6.875J/18.425J Cryptography and Cryptanalysis	October 1, 2001
Lecture 7	
Instructor: Anna Lysyanskaya	*Scribed by: Nick Feamster*

1 Recap

1.1 Definition of GM-security for l-bit messages

Definition 1 *A cryptosystem (G, E, D) for message space $\{0,1\}^l$ is* GM-secure *if*

$$\begin{aligned}
&\forall PPTF\{Eve_k\}\forall m_0, m_1 \in \{0,1\}^l \exists f(k) \in negl \exists K \text{ such that for } k > K: \\
&\quad \Pr[\ (PK, SK) \leftarrow G(1^k); \\
&\qquad b \leftarrow \{0,1\}; \\
&\qquad c \leftarrow E(PK, m_b); \\
&\qquad b' \leftarrow Eve(PK, c) \quad : \quad b = b'] \leq 1/2 + f(k)
\end{aligned}$$

1.2 GM-Security for l-bits implies GM-security for strings

Theorem 2 *Suppose a cryptosystem (G', E', D') is GM-secure for one-bit messages. Then the following one-to-many transformation produces a cryptosystem (G, E, D) that is secure for l-bit messages:*

Key Generation G *: Same as G'.*

Encryption E *: Use E' for each bit of m, individually.*

Decryption D *: Using D', decrypt each individual bit.*

Sketch of Proof Suppose, for contradiction, that (G, E, D) is not secure. We wish to show that this implies that (G', E', D') is not secure either.

What does it mean for (G, E, D) not to be GM-secure? It means that we can find an adversary $\{A_k\}$ which, for some polynomial $p(k)$ and infinitely many k, can distinguish between messages m_0 and m_1 with probability non-negligibly better than random guessing. More precisely, if (G, E, D) is not GM-secure, then

$$\begin{gathered}
\Pr[(PK, SK) \leftarrow G(1^k); b \leftarrow \{0,1\}; \\
c \leftarrow E(PK, m_{k,b}); b' \leftarrow A_k(PK, c, m_0, m_1) \quad : \quad b = b'] \geq 1/2 + 1/p(k)
\end{gathered}$$

Now, let us define the $l+1$ *hybrid* messages as follows: for $0 \leq i \leq l$, m^i is obtained by taking the first i bits of m_1 and concatenating them with the last $l - i$ bits of m_0. Note that $m^0 \equiv m_0$ and $m^l \equiv m_1$.

The proof rests on the following claim:

Claim 3 *If there exists an adversary $\{A_k\}$ which can distinguish between m_0 and m_1 with probability polynomially better than 1/2, then for some $0 \leq i < l$,*

$$\Pr[(PK, SK) \leftarrow G(1^k); b \leftarrow \{0,1\}; \\ c \leftarrow E(PK, m^{i+b}); b' \leftarrow A_k(PK, c, m_0, m_1) \quad : \quad b = b'] \geq 1/2 + 1/(l * p(k))$$

Moreover, given A_k, m_0 and m_1, one can efficiently determine i.

Our reduction R receives as input a public key PK' from the cryptosystem (G', E', D') and a challenge ciphertext c'. Our goal is to determine whether c is an encryption of 0 or an encryption of 1.

Now, we create the challenge ciphertext for our adversary A_k, as follows: the first i blocks of c will be encryptions of the first i bits of m_1. The $i+1$st block will be the input ciphertext c'. The last $l-i-1$ blocks will be the encryptions of the last $l-i-1$ bits of m_0. Note that, in case c' is the encryption of 0, c will be the encryption of m_k^i; otherwise c will be the encryption of m_k^{i+1}. Our reduction will simply output whatever A_k outputs. Let us estimate the success probability of our reduction R_k:

$$\begin{aligned} \Pr[(PK', SK') \leftarrow G'; b \leftarrow \{0,1\}; c' \leftarrow E'(PK', b); b' \leftarrow R_k(PK', c') : b = b'] &= \\ \Pr[(PK', SK') \leftarrow G'; b \leftarrow \{0,1\}; c' \leftarrow E'(PK', b); & \\ c \leftarrow (E'(PK', m_{k,0}^1), \ldots, E'(PK', m_{k,0}^i), c', E'(PK', m_{k,1}^{i+2}), \ldots, E'(PK', m_{k,1}^l)); & \\ b' \leftarrow A_k(PK', c) : b = b'] &= \\ \Pr[(PK, SK) \leftarrow G; b \leftarrow \{0,1\}; c \leftarrow E(m_k^{i+b}); b' \leftarrow A_k(PK, c) : b = b'] &\geq 1/2 + 1/(l * p(k)) \end{aligned}$$

