



Introduction to mathematical cryptography
by Péter Maga

Contents

Preface	v
1 Introduction	1
1.1 The principal goal of cryptography, Kerckhoff's principle	1
1.2 Cryptanalysis	2
1.3 Mathematics enters cryptography	2
1.3.1 Transform texts to numbers	2
1.3.2 The mathematical formulation of symmetric ciphers	3
1.3.3 The mathematical formulation of asymmetric ciphers	3
1.4 Computability	4
1.5 The XOR cipher and pseudorandom sequences	5
2 Discrete logarithms	7
2.1 The discrete logarithm problem	7
2.2 The Diffie-Hellman key exchange	8
2.3 The ElGamal cryptosystem	8
3 Integer factorization and RSA	11
3.1 The RSA cryptosystem	11
3.2 Primality testing	11
3.2.1 Fermat's little theorem and Carmichael numbers	12
3.2.2 The Miller-Rabin test	12
3.2.3 The Agrawal-Kayal-Saxena polynomial test	13
3.3 Multiplayer RSA – the bad way	15
4 Probability and information theory	17
4.1 The Vigenère cipher and its cryptanalysis	17
4.2 Collision and meet-in-the-middle attacks	18
4.3 Perfect secrecy and entropy	19
4.4 The redundancy of natural languages	24
5 Elliptic curves and cryptography	27
5.1 Elliptic curves and their abelian group structure	27
5.2 A sketch of the proof of Theorem 5.1.1	28
5.2.1 The resultant and Bézout's theorem	28
5.2.2 The Cayley-Bacharach theorem	29
5.2.3 Completion of the sketch	30
5.3 The elliptic curve discrete logarithm problem	30
5.4 Elliptic curve cryptography	31
5.4.1 The elliptic curve Diffie-Hellman	31
5.4.2 The elliptic curve ElGamal	31

6	Attacking the underlying problems	33
6.1	The discrete logarithm problem	33
6.1.1	A babystep-giantstep algorithm	33
6.1.2	The Pohlig-Hellman algorithm	33
6.1.3	The index calculus method	34
6.2	Factorization algorithms	35
6.2.1	Smooth numbers	35
6.2.2	Pollard's $p - 1$ method	35
6.2.3	Factorization via difference of squares	36
6.2.4	Lenstra's elliptic curve factorization	36
7	Additional topics	39
7.1	Interactive proofs	39
7.1.1	How to store the last move in chess?	39
7.1.2	A zero-knowledge proof of that a certain number is square modulo N	39
7.1.3	Using our password	41
7.2	Identification	41
7.2.1	An RSA-based digital signature	41
7.2.2	Multiplayer RSA – the good way	41
7.3	Secret sharing	42
	Appendices	45
A	Algebraic structures	47
A.1	Groups	47
A.2	Rings and fields	49
A.3	Vector spaces and projective spaces	50
B	Number theory	53
B.1	The fundamental theorem of arithmetic	53
B.2	Residue classes	54
B.3	The multiplicative group of the residue class rings	55
C	Probability theory	57
C.1	Discrete probability spaces	57

Preface

These lecture notes are written to provide a text to my Introduction to Mathematical Cryptography course at Budapest Semesters in Mathematics. The main source is [1], even the structure is borrowed from there. Note also that in [1], both the material and the collection of examples are much more extended.

As for prerequisites, needless to say, a general mathematical experience is indispensable. Also, even for such an introductory material, some familiarity with algebra, number theory and probability theory is necessary. For an overview of these subjects, I decided to include a detailed Appendix. Not so much inherently, yet for a good understanding of the notion of entropy introduced in Chapter 4, classical mathematical analysis might be useful, but certainly not beyond the level of [3].

Chapter 1

Introduction

1.1 The principal goal of cryptography, Kerckhoff's principle

The principal goal of cryptography is to **allow two people to exchange confidential information, even if they can only communicate via a channel monitored by an adversary.**

Assume for example that Bob wants to send a message to Alice in such a way that Eve – who reads/listens/spies the communication of Alice and Bob – cannot understand the message (Alice, Bob and Eve are the usual participants of the cryptographic setup).

The scheme of the solution is the following. Bob sends through the communication *something else* than his original message. Eve can read only this *something else*. Alice knows how this *something else* should be understood to get the original message.

Example 1.1.1 (an ancient method). We are in ancient times. Bob shaves the head of a slave. Then tattoos the message on the bald head. After hair has regrown, he sends the slave to Alice. Alice shaves the slave's head again and reads the message. (Here the assumption on the monitoring of the channel is that if Bob sent the slave with a letter to Alice, then Eve would steal the message from the slave and would read it.)

Example 1.1.2 (another ancient method – Caesar cipher). Bob shifts the alphabet the following way: he replaces each letter in his message with the letter which follows three later in the alphabet, e.g.

We attack the castle tomorrow.
Zh dwwdfn wkh fdvwoh wrpruurz.

which gives rise to the ciphertext (it is usual to rewrite the cipher in five-letter blocks, and leaving punctuation):

ZHDWW DFNWK HFDVW OHWRP RUURZ.

He writes down only the message 'ZHDWW DFNWK HFDVW OHWRP RUURZ', and sends this to Alice. Now Alice knows that she has to count three characters back and gets the original message, at least, in the form 'WEATT ACKTH ECAST LETOM ORROW', but it is obvious where words start and end. However, even if Eve can read the ciphertext, it is just meaningless for her.

Example 1.1.3 (Caesar's cipher improved – simple substitution ciphers). Set $\mathfrak{A} = \{a, b, c, \dots, z\}$ for the alphabet. Then assume Alice and Bob agree on a bijection $\pi : \mathfrak{A} \rightarrow \mathfrak{A}$ (this is called a permutation of \mathfrak{A}). Now if Bob wants to send a message to Alice, simply applies π to each letter, and sends the resulting text to Alice. Again, this text is a gibberish for Eve, while Alice knows how to recover the original message. Caesar cipher in Example 1.1.2 is just a special case of this when π is a shift of the letters.

What happens if Eve knows *what* the basic principle behind the encryption is? In Example 1.1.1, she simply captures and shaves the slave, and the cryptosystem is broken. However, in Example 1.1.2, the information that the encoding procedure is *some* shift of the alphabet still leaves 26 possibilities for the size of the shift (and hence for the original message). Even worse, in Example 1.1.3, the number of possibilities is $26! = 26 \cdot 25 \cdot \dots \cdot 1 > 10^{26}$.

Motivated by this phenomenon, Kerckhoff's principle says that the security of a cryptosystem should depend only on the secrecy of the *key* (this is the actual shift vector in Example 1.1.2 or the actual permutation in Example 1.1.3), not on the secrecy of the encryption *algorithm* (this is the fact that Alice and Bob use a shift in Example 1.1.2 or a permutation 1.1.3 of the alphabet to encode messages).

Already on this point, it is convenient to agree on the following ability of Eve (which is very natural to assume): she is able to recognize the message when she sees that. For example, in the Caesar cipher (Example 1.1.2), when she tries all the 26 possibilities for the translation vector, she will observe when she hits it. Although in principle it could happen that some other key gives another meaningful message, this is extremely unlikely (and its probability decreases very fast when the length of the message grows).

1.2 Cryptanalysis

In fact, Caesar's cipher can be easily broken: it is just 26 trials for the shift vector, and this can be done even by a human quite fast. Note that Eve has to try the first few characters in each possible shift, for an incorrect guess, most likely the first few letters will give rise to a gibberish.

This is not the case with the simple substitution cipher, the number of possible keys is quite large, $26! > 10^{26}$, which is too many even for a modern PC. However, even such a ciphertext can be revealed relatively easily. The point is that simple substitutions do not alter the characteristic of the underlying natural language, say, English in our case.

So Eve can argue as follows. In a normal English text (which is supposed the message to be), the most frequent characters are 'e', 't', 'a', 'o', 'n'. Even their frequency can be easily computed (you can find online tables telling them, or, by "bare" hand, you can open a long pdf file and count the occurrences of these letters). Also, you can consider the bigrams 'th', 'he', 'an', 're', 'er', and the trigrams 'the', 'and', 'ing'. Now given the ciphertext, if we count the often occurring letters, bigrams and trigrams, we may have reasonable guesses on the substitution: the point is that, for example, no matter which letter stands for 'e', it will stand for 'e' always, and since 'e' is the most often letter in English, this unknown character will be somewhat often in the cipher (and almost surely the most often, if the cipher is long enough). The same is true for other frequent letters, and also for bigrams and trigrams. As soon as some parts of the text are revealed, we can figure out further substitutions. After some trial and error, we can recover the text. Although this might seem complicated, it works surprisingly fast, read [1, Section 1.1.1] to see this in action.

As a historical note, we remark that cryptanalysis, letter frequency counts were known by Arab scholars in the 14th and the 15th century. The same time, in Italian states, more complicated cryptosystems were used than simple substitution ciphers, which suggests that cryptanalysis via frequency analysis was known there as well.

1.3 Mathematics enters cryptography

1.3.1 Transform texts to numbers

Up to this point, messages and ciphers were both textual objects. This is natural, our communication is textual, so most messages are simply texts. However, the communication of our computers (which deliver the messages nowadays) are rather based on numbers, and – more importantly – fancy properties of numbers provide us with extremely ingenious cryptosystems. For this reason, given any text, we – more precisely: our softwares – transform it to numbers. Our computers use encoding schemes, e.g. ASCII or UNICODE to convert characters to bytes. A byte consists of eight bits, where a bit is a 0 or a 1 (it is the abbreviation of the binary digit). For example, 01000001 is a byte, and in ASCII, it stands for the character 'A'. This encoding scheme is completely public, its purpose is not to hide information but on the contrary: it guarantees that if we send an e-mail to somebody, our computer transforms it to a number which can be transformed back to the same e-mail by his or her computer.

Of course, when Bob's e-mail is encoded in ASCII, and is sent to Alice, when Eve captures it, she is also able to read it. For this reason, Bob, after making the ASCII code, makes the message to a cipher (both are numbers this time). Ideally, when Eve reads the cipher, and decodes its ASCII, she will see a gibberish. Alice, who knows what to do, transforms the cipher back to the ASCII encoded version of the original message, then decodes it.

1.3.2 The mathematical formulation of symmetric ciphers

Definition 1.3.1 (symmetric cryptosystem). By a symmetric cryptosystem, we mean a 5-tuple

$$(\mathcal{K}, \mathcal{M}, \mathcal{C}, e, d),$$

where $\mathcal{K}, \mathcal{M}, \mathcal{C}$ are sets, and $e : \mathcal{K} \times \mathcal{M} \rightarrow \mathcal{C}$, $d : \mathcal{K} \times \mathcal{C} \rightarrow \mathcal{M}$ are functions satisfying that for any $k \in \mathcal{K}$ and $m \in \mathcal{M}$,

$$d(k, e(k, m)) = m.$$

Here, \mathcal{K} is the set of possible keys, \mathcal{M} is the set of possible messages, \mathcal{C} is the set of possible ciphers. The functions e and d are the encrypting and decrypting algorithms, respectively: for a given key $k \in \mathcal{K}$, e computes $c \in \mathcal{C}$ from $m \in \mathcal{M}$, while d computes $m \in \mathcal{M}$ from $c \in \mathcal{C}$. In practice, for a fixed $k \in \mathcal{K}$, $e(k, \cdot)$ and $d(k, \cdot)$ are often denoted by e_k and d_k , respectively.

In the realization, $\mathcal{K}, \mathcal{M}, \mathcal{C}, e, d$ are known to everyone (including Eve), while the particular $k \in \mathcal{K}$ in use is known only to Alice and Bob.

For example, in Example 1.1.2, the set of messages and ciphers are both the possible (finite) texts. The set of keys are the positive integers up to 26: $\mathcal{K} = \{1, 2, 3, \dots, 26\}$. Now e_k , for an integer $k \in \mathcal{K}$, shifts each letter by k , while d_k shifts it back.

In Example 1.1.3, the set of messages and ciphers are the same, but this time \mathcal{K} is the set of all permutations of the alphabet, a much bigger set than in Example 1.1.2.

We impose a few informal requirements on cryptosystems:

- (1) knowing $k \in \mathcal{K}$ and $m \in \mathcal{M}$, $e_k(m) \in \mathcal{C}$ must be easy to compute;
- (2) knowing $k \in \mathcal{K}$ and $c \in \mathcal{C}$, $d_k(c) \in \mathcal{M}$ must be easy to compute;
- (3) given one or more encoded ciphertexts $c_1, \dots, c_n \in \mathcal{C}$ encrypted using the key $k \in \mathcal{K}$, without knowing k , it must be hard to compute any of the plaintexts $d_k(c_1), \dots, d_k(c_n)$.

These are just natural conditions: the first two say that if Alice and Bob agree on the key, their algorithms to encrypt and decrypt texts are fast; the third one says that Eve – not knowing the particular key – cannot compute easily the original message.

A fourth condition is somewhat more complicated:

- (4) for any fixed $k \in \mathcal{K}$, given any number of pairs $(m_1, c_1), \dots, (m_n, c_n)$ (i.e. for any $1 \leq j \leq n$, $m_j \in \mathcal{M}$, $c_j \in \mathcal{C}$, $e_k(m_j) = c_j$, $d_k(c_j) = m_j$), if $c \in \mathcal{C} \setminus \{c_1, \dots, c_n\}$, then without knowing k , it must be hard to compute $d_k(c)$. This is security against a *chosen plaintext attack*.

For example, the Caesar cipher and even its generalized version, the simple substitution cipher (recall Examples 1.1.2, 1.1.3) satisfy the conditions (1), (2), but one can easily see that they do not satisfy (4). Indeed, in the Caesar cipher, even a single given pair (m_1, c_1) determines k , while in the simple substitution cipher, the pairs $(\text{'a'}, k(\text{'a'})), \dots, (\text{'z'}, k(\text{'z'}))$ determine k . It is not that obvious what we have seen in Section 1.2: even (3) does not hold for them, meaning that they are not secure in the sense that Eve does not have to organize any ingenious attack on the cryptosystem (like the chosen plaintext attack).

1.3.3 The mathematical formulation of asymmetric ciphers

In Examples 1.1.2 and 1.1.3, the role of Alice and Bob are symmetric. There is the key set \mathcal{K} , and they know which element $k \in \mathcal{K}$ they use (this has to be kept in secret). Such setups are the symmetric ciphers.

However, the relation of Alice and Bob is asymmetric (Bob sends, Alice reads the message), which might be utilized.

Definition 1.3.2 (asymmetric cryptosystem). By an asymmetric cryptosystem, we mean a 7-tuple

$$(\mathcal{K}, \mathcal{K}_{\text{pub}}, \mathcal{K}_{\text{priv}}, \mathcal{M}, \mathcal{C}, e, d),$$

where $\mathcal{K}, \mathcal{K}_{\text{pub}}, \mathcal{K}_{\text{priv}}, \mathcal{M}, \mathcal{C}$ are sets satisfying

$$\mathcal{K} \subseteq \mathcal{K}_{\text{pub}} \times \mathcal{K}_{\text{priv}},$$

and $e : \mathcal{K}_{\text{pub}} \times \mathcal{M} \rightarrow \mathcal{C}$, $d : \mathcal{K}_{\text{priv}} \times \mathcal{C} \rightarrow \mathcal{M}$ are functions satisfying that for any $(k_{\text{pub}}, k_{\text{priv}}) \in \mathcal{K}$ and $m \in \mathcal{M}$,

$$d(k_{\text{priv}}, e(k_{\text{pub}}, m)) = m.$$

Here, \mathcal{K} is the set of possible keys, \mathcal{M} is the set of possible messages, \mathcal{C} is the set of possible ciphers. The functions e and d are the encrypting and decrypting algorithms, respectively: for a given key $(k_{\text{pub}}, k_{\text{priv}}) \in \mathcal{K}$, e computes $c \in \mathcal{C}$ from $m \in \mathcal{M}$, while d computes $m \in \mathcal{M}$ from $c \in \mathcal{C}$. In practice, for a fixed $(k_{\text{pub}}, k_{\text{priv}}) \in \mathcal{K}$, $e(k_{\text{pub}}, \cdot)$ and $d(k_{\text{priv}}, \cdot)$ are often denoted by $e_{k_{\text{pub}}}$ and $d_{k_{\text{priv}}}$, respectively.

In the realization, $\mathcal{K}_{\text{pub}}, \mathcal{K}_{\text{priv}}, \mathcal{M}, \mathcal{C}, e, d$ are known to everyone (including Eve). It is only Alice who knows \mathcal{K} , and she picks a particular key $(k_{\text{pub}}, k_{\text{priv}}) \in \mathcal{K}$. Now she makes k_{pub} public, and keeps k_{priv} in secret. If Bob wants to send a message $m \in \mathcal{M}$ to Alice, then he applies $e_{k_{\text{pub}}}$ to m , and sends $e_{k_{\text{pub}}}(m)$ to Alice. Now Alice applies $d_{k_{\text{priv}}}$ to the incoming cipher, obtaining

$$d_{k_{\text{priv}}}(e_{k_{\text{pub}}}(m)) = m.$$

The requirements are modified the obvious way.

In this setup, there is nothing special about Bob, he has the same information as Eve has. Yet, if the system works well, Eve is unable to read Bob's messages to Alice. Although this seems more complicated and somewhat artificial, this scheme is commonly used. Think for example of your e-mail address: basically everyone knows it, none of your friends has a particular role, yet each of them can send you an e-mail which is hidden from anyone else¹.

1.4 Computability

In this section, we start to investigate the length of computations. Although a precise definition could be as well given using the notion of Turing machine, for our purpose an informal description will be completely sufficient.

Assume given an algorithm solving a problem. For example, let the problem be the addition of positive integers (given in base 10), and the algorithm is what we learned in the elementary school. Basically, the algorithm consists of two things: one is the addition table

+	0	1	2	3	4	5	6	7	8	9
0	0	1	2	3	4	5	6	7	8	9
1	1	2	3	4	5	6	7	8	9	10
2	2	3	4	5	6	7	8	9	10	11
3	3	4	5	6	7	8	9	10	11	12
4	4	5	6	7	8	9	10	11	12	13
5	5	6	7	8	9	10	11	12	13	14
6	6	7	8	9	10	11	12	13	14	15
7	7	8	9	10	11	12	13	14	15	16
8	8	9	10	11	12	13	14	15	16	17
9	9	10	11	12	13	14	15	16	17	18

and the other one is how to use this: the manipulation with the digits learned in the school. During this manipulation, we perform the addition of digits using the table (call them basic additions). If there are two numbers of n digits, then we have to use n basic additions (one at each digit). Also, there might be 1's taken on, at most n times, so the number of basic additions is at most $2n$. This means that given the basic addition table and the school algorithm, they can add up two numbers of n digits in $2n$ many steps, where by one step, we mean one reference to the basic addition table.

