

Lecture 8.a

*Lecturer: Victor Shoup**Scribe: Joël Alwen*

The first part of this lecture covers relevant building blocks and definitions in order to construct semantically secure *public key* encryption schemes. The second consists of two constructions which are proven secure under (relatively) standard complexity theoretic assumptions which can be instantiated via well known number theoretic assumptions.

1 Notation and Primitives

1.1 Trapdoor Functions

We begin with the definition of (and notation for) a classic cryptographic primitive: the trapdoor one-way function (TOWF). There existence is one of the most general (weakest) complexity theoretic assumption that can be made (as apposed to information or number theoretic ones) subsided only by the assumption of the existence of plain one-way functions and permutations. First introduced in 1976 by Diffie and Hellman in their seminal work [DH], various important tools can be constructed from TOWF such as semantically secure public key cryptosystems.

Intuitively the one-wayness property means that (except for very small security parameters) no efficient adversary has any significant chance of finding a pre-image to y when seeing only y and the public key.

Definition 1. Let $T = (G, F, I)$ be a triple of functions with security parameter α such that:

- $G(1^\alpha) \rightarrow (pk, sk)$ is a generator for T
- $F(pk, x) = y$ with $x \in X$ and $y \in Y$ evaluates the function
- $I(sk, y) = x$ inverts the function

Then T is called a **trapdoor one-way function** if the following 2 conditions hold.

Correctness : $F(pk, \cdot)$ is injective (a.k.a. onto 1-1) and
 $\forall x \in X \ I(sk, F(pk, x)) = x$

One-wayness : Let $\text{negl}()$ be a negligible function. Then $\exists n \in \mathbb{N}$ such that $\forall \alpha > n$ and $\forall A \in PPT$ efficient adversaries:

$$\Pr [F(pk, \tilde{x}) = y | (pk, sk) \leftarrow G(1^\alpha), x \leftarrow X, y = F(pk, x), \tilde{x} = A(pk, y)] \leq \text{negl}(k)$$

It turns out that candidate functions fitting this definition are harder to find than first thought. (The originally proposed subset-sum problem and several later candidates derived from various NP-complete problems have been demonstrated unsuitable. The main problem arises from the requirement that one-wayness hold for $x \leftarrow X$ a random and uniformly distributed x not just for a worst case value of x . Luckily there seems to be at least 1 good candidate, namely the RSA function (with the slight wrinkle that the sets X and Y actually depend on the specific public key chosen). Another (less used) candidate is the Rabin function ($R(n = pq, x) = x^2 \bmod n$) for which one-wayness can even be reduced to the hardness of factoring. Note that this has not been shown for RSA.

1.2 Semantically Secure Symmetric-Key Encryption

The first of the two constructions of SS public-key cryptosystems we will discuss is based on the symmetric-key variant of SS cryptosystems. The proof of security of the public-key scheme will then reduce to breaking the security of the symmetric scheme so we now give a formal definition.

Definition 2. Let $\text{negl}()$ be a negligible function. Then a symmetric key encryption scheme $\mathcal{E}_s = (E_s, D_s)$ with key-space \mathcal{K}_s , message-space \mathcal{M} , ciphertext-space \mathcal{C} and security parameter α is called a **semantically secure** if $\exists n \in \mathbb{N}$ such that $\forall \alpha > n$ and $\forall A_{SSS} \in PPT$ efficient adversaries it holds that:

$$\Pr[\tilde{b} = b \mid k \leftarrow \mathcal{K}_s, (m_0, m_1) \leftarrow A, b \leftarrow \{0, 1\}, c \leftarrow E_s(k, m_b), \tilde{b} \leftarrow A(c, m_0, m_1)] \leq \text{negl}(k)$$

As usual an alternative formulation of this definition looks at the outcome of a game between a symmetric-key semantic security attacker A_{SSS} and a challenger C_s .

