YALE UNIVERSITY
DEPARTMENT OF COMPUTER SCIENCE

CPSC 467a: Cryptography and Computer Security

Notes 7 (rev. 1)
Professor M. J. Fischer *September 22, 2005*

Lecture Notes 7

1 Message authentication codes (MACs)

So far in the course, we have been discussing only a single cryptographic application, namely, secret message transmission from Alice to Bob over a a publicly-readable channel. Our goal has been to maintain privacy in the face of an eavesdropper Eve.

Once we assume that Eve has the power to modify messages and generate her own messages as well as eavesdrop, life becomes more difficult. We usually call that kind of a malicious active adversary "Mallory" (to distinguish him from Eve, who only eavesdrops).

1.1 Attacks by malicious active adversaries

Encryption alone no longer solves Alice and Bob's problem. Alice sends $c = E_k(m)$, but Bob may receive a corrupted or forged $c' \neq c$. How then does Bob know that the message he receives really was sent by Alice?

The naive answer is that Bob computes $m' = D_k(c')$, and if m' "looks like" a valid message, then Bob accepts it as having come from Alice. The reasoning here is that Mallory, not knowing k, could not possibly have produced a valid-looking message.

For any particular cipher such as DES, that assumption may or may not be valid, but here are two things to watch out for:

1. There are three successively easier possible attacks in which Mallory might produce fraudulent messages:

   (a) He might produce $c' = E_k(m')$ for a message m' of his choosing.

   (b) He might produce a message c' for which the corresponding plaintext m' is a valid message, even though he could not choose m' in advance, nor perhaps he does not even know what m' is.

   (c) He might be able to alter a legitimate message c from Alice to produce a new message c' that corresponds to an altered form m' of the true message m. For example, if m represents an amount of money, it is conceivable that Mallory could find the encryption of $m + 1$ given the encryption of m, without knowing either m or $m + 1$.

   Attack (1a) is similar to, but not the same as, the notion of breaking the cryptosystem that we have been studying. We have been asking that it be hard for Eve to compute $m = D_k(c)$ knowing c but not k. To carry out attack (1a) requires that Mallory compute $E_k(m)$ knowing m but not k. It's conceivable that he could do the latter without being able to do the former.

   One form of attack (1b) clearly *is* possible, the so-called *replay* attack. This is when Mallory substitutes a legitimate old encrypted message c' for the current message c. It can be thwarted by adding timestamps and/or sequence numbers to the messages, so that Bob can recognize when old messages are being received. Of course, this only works if Alice and Bob anticipate the attack and incorporate appropriate countermeasures into the protocol they are using.

However, even if replay attacks are ruled out, a cryptosystem that is secure against attack (1a) might still permit attack (1b). There are all sorts of ways that Mallory can generate values c'. What gives us confidence that Bob won't accept one of them as being valid?

Attack (1c) might be possible even in a cryptosystem that is free from attacks (1a) and (1b). For example, if c_1 and c_2 are encryptions of valid messages, perhaps so is $c_1 \oplus c_2$. Whether or not it is depends entirely on particular properties of E_k. It does not follow in general from the difficulty of decrypting a given ciphertext. We will see some cryptosystems later which do have the property of being vulnerable to attack (1c). In some contexts, this can actually be a useful property, as we will see.

2. Cryptosystems are not always used to send natural language or other highly-redundant messages. For example, suppose Alice wants to send Bob her password to a web site. Knowing full well the dangers of sending passwords in the clear over the internet, she chooses to encrypt it instead. Since passwords are supposed to look like random strings of characters, Bob will likely accept anything he gets from Alice. He could be quite embarrassed (or worse) claiming he knew the correct password when in fact the password he thought was from Alice was actually a fraudulent one derived from a random ciphertext c' produced by Mallory.

What Alice and Bob need to solve their problem is called a *Message Authentication Code* or *MAC*. A MAC is generated by a function function $C_k(m)$ that can be computed by anyone knowing a secret key k. However, it should be hard for an attacker to find any pair (m, ξ) such that $\xi = C_k(m)$ without knowing k. In fact, this should remain hard even if the attacker knows a set of valid MAC pairs $\{(m_1, \xi_1), \ldots, (m_t, \xi_t)\}$, so long as m itself is not the message in one of the known pairs.

Using a MAC, Alice can send both $c = E_k(m)$ and also $\xi = C_k(m)$. Bob receives c' and ξ', possibly different from what Alice sent. Bob computes $m' = D_k(c')$ and then checks that $\xi' = C_k(m')$. If the check succeeds, then Bob accepts m' as a valid message from Alice. We say that Mallory successfully cheats if Bob accepts m' as valid but either $m' \neq m$ or Alice sent no message at all.