This is the general scheme: we can call *anything* a step, but that must be fixed once for all; also, we give an algorithm which can use the step as many times as it is needed. Then the input comes, and the algorithm computes the output. The running time (or time cost) of the algorithm (on the given input) is the number of steps used to compute the output.

Assume that the input is the positive integer n . We say that an algorithm computes the output in polynomial time, if the running time can be estimated from above by a polynomial of $\log n$. Why $\log n$?

¹At least would be, in a world where webmail providers do not read clients' mails – in ours they do.

Simply because the length of the input is just $\log_{10} n$ (and in other bases, just a constant times $\log_{10} n$), and not n . Similarly, the algorithm is exponential, if it is exponential in $\log n$. If the input consists of more than one numbers, say, n_1, \dots, n_k , then a polynomial algorithm is an algorithm whose running time can be estimated by a polynomial of $\max(\log n_1, \dots, \log n_k)$.

From the school, we know that addition, subtraction, multiplication and euclidean division can be performed by polynomial algorithms.

Proposition 1.4.1. *With input a, b , the number $\gcd(a, b)$ can be computed in polynomial time.*

Proof. Recall the notation of the proof of Proposition B.1.2: assume we have two numbers, $a \geq b > 0$, and then the sequence $(a, b, d_1, \dots, d_n, 0)$, where each d_i is defined via the euclidean division $d_{i-2} = c_{i-1}d_{i-1} + d_i$ (with $d_{-1} = a, d_0 = b, d_{n+1} = 0$) reproduces $d_n = \gcd(a, b)$. If we take the remainder d_i in each step such that $|d_i| \leq |d_{i-1}|/2$ (which can be done), then it is clear that the number of d_i 's is at most $1 + \log_2 a$. Since each euclidean divisions can be performed in polynomial time, the time cost altogether is

$$\text{Pol}(\log_2 a) + \text{Pol}(\log_2 b) + \text{Pol}(\log_2 d_1) + \dots + \text{Pol}(\log_2 d_{n-1}) \leq (1 + \log_2 a) \text{Pol}(\log_2 a) = \text{Pol}(\log_2 a).$$

The proof is complete. \square

In this proof, we spelled out that there are $\text{Pol}(\log)$ many euclidean divisions, and each of them can be computed in $\text{Pol}(\log)$ many steps, hence the algorithm runs in $\text{Pol}(\log)$ time. From now on, for brevity, when we have proved that *something* can be computed in polynomial time, we will use that *something* as a *step* in more complex algorithms: since the product of polynomials is a polynomial again, this simplification does not affect whether an algorithm is polynomial or not. For example, computing the gcd can be handled as a single step in a complex algorithm. Note that this simplification refers only to the polynomial being of the running time: the degree of the polynomial of course varies.

Proposition 1.4.2. *With input $a, b, N \in \mathbf{N}$, the residue class $a^b \bmod N$ can be computed in polynomial time.*

Proof. First of all, write b as

$$b = \sum_{0 \leq j \leq 1 + \log_2 b} \varepsilon_j 2^j,$$

where each ε_j is 0 or 1. (This can be done in polynomial time using only euclidean divisions: divide b by $2^0, 2^1, 2^2, 2^3, \dots, 2^{1 + \log_2 b}$. Alternatively, you may think of the input as numbers already given in base 2.)

Now

$$a^b = a^{\sum_{0 \leq j \leq 1 + \log_2 b} \varepsilon_j 2^j} = \prod_{0 \leq j \leq 1 + \log_2 b} a^{\varepsilon_j 2^j}.$$

Here, each factor with $\varepsilon_j = 0$ is just 1. When $\varepsilon_j = 1$, then

$$a^{\varepsilon_j 2^j} = a^{2^j} = \underbrace{((a^2)^2) \dots)^2}_{j \text{ many } 2\text{'s}}.$$

Now raise a to square (this is polynomial, since it is just a multiplication), then reduce it modulo N (a euclidean division). Then raise the result to square again, and reduce it modulo N . Doing this j many times, we get $a^{\varepsilon_j 2^j} \bmod N$ in polynomial time. Now we have to multiply together the factors $a^{\varepsilon_j 2^j}$ for all the j 's, this is polynomial many multiplications and modulo N reductions again. \square

1.5 The XOR cipher and pseudorandom sequences

Assume, the message is a number $0 \leq m \leq 2^t - 1$, i.e. a binary number on t bits. Now Alice and Bob agree on a binary number k also on t bits. So in this case,

$$\mathcal{M} = \mathcal{C} = \mathcal{K} = \{0 - 1 \text{ sequences of length } t\}.$$

Define now the \oplus operation as the bitwise addition, i.e. if $a = \sum_{j=0}^{t-1} a_j 2^j$, $b = \sum_{j=0}^{t-1} b_j 2^j$ (where a_j, b_j 's are binary digits, 0 or 1 each), then let

$$a \oplus b = \sum_{j=0}^{t-1} c_j 2^j,$$

where $c_j = 0$ if $a_j = b_j$, and $c_j = 1$ if $a_j \neq b_j$.

Given m and k , let $e_k(m) = m \oplus k$. One can easily see that then $d_k = e_k$:

$$d_k(e_k(m)) = (m \oplus k) \oplus k = m \oplus (k \oplus k) = m \oplus \underbrace{0 \dots 0}_{t \text{ zeros}} = m.$$

Since XOR addition can be computed fast, Bob can easily encrypt his message, and Alice can easily decrypt it. However, since Eve does not know k , she essentially has to check all possible k 's between 0 and $2^t - 1$, which is hopeless, if t is large enough.

In some sense, this cryptosystem is as perfect as can be. However, there are a few problems. First, it is vulnerable against a chosen plaintext attack, since even a single pair (m, c) reveals $k = m \oplus c$.

Another problem is that if Alice sends to message with the same key, say $c_1 = e_k(m_1)$, $c_2 = e_k(m_2)$, then Eve can XOR them to get

$$c_1 \oplus c_2 = (m_1 \oplus k) \oplus (m_2 \oplus k) = m_1 \oplus m_2.$$

It is not obvious how Eve can use this information, yet she has managed to get rid of k , and this is something that Alice and Bob would like to avoid.

Third, the key must be as long as the message is, which requires very long key sequences. And since Alice and Bob would like to use different keys for each encryption, they have to generate many k 's. Their security is the best possible, if they generate random k 's, for example, tossing a 0 – 1 coin t times gives a k , and they can repeat this as many times as they wish. But the setup is that they are at different places, their t -long tosses will almost certainly result differently. How can they solve this problem? Is it possible to securely and efficiently send long messages using only a single short key?

Assume there exists a function $R : \mathcal{K} \times \mathbf{N} \rightarrow \{0, 1\}$ satisfying the following properties:

- (1) for any $k \in \mathcal{K}$, $j \in \mathbf{N}$, it is easy to compute $R(k, j)$;
- (2) from any j_1, \dots, j_n and corresponding $R(k, j_1), \dots, R(k, j_n)$, it is hard to figure out k ;
- (3) from any j_1, \dots, j_n and corresponding $R(k, j_1), \dots, R(k, j_n)$, it is hard to guess the value of $R(k, j)$ with better than a 50% chance of success, if $j \notin \{j_1, \dots, j_n\}$.

Now Alice and Bob agree on a key k (a number on 200 bits, say), and they agree that for the first message, they will use $R(k, 1), \dots, R(k, 1000000)$, for the second one, they will use $R(k, 1000001), \dots, R(k, 2000000)$, and so on.

Such functions R are called pseudorandom number generators, however, it is unknown at the moment whether pseudorandom number generators exist.

Chapter 2

Discrete logarithms

Let us turn our attention back to asymmetric ciphers, those having a public and a private part. Their significance in computer age is bigger than ever, since it may easily happen that Alice and Bob would like to communicate although they have never met before (so they cannot be sure that any piece of information exchanged by them is handled confidentially).

How can Alice build up her cryptosystem? Given the sets \mathcal{M} and \mathcal{C} , she would like to publish a function $k_{\text{pub}} : \mathcal{M} \rightarrow \mathcal{C}$ in such a way that $k_{\text{pub}}^{-1} = k_{\text{priv}}$ can not be easily computed. For the first sight, this does not make too much sense: if it is difficult to invert k_{pub} for Eve (and Bob, and anyone else), then so is for Alice herself.

However, this is not the case. When Alice constructs k_{pub} , she uses some additional information, a trapdoor. This trapdoor plays an essential role in the construction of k_{pub} , and – ideally – has the following two properties: without the trapdoor, it is difficult, while using the trapdoor, it is easy to invert k_{pub} (and hence to get k_{priv}).

The second condition is the simpler one. To know that something is computationally easy, it suffices to have a good algorithm in hand, which solves the problem in little time. The critical point is the first condition. How to ensure that something is computationally difficult? Possibly there is no available method to invert a certain function but tomorrow there will be.

The cryptosystems presented below all depend on computational problems. Those which – according to the current state of art – are computationally difficult (i.e. at the moment, there are no fast algorithms to solve them) are considered to be secure.

2.1 The discrete logarithm problem

One of such computational problems is the discrete logarithm problem (DLP from now on). It is basically the following: given a finite abelian group G , and an element g generating it, for any $a \in G$, compute the smallest positive $k \in \mathbf{N}$ such that

$$\underbrace{g \cdot \dots \cdot g}_{k \text{ times}} = a.$$

This k is called the discrete logarithm of a (with base g), and we will denote it by $\log_g a$. The basic principle is the following: if we can perform the group operation fast, then from k , it is easy to get a (think of the repetitive squaring used in the proof of Proposition 1.4.2), while from a , it is not obvious to get k . This asymmetry gives us some chance to come up with a public key cryptosystem.

Following this principle, we will see below that using a finite abelian group, we can build up certain cryptosystems whose security rely a lot on the difficulty of the DLP. But is the DLP computationally difficult? We do not know it in general, but certainly it depends very much on the group.

For example, assume that the group is the additive group of the prime field \mathbf{F}_p . Then any $0 \neq g \in \mathbf{F}_p$ generates \mathbf{F}_p . Given a one may compute its logarithm as follows: if $\log_g a = k$, then

$$a \equiv \underbrace{g + \dots + g}_{k \text{ times}} \equiv gk \pmod{p}.$$

Now we see that k is nothing else but the multiplicative inverse of g multiplied by a (all understood modulo p). And the multiplicative inverse of g can be computed in polynomial time: either we write $1 = gu + pv$ and then $u \bmod g^{-1} \bmod p$ (and this can be done in polynomial time, since u is just produced by a euclidean algorithm, recall the proof of B.1.5) or by Euler-Fermat (Corollary B.3.7), $g^{p-2} \bmod p$ is the inverse of g modulo p (and this can also be computed in polynomial time, since this is just raising to a power with a modulus), recall Proposition 1.4.1 and Proposition 1.4.2. This altogether shows that the DLP in the additive group of prime fields is an easy problem, so certainly the methods below are useless with this group.

A much better choice is the multiplicative group \mathbf{F}_p^\times of the same field. We already know that \mathbf{F}_p^\times is cyclic, therefore, it is generated by some element. However, for our methods below, this is not really necessary, we can take an element g whose order modulo p is not too small, and can restrict to the subgroup generated by g . Then the DLP is the following: given g and a , compute the smallest possible k satisfying

$$g^k \equiv a \pmod{p}.$$

Current methods cannot solve this problem in polynomial time. This will be the starting point of the Diffie-Hellman key exchange and the ElGamal cryptosystem.

Later, we will learn how group structure can be attached to elliptic curves. It is also unknown but widely believed that the DLP is even more difficult in those groups. This gives us the possibility to build a cryptosystem on elliptic curves.

2.2 The Diffie-Hellman key exchange

In this section, we show how Alice and Bob can agree on a fixed number which is known only to them, even if they can only use a completely public channel for their whole communication. This is not as ambitious as sending complete encrypted messages, yet it is something: their common secret number can be for example used for a symmetric cipher.

First of all, they agree on a prime p and an element g whose order modulo p is sufficiently large. Now Alice takes a number a (and keeps it in secret), and Bob also takes a number b (and keeps it in secret). Then Alice computes $g^a \bmod p$, and sends it to Bob. Meanwhile, Bob computes $g^b \bmod p$, and sends it to Alice. Finally, Alice raises the incoming $g^b \bmod p$ to power a , while Bob raises the incoming $g^a \bmod p$ to power b . Both of them obtains the number $g^{ab} \bmod p$, so they have the common number (recall Proposition 1.4.2).

What happens to Eve? She just observes g^a and g^b (besides p, g). From this information, she should compute g^{ab} (this is called the Diffie-Hellman problem (DHP)). If she could solve effectively the DLP, she would immediately figure out the values a, b , and they lead to g^{ab} . Therefore, the DHP is not more difficult than the DLP. The converse is unknown (in principle, there can be a polynomial algorithm solving the DHP even if there is no algorithm solving the DLP in polynomial time).

2.3 The ElGamal cryptosystem

In this section, we describe a public key cryptosystem relying on the difficulty of the DLP for multiplicative groups of prime fields.

Again, Alice chooses a prime number p , and an element g of sufficiently large order modulo p . Also, she chooses a number a . Then she computes $A \equiv g^a \pmod{p}$, and publishes p, g, A (and keeps a in secret).

Then Bob chooses a number k (and keeps it in secret), and sends Alice the following two values:

$$c_1 \equiv g^k \pmod{p}, \quad c_2 \equiv mA^k \pmod{p},$$

where $1 \leq m \leq p - 1$ is the message he wants to send.

Now Alice gets the pair (c_1, c_2) , and computes $c_2(c_1^a)^{-1}$ modulo p . The result is

$$c_2(c_1^a)^{-1} \equiv mA^k((g^k)^a)^{-1} \equiv mg^{ak}g^{-ak} \equiv m \pmod{p},$$

so the message is recovered.

This is called the ElGamal cryptosystem. It is clear what k_{pub} is (note that it uses an external parameter k chosen by Bob). Now k_{priv} is (with current methods) is difficult to compute from k_{pub} , however, using the trapdoor a , it can be done fast.

It is clear that solving the DLP breaks the ElGamal, since Eve can easily compute a from A , then she can compute the same way m from (c_1, c_2) as Alice does.

Now we prove that breaking the ElGamal in general is essentially as difficult as solving the DHP.

Proposition 2.3.1. *Assume $p, g, A \equiv g^a \pmod p$ are fixed. If there is an algorithm, which computes m from (c_1, c_2) in polynomial time for any input (c_1, c_2) , then there is an algorithm solving the DHP with inputs g^a, g^b in polynomial time.*

Conversely, if there is an algorithm solving the DHP in polynomial time for any input g^a, g^b , then there is an algorithm computing the m -value from c_1, c_2 in polynomial time.

Proof. Assume we have to compute g^{ab} from g^a, g^b (all modulo p) in DHP. Set $c_1 \equiv g^b \pmod p$ and $c_2 = 1$. The computed m -value is g^{-ab} modulo p , its multiplicative inverse (computed in polynomial time) is g^{ab} modulo p .

As for the converse, from $A \equiv g^a \pmod p$ and $c_1 \equiv g^k \pmod p$, we can compute $c_1 \equiv g^{ak} \pmod p$ by our supposed algorithm. Now $m \equiv c_2 g^{-ak} \pmod p$. \square

Chapter 3

Integer factorization and RSA

In Chapter 2, the invertibility of k_{pub} reached by choosing the computationally complicated DLP: it is easy to raise a certain power, but it is hard to figure out the power from the result. In this chapter, we will use the following: it is easy to compute the product of prime powers, but it is hard to figure out the prime factorization of a given number.

3.1 The RSA cryptosystem

Alice chooses two large prime numbers p, q , they will serve for her as the trapdoor. Now she computes the product $N = pq$, and chooses a further number e coprime to $\varphi(N)$. Then she publishes N, e . Alice also computes the multiplicative inverse d of e modulo $\varphi(N)$ (that is, $de = \varphi(N)u + 1$). For her, this is easy to do: $\varphi(N) = (p - 1)(q - 1)$.

If Bob wants to send the message $1 \leq m \leq N$ to Alice, then he computes $c \equiv m^e \pmod{N}$ and sends it to Alice.

Then Alice simply raises $c \pmod{N}$ to power d , obtaining

$$c^d \equiv (m^e)^d \equiv m^{ed} \equiv m^{\varphi(N)u+1} \equiv m \cdot (m^{\varphi(N)})^u \equiv m \pmod{N},$$

at least when $\gcd(m, N) = 1$. For a randomly chosen message $1 \leq m \leq N$

$$\Pr(\gcd(m, N) = 1) = \frac{\varphi(N)}{N} = \frac{N - p - q + 1}{N} > 1 - \frac{1}{p} - \frac{1}{q},$$

which is very close to 1, if p, q are large.

Eve's problem is the following: she knows only N and e . If she could factorize N to pq , then it would immediately give $\varphi(N)$, leading to d fast. But there is no known method to factorize large numbers fast.

3.2 Primality testing

To implement RSA, one needs large prime numbers, but how can we get them? To answer this question, we first cite the prime number theorem:

$$\lim_{x \rightarrow \infty} \frac{\pi(x)}{x/\log x} = 1,$$

where $\pi(x)$ stands for the number of primes not exceeding x . Informally, we may say that up to x , there are approximately $x/\log x$ prime numbers. Reformulating this, we may say that if we pick a random large integer n , then we expect a prime in the set

$$\{n, n + 1, \dots, n + 2\lceil \log n \rceil\}.$$

So we choose some n , and then we take the integers $n, n + 1, \dots$ until we bump into a prime. The number of steps we have to take is estimated from above by the gap of consecutive primes. The prime number

theorem suggests that this is something like $\log n$ on average. However, it is known that the prime gap around X can be larger than any constant multiple of $\log X$ (it can be bigger than $100 \log X$, say). Conjecturally, the gap is always smaller than a constant multiple of $\log^2 X$ and the truth of this would suffice for a polynomial prime-generating algorithm.

