

## Lecture 1

*Lecture date: January 8, 2024*

*Scribes: Nakul Khambhati, Sasha Kononova, Nathan Leung*

# 1 Course Administration

This is the Winter 2023 edition of CS 282A/MATH 209A, Foundations of Cryptography. The text for this class is Prof. Ostrovsky's 2010 lecture notes, which are available on Prof. Ostrovsky's website.

## 1.1 Grading

The course grade will be weighted as follows: 40% midterm exam, 50% final exam, and 10% for scribe notes (these notes are one such example).

## 1.2 Exams

Exams will be closed book. The exact date of the midterm exam will depend on how fast Prof. Ostrovsky can lecture (this will be affected by the number of questions in lecture, etc.), but it will occur shortly after the lectures discussing zero-knowledge proofs and digital signatures.

## 1.3 Undergraduate PTEs

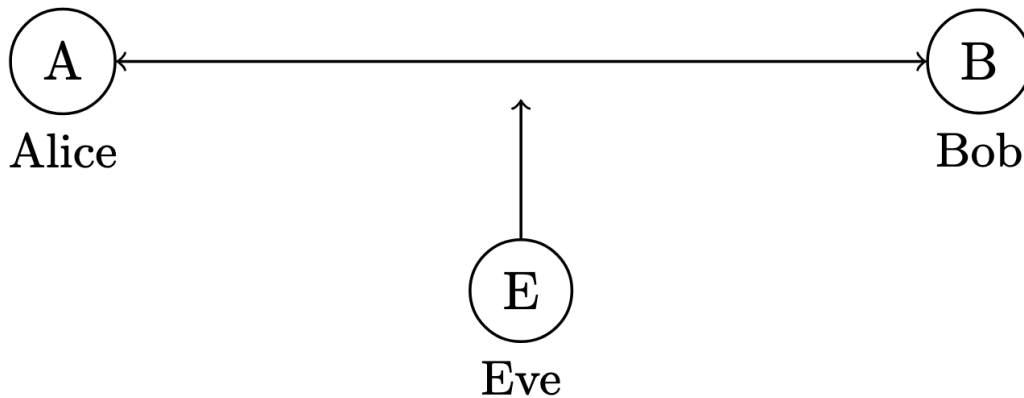
Prof. Ostrovsky is working on getting PTEs for undergraduate students who wish to take the course and will have them by Wednesday, January 10th.

For more details on course administration, see the syllabus on Prof. Ostrovsky's website.

# 2 Introduction to Encryption

To begin our study of cryptography, we consider the classic example of Alice, Bob, and Eve. Alice and Bob want to communicate secretly, while Eve is trying to listen in (or *eavesdrop*)

on their communications. Alice and Bob do not want Eve to know what they are talking about.



**Figure 1:** Communication diagram between Alice, Bob, and Eve.

So, what can Alice and Bob do to prevent Eve from listening in? One idea is to use a one-time pad.

## 2.1 One-Time Pads

A one-time pad is essentially a huge random string that Alice and Bob securely share. When Alice wants to send a message, she can encrypt her message by XORing the bits of her message with bits of the one-time pad and then sending them to Bob. Then Bob can XOR Alice’s encrypted message with his (identical) one-time pad to recover the original message (since XOR with one fixed parameter is an involution).

As long as the one-time pad is sufficiently random (this can be achieved using a psuedo-random generator, which allows the generation of massive amounts of data that are functionally indistinguishable from true randomness) and not leaked or reused (hence “one-time”), this encryption is unbreakable.

## 2.2 Public-Key Encryption

One-time pads, however, have the limitation that Alice and Bob must securely share the one-time pad in advance. What if they can’t do that? Public-key encryption solves this problem. Invented in 1976 by Diffie and Hellman, the concept is as follows:

1. Alice has a public key (shared) and a private key (never shared)
2. Bob has a public key (shared) and a private key (never shared)
3. When Alice wants to send a message to Bob, she can encrypt her message using Bob's **public** key
4. The mathematics of public-key encryption make it so that the only way to decrypt Alice's message is with Bob's **private** key

Thus, no secure secret exchange (as with one-time pads) is required; parties can communicate securely immediately. Public-key encryption can be implemented in various ways; one popular way today is RSA — a very efficient public-key encryption system based on the hardness of factoring numbers. One limitation of public-key encryption, however, is that public-key encryption is more mathematically complex than secret-key encryption (i.e. with a shared secret), and hence is computationally slower.