1.2 Computing MACs

A block cipher such as DES can be used to compute a MAC by making use of one of the ciphertext chaining modes, CBC or CFB. (See Lecture notes 4.) In these modes, the last ciphertext block c_t depends on all t message blocks $m_1, \ldots, m_t$. Therefore, we define $C_k(m) = c_t$. The result of this process is reputed to be a good MAC generation function. Note that the MAC is only a single block long, which in general is much shorter than the message. A MAC acts like a checksum for preserving data integrity, but it has the advantage that an adversary cannot compute a valid MAC for an altered message.

2 Asymmetric cryptosystems

A major advance in cryptography is the modern development of *asymmetric* (also called *2-key* or *public key*) cryptosystems. The idea is simple. Instead of having a single key k that is used by both Alice and Bob, an asymmetric cryptosystem has a pair of related keys, an *encryption key* e and a *decryption key* d. Alice encrypts a message m by computing $c = E_e(m)$. Bob decrypts by

computing $m = D_d(c)$.[1] As always, the decryption function inverts the encryption function, so the following is always satisfied:

$$m = D_d(E_e(m)).$$

What makes asymmetric systems useful is the additional requirement that it be difficult to find d from e, and more generally, that it be difficult to find m given both e and $c = E_e(m)$. Thus, the system remains secure even if the encryption key e is made public!

There are several reasons to make e public. One is that it is then possible for anybody to send a private message to Bob. Sandra need only obtain Bob's public key e and send Bob $E_e(m)$. Bob recovers m by applying D_d to the received ciphertext. This greatly simplifies the key management problem, for there is no longer the need for a secure channel between Alice and Bob for the initial key distribution (which I have carefully avoided talking about so far).

Of course, an active adversary Mallory can still do some nasty things. For example, he could send his own encryption key to Sandra when she attempts to obtain Bob's key. Not knowing she has been duped, Sandra would then encrypt her private data in a way that Bob could not read but Mallory could! If Mallory wanted to carry out this charade for a longer time without being discovered, he could intercept each message from Sandra to Mallory, decrypt the message using his own decryption key, then re-encrypt it using Bob's public encryption key and send it on its way to Bob. Bob, receiving a validly encrypted message, will be none the wiser to Mallory's shenanigans. This is an example of a *man-in-the-middle* attack.

The security requirements for an asymmetric cryptosystem are more stringent than for a symmetric cryptosystem. For example, the system must be secure against large-scale chosen plaintext attacks, for Eve can generate as many plaintext-ciphertext pairs as she wishes using the public encryption function $E_e()$.

The public encryption function also gives Eve a way to check the validity of a potential decryption. Namely, if Eve guesses that $D_d(c) = m_0$ for some candidate message m_0, she can check her guess by testing if $c = E_e(m_0)$. Thus, whether or not there is redundancy in the set of meaningful messages is of no consequence to her since she now has an independent way of determining a successful decryption.

Designing a good asymmetric cryptosystem is much harder than designing a good symmetric cryptosystem. All of the known asymmetric systems are orders of magnitude slower than corresponding symmetric systems. Therefore, they are often used in conjunction with a symmetric cipher to form a *hybrid* system. Here's how this works. Suppose (E^2, D^2) is a 2-key cryptosystem and (E^1, D^1) is a 1-key cryptosystem. To send a secret message m to Bob, Alice first generates a random *session key* k for use with the 1-key system. which she then uses to encrypt m. She then encrypts the session key using Bob's public key for the 2-key system and sends Bob both ciphertexts. In formulas, she sends Bob $c_1 = E^1_k(m)$ and $c_2 = E^2_e(k)$. Bob decrypts c_2 using $D^2_d()$ to obtain k and then decrypts c_1 using $D^1_k()$ to obtain m. This is much more efficient than simply sending $D^2_e(m)$ in the common case that m is much longer than k.

[1] We often get sloppy with notation when discussing asymmetric cryptosystems. Let $k = (e, d)$ be a key pair. We sometimes write k_e and k_d for the first and second keys, respectively, so we might use the rather cumbersome notation $E_{k_e}(m)$ and $D_{k_d}(c)$. But then we might simplify this by dropping the second-level subscripts to get the same notation we use for symmetric cryptosystems, namely $E_k(m)$ and $D_k(c)$. Nevertheless, it should still be understood that the "k" in E_k refers to the first element of the key pair, whereas the "k" in D_k refers to the second. In practice, it isn't generally as confusing as all this. but the potential for misunderstanding is there.