But this is just one part of the problem. Okay, assume there *is* a prime between n and $n + 100 \log^2 n$. How will we recognize it? Given a large number, how fast can we decide if it is prime or not?

Of course, we can try any number up to \sqrt{n} and if none of them divides n (except for 1 of course), then n is a prime. If n is large, this is awfully slow, so we need a better algorithm.

3.2.1 Fermat's little theorem and Carmichael numbers

Assume a number n is given, and we have to decide if n is a prime or not.

First assume n is a prime. Recall Fermat's little theorem (Corollary B.3.8), which tells us then that

$$a^n \equiv a \pmod{n}.$$

So if we pick a number a such that

$$a^n \not\equiv a \pmod{n},$$

then we can be sure that n is not a prime. This gives us a primality test: take some a , and if $a^n \not\equiv a \pmod{n}$, then n is composite.

Unfortunately, the converse is not true at all, for example,

$$2^{341} \equiv 2 \pmod{341},$$

but $341 = 11 \cdot 31$ is not a prime. Even worse, there exist composite integers n such that for any $a \in \mathbf{N}$,

$$a^n \equiv a \pmod{n}.$$

Such numbers are called Carmichael numbers, and although they are rare compared to primes, there are infinitely many of them (the smallest one is $561 = 3 \cdot 11 \cdot 17$).

So if our randomly found number n is a Carmichael number, then it is not a prime, yet we have no chance to prove its compositeness via Fermat's little theorem.

3.2.2 The Miller-Rabin test

The fundament of the Miller-Rabin test is the following observation.

Proposition 3.2.1. *Assume p is an odd prime, and write $p - 1$ as $2^k q$, where q is an odd number. Then for any integer a coprime to p , one of the following two alternatives hold. Either $a^q \equiv 1 \pmod{p}$ or one of $a^q, a^{2q}, a^{4q}, \dots, a^{2^{k-1}q}$ is -1 modulo p .*

Proof. Fix an a as in the statement. If $a^q \equiv 1 \pmod{p}$, then we are done. If not then consider the number $a^{2^k q} \equiv a^{p-1}$, which is 1 modulo p by Euler-Fermat (Corollary B.3.7). Since \mathbf{F}_p is a field, the only residue classes x satisfying $x^2 \equiv 1 \pmod{p}$ are ± 1 . Therefore since the square of $a^{2^{k-1}q}$ is 1 modulo p , it must be either 1 or -1 modulo p . If it is -1 , then we are done. If it is 1, we can go further and take $a^{2^{k-2}q}$: again, since its square is 1, it is ± 1 modulo p . We continue this, and since $a^q \not\equiv 1 \pmod{p}$, at some point, we must get -1 modulo p . \square

Like Fermat's little theorem, this immediately gives us a prime test. Assume n is given. If it is even, then it is composite (unless it is 2). If it is odd, take $n - 1$ and halve it as many times as possible, altogether getting $n - 1 = 2^k q$ for some odd number q and positive integer k .

Choose a random $1 \leq a \leq n - 1$. Check first if it is coprime to n (this can be done fast by Proposition 1.4.1). If $\gcd(a, n) > 1$, then n is composite. If $\gcd(a, n) = 1$, then compute $a^q, a^{2q}, \dots, a^{2^k q}$ modulo n . They are at most $1 + \log_2 n$ many numbers, and they can be computed in polynomial time by Proposition 1.4.2. If a^q is not 1, and none of them is -1 modulo n , then we can be certain that n is composite.

But we have already experienced that things might go wrong: it can happen in principle that something is a fake-prime, by which we mean that it is composite yet either $a^q \equiv 1 \pmod{n}$ or one of $a^q, a^{2q}, \dots, a^{2^k q}$ is $-1 \pmod{n}$ for all possible choices for a . Why this method is better is the fact that there are no fake-primes (i.e. there are no analogues for Carmichael numbers here).

Proposition 3.2.2. *If n is an odd composite number, then there exists a coprime residue class a such that $a^q \not\equiv 1 \pmod{n}$ and $a^q, a^{2q}, a^{4q}, \dots, a^{2^{k-1}q} \not\equiv -1 \pmod{n}$.*

Proof. First, when n is a power of a prime, say, p^α , then take a generator a of the cyclic group \mathbf{Z}_n^\times . Then, since

$$\gcd(\varphi(n), n-1) = \gcd(p^\alpha - p^{\alpha-1}, p^\alpha - 1) = (p-1) \gcd(p^{\alpha-1}, p^{\alpha-1} + \dots + p + 1) = p-1,$$

$a^{n-1} \not\equiv 1 \pmod{n}$. Then obviously none of $a^q, a^{2q}, \dots, a^{2^{k-1}q}$ is ± 1 modulo n .

Now assume n is not a power of a prime, then write $n = n_1 n_2$, where $1 < n_1, n_2 < n$ and $\gcd(n_1, n_2) = 1$. By the Chinese remainder theorem, there exists a residue class a modulo n such that $a \equiv 1 \pmod{n_1}$ and $a \equiv -1 \pmod{n_2}$. This clearly implies $a \not\equiv \pm 1 \pmod{n}$, and then $a^q \not\equiv \pm 1 \pmod{n}$. On the other hand, $a^2 \equiv 1 \pmod{n_1}$ and $a^2 \equiv 1 \pmod{n_2}$, implying $a^2 \equiv 1 \pmod{n}$. Then $a^q \not\equiv \pm 1 \pmod{n}$, and $a^{2q}, a^{4q}, \dots, a^{2^{k-1}q} \equiv 1 \pmod{n}$. \square

In fact, it can be proved that at least 75% of the residue classes do this job. This gives rise to a random primality testing: take a convincing number of residue classes $a_1, \dots, a_h \pmod{n}$, and run the Miller-Rabin test with them. If any of them shows n is composite, then we know n is composite. If all of them allows n to be prime, then we can say n is a prime with a high confidence. How high? Let us estimate this using Bayes's theorem.

Choose a number at random which is approximately n . Then denote by X the event that the chosen number is composite, by X' that it is a prime, and by Y that it survives the Miller-Rabin test with h numbers. Then

$$\Pr(X | Y) = \frac{\Pr(Y | X)\Pr(X)}{\Pr(Y | X)\Pr(X) + \Pr(Y | X')\Pr(X')}.$$

Here, from the prime number theorem, we know that $\Pr(X') \approx \log n/n$ and $\Pr(X) \approx 1 - \log n/n$. Also, if the number is composite, it survives Miller-Rabin by probability at most 4^{-h} , so $\Pr(Y | X) \leq 4^{-h}$, while primes certainly survive Miller-Rabin, so $\Pr(Y | X') = 1$. Then

$$\Pr(X | Y) \lesssim \frac{4^{-h}(1 - \log n/n)}{4^{-h}(1 - \log n/n) + \log n/n}.$$

We see that this converges to 0 very fast: even if h is chosen to be $\lfloor \log_4 n \rfloor$ (which keeps the whole test polynomial), the probability of judging a composite number to be a prime tends to 0 (as n tends to infinity).

Also, it is known that if the Generalized Riemann Hypothesis is true, then for any composite number n there is an $a \leq 2 \log^2 n$ which proves the compositeness of n in the Miller-Rabin test. Hence assuming the Generalized Riemann Hypothesis, the Miller-Rabin test is not only a probabilistic, but a deterministic test which decides the primality of a number in polynomial time.

3.2.3 The Agrawal-Kayal-Saxena polynomial test

However, even the classical Riemann Hypothesis is open, especially so are the generalized versions. In this section, we present the AKS primality test, which decides primality in polynomial time *unconditionally*. Note also that this has only theoretical significance: first, as we will see, the AKS test is very complicated, while Miller-Rabin is easy to implement; second, almost all mathematicians strongly believe that the Riemann Hypothesis (and all its reasonable generalizations) are true – all in all, for practical purposes, Miller-Rabin is totally fine.

Assume n is a large number, fixed once for all.

Proposition 3.2.3. *There exists a number $1 \leq r = O(\log^{O(1)} n)$ such that the order of n in \mathbf{Z}_r^\times is greater than $\log_2^2 n$.*

Proof. For any $1 \leq i \leq \log_2^2 n$, consider the number $n^i - 1$:

$$n^i - 1 < n^i \leq (2^{\log_2 n})^{\log_2^2 n} = 2^{\log_2^3 n}.$$

Therefore the number of prime divisors of $n^i - 1$ for any $1 \leq i \leq \log_2^2 n$ is $O(\log^{O(1)} n)$.

Then picking the first prime r which does not appear as a prime divisor of

$$\prod_{1 \leq i \leq \log_2^2 n} (n^i - 1),$$

$r = O(\log^{O(1)} n)$ and the order of n modulo r is bigger than $\log_2^2 n$. \square

Now with this r in hand, for any $1 \leq a \leq n$, we can consider the polynomial $(X + a)^n$ in variable X . If n is prime, then

$$(X + a)^n \equiv X^n + a^n \equiv X^n + a \pmod{n},$$

the second congruence holds because $a^n \equiv a \pmod{n}$ by Fermat's little theorem B.3.8, the first congruence holds because in the binomial expression, all intermediate terms have a factor n .

This implies, in particular, that if n is a prime, then

$$(X + a)^n \equiv X^n + a \pmod{(n, X^r - 1)}. \quad (3.2.1)$$

The point is that although the polynomial $(X + a)^n$ takes a long time to compute, its modulo $X^r - 1$ reduction can be computed in polynomial time. Indeed, write n in base 2, and apply repeated squarings modulo $X^r - 1$ (like in the proof of Proposition 1.4.2).

Proposition 3.2.4. *If n is not a prime or a power of a prime, then there exists a number $1 \leq a = O(\log^{O(1)} n)$ such that (3.2.1) fails or $\gcd(a, n) > 1$.*

Proof. We prove by contradiction. Assume n is not a prime, nor a power of a prime, yet $\gcd(a, n) = 1$ and (3.2.1) are true for any $a = O(\log^{O(1)} n)$. Then n has a prime divisor $p > \log^C n$ with some number C (bigger than the $O(1)$ in the bound for a , in particular, we may assume $p > r$).

Now consider the field $\mathbf{F} = \mathbf{F}_p(X)$, where X is a primitive r th root of unity, and from now on, we compute in \mathbf{F} . The hypothesis (3.2.1) gives

$$(X + a)^n = X^n + a,$$

for any $1 \leq a \leq A$, where $A = O(\log^{O(1)} n)$ (and we choose A to be bigger than $2r$). Also, we have

$$(X + a)^p = X^p + a,$$

for $1 \leq a \leq A$, as above, since p is a prime. Then, for $1 \leq a \leq A$,

$$(X^p)^{n/p} + a = X^n + a = (X + a)^n = ((X + a)^p)^{n/p} = (X^p + a)^{n/p}.$$

Since X^p is a primitive r th root of unity again (and every primitive r th of unity is the p th power of a primitive r th root of unity), this implies

$$X^{n/p} + a = (X + a)^{n/p},$$

for any $1 \leq a \leq A$.

For any $1 \leq m \leq n$ coprime to r , we define the ring homomorphism $\tilde{\phi}_m : \mathbf{F} \rightarrow \mathbf{F}$ sending X to X^m , and we say that m is *nice*, if

$$\left(\tilde{X} + a\right)^m = \tilde{\phi}_m\left(\tilde{X} + a\right) = \tilde{X}^m + a$$

holds for all $1 \leq a \leq A$ and for any $\tilde{X} = X^t$, where t is coprime to r (i.e. \tilde{X} runs through all the primitive r th roots of unity).

Obviously, 1 is nice, and we have already seen that $n, p, n/p$ are all nice. Also, if m and m' are both nice, then so is their product:

$$\left(\tilde{X} + a\right)^{mm'} = \left(\left(\tilde{X} + a\right)^m\right)^{m'} = \left(\tilde{X}^m + a\right)^{m'} = \tilde{X}^{mm'} + a.$$

Let $G \subseteq \mathbf{F}^\times$ denote the multiplicative group generated by the elements $X + a$ with $1 \leq a \leq A$. Obviously, for any $z \in G$, $\phi_m(z) = z^m$.

Lemma 3.2.5. *Assume that there are exactly t nice residue classes modulo r of the form $p^i(n/p)^j$ for $i, j \geq 0$. Then $\#G \leq n^{\sqrt{t}}$.*

Proof of this lemma. For some $0 \leq i, j, i', j' \leq \sqrt{t}$ with $(i, j) \neq (i', j')$, we have

$$m = p^i(n/p)^j \equiv p^{i'}(n/p)^{j'} = m' \pmod{r},$$

implying $\phi_m(z) = \phi_{m'}(z)$ for any $z \in G$. The numbers m, m' here are distinct positive integers not exceeding $n^{\sqrt{t}}$. Then for any $z \in G$, $z^m - z^{m'} = 0$, which means that each $z \in G$ is a root of the polynomial $Z^m - Z^{m'}$. Since a nonzero polynomial over a field cannot have more roots than its degree, $\#G \leq \max(m, m') \leq n^{\sqrt{t}}$. \square

Lemma 3.2.6. *Assume that there are at least t nice residue classes modulo r . Then $\#G \geq 2^t$.*

Proof of this lemma. Make all the products with less than t factors from the elements $X + a$ (with $1 \leq a \leq A$, and allowing repetitions). Since $A > 2r \geq 2t$, there are at least 2^t many such products, and we claim that they are distinct.

By contradiction, let $P(X) = Q(X)$ for two such products, where $P(Z), Q(Z)$ are distinct polynomials of degree less than t . Since $P(X), Q(X) \in G$, for any nice m ,

$$P(X^m) = \phi_m(P(X)) = \phi_m(Q(X)) = Q(X^m).$$

If m_1, \dots, m_t are nice numbers of distinct remainders modulo r , then this means that X^{m_1}, \dots, X^{m_t} are all distinct roots of the polynomial $P(Z) - Q(Z)$, which is a contradiction: the polynomial is nonzero and has degree less than t , while it has at least t roots. \square

Since the order of n modulo r is at least $\log_2^2 n$, we see that the numbers $p^i(n/p)^j$ fall into at least $\log_2^2 n$ residue classes modulo n . Then $t > \log_2^2 n$ in the notation of the lemmas above. Then

$$2^t = 2^{t - \log_2^2 n} 2^{\log_2^2 n} = 2^{t - \log_2^2 n} 2^{\log_2^2 n} = 2^{t - \log_2^2 n} n^{\log_2 n},$$

and

$$n^{\sqrt{t}} = n^{\sqrt{t} - \log_2 n} n^{\log_2 n} = 2^{\sqrt{t} \log_2 n - \log_2^2 n} n^{\log_2 n},$$

so

$$2^t > n^{\sqrt{t}},$$

the lower bound beating the upper bound, which is obviously a contradiction. \square

Now the AKS test is the following. Given n , find r (as in Proposition 3.2.3) in polynomial time. Then for each of the polynomially many $1 \leq a \leq A$ (as in Proposition 3.2.4), check in polynomial time if it is coprime to n and then if the identity (3.2.1) holds. If they always hold, then n is a prime power (otherwise, it is composite). Now it is easy to check in polynomial time if n is in fact a prime: all the possible square roots, cube roots etc. can be found by dyadic splits of $[1, n]$ (this is polynomially many steps), and the number of possible exponents is trivially at most $\log_2 n$.

3.3 Multiplayer RSA – the bad way

In this section, we make our first attempt to construct a secure way of communication between n participants – based on RSA. Assume that a Trusted Authority takes two large prime numbers p, q , and computes their product $N = pq$. Then also computes, for any $1 \leq i \leq n$, a pair (e_i, d_i) such that $e_i d_i \equiv 1 \pmod{\varphi(N)}$. Now each e_i is made public, but the Trusted Authority sends only each decrypting exponent d_i to the i th participant. Now we see that the participants can communicate as we have seen above.

Unfortunately, there are several problems with this setup. First of all, the Trusted Authority is able to read any message, but let us put this aside. An even bigger problem is that basically any participant can read any message, even that ones which have been sent to other than her/him.

Assume the j th participant wants to decrypt a message sent to the k th participant. First, (s)he computes a factorization uv of $e_j d_j - 1$ such that $\gcd(u, e_k) = 1$, and each prime factor of v divides e_k .

This is easy to do: first (s)he computes $v_1 = \gcd(e_k, e_j d_j - 1)$, then $v_2 = \gcd(e_k, (e_j d_j - 1)/v_1)$, then $v_3 = \gcd(e_k, (e_j d_j - 1)/v_1 v_2)$ and so on, until the gcd becomes 1. Then $v = v_1 \cdot \dots \cdot v_t$ and $u = (e_j d_j - 1)/v$, and it is easy to see that both conditions hold.

Now since $\gcd(\varphi(N), e_k) = 1$, $\gcd(\varphi(N), v) = 1$. Since $\varphi(N) \mid uv$, this implies $\varphi(N) \mid u$. Then (s)he can compute a d'_k such that $e_k d'_k \equiv 1 \pmod{u}$, which in particular implies $e_k d'_k \equiv 1 \pmod{\varphi(N)}$. Maybe this d'_k is not the same as d_k but serves as well as d_k for decryptions.

But even if the participants are honest, an external eavesdropper can make use of the fact that there are several exponents in the setup. It is quite likely that in a long communication, it happens that a message is sent to two different recipients, i.e. the eavesdropper can intercept a pair $c_1 \equiv m^{e_1}, c_2 \equiv m^{e_2} \pmod{N}$. With a little luck, the exponents are coprime, and then $u, v \in \mathbf{Z}$ satisfying $e_1 u + e_2 v = 1$ are computable. But then

$$c_1^u c_2^v \equiv m^{e_1 u} m^{e_2 v} \equiv m^{e_1 u + e_2 v} = m \pmod{N}.$$

So messages sent to many participants are simply readable.

Chapter 4

Probability and information theory

In this chapter, a little background in mathematical analysis can be useful. We advise the reader to consult [3].

4.1 The Vigenère cipher and its cryptanalysis

For awhile, we return to alphabetic encryptions. In Chapter 1, we have already seen such ciphers: these replace each character with some other *fixed* character – these are called *monoalphabetic* ciphers. Now we turn to the Vigenère cipher, an example of *polyalphabetic* ciphers, where each letter is replaced with some other one, but this cipher letter may vary.