So, we have proved that GM-security for bits implies GM-security for strings by showing that (G, E, D) is not l-GM-secure $\Rightarrow$ (G', E', D') is not GM-secure. ■

2 One-pass security

Let us revise the definition of security. So far, the scenario we have envisioned is like this:

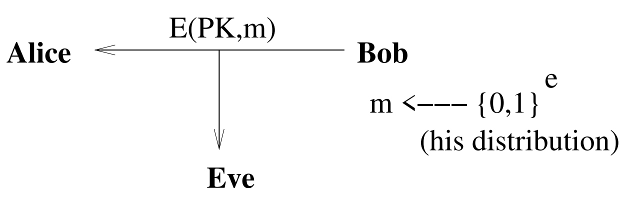

Figure 1: Scenario for traditional definition of security. Eve never talks to anyone, Bob decides what to send independently of any other party. The contents of Bob's message stays secret even when Eve has a good guess about the contents of the message.

But what if Eve can actually set up Bob's vocabulary? For example, in case of a love letter, we can imagine that Eve is Bob's ex, and she's the one who taught him how to express "I love you" and "I hate you." So, instead, the scenario looks like this:

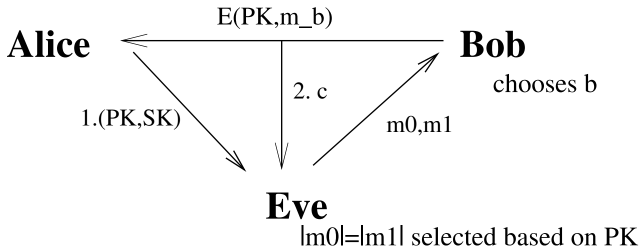

Figure 2: In this scenario, we consider the case where Eve is Bob's "ex-girlfriend". That is, Eve taught Bob some vocabulary, such as how to express messages such as "I love you" and "I hate you." We also assume that Eve was jealous of Alice from the start, so the vocabulary that Eve gave to Bob was dependent on Alice's public key. Note also that Eve has an advantage in this case because she restricts Bob to sending messages that she knows.

What happens is that: first, Alice publishes a public key. Then Eve figures out which two messages she is best at distinguishing. Then one of these messages is encrypted. Eve should have no idea which one it was.

Let us write this, stronger definition, in our notation. First, we observe that the adversary Eve_k is invoked twice, first with Alice's public key PK, and then with the challenge. So we let α denote some useful information that Eve_k may save for itself for later.

Definition 4 *(One-pass GM-security) A cryptosystem (G, E, D) for is* one-pass GM-secure *if*

$$\begin{aligned}
&\forall PPTF\{Eve_k\}\, \exists f(k) \in negl\, \exists K \text{ such that for } k > K : \\
&\quad \Pr[\ (PK, SK) \leftarrow G(1^k); \\
&\qquad (\alpha, m_0, m_1, |m_0 = m_1|) \leftarrow Eve_k(PK), \\
&\qquad b \leftarrow \{0,1\}; \\
&\qquad c \leftarrow E(PK, m_b); \\
&\qquad b' \leftarrow Eve(\alpha, c) \quad : \quad b = b'] \leq 1/2 + f(k)
\end{aligned}$$

Note that in this definition, we are not limited to messages of any particular length: the adversary decides on the length of the messages. But the price we pay is that we allow him to distinguish between encryptions of messages of different lengths for free; only encryptions of same-length messages should be indistinguishable.

Note also that the use of the parameter α that is passed as input to Eve is a piece of state that implies that she also knows m_0, m_1, and PK.

Is this definition stronger than the other one? Yes.

Suppose a cryptosystem is not GM-secure. Then there exist some l-bit messages m_0, m_1 and an adversary Eve who distinguishes between their encryptions. Then the one-pass adversary can simply always ask for m_0 and m_1 to be encrypted.

