NYU COMPUTER SCIENCE FALL 2026

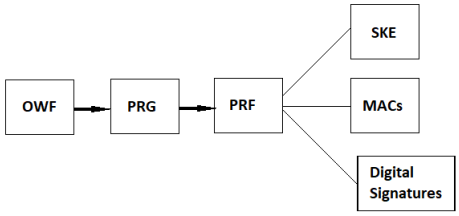
Introduction to Cryptography

CSCI-GA.3210-001 Instructor: Nir Bitansky

Lecture 8: Public-Key Encryption

1 Previously on Introduction to Crypto

We've seen that starting from one-way functions, we can go a long way. We've constructed PRGs, PRFs, secret-key encryption, authentication and even digital signatures.



Today, we will move beyond into the world of public-key encryption.

2 Privately Communicating without a Secret Key?

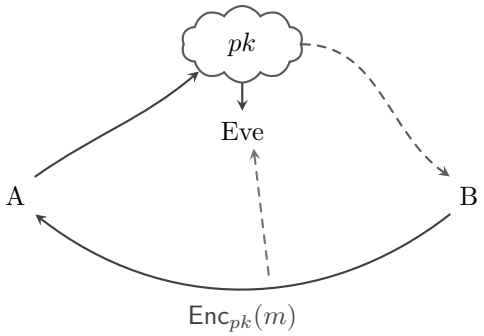
From the dawn of cryptography it was evident that *to communicate secretly, parties have to share a secret*, and this perception dominated for centuries. In the 70's, People had begun to challenge this perception, suggesting that secret communication may be achievable without the parties ever meeting! Curiously, these ideas emerged both publicly within the scientific community, and (a few years before that) behind closed doors within the British intelligence agency (you can read more about this fascinating story in Barak's notes).

One of the first to challenge the need for a secret key was Ralph Merkle, then an undergrad at Berkeley. He both conjectured that this is possible, and gave some evidence for this possibility (he demonstrated a public key encryption system where the attacker has to work quadratically, rather than super-polynomially, harder than the honest parties). The paper that officially started the age of public-key cryptography was published in '76 by Diffie and Hellman, and was titled "New Directions in Cryptography" [DH76]. It describes the concepts of public key-exchange, public-key encryption, and digital signatures, and suggested generic approaches and concrete constructions.

We will now start to explore these concepts.

3 Public-Key Encryption

In the setting of public key encryption, Alice shares (say, on her webpage) with the world a *public encryption key* pk and keeps a corresponding secret key sk. Then any person, Bob, in hold of pk can encrypt messages to Alice. While she can decrypt, any eavesdropper Eve that doesn't know the secret key (but does know the public key) learns nothing.



Definition 3.1 (Public-Key Encryption (PKE) against Chosen Plaintext Attacks (CPA))**.** *A public-key encryption scheme consists of polynomial-time algorithms* (Gen, Enc, Dec) *with the following properties.*

- **Syntax:**
  - $Gen(1^n)$ *is a randomized algorithm that takes as input a security parameter* n *and outputs a secret key* sk *and a public encryption key* pk*.*
  - $Enc_{pk}(m)$ *is a randomized algorithm that takes as input the public key* pk *and message* $m \in \{0,1\}^*$*, and outputs a ciphertext* ct*.*
  - $Dec_{sk}(ct)$ *is a deterministic algorithm that takes as input the secret key* sk *and ciphertext* ct*, and outputs a decrypted message.*
- **Correctness:** *for any message* $m \in \{0,1\}^*$*,*

$$\Pr\left[Dec_{sk}(Enc_{pk}(m)) = m \mid (sk, pk) \leftarrow Gen(1^n)\right] = 1 \ .$$

- **Indistinguishability against Chosen Plaintext Attack:** *For any (interactive) n.u. PPT* $A = \{A_n\}$ *there is a negligible* μ *such that* A *wins the following CPA game with probability at most* $1/2+\mu(n)$*.*