First Alice and Bob agree on a keyword. Then Bob encrypts his message as follows: he determines the shift he uses on each letter of the message by using the letters of the keyword one by one. The keyword letter **a** stands for no shift, **b** stands for a shift by one, and so on.

Example 4.1.1. Assume the keyword is 'CHESS'. Encrypt the message 'Learning mathematics is fun'. Then

```
LEARN INGMA THEMA TICSI SFUN
CHESS CHESS CHESS CHESS CHES
-----
NLEJF KUKES VOIES VPGKA UMYF,
```

so the ciphertext is NLEJF KUKES VOIES VPGKA UMYF.

The principle behind is the following. The letters of the keyword fix a variety of permutations of the alphabet, so breaking the cipher via counting common characters, bigrams, trigrams (which breaks simple substitution ciphers) is more difficult to perform in this case. Even if the length of the keyword has only five characters (four different, one is doubled), the common letter 'e' implies four common letters 'g', 'l', 'i', 'w' ('w' appearing with double frequency because of the double 's' in the keyword), and the common trigram 'the' in a long text implies four common trigrams 'voi', 'alw', 'xzw', 'lzg', 'ljl'. All in all, the characteristic of the language is totally mixed up.

For some time in history, the Vigenère cipher was considered to be unbreakable. Then in the 19th century, when basic statistical tools were developed, this view turned out to be very far from truth. In this section, we only describe the main methods to attack the cipher and suggest read [1, Section 4.2.2] to see how these approaches work in action.

The first goal is to determine the length of the keyword. We present two methods to do this.

Consider the repeated fragments (bigrams and trigrams) in the ciphertext, and list all the distances between the repetitions. Some of the repetitions may occur only by chance, but the length of the keyword is likely to divide many of the distances. This method depends on the fact that certain fragments are very frequent in natural languages (think of trigrams 'the' or 'ing' in English). In a not too short cipher, there is a considerable amount of occurrences of 'the', and if the key is not too long, some of them will be encrypted the same way. This method is known as the Kasiski test.

Another method is the index of coincidence test, which relies on individual letters rather than on two- or three-letter fragments. For each letter of the alphabet $i \in \mathfrak{A} = \{a, b, c, \dots, z\}$, let $F_i(\mathbf{s})$ be its frequency, i.e. the number of occurrences of i in a given string \mathbf{s} . With this notation, define, for any string \mathbf{s} , the index of coincidence

$$\text{IndCo}(\mathbf{s}) = \frac{1}{\text{length}(\mathbf{s})(\text{length}(\mathbf{s}) - 1)} \sum_{i \in \mathfrak{A}} F_i(\mathbf{s})(F_i(\mathbf{s}) - 1).$$

If a very long string is completely random, then all $F_i(\mathbf{s})$ are essentially $\text{length}(\mathbf{s})/26$, so we expect

$$\text{IndCo}(\mathbf{s}) \approx \frac{1}{26} \approx 0.0385.$$

However, if you consider a long text \mathbf{t} written in English, then you will find that

$$\text{IndCo}(\mathbf{t}) \approx 0.0685.$$

Although the difference of these two numbers seems small, it can be applied in practice. Assume there is a ciphertext c_1, \dots, c_n . Now for some $k \in \mathbf{N}$, split the ciphertext into k pieces $\mathbf{s}_0, \dots, \mathbf{s}_{k-1}$:

$$\mathbf{s}_i = c_i c_{i+k} c_{i+2k} \dots c_{i+\lfloor n/k \rfloor k}.$$

For each $0 \leq i \leq k-1$, compute $\text{IndCo}(\mathbf{s}_i)$. If k is the length of the keyword, then characters in \mathbf{s}_i are the encrypted versions of the plaintext using the same encryption, i.e. $\text{IndCo}(\mathbf{s}_i)$ is close to the index of coincidence 0.0685 of a natural English text (provided that \mathbf{s}_i is not too short). On the other hand, if k is not the right guess for the keyword, then $\text{IndCo}(\mathbf{s}_i)$'s are expected to be close to the index of coincidence 0.0385 of a random text. All in all, if the cipher is long enough, and the key is short enough, this method is expected to tell us the right key length.

Having figured out the key length, we introduce another coefficient which can help to figure out the difference between the shifts applied to get the \mathbf{s}_i 's. This is the mutual index of coincidence, defined as

$$\text{MutIndCo}(\mathbf{s}, \mathbf{t}) = \frac{1}{\text{length}(\mathbf{s}) \text{length}(\mathbf{t})} \sum_{i \in \mathfrak{A}} F_i(\mathbf{s}) F_i(\mathbf{t}),$$

for any strings \mathbf{s} and \mathbf{t} . Again, if \mathbf{s} and \mathbf{t} are encrypted with the same shift, then their mutual index of coincidence is large, otherwise, it is small. For any \mathbf{s}_i created above and any $\sigma \in \mathfrak{A}$, we define $\mathbf{s}_i + \sigma$, which is a further shift of each character of \mathbf{s}_i by σ (as in the Vigenère cipher: **a** stands for no shift, **b** stands for a shift by one, and so on). For any $0 \leq i < j \leq k-1$, and any $\sigma \in \mathfrak{A}$, compute

$$\text{MutIndCo}(\mathbf{s}_i, \mathbf{s}_j + \sigma).$$

Denoting the keyword by $\beta_1 \dots \beta_k$, and singling out the large values of $\text{MutIndCo}(\mathbf{s}_i, \mathbf{s}_j + \sigma)$, we find a linear system of equations of the form $\beta_i - \beta_j = \sigma$. This system can easily be overdetermined, when we have to leave out a few equations, but after some trial and error, we can determine $\beta_1 \dots \beta_k$ up to a shift which is *uniform* at all the k places. This leaves 26 possibilities, and trying them all, we arrive at the original plaintext.

4.2 Collision and meet-in-the-middle attacks

Proposition 4.2.1. *Assume A_1, \dots, A_n are independent events such that*

$$\Pr(A_1), \dots, \Pr(A_n) \geq p.$$

Then

$$\Pr((A_1 \cup \dots \cup A_n)^c) \leq e^{-np}.$$

Proof. First, we claim, for any $x \in [0, 1]$,

$$1 - x \leq e^{-x}.$$

Indeed, equality holds for $x = 0$, and the derivative (with respect to x) of the left-hand side is -1 , while that of the right-hand side is not less than -1 for any $x \in [0, 1]$. (Alternatively, writing $-x$ in place of x in the Taylor expansion $e^x = \sum_{j=0}^{\infty} x^j/j!$,

$$e^{-x} = 1 - x + \frac{x^2}{2!} - \frac{x^3}{3!} \pm \dots,$$

and for $0 \leq x \leq 1$, this is indeed at least $1 - x$, since for any $j \geq 1$, $x^{2j}/(2j)! - x^{2j+1}/(2j+1)! \geq 0$.) Then from the independency of A_1, \dots, A_n ,

$$\Pr((A_1 \cup \dots \cup A_n)^c) = \Pr(A_1^c \cap \dots \cap A_n^c) = \prod_{i=1}^n \Pr(A_i^c) \leq (1-p)^n \leq (e^{-p})^n = e^{-np}.$$

The proof is complete. \square

This statement enables us to attack certain problems with probabilistic algorithms. We illustrate this phenomenon via the DLP.

Proposition 4.2.2. *Assume G is a finite abelian group on N elements, and g is a generator. Let $a \in G$ be given. Then the DLP $g^x = a$ can be solved in $O(\sqrt{N})$ steps with high probability ($O(C\sqrt{N})$ steps lead to probability $1 - e^{-C^2}$ for any $C > 0$), where a step is an exponentiation in the group (which can be computed in $O(\log N)$ group multiplications, recall the repetitive squaring from Proposition 1.4.2).*

Proof. For some fixed $1 \leq n \leq N$, pick the elements y_1, \dots, y_n and z_1, \dots, z_n , and compute

$$g^{y_1}, \dots, g^{y_n}, ag^{z_1}, \dots, ag^{z_n}.$$

If there is a match, i.e.

$$g^{y_i} = ag^{z_j}$$

for some $1 \leq i, j \leq n$, then $a = g^{y_i - z_j}$. We claim that such a match occurs with high probability, if n is on the magnitude \sqrt{N} , the y_i 's are chosen independently and the z_j 's are chosen to be distinct.

As soon as z_1, \dots, z_n are fixed, each y_i gives a match with probability n/N . Since the y_i 's are independent, by Proposition 4.2.1,

$$\Pr(\text{there is no match}) \leq e^{-n^2/N}.$$

Here, if n is chosen to be $\lceil C\sqrt{N} \rceil$, the probability of not getting a match is at most e^{-C^2} . \square

4.3 Perfect secrecy and entropy

Now we try to understand how much information a given ciphertext reveals about the plaintext. For a proper formalization, we introduce the random variables M, K, C : M is a random plaintext ($M : \Omega \rightarrow \mathcal{M}$), K is a random key ($K : \Omega \rightarrow \mathcal{K}$), and C is a random ciphertext ($C : \Omega \rightarrow \mathcal{C}$). We denote by f_M, f_K, f_C their density functions, i.e.

$$f_M(m) = \Pr(M = m), \quad f_K(k) = \Pr(K = k), \quad f_C(c) = \Pr(C = c).$$

Assume that each $f_M(m), f_K(k), f_C(c)$ is positive, i.e. drop the possible messages, keys and ciphers of zero probability. (Although theoretically, impossibility does not mean the same as zero probability.)

Also, we may introduce the conditional probabilities

$$\begin{aligned} f_{M|K}(m | k) &= \Pr(M = m | K = k), & f_{M|C}(m | c) &= \Pr(M = m | C = c), \\ f_{K|C}(k | c) &= \Pr(K = k | C = c), & f_{K|M}(k | m) &= \Pr(K = k | M = m), \\ f_{C|M}(c | m) &= \Pr(C = c | M = m), & f_{C|K}(c | k) &= \Pr(C = c | K = k). \end{aligned}$$

Definition 4.3.1 (perfect secrecy). We say that the cryptosystem has perfect secrecy, if for all $m \in \mathcal{M}$ and $c \in \mathcal{C}$,

$$f_{M|C}(m | c) = f_M(m).$$

This definition indeed grabs the notion that c does not reveal information on m : intercepting the ciphertext c , Eve has the same chance when guessing m a priori.

Proposition 4.3.2. *If a cryptosystem has perfect secrecy, then $\#\mathcal{K} \geq \#\mathcal{M}$.*

Proof. Assume $\#\mathcal{K} < \#\mathcal{M}$. Take any $m \in \mathcal{M}$. Since for each $k \in \mathcal{K}$, $e_k(m)$ is a well-defined element of \mathcal{C} , we see that by assumption, there exists an element $c \in \mathcal{C}$ which does not appear as $e_k(m)$ when k runs through \mathcal{K} . This means that

$$f_{M|C}(m | c) = 0.$$

On the other hand, $f_M(m) > 0$, which excludes perfect secrecy. \square

In efficient cryptosystems, however, a single key is used many times and for many plaintexts, therefore, perfect secrecy is impossible to reach. Alice and Bob still have the endeavor to build a cryptosystem which gives Eve the least possible information.

To measure information, Shannon introduced the notion of entropy in communication theory.

Definition 4.3.3 (entropy). Given a random variable X taking the values x_1, \dots, x_n with probabilities p_1, \dots, p_n , respectively, the entropy is defined as

$$H(X) = H(p_1, \dots, p_n) = - \sum_{i=1}^n p_i \log_2 p_i,$$

with the convention that $0 \log_2 0$ is understood as 0 (note that this is in perfect accordance with $\lim_{x \rightarrow 0^+} x \log_2 x = 0$, which can be easily seen from l'Hôpital's rule).

For any $n \in \mathbf{N}$, set

$$U_n = \{(x_1, \dots, x_n) \in \mathbf{R}^n : x_1, \dots, x_n \geq 0 \text{ and } x_1 + \dots + x_n = 1\},$$

and

$$U = \bigcup_{n=1}^{\infty} U_n.$$

With this notation, we may think of the entropy function H as a function from U to \mathbf{R} . It is also symmetric, by which we mean that restricted to any U_n , it is invariant under any permutation of the n coordinates.

Proposition 4.3.4. *The function H satisfies that*

(H1) *H is continuous on any U_n ;*

(H2) *H is monotonically increasing on the uniform distributions, i.e. for any $m, n \in \mathbf{N}$, if $m > n$, then*

$$\log_2 m = H \left(\underbrace{\frac{1}{m}, \dots, \frac{1}{m}}_{m \text{ many}} \right) > H \left(\underbrace{\frac{1}{n}, \dots, \frac{1}{n}}_{n \text{ many}} \right) = \log_2 n;$$

(H3) *H is decomposable, i.e. if $(p_0, p_1, \dots, p_n) \in U_{n+1}$, and $q_1, \dots, q_m \in U_m$, then*

$$H(q_1 p_0, \dots, q_m p_0, p_1, \dots, p_n) = H(p_0, p_1, \dots, p_n) + p_0 H(q_1, \dots, q_m)$$

(obviously $(q_1 p_0, \dots, q_m p_0, p_1, \dots, p_n) \in U_{n+m}$).

Proof. All are straight-forward. \square

Proposition 4.3.5 (characterization of the entropy). *Assume that H' is a symmetric function from U to \mathbf{R} such that*

(H1) *H' is continuous on any U_n ;*

(H2) H' is monotonically increasing on the uniform distributions, i.e. for any $m, n \in \mathbf{N}$, if $m > n$, then

$$H' \left(\underbrace{\frac{1}{m}, \dots, \frac{1}{m}}_{m \text{ many}} \right) > H' \left(\underbrace{\frac{1}{n}, \dots, \frac{1}{n}}_{n \text{ many}} \right);$$

(H3) H' is decomposable, i.e. if $(p_0, p_1, \dots, p_n) \in U_{n+1}$, and $q_1, \dots, q_m \in U_m$, then

$$H'(q_1 p_0, \dots, q_m p_0, p_1, \dots, p_n) = H'(p_0, p_1, \dots, p_n) + p_0 H'(q_1, \dots, q_m)$$

(obviously $(q_1 p_0, \dots, q_m p_0, p_1, \dots, p_n) \in U_{n+m}$).

Then $H' = cH$ for some $c > 0$.

Proof. For any $n \in \mathbf{N}$, set

$$f(n) = H' \left(\underbrace{\frac{1}{n}, \dots, \frac{1}{n}}_{n \text{ many}} \right).$$

We claim that for any $m, n \in \mathbf{N}$, $f(mn) = f(m) + f(n)$. Indeed, applying (H3) m times:

$$\begin{aligned} H' \left(\underbrace{\frac{1}{mn}, \dots, \frac{1}{mn}}_{mn \text{ many}} \right) &= H' \left(\underbrace{\frac{1}{m}, \frac{1}{mn}, \dots, \frac{1}{mn}}_{(m-1)n \text{ many}} \right) + \frac{1}{m} H' \left(\underbrace{\frac{1}{n}, \dots, \frac{1}{n}}_{n \text{ many}} \right) \\ &= H' \left(\frac{1}{m}, \frac{1}{m}, \underbrace{\frac{1}{mn}, \dots, \frac{1}{mn}}_{(m-2)n \text{ many}} \right) + \frac{2}{m} H' \left(\underbrace{\frac{1}{n}, \dots, \frac{1}{n}}_{n \text{ many}} \right) \\ &\quad \dots \\ &= H' \left(\underbrace{\frac{1}{m}, \dots, \frac{1}{m}}_{m \text{ many}} \right) + H' \left(\underbrace{\frac{1}{n}, \dots, \frac{1}{n}}_{n \text{ many}} \right), \end{aligned}$$

which gives $f(mn) = f(m) + f(n)$.

Writing here $m = n = 1$, we obtain $f(1) = H'(1) = 0$ immediately. Then $0 < f(2) < f(3) < \dots$ by (H2). We claim that $f(n)/\log n$ is constant. To this aim, assume $p > 2$ is a prime number. For a large number $k \in \mathbf{N}$, sandwich p^k between two consecutive powers of 2:

$$2^{\lfloor \log_2 p^k \rfloor} < p^k < 2^{\lceil \log_2 p^k \rceil}.$$

Applying f and $f(mn) = f(m) + f(n)$, we obtain

$$\lfloor \log_2 p^k \rfloor f(2) < k f(p) < \lceil \log_2 p^k \rceil f(2).$$

Then

$$(\log_2 p^k - 1) f(2) < k f(p) < (\log_2 p^k + 1) f(2),$$

which gives

$$(k \log_2 p - 1) f(2) < k f(p) < (k \log_2 p + 1) f(2).$$

Dividing by $k \log p$,

$$\frac{f(2)}{\log 2} - \frac{f(2)}{k \log p} < \frac{f(p)}{\log p} < \frac{f(2)}{\log 2} + \frac{f(2)}{k \log p}.$$

Letting $k \rightarrow \infty$, this implies that $f(p)/\log p$ is the same positive number for each prime p . This clearly implies that $f(n)/\log n$ is a fixed positive constant. From this, it is clear that for some $c > 0$,

$$H' \left(\underbrace{\frac{1}{n}, \dots, \frac{1}{n}}_{n \text{ many}} \right) = cH \left(\underbrace{\frac{1}{n}, \dots, \frac{1}{n}}_{n \text{ many}} \right)$$

holds for all $n \in \mathbf{N}$.

