CS255: Cryptography and Computer Security **Winter 2015**

Assignment #1

Due: Wednesday, Jan. 28, 2015, in class.

Problem 0 In class we explained that the one time pad is malleable. Let's see a concrete example. Suppose you are told that the one time pad encryption of the message "attack at dawn" is 09e1c5f70a65ac51626bc3d25f17 (the plaintext letters are encoded as 8-bit ASCII and the given ciphertext is written in hex). What would be the one time pad encryption of the message "attack at dusk" under the same OTP key?

Problem 1 Simple secret sharing.

a. Suppose Alice shares a secret block cipher key, K_{AB} with Bob, and a different secret block cipher key, K_{AC} with Charlie. Describe a method for Alice to encrypt an m-block message such that it can only be decrypted with the cooperation of both Bob and Charlie. The ciphertext should only be a constant size greater than m blocks. You may assume that Bob and Charlie have a pre-established secret channel on which to communicate.

b. Now, suppose Alice shares a block cipher key, K_{AB} with Bob, a block cipher key K_{AC} with Charlie, and a block cipher key K_{AD} with David. Describe a method for Alice to encrypt an m-block message such that any two of Bob, Charlie, and David can decrypt (for example, Bob and Charlie can decrypt), but none of them can decrypt the message themselves. Again, the ciphertext should only be a constant size greater than m blocks. **Hint:** Pick a random message encryption key to encrypt the message with. Then add three ciphertext blocks to the ciphertext header.

c. How does your solution from part (b) scale as we increase the number of recipients? In other words, suppose Alice has a secret key with each of n recipients and wants to encrypt so that any k out of n recipients can decrypt, but any $k-1$ cannot. What would be the length of the header as a function of n and k?
Your answer shows that this solution scales poorly. We will discuss a far more efficient solution later on in the class.

Problem 2 The movie industry wants to protect digital content distributed on DVD's. We study one possible approach. Suppose there are at most a total of n DVD players in the world (e.g. $n = 2^{32}$). We view these n players as the leaves of a binary tree of height $\log_2 n$. Each node v_i in this binary tree contains an AES key K_i. These keys are kept secret from consumers and are fixed for all time. At manufacturing time each DVD player is assigned a serial number $i \in [0, n-1]$. Consider the set S_i of $1 + \log_2 n$ nodes along the path from the root to leaf number i in the binary tree. The manufacturer of the DVD player embeds in player number i the $1 + \log_2 n$ keys associated with the nodes in S_i. In this way each DVD player ships with $1 + \log_2 n$ keys embedded in it (these keys are supposedly inaccessible to consumers). A DVD movie M is encrypted as

$$DVD = \underbrace{E_{K_{root}}(K)}_{\text{header}} \Big\| \underbrace{E_K(M)}_{\text{body}}$$

where K is some random AES key called a content-key. Since all DVD players have the key K_{root} all players can decrypt the movie M. We refer to $E_{K_{root}}(K)$ as the header and $E_K(M)$ as the body. In what follows the DVD header may contain multiple ciphertexts where each ciphertext is the encryption of the content-key K under some key K_i in the binary tree.

a. Suppose the $1 + \log_2 n$ keys embedded in DVD player number r are exposed by hackers and published on the Internet (say in a program like DeCSS). Show that when the movie industry is about to distribute a new DVD movie they can encrypt the contents of the DVD using a header of size $\log_2 n$ so that all DVD players can decrypt the movie except for player number r. In effect, the movie industry disables player number r.
Hint: the header will contain $\log_2 n$ ciphertexts where each ciphertext is the encryption of the content-key K under certain $\log_2 n$ keys from the binary tree.

b. Suppose the keys embedded in k DVD players $R = \{r_1, \ldots, r_k\}$ are exposed by hackers. Show that the movie industry can encrypt the contents of a new DVD using a header of size $O(k \log n)$ so that all players can decrypt the movie except for the players in R. You have just shown that all hacked players can be disabled without affecting other consumers.

Side note: the AACS system used to encrypt Blu-ray and HD-DVD disks uses a related system. It was quickly discovered that bored hackers can expose player secret keys faster than the MPAA can revoke them.

Problem 3. Ciphertext expansion vs. security. Let (E, D) be a symmetric encryption scheme encrypting bit strings.

a. Suppose that for all keys and all messages m, the encryption of m is the exact same length as m. Show that (E, D) cannot be CPA-secure (i.e. cannot be semantically secure under a chosen plaintext attack).

b. Suppose that for all keys and all messages m, the encryption of m is exactly ℓ bits longer than the length of m. Show an attacker that can win the CPA security game using $2^{\ell/2}$ queries and non-negligible advantage (in fact, advantage close to $1/2$). Consequently the cipher becomes insecure if a key is used to encrypt $2^{\ell/2}$ messages.
Hint: for simplicity you may assume that every message m can be mapped to exactly 2^ℓ ciphertexts. Note that a similar statement can be shown to hold without this assumption. You may also assume that the message space contains more than 2^ℓ messages.