Public-key encryption is the basis for all secure internet communication today (think SSH, HTTPS), and as a result Diffie and Hellman received a Turing Award for their work. In the HTTPS example, however, (and in other cryptographic protocols, too), public-key encryption is only used to exchange an initial shared secret key. The shared secret key is used for all further communication, because secret-key encryption is much faster than public-key encryption.

### 2.3 Digital Signatures

Another concern is what happens if Eve, despite being unable to view the message between Alice and Bob, is able to tamper with it, for instance by flipping certain bits. How would Bob know that the message he is receiving and decrypting is actually what Alice initially sent?

This problem can be resolved using digital signatures, implemented with public-key encryption. In essence, Alice can create a signature attesting to the contents of her message with her private key, and then Bob can verify that the signature is authentic with Alice's public key and the contents of the received message. If the signature verification step fails, Bob will know that the message was tampered with.

### 2.4 Cryptographic Proofs and “Hardness”

Above, we used the term “hardness” — for instance, we said that RSA is based on the “hardness” of factoring. What does this mean? How can we show “hardness”?

In academic cryptography, we typically show “hardness” by reducing our problem of interest to another problem known to be “hard”, meaning that the problem can only be solved, as far as we know, by a superpolynomial algorithm. In other words, we show that an attack on a cryptosystem implies a solution to a famously difficult math problem.

This allows cryptographers to delegate the mathematical security of their protocols to mathematicians. The idea is that if a system is broken, it isn’t because of something in the domain of cryptography, but because of a transformative result in the domain of mathematics (e.g.  $\mathbf{P}=\mathbf{NP}$ ).

Another important aspect of cryptographic proofs is that there is no “security through obscurity” — we assume that the attacker (Eve) knows all the algorithms, functions, etc. being used in the cryptosystem.

This assumption is one reason why provably secure public-key encryption needs randomness — a result of Goldwasser and Micali. Otherwise, the same message would yield the same ciphertext every time; the ciphertext would be revealing something about the contents of the message. By incorporating randomness, the ciphertext reveals nothing about the underlying message. As a sidenote, lack of randomness is one reason why Turing was able to break the Nazi Enigma cipher — every Enigma message started with the same standard header, and Turing was able to take advantage of that regularity.

### 3 Randomness Used in Encryption: Average Salary

Suppose there are three people ( $P_1, P_2, P_3$ ) who wish to find their average salary, but do not wish to share their own salaries with each other. There are several ways to do this by utilizing randomness. Suppose there is some number  $M$  that is definitely larger than any of the three people’s salaries. Let each person choose two random numbers between 0 and  $M$ , and tell each of the other people one of those numbers. The information is then shared as in Figure 2; an arrow indicates a piece of information being communicated in private between the two people involved.


For example,  $P_1$  generates the random numbers  $A_1, A_2$  and communicates them to  $P_2, P_3$  respectively; in turn,  $P_1$  receives the random numbers  $R_1, B_1$  generated by  $P_2, P_3$  respectively.

Now, let each person report to both other people (i.e., publicly) the following:

$$\text{their salary} + \text{sum of incoming arrows} - \text{sum of outgoing arrows}$$

For example,  $P_1$  would report

$$\text{their salary} + (R_1 + B_1) - (A_1 + A_2)$$



**Figure 2:** Information between  $P_1$ ,  $P_2$ , and  $P_3$ .

The sum of these three reported values is equal to the sum of the actual salaries, since each of the random variables is added to one reported value (from the person who received the random number) and subtracted from another reported value (from the person who generated the random number). Thus, their average is equal to the true average salary. Note that no single person can figure out how much the other people make, as long as the lines of communication are indeed secure (i.e.,  $P_3$  does not know the values of  $A_1$  and  $R_1$ ).

While finding the average of three values is easy mathematically, it can be shown that any function of several variables that can be computed in polynomial time can be encrypted with a similar strategy. That is, using randomness in a clever way can solve a relatively complex mathematical problem correctly and without “leaking” information.

#### 4 NP-Completeness: “Decision = Search”

Recall that a problem lies in complexity class  $\mathbf{P}$  if, for an input that can be represented in  $n$  bits, it can be solved in time  $O(n^c)$  for some constant  $c$ . Problems in class  $\mathbf{P}$  will typically have an answer beyond “yes” or “no” – that is, they are search problems, where one needs to actually present a solution. A problem that does not appear to be in class  $\mathbf{P}$  is the problem of factoring: given the product of two (large) primes  $P_1, P_2$ , what are the

values of  $P_1$  and  $P_2$ ? The best known way to compute this runs in  $O\left(2^{\left(\frac{N}{\log \log N}\right)^{1/3}}\right)$ , so the problem does not lie in  $\mathbf{P}$  as far as we know.