In the next step, we would like to extend this $H' = cH$ from uniform distributions via (H3) to all distributions in $U_n \cap \mathbf{Q}^n$ for any $n \in \mathbf{N}$. First, for any $1 \leq a \leq n-1$, we have

$$H' \left(\underbrace{\frac{1}{n}, \dots, \frac{1}{n}}_{n \text{ many}} \right) = H' \left(\underbrace{\frac{a}{n}, \frac{1}{n}, \dots, \frac{1}{n}}_{n-a \text{ many}} \right) + \frac{a}{n} H' \left(\underbrace{\frac{1}{a}, \dots, \frac{1}{a}}_{a \text{ many}} \right).$$

The left-hand side and the second term on the right-hand side equals cH on the corresponding distributions, so is the first term on the right-hand side, i.e. $H' = cH$ for distributions of the form $(a/n, 1/n, \dots, 1/n)$:

$$\begin{aligned} H' \left(\underbrace{\frac{a}{n}, \frac{1}{n}, \dots, \frac{1}{n}}_{n-a \text{ many}} \right) &= H' \left(\underbrace{\frac{1}{n}, \dots, \frac{1}{n}}_{n \text{ many}} \right) - \frac{a}{n} H' \left(\underbrace{\frac{1}{a}, \dots, \frac{1}{a}}_{a \text{ many}} \right) \\ &= cH \left(\underbrace{\frac{1}{n}, \dots, \frac{1}{n}}_{n \text{ many}} \right) - \frac{a}{n} cH \left(\underbrace{\frac{1}{a}, \dots, \frac{1}{a}}_{a \text{ many}} \right) \\ &= c \left(H \left(\underbrace{\frac{1}{n}, \dots, \frac{1}{n}}_{n \text{ many}} \right) - \frac{a}{n} H \left(\underbrace{\frac{1}{a}, \dots, \frac{1}{a}}_{a \text{ many}} \right) \right) \\ &= cH \left(\underbrace{\frac{a}{n}, \frac{1}{n}, \dots, \frac{1}{n}}_{n-a \text{ many}} \right), \end{aligned}$$

where we used Proposition 4.3.4. Secondly, for any $1 \leq a \leq n-1$, we have

$$H' \left(\underbrace{\frac{a}{n}, \frac{1}{n}, \dots, \frac{1}{n}}_{n-a \text{ many}} \right) = H' \left(\frac{a}{n}, \frac{n-a}{n} \right) + \frac{n-a}{n} H' \left(\underbrace{\frac{1}{n-a}, \dots, \frac{1}{n-a}}_{n-a \text{ many}} \right),$$

and again $H' = cH$ for the left-hand side and the second term on the right-hand side implies the same for first term of the right-hand side, i.e. $H' = cH$ for $U_2 \cap \mathbf{Q}^2$. Now for any $(p_1, \dots, p_n) \in U_n \cap \mathbf{Q}^n$ with $n \geq 3$,

$$H'(p_1, \dots, p_n) = H'(p_1 + p_2, p_3, \dots, p_n) + (p_1 + p_2) H' \left(\frac{p_1}{p_1 + p_2}, \frac{p_2}{p_1 + p_2} \right),$$

and if $H' = cH$ on the right-hand side, then so is on the left, hence induction completes the proof of the fact that $H' = cH$ on $U_n \cap \mathbf{Q}^n$ for any $n \in \mathbf{N}$.

Now since $U_n \cap \mathbf{Q}^n$ is dense in U_n for any $n \in \mathbf{N}$, (H1) completes the proof. \square

Definition 4.3.6 (equivocation (conditional entropy)). Assume there are two random variables X, Y with values x_1, \dots, x_m and y_1, \dots, y_n , respectively. Their equivocation (conditional entropy) is defined as

$$H(X | Y) = - \sum_{i=1}^m \sum_{j=1}^n \Pr(Y = y_j) \Pr(X = x_i | Y = y_j) \log_2 \Pr(X = x_i | Y = y_j).$$

Proposition 4.3.7. *We have*

$$H(X, Y) = H(Y) + H(X | Y), \quad (4.3.1)$$

where by X, Y , we mean the random variable $X \times Y$, which takes the value (x, y) with probability $\Pr(X = x, Y = y)$ (and analogously to more variables). Also, if X and Y are independent, then

$$H(X, Y) = H(X) + H(Y). \quad (4.3.2)$$

Proof. We prove (4.3.1) first. Assuming the values of X are x_1, \dots, x_m and those of Y are y_1, \dots, y_n , we introduce the notation $p_{i,j} = \Pr(X = x_i, Y = y_j)$ for any $1 \leq i \leq m, 1 \leq j \leq n$. Then for any $1 \leq j \leq n$,

$$\Pr(Y = y_j) = \sum_{u=1}^m p_{u,j},$$

and for any $1 \leq i \leq m, 1 \leq j \leq n$,

$$\Pr(X = x_i | Y = y_j) = \frac{p_{i,j}}{\Pr(Y = y_j)} = \frac{p_{i,j}}{\sum_{u=1}^m p_{u,j}}$$

The left-hand side of (4.3.1) is, by definition,

$$H(X, Y) = - \sum_{i=1}^m \sum_{j=1}^n p_{i,j} \log_2 p_{i,j}.$$

On the right-hand side

$$H(X | Y) = - \sum_{i=1}^m \sum_{j=1}^n p_{i,j} \log_2 \left(\frac{p_{i,j}}{\sum_{u=1}^m p_{u,j}} \right).$$

Then

$$H(X, Y) - H(X | Y) = - \sum_{i=1}^m \sum_{j=1}^n p_{i,j} \log_2 \left(\sum_{u=1}^m p_{u,j} \right) = - \sum_{j=1}^n \Pr(Y = y_j) \log_2 \Pr(Y = y_j) = H(Y),$$

and the proof of (4.3.1) is complete.

Now assume X and Y are independent random variables with values $x_1, \dots, x_m, y_1, \dots, y_n$, respectively. Let $\Pr(X = x_i) = p_i, \Pr(Y = y_j) = q_j$, by independency, $\Pr(X = x_i, Y = y_j) = p_i q_j$ then. Now

$$H(X | Y) = - \sum_{i=1}^m \sum_{j=1}^n p_i q_j \log_2 p_i = - \sum_{i=1}^m p_i \log_2 p_i = H(X),$$

and then (4.3.2) follows from (4.3.1). \square

Proposition 4.3.8. *We have*

$$H(X | Y) \leq H(X).$$

Proof. Assuming the values of X are x_1, \dots, x_m and those of Y are y_1, \dots, y_n , we introduce the notation $p_{i,j} = \Pr(X = x_i, Y = y_j)$ for any $1 \leq i \leq m, 1 \leq j \leq n$. The difference $H(X) - H(X | Y)$ can be written as

$$\begin{aligned} H(X) - H(X | Y) &= - \sum_{i=1}^m \left(\sum_{j=1}^n p_{i,j} \right) \log_2 \left(\sum_{v=1}^n p_{i,v} \right) - \left(- \sum_{i=1}^m \sum_{j=1}^n p_{i,j} \log_2 \left(\frac{p_{i,j}}{\sum_{u=1}^m p_{u,j}} \right) \right) \\ &= - \sum_{i=1}^m \sum_{j=1}^n p_{i,j} \log_2 \left(\frac{\sum_{u=1}^m p_{u,j} \sum_{v=1}^n p_{i,v}}{p_{i,j}} \right) \\ &\geq - \log_2 \left(\sum_{i=1}^m \sum_{j=1}^n p_{i,j} \cdot \frac{\sum_{u=1}^m p_{u,j} \sum_{v=1}^n p_{i,v}}{p_{i,j}} \right) \\ &= - \log_2 \left(\sum_{j=1}^n \sum_{u=1}^m p_{u,j} \sum_{i=1}^m \sum_{v=1}^n p_{i,v} \right) = - \log_2 1 = 0, \end{aligned}$$

by applying Jensen's inequality in the penultimate line, combined with the concavity of $t \mapsto \log_2 t$. \square

Definition 4.3.9 (mutual information). Given random variables X and Y , their mutual information is defined as

$$I(X, Y) = H(X) - H(X | Y).$$

Proposition 4.3.10. *The mutual information is symmetric and nonnegative, i.e. for any random variables X and Y , $I(X, Y) = I(Y, X) \geq 0$.*

Proof. Nonnegativity is clear from Proposition 4.3.8. As for symmetry, by Proposition 4.3.7,

$$I(X, Y) = H(X) - H(X | Y) = H(X) - (H(X, Y) - H(Y)) = H(X) + H(Y) - H(X, Y),$$

which clearly implies the statement. \square

In the cryptographic setup, we would like to measure the information that the cipher gives on the key, and define $H(K | C)$ to be the key equivocation. The higher its value is, the less information the cipher reveals about the key.

Proposition 4.3.11. *If M and K are independent, then*

$$H(K | C) = H(M) + H(K) - H(C).$$

Proof. First, by (4.3.1), we have

$$H(M, K, C) = H(C | M, K) + H(M, K).$$

Here, since C is determined by M and K , $H(C | M, K) = 0$, and by (4.3.2),

$$H(M, K) = H(M) + H(K).$$

Similarly, since M is determined by K and C ,

$$H(M, K, C) = H(M | K, C) + H(K, C) = H(K, C).$$

Then

$$H(K | C) = H(K, C) - H(C) = H(M, K, C) - H(C) = H(M, K) - H(C) = H(M) + H(K) - H(C).$$

The proof is complete. \square

4.4 The redundancy of natural languages

If \mathbf{L} is a language, denote by L the random variable where the values are the letters of the language, and the probabilities are the relative frequencies. Set L^2 for the bigrams, L^3 for the trigrams and so on. Then the entropy of the language is defined to be

$$H(\mathbf{L}) = \lim_{n \rightarrow \infty} \frac{H(L^n)}{n}.$$

Proposition 4.4.1. *For any \mathbf{L} , $H(\mathbf{L})$ exists.*

Proof. For any $m, n \in \mathbf{N}$, we have

$$\frac{H(L^{nm})}{nm} \leq \frac{H(\overbrace{L^n, \dots, L^n}^{m \text{ many}})}{nm} = \frac{H(L^n)}{n}.$$

Now for a fixed n , and letting $m \rightarrow \infty$, introducing $m = nc_m + d_m$, where $0 \leq d_m < n$,

$$\frac{H(L^m)}{m} = \frac{H(L^{nc_m + d_m})}{m} \leq \frac{H(L^n)}{n} + o(1).$$

Clearly the sequence $(H(L^n)/n)_{n \in \mathbf{N}}$ is bounded (each element is between 0 and $H(L)$). Assume

$$S = \limsup_{n \rightarrow \infty} \frac{H(L^n)}{n}, \quad I = \liminf_{n \rightarrow \infty} \frac{H(L^n)}{n}$$

are different. Fix n such that $H(L^n)/n < I + (S - I)/3$. Now take m 's tending to infinity such that $H(L^m)/m > S - (S - I)/3$. This obviously gives a contradiction, when $o(1)$ in the above calculation gets smaller than $(S - I)/3$. \square

Experimentally, it seems that

$$1 \leq H(\text{English}) \leq 1.5.$$

Note also that

$$\log_2 26 \approx 4.7.$$

This means that the English language (like any other natural language) admits a high redundancy: everything which is expressed, could be expressed on less than one third of the used space. This explains the phenomenon that when you apply cryptanalysis to a simple substitution cipher (recall Section 1.2), you do not have to decrypt all the characters, the plaintext becomes readable earlier.

Chapter 5

Elliptic curves and cryptography

5.1 Elliptic curves and their abelian group structure

Given a field \mathbf{F} of characteristic bigger than 3, and elements $A, B \in \mathbf{F}$ such that $4A^3 + 27B^2 \neq 0$, the elliptic curve defined by them is defined on the projective plane is

$$E = \{[X, Y, Z] \in \mathbf{P}^2(\mathbf{F}) : Y^2Z = X^3 + AXZ^2 + BZ^3\}.$$

On the subset $\{Z \neq 0\}$, this is the set (with affine coordinates $x = XZ^{-1}$, $y = YZ^{-1}$)

$$E_{Z \neq 0} = \{(x, y) \in \mathbf{F}^2 : y^2 = x^3 + Ax + B\},$$

and when $Z = 0$, then automatically $X = 0$, giving

$$E_{Z=0} = \{[0, 1, 0]\},$$

from now on, this point $[0, 1, 0]$ will be denoted by \mathcal{O} .

Given two different points $P = [X_P, Y_P, Z_P]$ and $Q = [X_Q, Y_Q, Z_Q]$ on E , they determine a line passing through both of them, namely, when $P \neq Q$, then

$$L_{P,Q} = \{\lambda_1[X_P, Y_P, Z_P] + \lambda_2[X_Q, Y_Q, Z_Q] : [\lambda_1, \lambda_2] \in \mathbf{P}^1(\mathbf{F})\},$$

while if $P = Q$, then it is the tangent line of E at P . To understand this tangent line, let us use the affine coordinates. On $E_{Z \neq 0}$, at a point $(x, y) \in E$ (assume for a moment that $y \neq 0$),

$$\frac{d}{dx}y^2 = \frac{d}{dx}(x^3 + Ax + B), \quad \frac{2ydy}{dx} = 3x^2 + A, \quad \frac{dy}{dx} = \frac{3x^2 + A}{2y}.$$

This gives that the tangent line at $(x_P, y_P) \in E_{Z \neq 0, Y \neq 0}$ is

$$\left\{ (x, y) \in \mathbf{F}^2 : y = \frac{3x_P^2 + A}{2y_P}x - \frac{3x_P^2 + A}{2y_P}x_P + y_P \right\}.$$

Taking the projectivization, we obtain the line

$$\{[X, Y, Z] \in \mathbf{P}^2(\mathbf{F}) : 2y_P Y = (3x_P^2 + A)X + ((-3x_P^2 - A)x_P + 2y_P^2)Z\},$$

and we see that this definition gives the tangent line even for $y_P = 0$: it contains the point $[X_P, Y_P, Z_P]$, and also any point with coordinates $[X_P, Y, Z_P]$, which is just the vertical line on the xy -plane (and this is the tangent line, when the slope is infinity, i.e. when we divide by $y_P = 0$ above).

At the point \mathcal{O} , we see that the tangent line is $\{Z = 0\}$: indeed, setting $Z = 0$, we obtain the equation $0 = X^3$, which has a triple root at $X = 0$, so \mathcal{O} is a triple intersection point of E and $L_{\mathcal{O}, \mathcal{O}}$.

All in all, given any two point $P, Q \in E$, we can attach a line $L_{P,Q}$ passing through them, and if $P = Q$, then further being tangent to E .

We claim this line intersects E in a further point of E . We go by cases.

If $P = Q = \mathcal{O}$, then \mathcal{O} is the third intersection point (recall that \mathcal{O} is a triple intersection of E and $L_{\mathcal{O},\mathcal{O}}$).

If $P = \mathcal{O}$, and $Q = [X_Q, Y_Q, Z_Q] \neq \mathcal{O}$, then we have two subcases. If $Y_Q \neq 0$, then $[X_Q, -Y_Q, Z_Q]$ is the third intersection point. If $Y_Q = 0$, then $L_{\mathcal{O},Q}$ is the tangent line at Q , so the third intersection point is Q itself.

If $P = [X_P, Y_P, Z_P] \neq \mathcal{O}$, $Q = [X_Q, Y_Q, Z_Q] \neq \mathcal{O}$, and $x_Q = x_P, y_Q = -y_P$ in the affine notation (including the case $Y_P = Y_Q = 0$, i.e. $P = Q$), then $L_{P,Q}$ is a vertical line on the xy -plane, which intersects E further in \mathcal{O} .

Now assume $P = [X_P, Y_P, Z_P] \neq \mathcal{O}$, $Q = [X_Q, Y_Q, Z_Q] \neq \mathcal{O}$, and their x -coordinates (on the xy -plane) are different. Then the line passing through them is $y = mx + c$ for some $m, c \in \mathbf{F}$. Then

$$(mx + c)^2 = x^3 + Ax + B$$

has two roots x_P, x_Q (counted with multiplicity, when $P = Q$). Then factoring them out, we obtain a third solution, giving rise to a third intersection point.

Then we may define an addition on E . Given two points $P, Q \in E$, take the third intersection point of $L_{P,Q}$ and E (call it R), and then the third intersection point of $L_{R,\mathcal{O}}$ and E . This will be the point $P + Q$.

Theorem 5.1.1. *Under this addition, the points of E form an abelian group with unit element \mathcal{O} .*

One can easily check that the third intersection point R is nothing else but $-P - Q$.

5.2 A sketch of the proof of Theorem 5.1.1

In this section, we are going to sketch the proof of Theorem 5.1.1. Commutativity, unit element and invertibility are easy to see, the difficult (and deep) part is the associativity. To this aim, we assume that the underlying field \mathbf{F} is algebraically closed: we are free to do this, since the points are always defined as third intersection points, and if two intersection points are in a certain subfield, then so is the third one.

5.2.1 The resultant and Bézout's theorem

For fixed $n \in \mathbf{N} \cup \{0\}$, denote by $\mathcal{P}_n \subset \mathbf{F}[Z]$ the vector space of polynomials of degree smaller than n (completed with the zero polynomial). Define the linear map

$$\varrho(A, B) = PA + QB, \quad \varrho : \mathcal{P}_n \times \mathcal{P}_m \rightarrow \mathcal{P}_{m+n},$$

where $P(Z) = a_m Z^m + \dots + a_1 Z + a_0$ and $Q(Z) = b_n Z^n + \dots + b_1 Z + b_0$ are fixed polynomials (with $a_m, b_n \neq 0$). Taking the basis $(Z^{n-1}, \dots, Z, 1, Z^{m-1}, \dots, Z, 1)$ in $\mathcal{P}_n \times \mathcal{P}_m$ and the basis $(Z^{m+n-1}, \dots, Z, 1)$ in \mathcal{P}_{m+n} , we see that the matrix of ϱ is

$$S = \begin{pmatrix} a_m & a_{m-1} & \dots & a_0 & 0 & \dots & 0 \\ 0 & a_m & \dots & a_1 & a_0 & \dots & 0 \\ & & & \dots & & & \\ 0 & \dots & 0 & a_m & a_{m-1} & \dots & a_0 \\ b_n & b_{n-1} & \dots & b_0 & 0 & \dots & 0 \\ 0 & b_n & \dots & b_1 & b_0 & \dots & 0 \\ & & & \dots & & & \\ 0 & \dots & 0 & b_n & b_{n-1} & \dots & b_0 \end{pmatrix}.$$

Now ϱ is singular if and only if $\det S = 0$. Singularity is also equivalent to that P and Q share a common root: if $\deg \gcd(P, Q) > 0$, then ϱ is not surjective (since $\deg(\gcd(P, Q))$ divides all elements of $\varrho(\mathcal{P}_n \times \mathcal{P}_m)$), while if $\deg \gcd(P, Q) = 0$, then ϱ is surjective (since all elements of \mathcal{P}_{m+n} is a polynomial combination of P and Q).