- System Parameters:
 - Encryption scheme \mathcal{E}_s
 - Message-space \mathcal{M}
 - Security parameter α
- The Game:
 1. $(m_0, m_1) \leftarrow A$ and sends them to the C_s
 2. C_s chooses $k \leftarrow \mathcal{K}_s$ and $b \leftarrow \{0, 1\}$ and sends $c = E_s(k, m_b)$ to A
 3. A guesses \tilde{b} the value of b and sends the guess to C_s

The advantage Adv_{SSS} is defined as $|\Pr[b = \tilde{b}] - \frac{1}{2}|$ and \mathcal{E}_s is said to be semantically secure if Adv_{SSS} is a negligible function in α .

A common candidate that fits the above definitions is the AES algorithm in cipher-block chaining (CBC) mode with key-space $\mathcal{K}_s = \{0, 1\}^{128}$ and message space $\{0, 1\}^*$.

1.3 Hash Function

As a third tool we will be needing a hash function $H : X \rightarrow \mathcal{K}_s$. However, at this point there is no need for more precise definition's of it's properties as, in the first proof of security, it will be modeled as a random oracle. Later we define a specific property of H which will enable the removal of random oracles from the proof.

2 Construction 1

We are now ready to begin with the first construction of a semantically secure public key-cryptosystem.

Intuitively the system works as follows. It uses a hash function (modeled as a random oracle) to generate a key for E_s . In order that the decryptor can recover this key, a TOWF is applied to the input of the random oracle. Thus by the one-wayness property of the TOWF, only the intended recipient (who holds the appropriate trap-door) can recover the correct input so as to query the random oracle at the right point. Without this input the probability of guessing the correct key for D_s is negligible by a standard information theoretic argument on the output distribution of random oracles. So if an attacker doesn't have the decryption key for the E_s algorithm then the construction is secure by the security properties of \mathcal{E}_s .

2.1 Description

We now give the details of the construction. Let $T = (G, F, I)$ be a TOWF, $H : X \rightarrow \mathcal{K}_s$ a hash function modeled as a random oracle, and \mathcal{E}_s a semantically secure symmetric-key encryption scheme as above. Then the semantically secure public-key encryption scheme is a triple of efficient algorithms $\mathcal{E} = (Gen, E, D)$ with message-space \mathcal{M} and ciphertext-space \mathcal{C} where:

| | | |
|--------------------------------------|-----------------------------|--------------------------|
| $Gen(1^\alpha) :$ | $E(pk, m) :$ | $D(sk, (y, c)) :$ |
| 1. $(pk, sk) \leftarrow G(1^\alpha)$ | 1. $x \leftarrow X$ | 1. $x = I(sk, y)$ |
| | 2. $k = H(x)$ | 2. $k = H(x)$ |
| | 3. $y = F(k, x)$ | 3. output($D_s(k, c)$) |
| | 4. $c \leftarrow E_s(k, m)$ | |
| | 5. output(y, c) | |

2.2 Proof of Security

The security properties of this scheme are summed up by the following theorem:

Theorem 3. *Assuming T is a TOWF, H is a random oracle and \mathcal{E}_s is semantically secure then \mathcal{E} is a semantically secure public-key encryption scheme.*

Proof. Let $L[\cdot]$ be an array, indexed by X , of elements in \mathcal{K}_s and let $A \in PPT$ be an efficient adversary attacking the semantic security of the scheme. The proof follows from the interpretation and comparison of two games between a challenger and A . The following are the descriptions of the challenger's roles in the games:

- Game 0:
 1. Compute $(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}(1^\alpha)$, $x \leftarrow X$ and $y = F(\mathbf{pk}, x)$
 2. Compute $k \leftarrow \mathcal{K}_s$ and $b \leftarrow \{0, 1\}$ and set $L[\tilde{x}] := \perp \forall \tilde{x} \in X$
 3. Set $L[x] := k$
 4. From this point on, upon receiving RO query $\tilde{x} \in X$ from A , if $L[\tilde{x}] = \perp$ then set $L[\tilde{x}] \leftarrow \mathcal{K}_s$. Respond with $L[\tilde{x}]$.
 5. Send \mathbf{pk} to A receiving encryption query (m_0, m_1)
 6. Send $c \leftarrow E_s(k, m_b)$ to A receiving adversaries guess \tilde{b} .
- Game 1:
 - Same as game 0 without step 3.

