CSci 5271: Introduction to Computer Security

Exercise Set 3 **due: March 27th, 2019**

Ground Rules. You may choose to complete these exercises in a group of up to two students. Each group should turn in **one** copy with the names of all group members on it. You may use any source you can find to help with this assignment but you **must** explicitly reference any source you use besides the lecture notes or textbook. An electronic (plain text or PDF) copy of your solution should be submitted on Canvas by 11:59pm on Wednesday, March 27th.

1. Caesar's block cipher. (40 pts) The Caesar cipher is a historical encryption method based on advancing letters circularly through the alphabet. To discuss it in a modern context on ASCII, we can consider it to be a block cipher with an 8-bit block and a 5-bit key k. The encryption function E_k is defined as:

$$E_k(b) = \begin{cases} \texttt{0x41} + ((b - \texttt{0x41} + k) \bmod 26) & \text{if } \texttt{0x41} \le b \le \texttt{0x5A} \\ \texttt{0x61} + ((b - \texttt{0x61} + k) \bmod 26) & \text{if } \texttt{0x61} \le b \le \texttt{0x7A} \\ b, & \text{otherwise} \end{cases}$$

Recall that `0x41` through `0x5A` are the ASCII codes for `A` through `Z`, and similarly `0x61` through `0x7A` are `a` through `z`. The inverse operation is just shifting by the same amount in the opposite direction, so $D_k = E_{26-k}$ (we use the convention that the result of mod is always positive when the modulus is). ROT-13 corresponds to the special case $E_{13} = D_{13}$.

Suppose we want to use this block cipher to encrypt the message "`VENI.`" using the key $k = 5$. (In hex, the plaintext is `0x56 0x45 0x4E 0x49 0x2E`.) Conveniently with an 8-bit block there is no need for padding, but we still need to choose a mode of operation. In (a)-(d), give the encryption of this message, as a sequence of hex bytes, under each of the following modes:

(a) ECB mode (the ancient Roman standard)

(b) CTR mode with an initial counter value of `0x45`

(c) CBC mode with an IV of `0x08`

(d) OFB mode with an IV of `0x66`

(e) CaesarCrypt S.p.A. is an Italian computer security company which builds on their national heritage to market modern block ciphers that also have an 8-bit block size, but they have taken the lesson that the original Caesar cipher had too small a key size. Their first flagship product CCEA1 was a 8-bit block cipher with a 2048-bit key size. Their new successor cryptosystem, named CCEA2, increases CCEA1's key size to 4096 bits. CaesarCrypt's marketing materials suggest that this yields an astronomical increase in security by a factor of 2^{2048}. What do you think of this security claim: can CCEA2 really be more secure than CCEA1?

(f) In fact there are some general problems that affect any block cipher with a small block size. Describe a chosen-plaintext attack that would easily break any 8-bit block cipher.

2. (Mis-)using message authentication codes. (35 pts) Armed with a copy of Schneier's *Applied Cryptography* from a used bookstore, Sly can't wait to design his own encrypted thinga-madoodad protocol. He starts off with a super-secure key exchange protocol that ends with Alice and Bob sharing secret keys for encryption (K_e) and authentication (K_a). Now he wants to design a secure symmetric channel using these keys.

(a) Sly decides at first that he wants to use a CBC-MAC based on AES with 128 bit blocks for integrity. He looks carefully at his key exchange protocol and realizes that an adversary can interfere to make Alice and Bob end up deciding on different keys. So the first message sent over by Alice will be $\tau_0 = \mathsf{cbcMAC}_{K_a}(0^{128}) = \mathsf{aesEncrypt}_{K_a}(0^{128})$. (The notation 0^n means n zero bits.) If Bob's local value doesn't check out, he aborts, otherwise the channel is usable. Afterwards, whenever Alice wants to send the message M over the secure channel, she'll compute $\tau_M \leftarrow \mathsf{cbcMAC}_{K_a}(M)$ and send the pair (M, τ_M) over the channel; Bob will check whether $\tau_M = \mathsf{cbcMAC}_{K_a}(M)$ and if so will conclude that Alice said M.

This is a pretty bad idea. Show how to use the values τ_0, M and τ_M to compose a message to Bob that will convince him Alice meant to say the two-block message (M, τ_M) instead of just M. Explain why your message will convince Bob that Alice meant to say (M, τ_M) rather than just M. Hint: try writing a recursive definition of CBC-MAC, and use the facts that for any string A, $A \oplus A = 0^{|A|}$ and $A \oplus 0^{|A|} = A$.

Since τ_M is just 128 random-looking bits, why is this a big deal?

(b) Sly's friend Sally notices the same attack on his scheme. She proposes a different method of authenticating (and encrypting) messages: ignore the key K_a. Instead, to authenticate and encrypt the message M, first compute $H(M)$ using SHA-256; then encrypt $(M, H(M))$ together, using AES-CTR encryption. So the message sent on the insecure channel would be $\mathsf{CTR\text{-}Encrypt}_{K_e}(M, H(M))$; Bob would decrypt the message using K_e, check that the last 256 bits of the plaintext are the hash of all of the previous bits, and accept the message if they are.

Show that this is also a bad idea: if Alice ever sends a ciphertext corresponding to the message M, where Eve knows M, Eve can generate a ciphertext corresponding to any message M', (of the same length as M) that Bob will accept. (For example, if Alice sends the message "ATTACK AT TEN AM" Eve can drop it and make Bob accept the message "GO BACK HOME BOB" instead.)

3. Hashing and Signing. (25 pts) Nearly every digital signature scheme works by first hashing a message to be signed (with a cryptographic hash function) and then performing some operation on the hash—so in essence, we are "signing the hash" and not the message. In particular, if Eve sees Alice's signature on the message M and can find a message $M' \neq M$ so that $H(M) = H(M')$, she can convince people that Alice signed M'. This is OK, since a good crypto hash function H will resist finding targeted collisions (second pre-images) like this.

Suppose our signature scheme uses a hash function H with an output length ℓ that is sufficient to resist second pre-images but NOT resistant to free collisions (e.g. the hash length is around 100-120 bits). Then it is possible that if Eve can get Alice to sign one of a pair of colliding messages, she can later claim that Alice signed the other.

The classic birthday attack works by hashing random messages until two have the same hash. This could already be a problem in some applications, but you might object that Alice is unlikely to agree to sign a random message. So let's think about how to create a collision with more specific messages.

Suppose that a message is "favorable" if it is something that Alice would sign, for example "I will pay $5 to McDonald's for my lunch." Suppose that a message is "undesirable" if it is something that Alice would not sign, like "I will pay $500,000 to Eve for her lunch." Notice that we can generate 256 different "favorable" messages from the example above, for instance by varying the number of space characters between words between 1 and 2. Extend this idea to show how to generate a pair of messages, one favorable and one undesirable, with the same hash. Your attack should compute about as many hashes as the birthday attack.

Then, describe how Eve completes the attack using the pair she generates to her advantage.