Therefore, $\det S$ is zero if and only if P and Q share a common root. This quantity $\det S$ is called the resultant of P and Q .

Proposition 5.2.1. *Let A be a unique factorization domain, and assume $f, g \in A[x]$ are nonzero polynomials. Then f and g have a common nonconstant factor if and only if the equation $uf + vg = 0$ has a nontrivial solution such that $\deg v < \deg f$ and $\deg u < \deg g$.*

Proof. First, if f and g share a nonconstant factor $h \in A[x]$. Then setting $(g/h)f + (-f/h)g = 0$, and obviously $\deg(f/h) < \deg f$, $\deg(-g/h) < \deg g$. For the converse, take an admissible pair (u, v) , then each irreducible factor of f divides vg . Since $\deg v < \deg f$, there must be an irreducible factor of f which appears as a divisor of g . \square

Theorem 5.2.2 (Bézout). *Assume $f, g \in \mathbf{F}[x, y]$ are polynomials such that their degrees are m, n , respectively. If f, g do not share common factors, then they have at most mn common zeros.*

Proof. Assume the corresponding projectivizations are $F(X, Y, Z) = a_0Z^m + \dots + a_{m-1}Z + a_m$, and $G(X, Y, Z) = b_0Z^n + \dots + b_{n-1}Z + b_n$, where a_j, b_j are the degree j part of f, g , respectively. It is easy to see that no common factor of f, g implies the same for F, G . Then there are no nontrivial solutions of $uF + vG = 0$ such that $\deg v < \deg F$, $\deg u < \deg G$. This means that the resultant S made of the coefficients $a_0, \dots, a_m, b_0, \dots, b_n$ (note that these are now polynomials) is not constant zero. Then $\deg \det S \leq mn$, and by the fundamental theorem of algebra, we may factorize $\det S$ as $\prod(\xi x - \zeta y)$, and the number of factors is at most mn . Each such factor gives at most $\max(m, n)$ possibilities for Z , so there are finitely many common zeros $[X : Y : Z]$. Changing coordinates, we may choose a base point which is not a common zero, neither lies on a line connecting to common zeros (since \mathbf{F} is infinite). With this base point as $[0 : 0 : 1]$, and repeating the above argument, each factor $(\xi x - \zeta y)$ gives at most one Z . \square

5.2.2 The Cayley-Bacharach theorem

If $F(X, Y, Z)$ is a homogeneous polynomial of degree 3 over the base field \mathbf{F} , then its zero set $\gamma = \{[X : Y : Z] \in P\mathbf{F}^2\}$ is called a cubic curve.

Theorem 5.2.3 (Cayley-Bacharach). *Assume γ_1 and γ_2 are two cubic curves sharing exactly nine intersection points P_1, \dots, P_9 . Then if γ is a further cubic curve passing through the first eight intersection points P_1, \dots, P_8 , then γ is a linear combination of γ_1 and γ_2 (by this we mean that if $\gamma_1, \gamma_2, \gamma$ are the vanishing sets of F_1, F_2, F , respectively, then for some $\lambda, \mu \in \mathbf{F}^2$, $\lambda F_1 + \mu F_2 = F$). In particular, γ passes through P_9 .*

Proof. We prove by contradiction, assume F_1, F_2 are given as in the statement, and F is a degree 3 homogeneous polynomial vanishing at P_1, \dots, P_8 , which is not a linear combination of F_1 and F_2 .

First we claim that no four of P_1, \dots, P_9 lie on any line l (the vanishing set of the homogeneous degree 1 polynomial L), since by Bézout's theorem, this would mean that this L is a factor of F_1, F_2 , leading to infinitely many common points of γ_1, γ_2 . Similarly, no seven of P_1, \dots, P_9 can lie on any conic q (the vanishing set of the homogeneous degree q polynomial Q), since by Bézout's theorem, this would mean that this Q is a factor of F_1, F_2 , leading to infinitely many common points of γ_1, γ_2 .

Then we claim that any five of P_1, \dots, P_9 determine a unique conic passing through them. First, for the existence, observe that a conic has an equation

$$aX^2 + bXY + cZ^2 + dXZ + eYZ + fZ^2 = 0,$$

with not all of a, b, c, d, e, f being zero. Observe that fixing a point is nothing else but a linear condition on the coefficients a, b, c, d, e, f . Giving five such conditions still leaves a nonzero solution. Now assume there are two conics q_1, q_2 passing through the same five points. By Bézout's theorem, the corresponding polynomials Q_1, Q_2 must share a common factor L (with vanishing set l , a line). This line l can contain at most three of the points, so the line containing the remaining two points must be another common factor of Q_1, Q_2 , leading to $q_1 = q_2$.

Case 1: three of the first eight points, say, P_1, P_2, P_3 are collinear. Let the line containing them be l . Let q be the conic containing P_4, \dots, P_8 . Now let X be a further point on l , and Y be a further point not on either l or q . Since F, F_1, F_2 are linearly independent, there is a nonzero linear combination of them $G = aF + bF_1 + cF_2$ which vanishes at both X and Y . Then the vanishing set of G passes through P_1, \dots, P_8, X , so it must be the union of l and q (because it has 4 common points with l ,

and Bézout applies: $4 > 3 \cdot 1$; so the vanishing set of G must be the union of l and a conic, and this conic must be q , because of 5 common points. However $Y \notin l \cup q$, which is a contradiction.

Case 2: Case 1 does not hold, and six of the first eight points, say, P_1, \dots, P_6 are on a conic. Let the conic containing them be q . Let l be the line containing P_7, P_8 . Now let X be a further point on q , and Y be a further point not on either q or l . Since F, F_1, F_2 are linearly independent, there is a nonzero linear combination of them $G = aF + bF_1 + cF_2$ which vanishes at both X and Y . Then the vanishing set of G passes through P_1, \dots, P_8, X , so it must be the union of l and q (because it has 7 common points with q , and Bézout applies: $7 > 3 \cdot 2$; so the vanishing set of G must be the union of q and a line, and this conic must be l , because of 2 common points). However $Y \notin q \cup l$, which is a contradiction.

Case 3: none of Case 1 and Case 2 holds. Let l be the line passing through P_1, P_2 , and q the conic passing through P_3, \dots, P_7 . In this case then, by assumption, $P_8 \notin l \cup q$. Let X and Y be further points on l but not on q . Since F, F_1, F_2 are linearly independent, there is a nonzero linear combination of them $G = aF + bF_1 + cF_2$ which vanishes at both X and Y . Then, as above, the vanishing set of G must be $l \cup q$. The contradiction follows from that on the one hand, G vanishes at P_8 (since it is a linear combination of F, F_1, F_2), but on the other hand, $P_8 \notin l \cup q$.

The proof is complete. □

5.2.3 Completion of the sketch

Assume P, Q, R are points on the elliptic curve, and assume the points $\mathcal{O}, P, Q, R, P+Q, Q+R, -(P+Q), -(Q+R), -((P+Q)+R)$ are all different. Now consider the following cubics c_1, c_2, c_3 . Let $c_1 = E$ itself. Let

$$\begin{aligned} c_2 = & \text{line passing through } P \text{ and } Q \text{ and } -(P+Q) \\ & \cup \text{line passing through } P+Q \text{ and } R \text{ and } -((P+Q)+R) \\ & \cup \text{line passing through } \mathcal{O} \text{ and } Q+R \text{ and } -(Q+R). \end{aligned}$$

Finally, let

$$\begin{aligned} c_3 = & \text{line passing through } Q \text{ and } R \text{ and } -(Q+R) \\ & \cup \text{line passing through } P \text{ and } Q+R \\ & \cup \text{line passing through } \mathcal{O} \text{ and } P+Q \text{ and } -(P+Q). \end{aligned}$$

One can prove that c_1 and c_2 have exactly 9 common points (by the condition, it follows that E cannot contain a line, so by Bézout's theorem, the number of intersection points is at most 9), and c_3 passes through 8 of them. Then it passes through $-((P+Q)+R)$ as well, so it has to be equal to $-(P+(Q+R))$.

Of course, there can be many coincidences among the points. Having some topology in hand (e.g. over the complexes), this can be handled by continuity. As soon as we have the theorem in hand for the complex field, we might say that the intersection (or tangent) points are computed from algebraic formulae, and the coincidence giving associativity must be a formal coincidence, which then must hold in all fields.

5.3 The elliptic curve discrete logarithm problem

Letting \mathbf{F}_p the prime field for some $p > 3$, we may consider an elliptic curve defined by the equation $y^2 = x^3 + Ax + B$ over it (with the discriminant condition $4A^3 + 27B^2 \neq 0$).

Let us estimate the number of points on E . There is one point at infinity, so restrict to the solutions of $y^2 = x^3 + Ax + B$ in \mathbf{F}^2 . Now for any x , $x^3 + Ax + B$ is either square or not, and it is a square by probability $\approx 50\%$. This is because

$$\mathbf{F}^\times \ni x \mapsto x^2 \in \mathbf{F}^{\times 2}$$

is a two-folded cover of the quadratic residues (indeed, $x^2 = (-x)^2$, and for any $a \in F$, $x^2 - a$ has at most two solutions). Since $x^3 + Ax + B = 0$ has at most three solutions, we may say that essentially for half of the possible x 's, there are two y 's, so the number of points on E should be around $\#F$.

A deep theorem of Hasse tells us that this is the truth.

Theorem 5.3.1 (Hasse). *We have*

$$|\#\mathbf{F} + 1 - \#E| \leq 2\sqrt{\#\mathbf{F}}.$$

Proof omitted.

As we have seen earlier, from any x , the value $x^3 + Ax + B \pmod p$ can be computed in polynomial time. To get a point on E , we should compute its square-root. Using random algorithms, this can be done fast with high probability, but if $p \equiv 3 \pmod 4$, then we also have a simple deterministic way. We will return to this question later.

Using this, we may find points of E by random methods: pick a random x , compute $x^3 + Ax + B$, compute its $(p+1)/4$ th power, and check if its square is $x^3 + Ax + B$ or not. With high probability, after a few trials, we find a point $P \in E$. Having some points P and Q in hand, it is easy to compute $P + Q$ (the coefficients determining $L_{P,Q}$ come from simple algebraic manipulations, leading to $R = -P - Q$, then $L_{R,O}$ and $P + Q$ come similarly).

Given a point $P \in E$ on an elliptic curve, we may consider

$$nP = \underbrace{P + \dots + P}_{n \text{ many}} \in E.$$

Then nP can be computed in polynomial time with the same trick we applied in the proof of Proposition 1.4.2. Indeed, write

$$n = \sum_{0 \leq j \leq 1 + \lceil \log_2 n \rceil} \varepsilon_j 2^j, \quad \varepsilon_j \in \{0, 1\}.$$

Then $P, 2P = P + P, 4P = 2P + 2P, \dots, 2^{1 + \lceil \log_2 n \rceil} P = 2^{\lceil \log_2 n \rceil} P + 2^{\lceil \log_2 n \rceil} P$ can be computed in polynomial time, so is their weighted sum (with weights ε_j).

The elliptic curve discrete logarithm problem (ECDLP from now on), is the following: given $P, Q \in E$, provided that $Q = nP$ for some $n \in \mathbf{N}$, compute the smallest such n . Since even the DLP for multiplicative groups over prime fields seems to be a computationally difficult problem, and the group operation in an elliptic curve is much more complicated than that of multiplicative groups of prime fields, we expect that the ECDLP is also computationally difficult. So far, this seems to be the truth, there is no known algorithm which would solve the problem in fewer than $O(\sqrt{p})$ steps (and this can be achieved by a meet-in-the-middle attack, using the collision phenomenon).

5.4 Elliptic curve cryptography

5.4.1 The elliptic curve Diffie-Hellman

Alice and Bob agree on a prime p , an elliptic curve E , and a point $P \in E$. They make them public (they can agree on it publicly). Now Alice chooses a number n_A and Bob chooses a number n_B , and they keep them in secret. Then Alice sends $Q_A = n_A P$ to Bob, and Bob sends $Q_B = n_B P$ to Alice. Now Alice computes $n_A Q_B$, and Bob computes $n_B Q_A$, and this will be their secret key. Observe that

$$n_A Q_B = n_A n_B P = n_B n_A P = n_B Q_A,$$

so they have the same point $n_A n_B P \in E$ in hand. However, Eve known only $n_A P$ and $n_B P$, and from this, she should compute somehow $n_A n_B P$, i.e. she should solve the Diffie-Hellman problem over elliptic curves.

5.4.2 The elliptic curve ElGamal

It starts similarly, $p, E, P \in E$ are fixed and public. Now Alice chooses a number n_A and keeps it in secret. Now she computes $Q = n_A P$ and makes it public. If Bob wants to send a message $M \in E$ to Alice, he chooses an ephemeral key $k \in \mathbf{N}$ and computes the following two points:

$$C_1 = kP, \quad C_2 = M + kQ.$$

Then Alice computes $C_2 - n_A C_1$, obtaining

$$C_2 - n_A C_1 = M + kQ - n_A kP = M + kQ - kQ = M.$$

For Eve, to break the cipher in general, should solve a problem not easier than the elliptic curve Diffie-Hellman problem (recall Proposition [2.3.1](#)).

Chapter 6

Attacking the underlying problems

6.1 The discrete logarithm problem

6.1.1 A babystep-giantstep algorithm

Proposition 6.1.1. *Assume g generates the abelian group G of order N . Then the DLP $g^x = h$ can be solved in $O(\sqrt{N}) \text{Pol}(\log N)$ steps.*

Proof. For $n \geq \sqrt{N} + 1$ Make the following two lists:

$$1, g, \dots, g^n, \quad h, hg^{-n}, \dots, hg^{-n^2}.$$

We claim that the two lists have a common element. Ineed, the solution x satisfies $x = nq + r$ for some $0 \leq q, r \leq n$. Then $h = g^{qn+r}$, which implies $h^{-qn} = h^r$.

Therefore, the two lists indeed intersect nontrivially. Assume hence $g^i = hg^{-jn}$ for some $0 \leq i, j \leq n$. Then $x = i + jn$ is a solution to the DLP $g^x = h$. \square

6.1.2 The Pohlig-Hellman algorithm

The following proposition reduces the DLP in a group to the DLP in its Sylow subgroups.

Proposition 6.1.2. *Assume G is a cyclic group of order N , where $N = \prod_{j=1}^r p_j^{\alpha_j}$. Assume we can solve the DLP in $O(S(p^\alpha))$ time for any element, whose order is p^α . Then for G , we can solve the DLP in $O(\sum_{j=1}^r S(p_j^{\alpha_j})) \text{Pol}(\log N)$ time.*

Proof. For any $1 \leq j \leq r$, set $g_j = g^{N/p_j^{\alpha_j}}$, and $h_j = h^{N/p_j^{\alpha_j}}$. Now for any $1 \leq j \leq r$, solve the problem

$$g_j^{y_j} = h_j$$

in $O(S(p_j^{\alpha_j}))$ time.

Then in $O(\log N)$ time, using the Chinese remainder theorem (Corollary B.2.8), we get

$$x \equiv y_j \pmod{p_j^{\alpha_j}} \quad \text{for each } 1 \leq j \leq r.$$

Then, for any $1 \leq j \leq r$, for some $z_j \in \mathbf{Z}$, $x = y_j + p_j^{\alpha_j} z_j$, and then

$$(g^x)^{N/p_j^{\alpha_j}} = \left(g^{y_j + p_j^{\alpha_j} z_j}\right)^{N/p_j^{\alpha_j}} = h_j = h^{N/p_j^{\alpha_j}}.$$

This means that

$$\frac{N}{p_j^{\alpha_j}} x \equiv \frac{N}{p_j^{\alpha_j}} \log_g(h) \pmod{N},$$

which implies

$$x \equiv \log_g(h) \pmod{p_j^{\alpha_j}}.$$

Since this holds for any $1 \leq j \leq r$, $x \equiv \log_g(h)$ modulo N . \square

Further, the prime power groups can be reduced to prime groups.

Proposition 6.1.3. *Assume G is a cyclic group on p^e elements. Assume we can solve the DLP in $O(S(p))$ time for any element, whose order is p . Then in G , we can solve the DLP in $O(eS(p)) \text{Pol}(e \log p)$ time.*

Proof. Assume we have to solve $g^x = h$. We know that for some $0 \leq x_0, x_1, \dots, x_{e-1} \leq p-1$

$$g^{x_0 + x_1 p + \dots + x_{e-1} p^{e-1}} = h.$$

Then raising to power p^{e-1} ,

$$(g^{p-1})^{x_0} = h^{p-1},$$

and we obtain x_0 in $O(S(p))$ steps. Recursively, assume x_0, x_1, \dots, x_{j-1} are already determined. Then

$$g^{x_j p^j + \dots + x_{e-1} p^{e-1}} = h g^{-x_0 - x_1 p - \dots - x_{j-1} p^{j-1}},$$

and raising to power p^{e-1-j} , we obtain

$$(g^{p-1})^{x_j} = \left(h g^{-x_0 - x_1 p - \dots - x_{j-1} p^{j-1}} \right)^{p-1},$$

and x_j is computed in $O(S(p))$ time. □

6.1.3 The index calculus method

In this section, we present the index calculus method, which solves the DLP in the multiplicative group of \mathbf{F}_p .

First of all, fix a parameter B , and solve the discrete logarithm problem for any prime $q \leq B$, i.e. compute $\log_g(q)$. How to do this? For some random numbers $1 \leq i \leq p-1$, compute

$$g_i \equiv g^i \pmod{p}.$$

If g_i has a prime factor bigger than B , then discard it, otherwise factorize it as

$$g_i = \prod_{\substack{q \leq B \\ q \text{ prime}}} q^{u_{q,i}}.$$

Then by Euler-Fermat (Corollary B.3.7),

$$i \equiv \sum_{\substack{q \leq B \\ q \text{ prime}}} u_{q,i} \log_g(q) \pmod{p-1}.$$

Here, if the number of congruences is large enough, then we may be able to solve this congruence system for $\log_g(q)$ (be careful: \mathbf{Z}_{p-1} is not a field).