  The CPA Game:
  - $Gen(1^n)$ *samples* sk, pk*.*
  - A_n *is given* pk*, and chooses two equal-length messages* m_0, m_1*.*
  - A_n *is then given an encryption* $Enc_{pk}(m_b)$ *for a random* $b \leftarrow \{0,1\}$*.*
  - A_n *wins if it guesses* b *correctly.*

Remarks:

- **Encryption Queries:** Typically, in a chosen plaintext attack, the adversary may also get oracle access to the encryption algorithm, allowing her to adaptively compute encryptions of messages of her choice and choose the challenge messages m_0, m_1 accordingly. In the case of public-key encryption, the above definition (without an explicit oracle) is equivalent (make sure you understand why).
- **Multi-Message Security:** The definition addresses a single challenge ciphertext. In the homework, you will show that it, in fact, implies multi-message security.
- **Public-Key Encryption vs Key Exchange:** A relaxation of the notion of public-key encryption is that of *key exchange*. Here two parties may run an *interactive protocol* at the end of which they agree on a secret key sk, such that an external observer that sees the transcript of the protocol cannot recover sk. It is not hard to show that PKE is equivalent to two-message key exchange. However, key exchange protocols with more than two rounds are not known to imply PKE.

Why Would PKE Exist? One has to ask what made Diffie, Hellman and the other co-discoverers of public-key encryption believe that such a thing is even possible. It seems that one intuition they had is the following. Imagine that parties can not only send messages but also *magic boxes*. A sender party can lock a computer program Π in such a box so that the receiver can provide inputs x to the program and get the outputs $\Pi(x)$, but the box is otherwise opaque, and does not allow to see the contents of Π itself.



In this case, public key encryption can be done as follows. The public key pk can be a "publicly-verifiable puzzle", say $f(x)$ for a OWF f, and the secret key sk would be the solution x. Then to encrypt m, we will simply create a magic box and lock inside it the program $\Pi_{f(x),m}$ that returns m only when given the solution x. Another natural idea is to publish such a box as the public key and put inside the encryption algorithm of a secret key encryption algorithm (and say derandomize the encryption algorithm using a PRF).

Diffie and Hellman believed that it might be possible to digitally realize such a magic box, by simply *rewriting* Π in a way that will make it hard to read, e.g. by writing it in some low-level assembly-like language (or asking an undergrad to code it). Indeed, in general, reverse engineering is a notoriously hard problem. This idea is known today as *program obfuscation*. Diffie and Hellman didn't give any formal definitions or constructions, but rather used it as intuition to why PKE is conceivable. (In time, it had turned out that obfuscation is a fascinating concept beyond the context of PKE. It is also very challenging to define and construct. In fact, currently, there is still no satisfying solution to the problem. Hopefully, we will return to it toward the end of the course.)

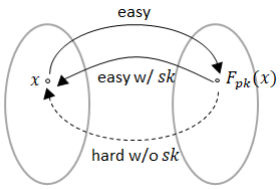
4 Basing PKE on Hard Problems

The question now is how to fulfill Diffie and Hellman's intuition based on reasonable computational assumptions. So far in the course we've been lucky enough to manage to construct everything from OWFs. It would be great if we could also prove:

Theorem 4.1 (In Our Dreams)**.** *If OWFs exist so do public-key encryption schemes.*

Unfortunately, we're very far from being able to prove such a theorem, and the common belief today is that public-key encryption requires considerably more structure than OWFs.

Trying to capture the structure required for PKE, Diffie and Hellman suggested the following generalization of one-way functions to trapdoor one-way functions. Here, a function is associated with a public key and a secret *trapdoor key*. Using the public key, anyone could compute the function forward, and using the secret key it can also be computed backwards (we'll assume for now that such functions are injective). However, for a random public key, and without the secret key, the function is one way.