A decision problem, on the other hand, has an answer that is simply “yes” or “no.”

There is a result that states that if one has an algorithm that can, in polynomial time, answer decision questions, it can be used to obtain an actual solution to an analogous problem. For example, consider a boolean circuit. A boolean circuit takes several inputs (either 1’s or 0’s), which then go through some series of logic gates (AND, OR, NOT, and so on), leading to a single output. The decision question attached to such a circuit is whether or not some set of inputs can cause the output to be 1. The search question is what set of inputs will cause the output to be 1.

For the boolean circuit, suppose we have an algorithm that answers the decision question. If it says that it is impossible to produce a 1, there is nothing to show for the search question. If it says a solution exists that produces a 1, we can set the first input to be 1, and ask the algorithm whether the modified circuit, with the first input fixed, can produce a 1. If it says yes, we know there exists a solution where the first input is a 1; otherwise, we know there exists a solution where the first input is a 0, and so we set the first input to be 0. We can continue doing this for each bit; thus, we get a solution to the search problem after running the algorithm  $n$  times. If our initial algorithm solved the decision question in polynomial time, then we have solved the search question in polynomial time as well.

## 5 Complexity classes

### 5.1 Uniform complexity classes

The modern approach to cryptography is based on certain complexity assumptions. As a result, some background on complexity classes will be helpful. We quickly recap some complexity classes  $\mathbf{P}$ ,  $\mathbf{NP}$ ,  $\mathbf{RP}$ ,  $\mathbf{co-RP}$  and  $\mathbf{BPP}$ .

**Definition 1 (Polytime (P))**  $P$  is the set of languages  $L \subset \{0,1\}^*$  that can be decided by a deterministic Turing machine running in polynomial time on the length of the input.

**Definition 2 NP** is the set of languages  $L \subset \{0,1\}^*$  that can be decided by a non-deterministic Turing machine running in polynomial time on the length of the input.

**Definition 3 (Randomized Polytime (RP))** A language  $L \subset \{0,1\}^*$  is in  $\mathbf{RP}$  if and only if there exists a probabilistic Turing machine  $A$  that runs on polynomial time for all inputs and

$$(1) \quad \forall x \in L : Pr[A(x) = 1] > \frac{2}{3}.$$

$$(2) \quad \forall x \notin L : Pr[A(x) = 1] = 0.$$

**Definition 4 (Co-RP)** A language  $L \subset \{0,1\}^*$  is in **co-RP** if and only if there exists a probabilistic Turing machine  $A$  that runs on polynomial time for all inputs and

$$(1) \quad \forall x \in L : Pr[A(x) = 1] = 1$$

$$(2) \quad \forall x \notin L : Pr[A(x) = 1] < \frac{1}{3}.$$

**Definition 5 (Bounded Probabilistic Polytime (BPP/PPT))** A language  $L \subset \{0,1\}^*$  is in **BPP** if and only if there exists a probabilistic Turing machine  $A$  that runs on polynomial time for all inputs and

$$(1) \quad \forall x \in L : Pr[A(x) = 1] > \frac{2}{3}.$$

$$(2) \quad \forall x \notin L : Pr[A(x) = 1] < \frac{1}{3}.$$

**Remark** We require that the probability of  $A$  making a mistake on any input is bounded away from  $1/2$ . We can reduce the probability that  $A$  makes an error by running it  $n$  times with fresh randomness each time and taking the majority output.

Next, we justify that this method works by showing that repeating the computation  $n$  times and taking majority gives us a low error probability – in fact one that decays exponentially as  $n$  increases.

**Proposition 6** Consider the Turing machine  $A'$  that executes  $A$  a total of  $n$  times with fresh randomness obtaining outputs  $b = b_1b_2 \cdots b_n$  and outputs  $MAJ(b)$ . For any input  $x \in \{0,1\}^*$  the probability that  $A'$  makes a mistake (i.e.  $A$  makes at least  $n/2$  mistakes in  $n$  independent trials) is less than  $e^{-n/24}$ .

