6.875J/18.425J Cryptography and Cryptanalysis	November 21, 2001

Handout 14: Problem Set 8

Instructor: Anna Lysyanskaya *Teaching Assistant: Moses Liskov*

This problem set is due on Wednesday, December 5.

Problem 1: Adaptive Attacks

(40 points) When we considered the security of public-key cryptosystems, we basically only considered allowing the adversary access to the public key. However, a stronger attack scenario may be considered: where the adversary has access to a decryption oracle. Informally, we want that the adversary cannot succeed in a one-pass GM security scenario even when the adversary has access to a decryption oracle under the relevant secret key, though of course querying the oracle on the actual challenge should be disallowed.

(a) (15 points) Give a formal definition of a public-key cryptosystem which is one-pass GM secure against adaptive chosen ciphertext attack, as described above.

(b) (15 points) Recall the ElGamal-A cryptosystem. Alice selects at random (1) a k-bit prime $p = 2q + 1$ where q is also prime, (2) an element g in Z_p^* of order q, and (3) an x in $\mathbb{Z}_q^*$. Then, she sets her public key to be (p, g, h), where $h = g^x \bmod p$, and her secret key to be x.

Alice's public key restricts the message space to consist of all elements in Z_p^* having order q. To send such a message m to Alice, one picks an $r \leftarrow \mathbb{Z}_q$ and computes his ciphertext $c = (g^r \bmod p, h^r m \bmod p)$.

To decrypt $c = (a, b)$, Alice computes $b/(a^x) \bmod p$.

In this problem we go beyond proving that ElGamal-A is not secure against adaptive chosen ciphertext attack. Suppose the adversary has access to a challenge ciphertext c, and may ask only *one* query of the decryption oracle, which much be different from c. Show how to design a c' from c such that given the decryption of c' it is easy to recover the decryption of c.

Can you strengthen this result so that if $c = (a, b)$ and $c' = (a', b')$ then $a \neq a'$ and $b \neq b'$? If so, show how.

(c) (10 points) Suppose the adversary has access to a ciphertext c of a message under the (1-bit) Goldwasser-Micali cryptosystem, and may ask one query of a decryption oracle, other than c. Show how to design a c' from c such that given the decryption of c' it is easy to recover the decryption of c. Also, generalize your attack on the 1-bit scheme to an attack on the ℓ-bit scheme.

Problem 2: Collision-resistant hashing

(30 points) Define a family of collision-resistant hash functions $\mathcal{H}$ as follows:

1. On input 1^k, $\mathcal{H}$ outputs the description of a key PK.
2. For any PK, it is efficient to compute H_{PK}.
3. H_{PK} is length decreasing: $\forall k, \forall PK \in \mathcal{H}(1^k) \forall x \in \{0,1\}^*, |H_{PK}(x)| \leq \min(k, |x|)$.
4. It is hard to find collisions:

$$\forall \{A_k\} PPTF\ \exists \nu \in negl\ \forall k$$
$$\Pr[PK \leftarrow \mathcal{H}(1^k); (x,y) \leftarrow A_k(PK) : H_{PK}(x) = H_{PK}(y) \wedge x \neq y] < \nu(k)$$

(a) Prove that the existence of a family of collision-resistant hash functions implies the existence of a one-way function.

(b) Suppose a family of collision-resistant hash functions $\mathcal{H}$ is given, and let $\mathcal{H}'$ be a family of hash functions that takes two security parameters, l and n, so that it has two keys $PK_l \leftarrow \mathcal{H}(1^l)$ and $PK_n \leftarrow \mathcal{H}(1^n)$. Define $H'_{PK_l,PK_n}(x) = H_{PK_l}(H_{PK_n}(x))$. Prove or disprove: there is an infinite set of security parameter pairs for which $\mathcal{H}'$ is not collision-resistant. Prove or disprove: for all pairs (k,k), $\mathcal{H}'$ is a family of collision-resistant hash functions.

(c) Suppose G is a PRG. Now suppose $\mathcal{H}^G$ is a family of hash functions defined by $\mathcal{H}^G(1^k) = \mathcal{H}(1^k)$ and $H^G_{PK}(x) = G(H_{PK}(x))$. Prove or disprove: $\mathcal{H}^G$ is a family of collision-resistant hash functions. Prove or disprove: $\mathcal{H}^G$ is a family of functions indistinguishable from random.

Problem 3: Signature Schemes in the Random Oracle Model

(30 points) In the "random oracle model," we make the assumption that some function behaves as if it were a random oracle $\mathcal{O}$. We assume the adversary, signer, and verifier have access to the random oracle. Furthermore, we make the assumption only that for every $x \in \{0,1\}^*, \mathcal{O}(x)$ is a truly random string of the same length as x.

Assume that a TDP family p_{PK} exists, and design a simple signature scheme using p and a random oracle. Prove that the scheme is non-existentially forgeable against an adaptive chosen message attack. (Hint: in order to do this, you will have to describe exactly how the oracle $\mathcal{O}$ works, but $\mathcal{O}$ must still be truly random.)