Definition 4.2 (Trapdoor Function (TDF))**.** *A trapdoor function is given by polynomial time algorithms* (Gen, F, F^{-1}) *with the following properties:*

- **Syntax:**
  - *$Gen(1^n)$ is a probabilistic algorithm that given a security parameter 1^n outputs a secret key (also called the trapdoor) and a public key (sk, pk).*
  - *$F_{pk}(x)$ is a deterministic algorithm that given a public key pk and input $x \in \{0,1\}^n$, outputs an image y.*
  - *$F_{sk}^{-1}(y)$ is a deterministic algorithm that given a secret key sk and $y \in \{0,1\}^*$, outputs $x \in \{0,1\}^n$.*
- **Correctness:** *for any $x \in \{0,1\}^n$,*

$$\Pr\left[F_{sk}^{-1}(F_{pk}(x)) = x \mid (sk, pk) \leftarrow Gen(1^n)\right] = 1 \ .$$

- **One-Wayness:** *for any n.u. PPT $A = \{A_n\}$, there exists a negligible μ such that for all $n \in \mathbb{N}$,*

$$\Pr\left[A_n(pk, F_{pk}(x)) = x \;\middle|\; \begin{array}{r}(sk, pk) \leftarrow Gen(1^n) \\ x \leftarrow \{0,1\}^n\end{array}\right] \leq \mu(n) \ .$$

On the surface, TDFs seem like a simpler object than PKE. In particular, they only guarantee that random messages x are hard to fully recover, but they certainly don't guarantee to fully hide any input x. In fact, for any two messages x, x' you can always easily tell $F_{pk}(x)$ and $F_{pk}(x')$, as F_{pk} is deterministic. Historically, this is still how people imagined using TDFs for several years, until the work of Goldwasser and Micali [GM82] who introduced *probabilistic encryption* and the requirement that encryption leaks nothing about plaintexts, which later became the gold standard.[1]

Formally, however, public-key encryption is not known to be stronger than TDFs — they are not known to imply TDFs (think why a PKE scheme isn't already a TDF by itself). In contrast, the other direction is true.

Theorem 4.3. *TDFs imply PKE.*

Proof Sketch. Let (Gen, F, F^{-1}) be a TDF. We define a PKE (Gen, E, D) for one-bit messages.

- The Gen algorithms are the same (as suggested by our notation).
- $E_{pk}(m)$ samples $x, r \leftarrow \{0,1\}^n$, and outputs

$$F_{pk}(x), r, \langle x, r\rangle \oplus m \ .$$

- $D_{sk}(y, r, b)$ outputs $\langle F_{sk}^{-1}(y), r\rangle \oplus b$.

The correctness of the scheme is obvious. As for security, by the Goldreich-Levin theorem, we know that

$$pk, F_{pk}(x), r, \langle x, r\rangle \approx_c pk, F_{pk}(x), r, u \ ,$$

where $pk \leftarrow Gen(1^n), x, r \leftarrow \{0,1\}^n, u \leftarrow \{0,1\}$.

The above already implies that

$$pk, E_{pk}(0) \approx_c pk, E_{pk}(1) \ .$$

In the homework, you will prove that in the case of PKE, the above is equivalent to CPA security:

Claim 4.4. *A bitwise PKE is CPA secure if and only if $pk, E_{pk}(0) \approx_c pk, E_{pk}(1)$.*

□

Diffie and Hellman defined trapdoor functions and demonstrated how to use them for encryption (although, in a deterministic fashion. This is before Goldwasser-Micali and Goldreich-Levin). However, they did not provide a candidate for such functions (they did provide a direct construction of PKE, which we'll touch in a bit).

[1] This requirement, which we have so far formalized as indistinguishability, also has an alternative formulation known as semantic security, which we might mention later on.

RSA. The first trapdoor function that (publicly) emerged to the world in 1977 was suggested by Rivest, Shamir, and Adleman [RSA78], which became to be known as the RSA function. The trapdoor function, or actually permutation (TDP), is based on a number-theoretic problem related to factoring. We'll briefly present it, just to get a sense of how it looks, without getting too much into details.