Problem 4 Advantage. The purpose of this problem is to clarify the concept of *advantage*. Consider the following two experiments EXP(0) and EXP(1):

- In EXP(0) the challenger flips a fair coin (probability 1/2 for HEADS and 1/2 for TAILS) and sends the result to the adversary $\mathcal{A}$.
- In EXP(1) the challenger always sends TAILS to the adversary.

The adversary's goal is to distinguish these two experiments: at the end of each experiment the adversary outputs a bit 0 or 1 for its guess for which experiment it is in. For $b = 0, 1$ let W_b be the event that in experiment b the adversary output 1. The adversary tries to maximize its distinguishing advantage, namely the quantity

$$\text{Adv} = \big| \Pr[W_0] - \Pr[W_1] \big| \quad \in [0, 1] \ .$$

The advantage Adv captures the adversary's ability to distinguish the two experiments. If the advantage is 0 then the adversary behaves exactly the same in both experiments and therefore does not distinguish between them. If the advantage is 1 then the adversary can tell perfectly what experiment it is in. If the advantage is negligible for all efficient adversaries (as defined in class) then we say that the two experiments are indistinguishable.

a. Calculate the advantage of each of the following adversaries:
   - $\mathcal{A}_1$: Always output 1.
   - $\mathcal{A}_2$: Ignore the result reported by the challenger, and randomly output 0 or 1 with even probability.
   - $\mathcal{A}_3$: Output 1 if HEADS was received from the challenger, else output 0.
   - $\mathcal{A}_4$: Output 0 if HEADS was received from the challenger, else output 1.
   - $\mathcal{A}_5$: If HEADS was received, output 1. If TAILS was received, randomly output 0 or 1 with even probability.

b. What is the maximum advantage possible in distinguishing these two experiments? Explain why.

Problem 5 Let us see why in CBC mode an unpredictable IV is necessary for CPA security.

a. Suppose a defective implementation of CBC encrypts a sequence of packets by always using the last ciphertext block of packet number i as the IV for packet number $i+1$ (up until a few years ago all web browsers implemented CBC this way). Construct an efficient adversary that wins the CPA game against this implementation with advantage close to 1. Recall that in the CPA game the attacker submits packets (a.k.a messages) to the challenger one by one and receives the encryption of those packets. The attacker then submits the semantic security challenge which the challenger treats as the next packet in the packet stream.

b. Can you suggest a simple fix to the problem from part (a) that does not add any additional bits to the ciphertext?

c. Suppose the block cipher (E, D) used for CBC encryption has a block size of n bits. Construct an attacker that wins the CPA game against CBC with a random IV (i.e. where the IV for each message is chosen independently at random) with advantage close to $1/2^n$.

Your answer for part (c) explains why CBC cannot be used with a block cipher that has a small block size (e.g. $n = 32$ bits). Note that there are many other problems with such a small block size, which is why AES has a block size of 128 bits.

Problem 6 PRFs. Let $F : K \times X \to Y$ be a secure PRF with $K = X = Y = \{0,1\}^n$.

a. Show that $F_1(k, x) = F(k, x)\|0$ is not a secure PRF. (for strings y and z we use $y\|z$ to denote the concatenation of y and z)

b. Prove that $F_2(k, x) = F(k,\ x \oplus 1^n)$ is a secure PRF. Here $x \oplus 1^n$ is the bit-wise complement of x. To prove security argue the contra-positive: a distinguisher $\mathcal{A}$ that breaks F_2 implies a distinguisher $\mathcal{B}$ that breaks F and whose running time is about the same as $\mathcal{A}$'s.

c. Let $K_3 = \{0,1\}^{n+1}$. Construct a new PRF $F_3 : K_3 \times X \to Y$ with the following property: the PRF F_3 is secure, however if the adversary learns the last bit of the key then the PRF is no longer secure. This shows that leaking even a *single* bit of the secret key can completely destroy the PRF security property.
Hint: Let $k_3 = k\|b$ where $k \in \{0,1\}^n$ and $b \in \{0,1\}$. Set $F_3(k_3, x)$ to be the same as $F(k, x)$ for all $x \neq 0^n$. Define $F_3(k_3, 0^n)$ so that F_3 is a secure PRF, but becomes easily distinguishable from a random function if the last bit of the secret key k_3 is known to the adversary. Prove that your F_3 is a secure PRF by arguing the contra-positive, as in part (b).

d. Construct a new PRF $F_4 : K^2 \times X \to Y$ that remains secure if the attacker learns any *single* bit of the key. Your function F_2 may only call F once. Briefly explain why your PRF remains secure if any single bit of the key is leaked.