

Homework is due by **noon of Dec 2**. Send by email to both “regev” and “tess” under the cs.nyu.edu domain with subject line “CSCI-GA 3210 Homework 11” and name the attachment “YOUR NAME HERE HW11.tex/pdf”. Please also bring a printed copy to class. Start early!

Instructions. Solutions must be typeset in L^AT_EX (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/hints used* for each problem.

1. (*Lossy Encryption*)¹ Let (Gen, E, D) be a public key encryption scheme. In this problem we define a new property for PKE schemes that we call “lossy encryption”. We say that a scheme (Gen, E, D) is *lossy* if there exists an algorithm $LossyGen(1^n)$ which generates a “lossy” public key PK' (without a secret key) such that the following two properties are satisfied:

1. A lossy public key is computationally indistinguishable from a public key generated by Gen : $PK \approx PK'$. More formally, for any PPT adversary A it holds:

$$|\Pr[A(PK) = 1 | (PK, SK) \leftarrow Gen(1^n)] - \Pr[A(PK') = 1 | PK' \leftarrow LossyGen(1^n)]| \leq \text{negl}(n)$$

2. For any lossy public key $PK' \leftarrow LossyGen(1^n)$, encrypting any message using PK' produces ciphertexts that have identical distribution. Namely, for any $PK' \leftarrow LossyGen(1^n)$, and any pair of messages $m_0, m_1 \in \mathcal{M}$, we have $(PK', E(PK', m_0)) \equiv (PK', E(PK', m_1))$.

Intuitively, notice that this second property is telling that encrypting using the lossy public key completely loses information about the original plaintext, and thus it is not possible to decrypt.

- (a) (5 points) Prove that if an encryption scheme is lossy according to the definition provided above, then the scheme is also IND-CPA-secure.

A hint for 1 point (ID 51588)

Consider the following scheme as a potential candidate for being a lossy public key encryption. $Gen(1^n)$ chooses a random n -bit large safe prime p (i.e., $p = 2q + 1$ for a large prime q) and chooses two random generators g_0, g_1 of $G = QR_p$ (recall that QR_p is the subgroup of quadratic residues in \mathbb{Z}_p^*). Next, it chooses two random (but distinct) values $x_0, x_1 \in \mathbb{Z}_q$, computes $h_0 = g_0^{x_0}$, $h_1 = g_1^{x_1}$, and outputs $PK = (p, g_0, g_1, h_0, h_1)$ and $SK = (x_0, x_1)$.

To encrypt a 1-bit message $m \in \{0, 1\}$, $E(PK, m)$ proceeds as follows: choose a random $r \in \mathbb{Z}_q$ and output $C = (g_m^r, h_m^r)$.

- (b) (3 points) Describe a decryption algorithm.
- (c) (8 points) Second, prove that the scheme described above (together with the decryption algorithm that you obtained from part (b)) is a lossy public key encryption based on the DDH assumption. Namely, first describe a lossy key generation algorithm $LossyGen(1^n)$ and then show that it satisfies both properties (1) and (2). Deduce that the scheme is CPA-secure. A hint for 2 points (ID 51589)
- (d) (5 points) Although the lossy property may be nice and useful in some contexts, this is not necessary to prove that the scheme is CPA-secure. Prove *directly* that this scheme is CPA-secure under the

¹From Dodis

DDH assumption; namely,

$$(g_0, g_1, h_0, h_1, g_0^{r_0}, h_0^{r_0}) \approx (g_0, g_1, h_0, h_1, g_1^{r_1}, h_1^{r_1})$$

A hint for 1 point (ID 51599)

2. (*Semantic Security*) Assume Alice sends Bob a random n -bit string $x \in \{0, 1\}^n$ using a PKC. Eve eavesdrops to the communication and gets to see $Enc_{pk}(x)$. Her goal is to correctly guess x .
 - (a) (3 points) Using the semantic security definition of CPA-security from class (see Dodis's Lecture 6 for a reminder), show that Eve's success probability is negligible. This should follow immediately from the definition.
 - (b) (3 points) Prove the same using the (equivalent) definition of IND-CPA security. This requires a bit more work.
3. (0 points) (*Expanding domain of PRF.*♣) Assume we have a PRF family $\{f_k : \{0, 1\}^n \rightarrow \{0, 1\}^n\}_k$. Let $H = \{h_k : \{0, 1\}^N \rightarrow \{0, 1\}^n\}$ be another family of functions for some large N , say $N = n^2$. What property does H need to satisfy so that the family $\{f_k(h_{k'}(\cdot)) : \{0, 1\}^N \rightarrow \{0, 1\}^n\}$ is a PRF family (where k and k' are chosen independently from the corresponding set of keys)? E.g., can we take H to consist of just the function that outputs the first n bits of its input?