- Let $p < q$ be two primes of magnitude $\Theta(2^n)$, and let $N = pq$.
- Let $\mathbb{Z}_N^*$ be the multiplicative group modulo N (these are all the numbers in $[N]$ that are coprime to N, and multiplication is done modulo N).
- The order of this group is $\phi(N) = (p-1)(q-1)$.
- A public key pk for the function consists of N and an integer $e \in \mathbb{Z}_{\phi(N)}^*$. The secret key sk consists of an integer $d \in \mathbb{Z}_{\phi(N)}^*$ such that $ed = 1 \mod \phi(N)$ (such an integer can be computed using the GCD algorithm, by those who know p and q).
- The function is defined (and in fact is a permutation) over $x \in \mathbb{Z}_N^*$ as follows

  $$F_{N,e}(x) = x^e \mod N \ .$$

  (This operation can be computed efficiently in time $\text{polylog}(N) = \text{poly}(n)$.)
- The inversion for $y \in \mathbb{Z}_N^*$ is given by

  $$F_d^{-1}(y) = y^d \mod N \ .$$

Assumption 4.5 (RSA Assumption, somewhat informally stated)**.** *For randomly chosen p, q and e, $F_{N,e}(x)$ is hard to invert for a random x.*

How hard is RSA? Today the best algorithms to solve this problem are those that factor N, but no formal reduction to factoring is known. (As we've already mentioned in previous lectures, factoring itself is not as hard as one may hope. It can be solved in subexponential time and also quantumly.)

Other TDFs and PKEs. In 1979, Rabin [Rab79] suggested another TDP based on the hardness of factoring N's of a special form (called Blum integers). For a very long time (roughly 30 years) these were the only known TDFs.[2] The most common way of constructing PKE schemes is directly based on different (or perhaps, as we'll later mention, not so different) computational problems. In fact, such a construction was already given by Diffie and Hellman. The construction was actually described as a two-message key exchange protocol. Here we'll present a slight variant of it as a PKE (which reflects the equivalence between the notions).

Diffie-Hellman PKE. The scheme comes from the *discrete log* problem in groups, which is perhaps one of the oldest and most investigated problems in computational number theory. Specifically, consider the multiplicative group $\mathbb{Z}_p^*$ for a prime $p \sim 2^n$. The discrete logarithm problem in this group is given $g, g^x \in \mathbb{Z}_p^*$, find x. The scheme is defined as follows (in what follows all group operations are $\mod p$).

- A public key pk consists of p, g, g^x where g is chosen at random from $\mathbb{Z}_p^*$, and $x \in \mathbb{Z}_{p-1}$ is a random exponent. The secret key sk consists x.
- The encryption algorithm is defined for bit messages as follows

  $$E_{p,g,g^x}(m; y, r) = g^y, \langle g^{xy}, r\rangle \oplus m, r \ ,$$

  where $y \leftarrow \mathbb{Z}_{p-1}$ and $r \leftarrow \{0,1\}^n$, are randomness used by the encryption algorithm.

[2] Today, we know of some additional ones, but these two constructions are still the only known permutations based on assumptions that are considered standard.

- Decryption for $h \in \mathbb{Z}_p^*$, $b \in \{0,1\}$, and $r \in \{0,1\}^n$ is given by

$$D_x(h,b,r) = \langle h^x, r\rangle \oplus b \ .$$

It is not hard to see that if one can solve the discrete log problem in $\mathbb{Z}_p$, then the scheme is completely insecure, but this is not enough. The security of the scheme is based on the following assumption.

Assumption 4.6 (Computational Diffie-Hellman Assumption (CDH), somewhat informally stated)**.** *For randomly chosen p and g, x, y, given p, g, g^x, g^y, it is hard to find g^{xy}.*

Claim 4.7. *Assuming CDH, the Diffie-Hellman PKE described above is CPA secure.*

Proof Sketch. By Claim 4.4, it is sufficient to show that

$$pk, E_{pk}(0) \approx_c pk, E_{pk}(1) \ .$$

For this it suffices to show that,

$$p, g, g^x, g^y, \langle g^{xy}, r\rangle, r \approx_c p, g, g^x, g^y, b, r$$

where b is a uniformly independent bit (make sure you understand why this is enough). Indeed, assume towards contradiction, there exists an efficient distinguisher A for the above distributions with advantage $\epsilon(n) = n^{-O(1)}$. Then, we can use A to break the CDH assumption. Specifically, we can construct from A an efficient predictor P such that

$$\Pr_{p,g,x,y,r}[P(p,g,g^x,g^y,r) = \langle g^{xy}, r\rangle] \geq \frac{1}{2} + \Omega(\varepsilon) \ .$$

