Lecture 9: Public Key Encryption III

MIT - 6.5620
Lecture 7 (October 7, 2026)

> ***Warning:*** This document is a rough draft and may contain bugs. Please feel free to email us with corrections.

Recap:

In the last lecture, we presented the El-Gamal encyrption scheme and looked at constructions of trapdoor one-way functions, namely RSA and Rabin.

Today:

Today, we will look at public-key encryption schemes based on trapdoor one-way functions. Note that these notes only cover up to slide 25 of the lecture slides; for the remainder of the material covered in lecture, please see the slides and/or the lecture video.

Number-Theory Review

The constructions we will see are based on number theory, but this time over the group

$$\mathbb{Z}_N^* = \{x : \ \mathsf{GCD}(x, N) = 1\}.$$

where N is a product of two large primes $p, q \leftarrow \{0,1\}^\lambda$. Note that

$$|\mathbb{Z}_N^*| = N - p - q + 1 = (p-1)(q-1).$$

What makes this group very different from $\mathbb{Z}_p^*$ where p is prime, is that $\mathbb{Z}_N^*$ is a group of unknown order (assuming the hardness of Factoring).

*The Chinese Remainder Theorem (*CRT*)* the CRT theorem asserts that for $N = p \cdot q$, the group $\mathbb{Z}_N^*$ is isomorphic to the group $\mathbb{Z}_p^* \times \mathbb{Z}_q^*$, with the isomorphism given by

$$x \to (x \bmod p, x \bmod q).$$

Namely, $c = a \cdot b \bmod N$ if and only if

$$c = a \cdot b \bmod p \quad \wedge \quad c = a \cdot b \bmod q.$$

Furthermore, for any $a \in \mathbb{Z}_p^*$ and $b \in \mathbb{Z}_q^*$, there is a unique *efficiently computable* $c \in \mathbb{Z}_N^*$ such that

$$c = a \bmod p \quad \wedge \quad c = b \bmod q.$$

Uniqueness follows from the fact that if $x_1, x_2 \in \mathbb{Z}_N^*$ satisfy that

$$x_1 \bmod p = x_2 \bmod p \quad \wedge \quad x_1 \bmod q = x_2 \bmod q$$

then

$$x_1 - x_2 \bmod p = 0 \quad \wedge \quad x_1 - x_2 \bmod q = 0$$

which by the fact that p, q are co-prime implies that

$$x_1 - x_2 = 0 \bmod N,$$

as desired.

Claim 1. Given any primes p, q and any $a, b \in \mathbb{N}$ that are co-prime to p and q, respectively, one can efficiently compute $c \in \mathbb{Z}_N^*$, for $N = p \cdot q$, such that

$$c \bmod p = a \quad \wedge \quad c \bmod q = b$$

Proof. Given any primes p, q and any $a, b \in \mathbb{N}$, co-prime to p and q, respectively, compute $c \in \mathbb{Z}_N^*$, as follows.

- Compute p^{-1} which is the inverse of p modulo q.
- Compute q^{-1} which is the inverse of q modulo p.
- Output $c = q \cdot q^{-1} \cdot a + p \cdot p^{-1} \cdot b \bmod N$.

.

Remark. The notation above is a bit misleading since $p \cdot p^{-1} \neq 1 \bmod N$, and similarly for q. Moreover, neither p nor q have an inverse modulo N. Crucially, the inverses p^{-1} and q^{-1} were defined over q and p, respectively.

It is easy to see that

$$c \bmod p = a \quad \wedge \quad c \bmod q = b,$$

as desired. □

Notation We denote by (a, b) the CRT representation of an element $c \in \mathbb{Z}_N^*$. Namely (a, b) refers to the unique $c \in \mathbb{Z}_N^*$ such that

$$c \bmod p = a \quad \wedge \quad c \bmod q = b,$$

The Goldwasser-Micali Public-Key Encryption Scheme

Recall Rabin's construction of a trapdoor one-way function from the previous lecture. This construction turned out to be extremely useful, and has led to the Goldwasser-Micali encryption scheme [1], which was the first public-key encryption scheme that was proven to be secure. We describe the encryption scheme as follows:

- $\mathsf{Gen}(1^\lambda)$**:**
  1. Sample random primes $p, q \leftarrow \{0,1\}^\lambda$, such that
  $$(p \bmod 4 = 3) \quad \wedge \quad (q \bmod 4 = 3).$$
  This choice of p, q ensures that (-1) is not a quadratic residue modulo p or modulo q.
  This is the case since let $x \in \mathbb{Z}_p^*$ be any generator then $-1 = x^{\frac{p-1}{2}}$. Note that this would be a quadratic residue only if $\frac{p-1}{2}$ was even modulo $p-1$, i.e., if there exists $a \in \mathbb{N}$ such that
  $$2a = \frac{p-1}{2} \bmod (p-1)$$
  which holds only if $p = 1 \bmod 4$.
  2. Let $N = p \cdot q$.
  3. Output $\mathsf{pk} = N$ and $\mathsf{sk} = (p, q)$.
- $\mathsf{Enc}(\mathsf{pk}, m)$:
  1. Sample a random $x \leftarrow \mathbb{Z}_N^*$.
  2. Compute $y = x^2 \bmod N$.
  3. Output $\mathsf{ct} = (-1)^m \cdot y$.
- $\mathsf{Dec}(\mathsf{sk}, \mathsf{ct})$:
  1. Output $m = 0$ if $y \in \mathbb{QR}_N$ and output $m = 1$ otherwise.

Security follows from the assumption that it is secure! Namely, the assumption is that it is hard to distinguish random elements in $\mathbb{QR}_N$ from random elements in $\mathbb{Z}_N^*$ that are not quadratic residues modulo p or q.

Remark. It turns out that it is easy to distinguish random elements in $\mathbb{QR}_N$ from random elements in $\mathbb{Z}_N^* \setminus \mathbb{QR}_N$. But we believe that it is hard to distinguish between random elements in $\mathbb{QR}_N$ from random elements in $\{x \in \mathbb{Z}_N^* : (x \bmod p) \notin \mathbb{QR}_p \quad \wedge \quad (x \bmod q) \notin \mathbb{QR}_q\}$. This is called the Quadratic Residue assumption and is precisely what is needed in order to prover the security of the Goldwasser-Micali encryption scheme.

Public-key Encryption from Trapdoor One-Way Functions

Intuitively, to build a public-key encryption scheme from a trapdoor one-way function, we should take the public key to be hk and the secret key to be td. However, it is not clear what to do next. We will rely on the Goldreich-Levin theorem that there exists a hardcore predicate $\mathsf{HCP}(x, r) = x \cdot r$ such that

$$(\mathsf{hk}, F_{\mathsf{hk}}(x), r, \mathsf{HCP}(x, r) \approx (\mathsf{hk}, F_{\mathsf{hk}}(x), r, U)$$

where $U \leftarrow \{0,1\}$.

Remark. There seems to be a mismatch here since on the one hand $x \in \mathcal{D}_{\mathsf{hk}}$ and on the other hand HCP assumes that $x \in \{0,1\}^\lambda$. However, we argue that this mismatch is only syntactic and can be corrected easily. Basically, all we need to do is embed $\mathcal{D}_{\mathsf{hk}}$ in $\{0,1\}^\ell$, for an appropriate $\ell(\lambda) = \mathrm{poly}(\lambda)$.

Alternatively, the fact that $x \leftarrow \mathcal{D}_{\mathsf{hk}}$ is efficiently sampleable implies that one can think of F_{hk} as having the domain $\{0,1\}^\lambda$, and on input $r \in \{0,1\}^\lambda$ it samples $x \leftarrow \mathcal{D}_{\mathsf{hk}}$ using randomness r and outputs $F_{\mathsf{hk}}(x)$.[1] Thus, from now on we assume that $\mathcal{D}_{\mathsf{hk}} = \{0,1\}^\lambda$.[2]

In what follows, given a trapdoor one-way function family, defined by Gen_F and $\{F_{\mathsf{hk}}\}$, we construct the following public-key encryption scheme:

- Gen**:** Takes as input 1^λ and does the following:
  1. Sample $(\mathsf{hk}, \mathsf{td}) \leftarrow \mathsf{Gen}_F(1^\lambda)$.
  2. Set $\mathsf{pk} = \mathsf{hk}$ and $\mathsf{sk} = \mathsf{td}$.
  3. Output $(\mathsf{pk}, \mathsf{sk})$.
- Enc**:** takes as input the public key hk and a message $m \in \{0,1\}$, and does the following:
  1. Sample a random $x \leftarrow \{0,1\}^\lambda$.
  2. Compute $y = F_{\mathsf{hk}}(x)$.
  3. Choose a random $r \leftarrow \{0,1\}^\lambda$, and let $b = \mathsf{HCP}(r, x)$.
  4. Output $(y, r, b \oplus m)$.
- Dec**:** On input a secret key td and a ciphertext $\mathsf{ct} = (y, r, c)$, does the following:
  1. Compute $x = \mathcal{I}(\mathsf{td}, y)$.
  2. Compute $b = \mathsf{HCP}(r, x)$.
  3. Output $m = b \oplus c$.

Theorem 1. *The public-key encryptions scheme above is* CPA*-secure*

[1] We will also need to assume that the transformation from $r \in \{0,1\}^\lambda$ to x is injective. This can be achieved by taking a shorter λ and using a PRG. We omit the details here.

[2] We note that we will only be interested in how F_{hk} behaves on random elements $x \leftarrow \mathcal{D}_{\mathsf{hk}}$ hence this change of domain does not have an affect.

Proof. Suppose for the sake of contradiction that there exists a poly-size adversary $\mathcal{A}$ and a non-negligible ϵ such that $\mathcal{A}$ wins in the CPA game with probability ϵ. We construct a poly-size adversary that inverts the one-way function. To this end, it suffices to construct a poly-size $\mathcal{B}$ for which there exists a non-negligible ϵ such that for every $\lambda \in \mathbb{N}$,

$$\Pr[\mathcal{B}(\mathsf{hk}, F_{\mathsf{hk}}(x), r) = \mathsf{HCP}(x, r)] \geq \epsilon(\lambda),$$

where the probability is over $(\mathsf{hk}, \mathsf{td}) \leftarrow \mathsf{Gen}_F(1^\lambda)$ and over $x, r \leftarrow \{0,1\}^\lambda$.

$\mathcal{B}(\mathsf{hk}, F_{\mathsf{hk}}(x), r)$ emulates the adversary $\mathcal{A}$ in the CPA game, as follows:

1. Send $\mathsf{pk} = \mathsf{hk}$ to $\mathcal{A}$.
2. Emulate $\mathcal{A}(\mathsf{hk})$ to obtain $m_0, m_1 \in \{0,1\}$.
3. Choose at random $b \leftarrow \{0,1\}$.
4. For each $d \in \{0,1\}$ (which should be thought of as a guess for $\mathsf{HCP}(x,r)$),
   (a) Let $\mathsf{ct}_d = (F_{\mathsf{hk}}(x), r, d \oplus m_b)$.
   (b) Emulate $\mathcal{A}$ on input ct_d to obtain b'_d.
5. If $b'_0 = b'_1$ then output a random guess $b' \leftarrow b$.
6. Otherwise output d such that $b'_d = b$.

Note that for $d = \mathsf{HCP}(x,r)$ it holds that

$$\Pr[b'_d = b] \geq \frac{1}{2} + \epsilon(\lambda)$$

whereas for a random $d \leftarrow \{0,1\}$,

$$\Pr[b'_d = b] = \frac{1}{2}.$$

therefore $\mathcal{B}$ outputs $\mathsf{HCP}(x,r)$ with a non-negligible advantage, as desired.

□

References

[1] Shafi Goldwasser and Silvio Micali. Probabilistic encryption and how to play mental poker keeping secret all partial information. In Harry R. Lewis, Barbara B. Simons, Walter A. Burkhard, and Lawrence H. Landweber, editors, *Proceedings of the 14th Annual ACM Symposium on Theory of Computing, May 5-7, 1982, San Francisco, California, USA*, pages 365–377. ACM, 1982.