What if a cryptosystem is not one-pass GM secure — can it still be GM-secure? Yes. For example, consider a GM-secure cryptosystem, modified such that $E(PK, PK)$ is always just PK, while everything else is the same. Then it is clear that a one-pass adversary can always distinguish the encryption of PK from the encryption of, say 0^l, but it is still GM-secure, since for any fixed m_0, m_1, the probability that $PK = m_i$ is negligible.

3 1-pass Security is stronger than 3-Pass Security

In this section, we show that 1-pass security is stronger than 3-pass security. We do so by showing that 1-pass security implies 3-pass security, but 3-pass security does not imply 1-pass security.

Theorem 5 *1-pass Security $\Rightarrow$ 3-pass Security.*

This proof is trivial. Suppose we have an Eve_k such that for some m_0, m_1 she can distinguish between the encryptions of the two messages (i.e., Eve can break 3-pass security).

Then Eve_k can break 1-pass security simply by sending the *same* m_0, m_1 to Bob – i.e., the ones for which she can break 3-pass security.

By the contrapositive, 1-pass security implies 3-pass security.

Theorem 6 *3-pass security does not imply 1-pass security.*

The proof is by counterexample. We assume that 3-pass security implies 1-pass security and then show that this cannot be true.

Let (G, E, D) be a GM-secure cryptosystem. Consider (G', E', D'):

$$\begin{aligned}
G' &= \text{same as before} \\
E'(PK, m) &= 0 \circ PK \text{if } m = PK \\
& 1 \circ E(PK, m) \ \text{ otherwise.} \\
D'(PK, SK, c) &= PK \text{ if } c = 0 \circ PK \\
& D(PK, SK, c') \ \text{ if } c = 1 \circ c'
\end{aligned}$$

Note that (G', E', D') is still 3-pass secure:

Suppose $\forall m_0, m_1, \forall Eve, \exists f(k) \in negl$ such that:

$$\begin{aligned}
&\forall PPTF\{Eve_k\}\, \forall m_0, m_1 \in \{0,1\}^l \exists f(k) \in \text{negl}\, \exists K \text{ such that for } k > K: \\
&\quad \Pr[\ (PK, SK) \leftarrow G(1^k); \\
&\qquad b \leftarrow \{0,1\}; \\
&\qquad c \leftarrow E(PK, m_b); \\
&\qquad b' \leftarrow \text{Eve}(PK, c) \quad : \quad b = b'] \leq 1/2 + f(k)
\end{aligned}$$

However, if the definition does not hold for (G', E', D'), then:

$$\begin{aligned}
&\exists PPTF\{Eve_k\}\, \exists m_0, m_1 \in \{0,1\}^l \exists f(k) \in \text{negl} \\
&\quad \Pr[\ (PK, SK) \leftarrow G(1^k); \\
&\qquad b \leftarrow \{0,1\}; \\
&\qquad c \leftarrow E(PK, m_b) \text{or} PK = m_b; \\
&\qquad b' \leftarrow \text{Eve}(PK, c) \quad : \quad b = b'] > 1/2 + 1/poly
\end{aligned}$$

Thus, (G', E', D') can only become insecure if $\exists m_0, m_1$ such that:

$$\begin{aligned}
&\Pr[\ (PK, SK) \leftarrow G(1^k); \\
&\qquad b \leftarrow \{0,1\}; \\
&\qquad PK = m_b] \geq 1/poly
\end{aligned}$$

If this were possible, then an Eve equipped with SK corresponding to m_b would break 3-pass security. This is a contradiction, since we have assumed 3-pass security. Thus, by contradiction, we have shown that 3-pass security does not imply 1-pass security.

Theorem 7 *1-bit GM-security $\Rightarrow$ one-pass GM-Security. (If you encrypt bit-by-bit.)*

Sketch of Proof This proof goes roughly the same as last time (when we showed 1-bit security implies l-bit security).

Suppose we have an Eve_k that breaks one-pass security. Then we can produce the following reduction:

$Reduction(PK, c)$
$(m_0, m_1, \alpha) \leftarrow Eve_k(PK)$
Determine $1 \leq i < |m_0| = |m_1|$ such that eve can distinguish this hybrid
$c' \leftarrow E(PK, (m_1^1 \circ \ldots \circ m_1^i) \circ c \circ E(m_0^{i+2} \circ \ldots \circ m_0^{|m_1|})$
$b' \leftarrow Eve_k(\alpha, c')$
Return b'

Note that this is exacrly how we broke the 1-bit security scheme from last time. ■

4 Semantic security

Intuition tells us that the right way to define security is: "Whatever can be computed about the message from the ciphertext, can be computed without the ciphertext." Let us try to capture this intuition.