**Proof** Let  $x \in L$  be arbitrary. Let  $E$  denote the bad event that  $A$  makes a mistake on input  $x$  i.e.  $x \in L$  and  $A(x) = 0$  or  $x \notin L$  and  $A(x) = 1$ . Since  $L \in \text{BPP}$ , in either case  $\Pr[E] < 1/3$ . Formally, let  $X_i$  denote the random variable that equals 1 if  $A$  makes a mistake on input  $x$  in trial  $i$  and 0 otherwise. Note that the  $X_i$  are independent and identically distributed (i.i.d.). Let

$$X = \sum_{i=1}^n X_i$$

be the total number of mistakes that  $A$  makes. Our goal is then to show that

$$\Pr[X \geq n/2] < e^{-n/24}.$$

We will do this by appealing to the Chernoff Bound.

**Theorem 7 (Chernoff bound)** *Let  $X_1, \dots, X_n$  be i.i.d. random variables and  $X = \sum_{i=1}^n X_i$ . Let  $E(X)$  denote the expectation of  $X$  and  $\beta < 1.5$ . Then*

$$\Pr[X \geq (1 + \beta)E(X)] < e^{-\beta^2 E(X)/2}.$$

**Proof** Omitted. ■

Setting  $\beta = 1/2$  and observing that  $E(X) = n/3$  we get

$$\Pr[X \geq \frac{3}{2} \cdot \frac{n}{3}] < e^{-n/24}$$

$$\Pr[X \geq n/2] < e^{-n/24}$$

as required ■

## 5.2 Non-uniform complexity classes

In the previous section, we looked at languages that can be decided by Turing machines. Next, we look at languages that can be decided by sequences of circuits  $\{C_n\}_{n \in \mathbb{N}}$  where each  $C_n$  has  $n$  inputs, one output, and a number of gates polynomial in  $n$ .

**Definition 8 (P/poly)** *P/poly is the set of languages that can be decided by polynomial-sized circuit families.*

As it turns out, **P/poly** circuits are so powerful that they do not require coin flips to make their decisions. We capture this idea in the following theorem due to Leonard Adelman.



**Theorem 9 (Adelman)  $BPP \subset P/Poly$**

**Proof** Let  $L \subset BPP$  i.e. it can be decided by a Turing machine that makes mistakes with low probability. We want to show that it can be decided *deterministically* by some family of circuits  $\{C_n\}_{n \in \mathbb{N}}$ . First, by Proposition 6, we can assume wlog that there exists a Turing machine  $A$  that runs in polynomial time on all inputs and letting  $n = |x|$  denote the length of string  $x$  we have

$$(1) \quad \forall x \in L : Pr[A(x) = 1] > 1 - 2^{-(n+1)}.$$

$$(2) \quad \forall x \notin L : Pr[A(x) = 1] < 2^{-(n+1)}.$$

Let  $r$  denote the total number of coins used by  $A$ . Fix some  $n \in \mathbb{N}$  (i.e. for the time being only consider strings of length  $n$ ) and consider the  $2^n \times 2^r$  matrix  $M$  below that represents all possible inputs (including the coins) to  $A$  on a string of length  $n$ .

	$\rho_1$	$\rho_2$	$\rho_3$	$\dots$	$\rho_{2^r}$
$x_1$	1	0	0	$\dots$	1
$x_2$	0	1	0	$\dots$	0
$x_3$	1	0	0	$\dots$	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
$x_{2^n}$	1	0	0	$\dots$	1

**Figure 3:** Caption

Let the  $(i, j)$ -th position of this matrix denote  $A$ 's output on input  $x_i$  where  $x_i \in \{0, 1\}^n$  and sequence of coin flips  $\rho_j \in \{0, 1\}^r$ . By assumption, on each input,  $A$  makes a mistake with probability less than  $2^{-(n+1)}$  which means that the fraction of 1's in each row must be less than  $2^{-(n+1)}$ . This implies that the total number of 1's in matrix  $M$  is  $< 2^n \cdot 2^r \cdot 2^{-(n+1)} = 2^r/2$ . However, there are only  $2^r$  many columns in  $M$ . By the pigeonhole principle, at least half of the columns have all zeroes. Interpreting this differently, at least half of the configurations of the coin flips will give the correct result on all  $2^n$  inputs of length  $n$ . We pick any one such configuration  $\rho = \rho_1\rho_2 \dots \rho_r$  and hardwire it into the **P/poly** circuit  $C_n$  for strings of length  $n$ . This circuit is of polynomial size as  $A$  runs in polynomial time and decides  $L$  for strings of length  $n$  without using any randomness. Since  $n$  was picked arbitrarily, we can construct the entire family  $\{C_n\}_{n \in \mathbb{N}}$  in this way which shows that  $L \subset P/poly$  as required. ■