CS195-7: Introduction to Cryptography DUE: September 10, 2002

Problem Set 1

Instructor: Anna Lysyanskaya

Problem 1

In Lecture 1, we gave a definition of a perfectly secure cryptosystem, and saw some limitations arising while using it. In this problem, we will look at perfectly secure message authentication schemes, and show their limitations.

Suppose we are given the following toy scenario: Alice and Bob are communicating over a channel controlled by Eve. Eve may delete or alter messages sent over the channel, and may inject messages of her own. Our goal is to program Alice and Bob in such a way that Bob will only accept a message if it indeed came from Alice.

To that end, we want to design algorithms (G, A, V) for message space M such that

1. Algorithm G is a randomized setup procedure that generates the public parameters P of the system, as well as private inputs to each party, denoted s_{name}, where "name" is the name of the player to which this input was given; for example s_{Alice}.

2. Algorithm A is a procedure that computes an *authentication code* on a message $m \in M$: $c = A(m, \text{"Alice"}, \text{"Bob"}, P, s_{Alice})$.

3. Algorithm V checks that the authentication code c corresponds to message m: for all $m \in M$, for all s_{Alice}, for all values of the public parameters P, if $c = A(m, \text{"Alice"}, \text{"Bob"}, P, s_{Alice})$, then $V(m, c, Alice, Bob, s_{Bob}) =$ "Accept".

4. ϵ-security for n uses: Even after Eve sees the authentication codes for at most n messages *of its own choice*, the probability that she can produce (using any amount of time and space) a pair (m, c), where m is a new message, and c is an authentication code for m, is bounded by ϵ. To make any sense, this is only defined when $n < |M|$, where M is the size of the message space.

(a) Show that $\epsilon = 0$ cannot be achieved for any $n \geq 0$, for any non-empty message space.

(b) Suppose that all messages in M can be represented as ℓ-bit strings. Consider the following authentication mechanism: the initialization procedure G picks an ℓ-bit shared secret s and gives s to Alice and to Bob. $A(m, Alice, Bob, s) = m \oplus s$. For what values of n and ϵ is this authentication mechanism secure? (In other words, how many messages can Alice send before Eve can authenticate a message on her own?)

(c) Suppose that all messages in M can be represented as ℓ-bit strings. Consider the following authentication mechanism:

The initialization procedure G picks an $\ell + 1$-bit prime number p, and two random integers a and b, $0 \leq a, b < p$. The pair of values a and b are given to Alice and Bob, that is their shared secret. The value p is a public parameter.

An authentication code on message m is $c = am + b \bmod p$, where m is treated as an ℓ-bit integer (and so $m < p$).

For what values of n and ϵ is this authentication mechanism secure?

(d) Show that if $\epsilon = 2^{-k}$ and an authentication mechanism is ϵ-secure for n uses, then s_{Alice} and s_{Bob} must both be at least $k(n+1)$ bits long.

Problem 2

(Moved to Problem Set 2 because we have not covered the discrete-logarithm problem in lecture yet.)

Suppose p is a prime and g is a generator modulo p.

Experiment 1: Pick x at random in $\{1, \ldots, p-1\}$. Output $g^x \bmod p$.

Experiment 2: Pick x, y at random in $\{1, \ldots, p-1\}$. Output $g^{xy} \bmod p$.

Prove or disprove: Experiment 1 and Experiment 2 produce identically distributed outputs.