed in class that definitions 1 and 2 are equivalent. Prove that all four definitions are equivalent.

2. (26 pts) In class, we saw that for any encryption scheme (E, D) for messages of length ℓ, with keys of length $n \leq \ell - 10$, if E is deterministic, there exist two messages m_0, m_1 and an inefficient A such that

$$\Pr\left[A(ct) = m_b \;\middle|\; \begin{array}{l} sk \leftarrow \{0,1\}^n \\ b \leftarrow \{0,1\} \\ ct \leftarrow E_{sk}(m_b) \end{array}\right] > 0.99 \ .$$

Show that the same holds even if E also tosses, say n, coins (on top of the key). That is, an encryption of any message $m \in \{0,1\}^\ell$ is drawn from a distribution $\{ct \mid sk \leftarrow \{0,1\}^n, r \leftarrow \{0,1\}^n, ct = E_{sk}(m;r)\}$.

3. (32 pts) In what follows, let X_0, X_1, Z be distributions over the same finite support S and let $A : S \to \{0,1\}$ be a function. Define:

$$\Delta_A(X_0, X_1) := |\Pr[A(X_0) = 1] - \Pr[A(X_1) = 1]| \ .$$

(a) (8 pts) Show that for any function A (efficient or not):

$$\left|\Pr\left[A(x) = b \;\middle|\; \begin{array}{l} b \leftarrow \{0,1\} \\ x \leftarrow X_b \end{array}\right] - \frac{1}{2}\right| = \frac{\Delta_A(X_0, X_1)}{2} \ .$$

(b) (8 pts) Show that

$$\max_{A \in \{0,1\}^S} \Delta_A(X_0, X_1) = \frac{1}{2}\sum_{x \in S} |\Pr[X_0 = x] - \Pr[X_1 = x]| = \max_{T \subseteq S} |\Pr[X_0 \in T] - \Pr[X_1 \in T]| \ .$$

(c) (8 pts) Show that

$$\Delta_A(X_0, X_1) \leq \Delta_A(X_0, Z) + \Delta_A(Z, X_1) \ .$$

(d) (8 pts) let A be a randomized distinguisher that uses n random coins. That is, $A(x;r)$ is given a sample x and a random string $r \leftarrow \{0,1\}^n$ (and as usually outputs a bit). Prove that there exists a fixed string $r \in \{0,1\}^n$ such that

$$\Delta_A((X_0;r),(X_1;r)) \geq \Delta_A((X_0;U_n),(X_1;U_n)) \ ,$$

where U_n is the uniform distribution on $\{0,1\}^n$.
(Interpretation: non-uniform distinguishers can often be assumed deterministic w.l.o.g.)