Now in the DLP $g^x \equiv h \pmod{p}$, compute, for $k = 1, 2, \dots$, the value $h g^{-k} \pmod{p}$ until we arrive at a value which has only prime factors not exceeding B . Then

$$h g^{-k} = \prod_{\substack{q \leq B \\ q \text{ prime}}} q^{e_q},$$

which immediately gives

$$\log_g(h) = k + \sum_{\substack{q \leq B \\ q \text{ prime}}} e_q \log_g(q).$$

Of course, the algorithm depends on the choice of the parameter B . For any $0 \leq \varepsilon \leq 1$, introduce the notation

$$L_\varepsilon(X) = e^{(\log X)^\varepsilon (\log \log X)^{1-\varepsilon}}.$$

It turns out that to find enough congruences, B has to run up to $L_{1/2}(p)^{1/\sqrt{2}}$ which altogether gives rise to a subexponential algorithm (which is still far worse than polynomial, but much better than exponential). With further improvements, the running time can be decreased to $L_{1/3}(p)$ (still subexponential).

6.2 Factorization algorithms

6.2.1 Smooth numbers

Definition 6.2.1 (*B*-smooth numbers). For a fixed number $B > 1$, we say that a positive integer is *B*-smooth, if none of its prime factors exceed B . Given $X, B > 1$, we denote by $\psi(X, B)$ the number of *B*-smooth positive integers not exceeding X .

Example 6.2.2. The 3-smooth numbers up to 20 are: 1, 2, 3, 4, 6, 8, 9, 12, 16, 18. Therefore, $\psi(20, 3) = 10$.

We introduce the notation

$$L(X) = e^{\sqrt{\log X \log \log X}}.$$

How large $L(X)$ is? On the one hand, it grows faster than any power of $\log X$:

$$(\log X)^c = e^{\log \log X^c} = e^{c \log \log X},$$

and clearly $c \log \log X < \sqrt{\log X \log \log X}$ (for any fixed $c > 0$), if X is large enough. On the other hand, it grows slower than any power of X :

$$X^c = e^{\log X^c} = e^{c \log X},$$

and clearly $c \log X > \sqrt{\log X \log \log X}$ (for any fixed $c > 0$), if X is large enough.

Theorem 6.2.3 (Canfield-Erdős-Pomerance). *For any $0 < c < 1$, we have*

$$\psi(X, L(X)^c) = XL(X)^{-\frac{1}{2c}(1+o(1))}.$$

Proof omitted.

6.2.2 Pollard's $p - 1$ method

When attacking RSA, there is a given large number $N = pq$, where we do not know the prime factors p and q .

Starting out from a small number a , say, $a = 2$, consider the following sequence

$$\gcd(a^{1!} - 1, N), \quad \gcd(a^{2!} - 1, N), \quad \gcd(a^{3!} - 1, N), \quad \dots$$

If $n!$ in the exponent is not too large, then these numbers can be computed in a reasonable time. Now assume that for some not too large n , $(p - 1) | n!$. Then, by Euler-Fermat (Corollary B.3.7),

$$a^{n!} = (a^{p-1})^{n!/(p-1)} \equiv 1^{n!/(p-1)} = 1 \pmod{p}.$$

If we are lucky enough that q does not divide $a^{n!} - 1$, then $\gcd(a^{n!} - 1, N)$ will be p , which gives the factorization of N .

How difficult it is to compute $\gcd(a^{n!} - 1, N)$? First, $a^{n!}$ is uncomputably large (with current tools) even for $n = 200$. However, we need it only modulo N , and the modulo N powering can be done fast, think of Proposition 1.4.2. Even better, for an actual calculation, we may get use of the identity

$$a^{(n+1)!} = (a^{n!})^{n+1},$$

so when we are interested in $a^{(n+1)!}$ (modulo N), we only have to raise $a^{n!}$ (modulo N) to power $n + 1$ (modulo N). Altogether, this means that even $n \approx \log N$ is an acceptable number for the number of trials in a polynomial algorithm.

This means that prime numbers p, q are insecure for the aim of RSA, if one of $p - 1$ and $q - 1$ has only small prime factors, i.e. one of them is *B*-smooth for a not too large B .

6.2.3 Factorization via difference of squares

Assume we would like to factorize a large odd integer N . In factorization via difference of squares, we try to find two squares, say, b^2 and a^2 , such that their difference is divisible by N . If none of $b - a$ and $b + a$ is divisible by N itself, then, since

$$kN = b^2 - a^2 = (b - a)(b + a),$$

both $b - a$ and $b + a$ share a common factor with N . Therefore, computing, say, $\gcd(N, b - a)$ (which can be done in polynomial time), we obtain a nontrivial factor of N .

If we find two squares $a^2 \equiv b^2 \pmod{N}$, then we expect that the other condition (i.e. $b \not\equiv \pm a \pmod{N}$) fails with probability $2^{1-\omega(N)}$, where $\omega(N)$ is the number of prime divisors of N (this is not completely obvious, but not too difficult as well). In RSA, when N has two different prime factors, so as soon as we find $b^2 \equiv a^2 \pmod{N}$, it has probability 50% that $\gcd(N, b - a)$ gives a nontrivial prime factor of N . The factorization algorithm consists of three phases.

In the first one, we build relations: take random numbers, until we find a_1, \dots, a_r such that $c \equiv a_j^2 \pmod{N}$ has only small prime factors. Letting p_1, \dots, p_t be the first t primes, we can compute the factorizations

$$\begin{aligned} c_1 &= p_1^{e_{11}} \cdots p_t^{e_{1t}} \\ &\quad \vdots \\ c_r &= p_1^{e_{r1}} \cdots p_t^{e_{rt}}. \end{aligned}$$

Now we arrive at the second step: we find a certain (nonempty) product $c_1^{u_1} \cdots c_r^{u_r}$, where each $u_j \in \{0, 1\}$. We solve the matrix equation

$$UE = 0$$

over the field \mathbf{F}_2 , where

$$U = (u_1 \quad \dots \quad u_r), \quad E = (e_{ij})_{1 \leq i \leq r, 1 \leq j \leq t},$$

and 0 stands for the $1 \times t$ zero matrix. Finding a nontrivial solution (by Gauss elimination over the field \mathbf{F}_2), we have

$$a^2 = c_1^{u_1} \cdots c_r^{u_r} = p_1^{2\alpha_1} \cdots p_t^{2\alpha_t},$$

and also

$$b^2 = a_1^{2u_1} \cdots a_r^{2u_r}.$$

Now obviously $b^2 \equiv a^2 \pmod{N}$, and because of the modulo N reductions (when computing $c_j \equiv a_j^2 \pmod{N}$), we have a good chance that none of $b - a$ and $b + a$ is divisible by N .

The final step is to compute $\gcd(N, b - a)$, which can be done polynomially. If by accident it is N , we take more a_j 's, find other solutions and do the same again.

How long does this procedure take (expectedly)? Letting $p_t = B$, for a singular system of equations, we need essentially $\pi(B)$ (or more precisely, $\pi(B) + 1$) equations. Of course, a random c_j is B -smooth with probability essentially $\psi(N, B)/N$. In order to apply Theorem 6.2.3, set $B = L(N)^c$, and optimize at the end. The number of numbers to be checked is then, expectedly, using also the prime number theorem,

$$\frac{L(N)^c / \log L(N)^c}{L(N)^{-1/2c}} = L(N)^{c+1/2c} \frac{1}{c \log N}.$$

The size of this is essentially determined by the first factor, and the exponent $c + 1/2c$ is minimized at $c = 1/\sqrt{2}$.

That is, we have to check

$$B = e^{\sqrt{\log N \log N/2}}$$

numbers to find $\pi(B)$ B -smooth numbers. This altogether gives a subexponential running time, i.e. a polynomial of $L(N)$.

6.2.4 Lenstra's elliptic curve factorization

Given N to be factorized, consider an elliptic curve $y^2 = x^3 + Ax + B$ modulo N . For the first sight, this seems nonsense, since N is not a prime, so \mathbf{Z}_N is not a field. Nevertheless, when computing the sum of two points, we just add, subtract, multiply and divide, and all these make sense, at least if the number we divide by is coprime to N .

For a point $P \in E$, its multiples

$$P, \quad 2P, \quad 3P, \quad \dots$$

can be computed in most cases. What happens, when we cannot compute nP for some $n \in \mathbf{N}$? It means that during its computation $(n-1)P + P$, we draw the line of slope $(y_{(n-1)P} - y_P)(x_{(n-1)P} - x_P)^{-1}$, and this slope does not make sense, meaning that $x_{(n-1)P} - x_P$ is not coprime to N . Assuming that it is divisible by p , this means that $nP = \mathcal{O}$ when considered over \mathbf{F}_p .

This gives us an algorithm: take a point P on E , and compute

$$P, \quad 2!P, \quad 3!P, \quad \dots$$

When $n!$ gets divisible by the order of P modulo p , the resulting point is \mathcal{O} over \mathbf{F}_p . If we cannot perform the calculation, it means that we tried to divide by a number (the difference of two x -coordinates) not coprime to N , that is, we have some $d > 1$ such that $d \mid N$. Computing $\gcd(d, N)$, we either find a proper divisor of N , or we find N itself. In the former case, we are done, in the latter one, we pick a new curve and a new point.

Again, an actual computation can be made faster by utilizing

$$(n+1)!P = (n+1) \cdot n!P.$$

There is a technical subtlety in the choice of the curve and the point on it. Earlier, we agreed that given a curve, a random trial gives a good x -coordinate by 50% chance, and then computing the square root is easy. However, this is not true now, since we do not know the factorization of N : if we compute $x^3 + Ax + B$, we cannot tell if it a square. Instead, choose the point P first, and adjust the elliptic curve: pick a random A , then set B such that $y^2 = x^3 + Ax + B$ holds for P .

Chapter 7

Additional topics

In this chapter, our main source is [2].

7.1 Interactive proofs

7.1.1 How to store the last move in chess?

Assume now Alice and Bob play chess via phone (they tell each other their moves using the chessboard coordinates). After some time, they decide to adjourn the game. In tournament chess, the player on move writes her/his move to a piece of paper, and gives it to the arbiter in a closed envelope. The next day the arbiter opens the envelope and makes the assigned move on the board, then the game continues. The goal is to deprive both players from clear advantage over each other¹.

Of course, Alice and Bob could go for a Trusted Authority and use the same scheme. Nevertheless, they can solve this problem on their own. Alice generates two large primes p, q such that the first few, say, ten digits of p (which is the smaller prime by their agreement) defines her move in a previously fixed way. Then she sends the product $N = pq$ to Bob. Bob is unable to read out Alice's move, but the next day, when they continue their game, Alice tells the factorization. Of course, Bob checks it, so that Alice cannot alter her move.

This way it is possible to create "electronic envelopes": Alice can put a certain piece of information in a deposit, which cannot be altered even by herself, but cannot be read by anyone else until she lets them.

7.1.2 A zero-knowledge proof of that a certain number is square modulo N

First assume p and q are large prime numbers, which Peggy (the prover) keeps in secret. She publishes their product $N = pq$, and also claims that a certain residue class $y \bmod N$ is a square modulo N , i.e. there exists a residue class $x \bmod N$ satisfying $x^2 \equiv y \bmod N$. Her goal is to prove this to Victor (the verifier) without telling him an actual square root. Of course, she could in principle tell the prime factors, but doing so would actually tell the square roots.

Proposition 7.1.1. *Assume p is an odd prime number, and $\gcd(a, p) = 1$. Then*

$$a^{\frac{p-1}{2}} \equiv \begin{cases} 1 \bmod p, & \text{if } a \text{ is a square modulo } p, \\ -1 \bmod p, & \text{if } a \text{ is not a square modulo } p. \end{cases}$$

Proof. Assume first that $a \equiv b^2 \bmod p$ for some residue class b modulo p . By Euler-Fermat (Corollary B.3.7), we have

$$a^{\frac{p-1}{2}} \equiv (b^2)^{\frac{p-1}{2}} \equiv b^{p-1} \equiv 1.$$

Assume then a is not a square modulo p . Again, by Euler-Fermat (Corollary B.3.7), we certainly have

$$\left(a^{\frac{p-1}{2}}\right)^2 \equiv 1 \bmod p,$$

¹In competitive chess, there are essentially no adjourned games any more, because of the fast development of engines in the last few decades.

and we see that $a^{(p-1)/2}$ is a root of the polynomial $x^2 - 1$ over \mathbf{F}_p , hence it must be ± 1 . Therefore, it suffices to exclude $a^{(p-1)/2} \equiv 1 \pmod p$ in this case. Observe that $x \mapsto x^2$ is a two-folded cover from \mathbf{F}_p^\times to the nonzero squares: indeed, if $x_1^2 \equiv x_2^2 \pmod p$, then

$$0 \equiv x_1^2 - x_2^2 \equiv (x_1 - x_2)(x_1 + x_2) \pmod p, \quad x_1 \equiv \pm x_2 \pmod p.$$

Then the number of nonzero square residue classes is exactly $(p-1)/2$. Then consider the polynomial $x^{(p-1)/2} - 1$ over \mathbf{F}_p . It has at most $(p-1)/2$ roots, since the number of roots cannot exceed the degree (over a field), and recall that we have already found that many roots: namely, the nonzero squares. This means that nonsquare residue classes are not roots, so $a^{(p-1)/2}$ cannot be 1 modulo p . \square

Proposition 7.1.2. *If p is a prime such that $p \equiv 3 \pmod 4$, and $a \in \mathbf{N}$ is coprime to p such that it is a square modulo p , then $a^{(p+1)/4}$ is a square root of a .*

Proof. We have the following simple calculation (using Proposition 7.1.1):

$$\left(a^{\frac{p+1}{4}}\right)^2 \equiv a^{\frac{p+1}{2}} \equiv a \cdot a^{\frac{p-1}{2}} \equiv a \pmod p,$$

and the proof is complete. \square

This already tells us how to compute the square root modulo a prime number congruent to 3 modulo 4. When a prime number is 1 modulo 4, we may use Cipolla's algorithm, which relies on the following proposition.

Proposition 7.1.3. *Assume p is an odd prime, and a is coprime to p such that it is a square modulo p . Choose $b \in \mathbf{F}_p^\times$ such that $b^2 - a$ is not a square in \mathbf{F}_p . Then*

$$\left(b + \sqrt{b^2 - a}\right)^{\frac{p+1}{2}}$$

is a square root of a , where the element $b + \sqrt{b^2 - a}$ is understood in the field \mathbf{F}_{p^2} .

Proof. Set $\omega = \sqrt{b^2 - a}$. We will calculate in the field $\mathbf{F}_{p^2} = \mathbf{F}_p(\omega)$. Then the Frobenius automorphism $\sigma : x \mapsto x^p$ maps ω to $-\omega$: indeed, $\sigma(\omega)^2 = \sigma(\omega^2) = \sigma(b^2 - a) = b^2 - a$ by Corollary B.3.8, therefore, $\sigma(\omega)^2 = \pm \omega$ (its square is fixed under σ); and the polynomial $x^p - x$ has at most p roots, so σ cannot fix anything outside \mathbf{F}_p . Then

$$\left((b + \omega)^{\frac{p+1}{2}}\right)^2 = (b + \omega)^{p+1} = (b + \omega)^p(b + \omega) = (b - \omega)(b + \omega) = b^2 - \omega^2 = b^2 - (b^2 - a) = a,$$

and the proof is complete. \square

So when $p \equiv 1 \pmod 4$, we may compute the square root as follows: we pick random b 's, apply Proposition 7.1.1 to decide if $b^2 - a$ is a square or not modulo p , until we arrive at a nonsquare (it is likely that this happens fast: note that we win by probability 1/2 in each trial). This nonsquare (or rather, its square root in the bigger field) gives us a square root of a modulo p .

All in all, if Victor learns p and q , he can easily compute a square root x of $y \equiv x^2 \pmod N$.

To solve this problem, Peggy generates another square, $s \equiv r^2 \pmod N$, and sends s to Victor. Now Victor chooses a random bit $\beta \in \{0, 1\}$ and sends it to Peggy. If $\beta = 0$, Peggy sends r to Victor, who checks $s \equiv r^2 \pmod N$. If $\beta = 1$, Peggy sends xr to Victor, who checks $ys \equiv (xr)^2 \pmod N$. Observe that in none of the cases Peggy tells too much information on x to Victor: indeed, when $\beta = 0$, then r has nothing to do with y ; while if $\beta = 1$, the original square root x is perturbed by a random square root. This means basically that Peggy does not inform Victor about x . However, Victor still learns something, at least with some probability: if one of y and s is not a square, then Peggy fails his check with probability 1/2. This is not very much convincing, but Victor can tell Peggy in advance: "I will challenge you one hundred times. In each challenge, give me another square s , and I will ask you to send me the square root of either s or ys , chosen at random. If you fail any of the challenges, I will refuse to believe that your number y is a square. If you pass all the challenges, I will believe that your number y is indeed a square, since the probability $1 - 2^{-100}$ is more than enough for me."

7.1.3 Using our password

When we withdraw cash from an ATM, we use our password, and the computer system of the bank checks it. There is a weak point in this procedure: if a hacker gets access to the files of the bank, (s)he can learn our password.

Using the complexity of factorization, we may find a solution to this weakness. Assume our password is a large prime p . We also generate another large prime number q , and tell the bank their product $N = pq$ (and, as usual, we keep p, q in secret). The computer system of the bank stores only N , and when withdrawing cash, we tell p to the ATM, and the system checks if p divides N , but does not store p .

Now even if the hacker manages to get access to the files of the bank, and learns our number N , (s)he is still unable to use our password p without factorizing N .

What happens if we cannot trust that the bank indeed does not store our prime number p ? Then after the first withdrawal, p is available for the bank – and also for the skillful hacker.

A version of the RSA can solve this problem, namely, instead of telling the bank the prime divisor p , we only prove to the bank that we know the prime factorization of N . To this aim, we tell the bank not only N , but also a number e , and at cash withdrawals, the following protocol can be applied. The ATM generates a random number $1 \leq x \leq N$, and tells us $y \equiv x^e \pmod{N}$. We can easily tell what the original x was: letting d be the multiplicative inverse of e (which is easily computable for us), $y^d \equiv x \pmod{N}$ (recall Section 3.1). Now the ATM checks if we really recovered x , and if yes, it accepts that we know the factorization of N .

7.2 Identification

There are certain situations when it is important for Alice to be able to check whether the incoming message have been really encrypted Bob. In this section we present a few situations where such issues are addressed.