What if Eve has even greater power over Bob? Let us generalize this definition even further. Suppose that Eve, upon examining Alice's public key, decides that the message should be selected according to some efficiently samplable distribution M. Further, suppose that Eve's goal is not to determine which specific message was sent, but rather, to compute some function of the message, for example, to determine whether this is a love letter or a hate letter. This scenario is summarized in Figure 3.

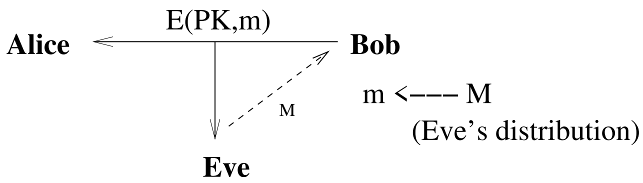

Figure 3: A generalized scenario for messages passed between Alice, Bob, and Eve.

Let us be even more general than that. Suppose that instead of computing a function on the encrypted message, the adversary is trying to compute a relation R. Figure 4 shows the distinction between a function and a relation. While a function maps one point in one space to *exactly* one point in another space, a relation can map one point in one space to *multiple* points in another space. In this sense, a function is a special case of a relation.

The following are examples of relations (*not* functions):

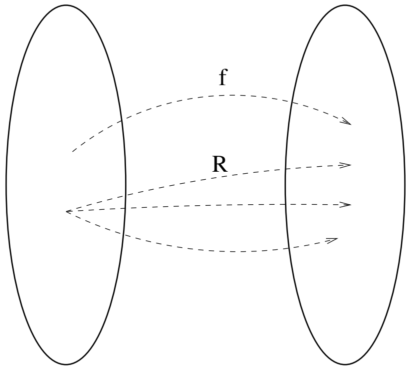

Figure 4: The distinction between a function and a relation. While a function maps one point in one space to *exactly* one point in another space, a relation can map one point in one space to *multiple* points in another space. In this sense, a function is a special case of a relation.

- $R(m_1, m_2)$ is true if m_1 and m_2 both contain the phrase "I love you". (For a fixed m_1, there may be many values for m_2 that satisfy this relation.)
- $R(m_1, m_2)$ is true if m_1 and m_2 overlap in 10 letters or more.

Eve's goal is now to compute some relation R, or some value z such that $R(m, z)$ is true.

Definition 8 *(One-pass semantic security) A cryptosystem* (G, E, D) *for is* one-pass semantically *if*

$$\forall PPTF\{Eve_k\}\,\exists PPTF\{S_k\}, f(k) \in negl\,\exists K \textit{ such that for } k > K,\ |p_0(k) - p_1(k)| \leq f(k), \textit{ where}$$

$$\begin{aligned} p_0(k) = \Pr[\ & (PK, SK) \leftarrow G(1^k); \\ & (\alpha, M, R) \leftarrow Eve_k(PK) \\ & m \leftarrow M; \\ & c \leftarrow E(PK, m); \\ & z \leftarrow Eve(\alpha, c) \quad : \quad R(m, z)]. \\ p_1(k) = \Pr[\ & (PK, SK) \leftarrow G(1^k); \\ & (\alpha, M, R) \leftarrow Eve_k(PK) \\ & m \leftarrow M; \\ & z \leftarrow S(\alpha, M, R) \quad : \quad R(m, z)]. \\ p_0 - p_1 \leq\ & f(k) \end{aligned}$$

Note that this is a very strong definition of security. Also note that if Eve receives PK, then we must restrict the possible relations to those that can be decided in polynomial time.

In the next lecture, we will prove that one-pass semantic security is equivalent to one-pass GM-security. This is very powerful, because it implies that if you can show that distinguishing between two messages is impossible, then you can show that the adversary can gain no information at all about the original messages (i.e., the adversary cannot even compute relations or functions of the original messages).

This essentially makes proving semantic security much easier, because proving one-bit GM-security implies semantic security. In other words, if a (G, E, D) cryptosystem is secure for 1-bit, then bit-by-bit encryption in that cryptosystem is semantically secure. Thus, if follows that the GM-cryptosystem is semantically secure.