The predictor P samples a random bit β and checks whether $A(p,g,g^x,g^y,\beta,r) = 1$. If so it returns β, and otherwise $1-\beta$. A standard calculation shows that this predictor works assuming $\Pr[A(p,g,g^x,g^y,b,r) = 1] - \Pr[A(p,g,g^x,g^y,b,r) = 1] \geq \varepsilon$. (If it is the other way around, we flip P's output.)

By averaging with probability $\Omega(\varepsilon)$ over p, g, x, y, it holds that

$$\Pr_{r}[P(p,g,g^x,g^y,r) = \langle g^{xy}, r\rangle] \geq \frac{1}{2} + \Omega(\varepsilon) \ .$$

In this case, we can apply to Goldreich-Levin algorithm to recover g^{xy}. □

A famous variant of this assumption is the Decision DH (DDH), that says g^{xy} is not only hard to compute, but also pseudorandom. It naturally leads to a simpler variant of the above PKE (without the GL hardcore bit) known as ElGamal encryption. The DDH assumption, in fact, does not hold in $\mathbb{Z}_p^*$, but it is conjectured to hold in appropriate groups.

A Scheme of a Different Nature Based on LPN. We now briefly describe yet another system by Aleknovich [Ale03] based on the learning parity with noise problem (which we've already encountered previously in the course).

- A public key pk consists of $\hat{A} = (A, As + e)$ where $A \leftarrow \{0,1\}^{m \times n}$ (where $m \approx 2n$), $s \leftarrow \{0,1\}^n$ and $e \in \{0,1\}^m$ is a noise vector, where each coordinate is one with probability δ. The secret key sk consists of e.
- The encryption algorithm is defined for bit messages as follows:
  - To encrypt a one, output a random vector $u \leftarrow \{0,1\}^m$.
  - To encrypt a zero, sample a random y in the kernel of $\hat{A}^T$ (namely a random codeword in the dual of $\hat{A}$) and an error vector e' (distributed as e) and output $y + e'$.

- To decrypt $c \in \{0,1\}^m$, output $\langle c, e \rangle$.

When choosing the noise rate δ to be small enough (roughly $1/\sqrt{n}$), it can be shown that zero encryptions are decrypted correctly with probability close to 1 and one encryption with probability close to 1/2. Correctness is then amplified by taking say n copies of the scheme (and using the fact that PKE remains secure under such repetition). Security is shown based on the LPN assumption, saying that

$$A, As + e \approx_c A, u \ .$$

On Basing PKE on Different Assumptions: All and all hard problems sufficient for PKE seem far more rare than those needed for secret-key encryption. In fact, we can essentially classify them into two groups of similar problems: 1) number/group theoretic problems (solved quantumly) and 2) lattice/decoding problems (believed to be quantum resistant). Coming up with essentially different PKE schemes is a major open problem.

References

[Ale03] Michael Alekhnovich. More on average case vs approximation complexity. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings*, pages 298–307, 2003.

[DH76] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Trans. Information Theory*, 22(6):644–654, 1976.

[GM82] Shafi Goldwasser and Silvio Micali. Probabilistic encryption and how to play mental poker keeping secret all partial information. In *Proceedings of the 14th Annual ACM Symposium on Theory of Computing, May 5-7, 1982, San Francisco, California, USA*, pages 365–377, 1982.

[Rab79] Michael O. Rabin. *Digitalized signatures and public-key functions as intractable as factorization.* MIT Laboratory for Computer Science, 1979. URL: `http://ncstrl.mit.edu/Dienst/UI/2.0/Describe/ncstrl.mit_lcs/MIT/LCS/TR-212`. Note: Technical Report 212.

[RSA78] Ronald L. Rivest, Adi Shamir, and Leonard M. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126, 1978.