Now we wish to show that the random variables in these two games have (almost) the same distribution (were "almost" means with at most negligible difference). Let Z be the event that A queries the RO at x in game 1. Observe that if A *doesn't* make a query at x then the two games proceed identically. Now for game $i \in \{0, 1\}$ let the event $\tilde{b} = b$ be denoted by W_i . Then the above observation can be formulated as $W_0 \wedge \bar{Z} \Leftrightarrow W_1 \wedge \bar{Z}$.

Lemma 4. "*Difference Lemma*"

If, for 3 events W_0, W_1 and Z it holds that $W_0 \wedge \bar{Z} \Leftrightarrow W_1 \wedge \bar{Z}$ then

$$|Pr[W_0] - Pr[W_1]| \leq Pr[Z]$$

The difference lemma tells us that if it can be shown that $Pr[Z]$ is negligible and $Pr[W_1]$ is negligibly far from $\frac{1}{2}$ then so is $Pr[W_0]$. (Note that \mathcal{E} is semantically secure iff $|Pr[W_0] - \frac{1}{2}|$ is negligible.)

Lemma 5. *$Pr[Z]$ is negligible in α .*

Proof. We construct an efficient algorithm $B_T \in PPT$ which breaks the one-wayness of T with advantage equal to $Pr[Z]$.

B_T is given a public key \mathbf{pk} and a challenge image $y = F(\mathbf{pk}, x)$ (for unknown x) as input. It simulates the challenger in game 1 to A exactly as the challenger would play except:

- In step 1 compute only $x \leftarrow X$.
- Upon receiving a RO query \tilde{x} from A check whether $y = F(\mathbf{pk}, \tilde{x})$. If so output y and terminate. Otherwise proceed as the challenger would.

- If A terminates then quit with no output.

The only way for B_T to win it's game against the challenger is for A to query at \tilde{x} equal to a pre-image of y . If this happens then B_T wins with probability 1 and if it doesn't happen then B_T wins with probability 0. In other words $Pr[B_T \text{ wins}] = Pr[Z]$. So by the one-wayness of T we have that $Pr[Z]$ is negligible in α . \square

Lemma 6. $|Pr[W_1] - \frac{1}{2}|$ is negligible in α .

Proof. We construct an efficient algorithm B_S which breaks the semantic security of \mathcal{E}_s with advantage equal to $|Pr[W_1] - \frac{1}{2}|$.

Conceptually B_S plays the game from definition 2 by sitting in between the semantic security challenger C_s (for symmetric-key schemes) and A (from game 1). B_S basically emulates the semantic security challenger (for public-key schemes) to A and wins iff A wins.

- Play round 1-4 of game 1 honestly
- Upon receiving encryption challenge (m_0, m_1) from A send it to C_s forwarding the response c back to A .
- When A sends it's guess \tilde{b} forward it to C_s and quit.

It is clear that A 's view of this interaction is identical to it's view in game 1. Therefor it's probability of guessing the value of b is identical and further if A guesses correctly then B_S wins against C_s . If the guess is bad though then B_S fails. Therefor the advantage of B_S is identical to that of A . So by the semantic security of \mathcal{E}_s we have that $|Pr[W_1] - \frac{1}{2}| \leq \text{negl}(\alpha)$. \square

By combining the Difference Lemma with lemmas 5 and 6 we have that $|Pr[W_0] - \frac{1}{2}| \leq \text{negl}(\alpha)$ thereby proving the theorem. \square

3 Construction 2

Now we will cover a second construction based on different primitives. Though these primitives are less common then the ones used by the previous construction, an advantage of this scheme is that it does not require the random oracle model to be shown secure.

Before we give the details of construction 2 we first cover three definitions: trapdoor function pair schemes, the notion of unpredictability and that of a secure key derivation function.

