Introduction to Cryptography
Courant, Fall 2014

Homework 8

Instructor: Oded Regev
Student: YOUR NAME HERE

Homework is due by **11am of Nov 10**. Send by email to both "regev" and "tess" under the cs.nyu.edu domain with subject line "CSCI-GA 3210 Homework 8" and name the attachment "YOUR NAME HERE HW8.tex/pdf". Please also bring a printed copy to class. Start early!

Instructions. Solutions must be typeset in LaTeX (a template for this homework is available on the course web page). Your work will be graded on *correctness*, *clarity*, and *conciseness*. You should only submit work that you believe to be correct; if you cannot solve a problem completely, you will get significantly more partial credit if you clearly identify the gap(s) in your solution. It is good practice to start any long solution with an informal (but accurate) "proof summary" that describes the main idea. You may collaborate with others on this problem set and consult external sources. However, you must ***write your own solutions*** and ***list your collaborators/sources*** for each problem.

1. *(Security definitions of SKE.)*[1]

   (a) (3 points) Recall that multi-message *non-adaptive* security for a symmetric-key encryption scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ says that for any $q = \mathrm{poly}(n)$ and any tuples $(m_1, \ldots, m_q), (m'_1, \ldots, m'_q) \in \mathcal{M}^q$, it should be the case that

   $$(\mathsf{Enc}_k(m_1), \ldots, \mathsf{Enc}_k(m_q)) \overset{c}{\approx} (\mathsf{Enc}_k(m'_1), \ldots, \mathsf{Enc}_k(m'_q)),$$

   where in both cases the distribution is over the choice of $k \leftarrow \mathsf{Gen}$ and the randomness in the encryption procedure. Show that non-adaptive security can be equivalently defined as saying that for any $q = \mathrm{poly}(n)$ and any $m_1, \ldots, m_q, m_0, m'_0 \in \mathcal{M}$, it should be the case that

   $$(\mathsf{Enc}_k(m_1), \ldots, \mathsf{Enc}_k(m_q), \mathsf{Enc}_k(m_0)) \overset{c}{\approx} (\mathsf{Enc}_k(m_1), \ldots, \mathsf{Enc}_k(m_q), \mathsf{Enc}_k(m'_0)),$$

   where in both cases the distribution is over the choice of $k \leftarrow \mathsf{Gen}$ and the randomness in the encryption procedure. I need a hint! (ID 17499)

   (b) (3 points) In class, we defined adaptive (or IND-CPA) security as the oracle indistinguishability

   $$(\mathsf{Enc}_k(\cdot), C^0_k(\cdot, \cdot)) \overset{c}{\approx} (\mathsf{Enc}_k(\cdot), C^1_k(\cdot, \cdot)),$$

   where $C^b_k(m_0, m_1)$ outputs $\mathsf{Enc}_k(m_b)$ on receiving the first query and then ignores all further queries, and $k \leftarrow \mathsf{Gen}$. Show that an equivalent definition is

   $$(\mathsf{Enc}^0_k(\cdot, \cdot)) \overset{c}{\approx} (\mathsf{Enc}^1_k(\cdot, \cdot)),$$

   where $\mathsf{Enc}^b_k(m_0, m_1)$ outputs $\mathsf{Enc}_k(m_b)$ and $k \leftarrow \mathsf{Gen}$. I need a hint! (ID 17499)

   (c) (4 points) Give a separation between the non-adaptive and the adaptive security definitions, i.e., construct a (possibly contrived) scheme and prove it secure according to the former definition (under some standard assumption), while showing that it is definitely insecure according to the latter definition. I need a hint! (ID 17495)

   (d) (2 points) (Extra credit)[2] Consider the weakening of the definition of multi-message non-adaptive security in which we take q to be some fixed polynomial, say, $q = n^2$. Show a separation between this definition and the original one.

[1] Based on a question from Peikert's class.

[2] A question asked in class by Konstantinos Vamvourellis