3.1 Trapdoor Function Pair Schemes

Definition 7. Let $P = (G, F_0, F_1, I)$ be a quadruple of efficiently computable functions with security parameter α such that:

- $G(1^\alpha) \rightarrow (\mathbf{pk}, \mathbf{sk})$
- $F_0 : X \rightarrow U$ with $F_0(\mathbf{pk}, x) = u$
- $F_1 : X \rightarrow V$ with $F_1(\mathbf{pk}, x) = v$
- $I : U \rightarrow V$ with $I(\mathbf{sk}, u) = v$

Then P is called a **trapdoor function pair scheme (TFP)** iff:

Correctness : $F_0(\mathbf{pk}, \cdot)$ and $F_1(\mathbf{pk}, \cdot)$ are injective (a.k.a. onto 1-1) and $\forall x \in X I(\mathbf{sk}, F_0(\mathbf{pk}, x)) = F_1(\mathbf{pk}, x)$

To exemplify this definition we give two special cases of TFPs.

1. Let $T = (G, F, I)$ be a TOWF. Then (G, F, Id, I) is a TFP where Id is the identity function and $V = X$.
2. Let F be a one-way permutation family (indexed by \mathbf{pk}), B be the associated hard-core predicates, and $I = B \circ F^{-1}$ where F^{-1} is the family of inverses to F indexed by \mathbf{sk} . Then (G, F, B, I) is TFP.

Definition 7 alone is not enough though. We need some intractability property much like the one-wayness of TOWF.

Definition 8. A TFP $P = (G, F_0, F_1, I)$ with sets X, U and V is called **unpredictable** iff:

Unpredictability : Let $\text{negl}(\cdot)$ be a negligible function. Then $\exists n \in \mathbb{N}$ such that $\forall \alpha > n$ and $\forall A \in \text{PPT}$ efficient adversaries:

$$\Pr[\tilde{v} = v | (\mathbf{pk}, \mathbf{sk}) \leftarrow G(1^\alpha), x \leftarrow X, u = F_0(\mathbf{pk}, x), v = F_1(\mathbf{pk}, x), \tilde{v} \leftarrow A(\mathbf{pk}, u)] \leq \text{negl}(\alpha)$$

As it turns out the above two special cases of TFPs are both also unpredictable (by the one-wayness of T for the first case, and by the hard-core property of B in the second case). An even more specific example of an unpredictable TFP is the following instantiation based on the CDH assumption.

- Let \mathbb{G} be a group of prime order q and let $g \in \mathbb{G}$ be a random generator
- Let $a \leftarrow \mathbb{Z}_q$, $h = g^a$, $\mathbf{pk} := (g, h)$ and $\mathbf{sk} := a$
- Then $F_0(\mathbf{pk}, b) = g^b$, $F_1(\mathbf{pk}, b) = h^b$ and $I(\mathbf{sk}, b) = b^a$ is a TFP with $X = \mathbb{Z}_q$ and $U = V = \mathbb{G}$

Finally we mention another useful property a function can have which will allow for a proof of security in the plain model (rather than the random oracle one).

Definition 9. Let \mathcal{E}_s be an encryption scheme with key-space \mathcal{K}_s and let $H : V \rightarrow \mathcal{K}_s$ be a function. Then H is called a **secure key derivation function** iff:

$$H(V) \sim_i \mathcal{K}_s$$

In other words if \mathcal{V} and \mathcal{K} are uniformly distributed random variables over V and K respectively then no efficient algorithm, given oracle (sampling) access to $H(\mathcal{V})$ and \mathcal{K} , can tell them apart with more than negligible advantage.

Two good examples of such functions are extractors [DRS04] and the random element of universal hash functions. (See the Leftover hash lemma in [ILL89].)

3.2 Construction Details

We are now ready to give the details of construction 2.

Let $P = (G, F_0, F_1, I)$ be a TFP, $H : X \rightarrow \mathcal{K}_s$ a hash function, and \mathcal{E}_s a semantically secure symmetric-key encryption scheme as above. Then the semantically secure public-key encryption scheme is a triple of efficient algorithms $\mathcal{E} = (Gen, E, D)$ with message-space \mathcal{M} and ciphertext-space \mathcal{C} where:

| | | |
|---|--|--|
| $Gen(1^\alpha) :$ 1. $(\mathbf{pk}, \mathbf{sk}) \leftarrow G(1^\alpha)$ | $E(\mathbf{pk}, m) :$ 1. $x \leftarrow X$ 2. $u = F_0(\mathbf{pk}, x)$ 3. $v = F_1(\mathbf{pk}, x)$ 4. $\mathbf{k} = H(v)$ 5. $c \leftarrow E_s(\mathbf{k}, m)$ 6. $\text{output}(u, c)$ | $D(\mathbf{sk}, (u, c)) :$ 1. $v = I(\mathbf{sk}, u)$ 2. $\mathbf{k} = H(v)$ 3. $\text{output}(D_s(\mathbf{k}, c))$ |
|---|--|--|

3.3 Proof of Security

The security of construction 2 can be summarized by the following theorem.

Theorem 10. *If P is unpredictable, \mathcal{E}_s is semantically secure and H is a random oracle then then construction 2 is a semantically secure public-key encryption scheme.*

Say we use the TFP instantiation based on the CDH assumption. Then g is a system parameter and the output of the encryption function consists of $(u = g^b, c)$. By the DDH assumption the output of $F_1(\mathbf{pk}, x) = h^b = g^{ab}$ is indistinguishable from a random element of \mathbb{G} . Therefore, if H is a secure key derivation function, $H(F_1(\mathbf{pk}, x))$ is a random key for \mathcal{E}_s (from the adversaries point of view). So no random oracle assumption need be made about H . In other words we can now restate the security of construction 2 as:

Theorem 11. *If*

1. P is unpredictable
2. F_1 is pseudo-random
3. \mathcal{E}_s is semantically secure
4. H is a secure key derivation function

then construction 2 is a semantically secure public-key encryption scheme.

Proof. As in the case of the previous construction we prove the theorem by comparing several attack games. This time though, we will not go into all the details concerning the construction of each reduction but merely provide a proof sketch. (Note: We use the same notation W_i for $i \in \{0, 2\}$ to denote the event that $\tilde{b} = b$ in game i .)

- 0) Game 0 is the standard attack game for semantic security of public-key cryptosystems.
- 1) Game 1 is as game 0 except that when an oracle query for v is made by the attacker v is replaced with a random value \tilde{v} for the query.
 - Note that Game 1 and 0 only differ if A actually queries v . Therefore $|Pr[W_0] - Pr[W_1]| \leq \epsilon_1$ where ϵ_1 is the advantage of A at breaking the unpredictability of P . (This can be shown with a reduction much like the algorithm B_T in the previous proof of security.)
- 2) Game 2 is as game 1 except that in the oracle query for the encryption procedure $H(v)$ is replaced by $k \leftarrow \mathcal{K}_s$.
 - Here again Game 2 and 1 only differ for an adversary that either breaks the pseud-randomness assumption or the secure key derivation assumption. That is $|Pr[W_1] - Pr[W_2]| \leq \epsilon_2 + \epsilon_4$

Via a straightforward reduction much as in the previous proof of security for B_S we have that $|Pr[W_2] - \frac{1}{2}| = \epsilon_2$ where ϵ_2 is the advantage at breaking the semantic security of \mathcal{E}_s . Therefore

$$|Pr[W_0] - \frac{1}{2}| \leq \sum_{i=1}^4 \epsilon_i \leq \text{negl}(\alpha)$$

and so construction 2 is semantically secure. \square

4 Next Time

In the following lecture we will begin our discussion of chosen ciphertext security and various constructions secure in both the random oracle model and the plain model.

References

- [DH] W. Diffie and M. E. Hellman. New directions in cryptography. *IEEE Trans. Info. Theory IT-22*.
- [DRS04] Y. Dodis, L. Reyzin, and A. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In *EUROCRYPT*, volume 3027 of *Lecture Notes in Computer Science*, pages 532–540. Springer, 2004.
- [ILL89] R. Impagliazzo, L. A. Levin, and M. Luby. Pseudo-random generation from one-way functions. In *STOC '89: Proceedings of the twenty-first annual ACM symposium on Theory of computing*, pages 12–24, New York, NY, USA, 1989. ACM Press.