7.2.1 An RSA-based digital signature

In the following setup, we turn the RSA upside down. Now Bob chooses two large primes p, q and exponents d, e such that $de \equiv 1 \pmod{\varphi(N)}$, where $N = pq$. Then publishes N and e and keeps the remaining numbers in secret. If his message is $1 \leq m \leq N$, a number coprime to N , he computes $m^d \pmod{N}$, and sends it to Alice. Then Alice raises the incoming message to power e , obtaining

$$(m^d)^e \equiv m \pmod{N},$$

recall Section 3.1. For a random number, its e th power is most likely a gibberish, so if Alice gets a meaningful message, she learns that it must have been sent by Bob indeed. This way Bob's m^d is not only a message, but it also serves as a personal identifier of him, an example of a digital signature.

7.2.2 Multiplayer RSA – the good way

In this section, we assume that there are n participants, each of them using a public key cryptosystem (let us index them by $1 \leq i \leq n$). To this aim, each of them chooses two large prime numbers p_i, q_i , and exponents d_i, e_i such that $d_i e_i \equiv 1 \pmod{\varphi(N_i)}$, where $N_i = p_i q_i$. Now each of them publishes N_i, e_i .

Since the moduli are different, this time they cannot use residue classes, so the possible messages are positive integers. Now assume the i th participant wants to send the message $m \in \mathbf{N}$ to the j th participant. First of all, (s)he writes m in base N_i :

$$m = \sum_{k=0}^t m(i, k) N_i^k.$$

Now (s)he applies her/his own decryption function d_i to each digit $m(i, k)$. They altogether give rise to a certain number, say,

$$D_i(m) = \sum_{k=0}^t (m(i, k)^{d_i} \pmod{N_i}) N_i^k.$$

Now (s)he rewrites this number $D_i(m)$ in base N_j :

$$D_i(m) = \sum_{k=0}^t D_i(m)(j, k) N_j^k,$$

and applies the addressee's encryption function e_j to each digit $D_i(m)(j, k)$, giving, say

$$E_j(D_i(m)) = \sum_{k=0}^t ((D_i(m)(j, k))^{e_j} \bmod N_j) N_j^k.$$

Then the (s)he sends the number $E_j(D_i(m))$ to the addressee.

Now the recipient applies her/his own decryption function d_j to each digit, this gives back $D_i(m)(j, k)$, as usual:

$$((D_i(m)(j, k))^{e_j})^{d_j} \equiv D_i(m)(j, k) \bmod N_j,$$

at least when $\gcd(D_i(m)(j, k), N_j) = 1$, but this holds with high probability (at least when t is much smaller than all the p_i, q_i 's, which we assume from now on). Now these $D_i(m)(j, k)$'s give rise to $D_i(m)$:

$$D_i(m) = \sum_{k=0}^t D_i(m)(j, k) N_j^k.$$

Finally, the recipient rewrites $D_i(m)$ in base N_i ,

$$D_i(m) = \sum_{k=0}^t (m(i, k)^{d_i} \bmod N_i) N_i^k.$$

then applies sender's encryption function to each digit, giving back, as usual:

$$(m(i, k)^{d_i})^{e_i} \equiv m(i, k) \bmod N_i,$$

at least when $\gcd(m(i, k), N_i) = 1$, but this holds with high probability. Then

$$m = \sum_{k=0}^t m(i, k) N_i^k,$$

so the reader obtains the original message. But this is not the only outcome of the method. Simultaneously, the j th participant can be almost certain that the sender was the i th participant, so the method provides digital signatures as a by-product.

7.3 Secret sharing

The basic problem of secret sharing is the following: given some information, share certain pieces of it among n people such that when they get together, they can read out the information, but no subgroup of $n - 1$ people can do so.

A simple solution is the following. Let the secret be a residue class S modulo m . Choose then $n - 1$ random values D_1, \dots, D_{n-1} modulo m , and set

$$D_n \equiv S - D_1 - \dots - D_{n-1} \bmod m.$$

Now the i th participant receives the value D_i for $1 \leq i \leq n$. It is clear that together they can reveal S , since

$$S \equiv D_1 + \dots + D_n \bmod m.$$

However, if $n - 1$ of them gets together, and share their knowledge, they still cannot have a good guess on S , which can be any residue class modulo m .

A more general problem is the following: there are n participants and the goal is to give each of them some piece of the given secret S such that any t of them can reveal the whole S , but no $t - 1$ of them can do so. The following idea is due to Shamir.

Let the secret S be a number in any fixed field \mathbf{F} . Set $a_0 = S$, and choose random numbers a_1, \dots, a_{t-1} . Construct the polynomial

$$f(x) = a_0 + a_1x + \dots + a_{t-1}x^{t-1} \in \mathbf{F}[x].$$

We see that $f(0) = S$. Take random nonzero elements $x_1, \dots, x_n \in \mathbf{F}^\times$ and let the i th participant receive the value $y_i = f(x_i)$.

Proposition 7.3.1. *Let \mathbf{F} be a field. Assume that $a_1, \dots, a_m, b_1, \dots, b_m \in \mathbf{F}$ such that the a_i 's are distinct. Then there exists a unique polynomial $f \in \mathbf{F}[x]$ such that it is either the zero polynomial or its degree is at most $m - 1$ which satisfies $f(a_i) = b_i$ for any $1 \leq i \leq m$.*

Proof of uniqueness. Assume $f(x)$ and $g(x)$ are two such polynomials. Then consider the polynomial $h(x) = f(x) - g(x)$. Then either $h(x) = 0$ or $\deg h(x) \leq m - 1$. Note on the other hand that a_1, \dots, a_m are all roots of $h(x)$, and the number of roots cannot exceed the degree of a polynomial over a field. Then the possibility $\deg h(x) \leq m - 1$ is excluded, so $h(x) = 0$, and then $f(x) = g(x)$. \square

Proof of existence via Lagrange's interpolation. Consider the polynomial

$$f(x) = \sum_{i=1}^m b_i \prod_{\substack{1 \leq j \leq m \\ j \neq i}} \frac{x - a_j}{a_i - a_j}.$$

It is clear that $f(x) = 0$ or $\deg f(x) \leq m - 1$. Also, for a single monomial,

$$\prod_{\substack{1 \leq j \leq m \\ j \neq i}} \frac{x - a_j}{a_i - a_j} \equiv \begin{cases} 1, & \text{if } x = a_i, \\ 0, & \text{if } x = a_{i'}, \text{ for some } 1 \leq i' \leq m \text{ with } i' \neq i. \end{cases}$$

From this, the statement is obvious. \square

Proof of existence via Newton's interpolation. We induct on $m \in \mathbf{N}$. The statement for $m = 1$ is obvious, since the polynomial can be chosen to be the constant b_1 . Now assume that $g(x)$ is an appropriate polynomial for the first $m - 1$ pairs in the input: $g(x) = 0$ or $\deg g(x) \leq m - 2$, and $g(a_i) = b_i$ for $1 \leq i \leq m - 1$. Then set

$$f(x) = g(x) + (b_m - g(a_m)) \prod_{j=1}^{m-1} \frac{x - a_j}{a_m - a_j}.$$

Clearly $f(x)$ satisfies the degree conditions, $f(a_i) = g(a_i) = b_i$ for $1 \leq i \leq m - 1$, while at a_m , we have

$$f(a_m) = g(a_m) + (b_m - g(a_m)) \prod_{j=1}^{m-1} \frac{a_m - a_j}{a_m - a_j} = g(a_m) + b_m - g(a_m) = b_m,$$

and the proof is complete. \square

Although both proofs provide an algorithm, Newton's one is faster in practice (although it is not as explicit as Lagrange's one).

With this tool in hand, it is clear that t people can figure out the shared secret: they simply interpolate the value $f(0) = S$. However, given only $t - 1$ pieces of the secret, the polynomial can still take any value of \mathbf{F} at 0.

Appendices

Chapter A

Algebraic structures

A.1 Groups

Definition A.1.1 (group). Assume G is a set, and there is a binary operation on it, e.g. a function \cdot from $G \times G$ to G , which satisfies the following conditions (for a better notation, we write $x \cdot y$ in place of $\cdot(x, y)$):

- for any $x, y, z \in G$, $(x \cdot y) \cdot z = x \cdot (y \cdot z)$ (associativity);
- there exists $e \in G$ satisfying that for any $x \in G$, $x \cdot e = e \cdot x = x$ (unit element);
- for any $x \in G$, there exists $y \in G$ satisfying $x \cdot y = y \cdot x = e$ (inverse).

Then we say that (G, \cdot) is a group (or G is a group, if the operation is clear).

Notational convention: the inverse of an element x is often denoted by x^{-1} .
In most applications below, our groups will be abelian.

Definition A.1.2 (abelian (commutative) group). Assume (G, \cdot) is a group such that for any $x, y \in G$,

$$x \cdot y = y \cdot x.$$

Then we say that G is an abelian (commutative) group.

Example A.1.3. The structures $(\mathbf{Z}, +)$, $(\mathbf{Q}, +)$, $(\mathbf{R}, +)$, $(\mathbf{C}, +)$, $(\mathbf{Q}^\times, \cdot)$, $(\mathbf{R}^\times, \cdot)$, $(\mathbf{C}^\times, \cdot)$, (\mathbf{Q}_+, \cdot) , (\mathbf{R}_+, \cdot) are all abelian groups.

Example A.1.4. Given any set S , denote by $\text{Perm}(S)$ the set of its permutations. Two permutations $\pi, \sigma \in \text{Perm}(S)$ can be concatenated by the composition

$$\pi \circ \sigma \in \text{Perm}(S) : \quad \forall x \in S : (\pi \circ \sigma)(x) = \pi(\sigma(x)).$$

Then $(\text{Perm}(S), \circ)$ is a group: composing functions is associative, the identical permutation is the unit element, and the inverse permutation of each permutation is its inverse. Note that this group is not abelian if $\#S \geq 3$.

Proposition A.1.5. Assume G is a group. Then its unit element is unique. Also, the inverse of each element of G is unique.

Proof. Assume e, f are unit elements. Then

$$\begin{aligned} e &= ef \text{ (since } f \text{ is a unit element)} \\ &= f \text{ (since } e \text{ is a unit element),} \end{aligned}$$

and the proof of the uniqueness of the unit element is complete.

Similarly, let $x \in G$, and assume y, z are both inverses of x . Then (denoting by e the unit element)

$$\begin{aligned} y &= ye \text{ (since } e \text{ is the unit element)} \\ &= y(xz) \text{ (since } z \text{ is the inverse of } x\text{)} \\ &= (yx)z \text{ (by associativity)} \\ &= ez \text{ (since } y \text{ is the inverse of } x\text{)} \\ &= z \text{ (since } e \text{ is the unit element),} \end{aligned}$$

and the proof of the uniqueness of the inverse is complete. \square

Proposition A.1.6. *Assume G is a group, and $a, b \in G$ are fixed elements. Then the equation $ax = b$ has a unique solution. The same holds for the equation $ya = b$.*

Proof. In the first equation, observe that $x = a^{-1}b$ is a solution: indeed, $a(a^{-1}b) = (aa^{-1})b = b$, where we used associativity and that aa^{-1} is the unit element. Also, if x is any solution, then multiplying both sides of the equation by a^{-1} on the left, we obtain $x = a^{-1}b$.

The proof is similar for the equation $ya = b$: $y = ba^{-1}$ is a solution, and it must be the only one, which can be seen by multiplying the both sides on the right by a^{-1} . \square

Proposition A.1.7. *Assume G is a finite group, and $x \in G$. Then for any $n \in \mathbf{N}$, $(x^n)^{-1} = (x^{-1})^n$.*

Proof. Assume the unit element is e . By associativity,

$$x^n(x^{-1})^n = \underbrace{x \cdot \dots \cdot x}_{n \text{ times}} \cdot \underbrace{x^{-1} \cdot \dots \cdot x^{-1}}_{n \text{ times}} = e,$$

and the proof is complete. \square

Therefore, the conventions $x^0 = e$ (the unit element) and $x^{-n} = (x^n)^{-1}$ for $n \in \mathbf{N}$ are completely consistent, so the notation x^n makes sense for $n \in \mathbf{Z}$.

Proposition A.1.8. *Assume G is a finite group with unit element e . Then for any element $x \in G$, there is a positive integer $n \in \mathbf{N}$ such that*

$$x^n = \underbrace{x \cdot \dots \cdot x}_{n \text{ times}} = e.$$

Proof. Consider the elements $x = x^1, x^2, x^3, \dots \in G$. Since G is finite, there exist positive integers $k < l$ such that $x^k = x^l$. Multiplying both sides by $(x^k)^{-1} = x^{-k}$, we obtain $e = x^{l-k}$, and here, $l-k \in \mathbf{N}$. \square

Definition A.1.9 (order of element). Assume G is a finite group with unit element e . For any element $x \in G$, we define its order ($\mathbf{o}(x)$ in notation) the smallest positive integer n satisfying $x^n = e$.

Proposition A.1.10. *Assume G is a finite group with unit element e , and $x \in G$. Then for any $k \in \mathbf{N}$, $x^k = e$ if and only if k is a multiple of $\mathbf{o}(x)$.*

Proof. First assume $\mathbf{o}(x) \mid k$, i.e. $k = l\mathbf{o}(x)$, where $l \in \mathbf{N}$. Then

$$x^k = (x^{\mathbf{o}(x)})^l = e^l = e.$$

Conversely, assume $\mathbf{o}(x) \nmid k$. Then by Proposition B.1.1, for some $c \in \mathbf{N} \cup \{0\}$ and $0 < d < \mathbf{o}(x)$, $k = c\mathbf{o}(x) + d$. Then

$$x^d = x^{k - c\mathbf{o}(x)} = x^k x^{-c\mathbf{o}(x)} = ee^{-1} = e,$$

which contradicts the minimal choice of $\mathbf{o}(x)$. \square

Definition A.1.11 (subgroup). Assume (G, \cdot) is a group. If $H \subseteq G$, and (H, \cdot) is a group, then we say that H is a subgroup of G . In notation: $H \leq G$.

Example A.1.12. We have $(\mathbf{Z}, +) \leq (\mathbf{Q}, +) \leq (\mathbf{R}, +) \leq (\mathbf{C}, +)$. Also $(\mathbf{Q}_+, \cdot) \leq (\mathbf{Q}^\times, \cdot)$.

Proposition A.1.13. *Assume G is a group, and H_i ($i \in I$) are subgroups of G . Then $H = \bigcap_{i \in I} H_i$ is also a subgroup of G .*

Proof. Associativity just follows from the associativity of G . We have to check that H is closed under multiplication and taking inverses. First, if $x, y \in H$, then for each $i \in I$, $x, y \in H_i$, and then $xy \in H_i$, since H_i is a subgroup. Therefore, $xy \in H_i$ for each $i \in I$, implying $xy \in H$. The proof for inverses goes similarly: if $x \in H$, then $x \in H_i$ for each $i \in I$, and since H_i is a subgroup, $x^{-1} \in H_i$, giving $x \in H$. \square

Definition A.1.14 (generated subgroup). Assume G is a group, and $A \subseteq G$. Then the subgroup generated by A is

$$\langle A \rangle = \bigcap_{A \subseteq H \leq G} H.$$

The intersection is nonempty, since $A \subseteq G$, and it is a subgroup, since it is an intersection of subgroups.

Proposition A.1.15. For a group G and $A \subseteq G$,

$$\langle A \rangle = \{a_1 \cdots a_n : \text{for each } 1 \leq i \leq n, a_i \in A \text{ or } a_i^{-1} \in A\}.$$

Proof. The right-hand side is clearly a subgroup: it is closed under multiplication and taking inverses, also it is containing A , so it is taken into account in the definition of $\langle A \rangle$. Also, any $H \leq G$ containing A contains the elements a_1, \dots, a_n and hence all the products on the right-hand side. \square

Corollary A.1.16. For a group G and $x \in G$, $\langle x \rangle = \langle \{x\} \rangle = \{x^n : n \in \mathbf{Z}\}$. \square

Theorem A.1.17 (Lagrange). Assume G is a finite group and H is a subgroup of G . Then $\#H \mid \#G$.

Proof. Introduce the following relation on the pair of elements of G : $x \sim y$, if for some $h \in H$, $xh = y$. This is an equivalence relation: $1 \in H$ implies $x \sim x$; if $xh \sim y$ for some $h \in H$, then $yh^{-1} = x$ and $h^{-1} \in H$; if $x \sim y \sim z$, then for some $h_1, h_2 \in H$, $xh_1 = y$, $yh_2 = z$, then $x(h_1h_2) = z$ and $h_1h_2 \in H$, yielding $x \sim z$. Then G is partitioned into equivalence classes, and we claim that each equivalence class has the same number of elements as H (this clearly implies the statement). Take any equivalence class C , let x be a representative of it. Then take the function $f(h) = xh$ for $h \in H$, obviously $f(H) \subseteq C$. Also, for any $y \in C$, there is h satisfying $xh = y$, then $f(h) = y$, thus $f(H) = C$. Then f surjects H onto C , it suffices to see that it also injects H into C . Assume that for h, h' , $f(h) = f(h')$. Then $xh = xh'$, so multiplying by x^{-1} on the left, $h = h'$. \square

Corollary A.1.18. For a finite group G and $x \in G$, $\mathfrak{o}(x) \mid \#G$.

Proof. The elements of $\langle x \rangle$ are exactly the elements $x = x^1, x^2, x^3, \dots, x^{\mathfrak{o}(x)}$ (and these elements are distinct: if any two of them coincide, then their quotient is the unit element e , contradicting the minimal choice of the order of x). \square

A.2 Rings and fields

Definition A.2.1 (ring). Assume R is a set with two binary operations $+, \cdot$ such that

- $(R, +)$ is an abelian group;
- for any $x, y, z \in R$, $(x \cdot y) \cdot z = x \cdot (y \cdot z)$ (associativity of \cdot);
- for any $x, y, z \in R$, $x \cdot (y + z) = x \cdot y + x \cdot z$ and $(x + y) \cdot z = x \cdot z + y \cdot z$ (distributivity).

Then we say that $(R, +, \cdot)$ is a ring.

Terminological and notational conventions: the unit element of $(R, +)$ is often called the zero element of the ring, and they will often be denoted by 0.

Our rings $(R, +, \cdot)$ will exclusively be commutative, and will have unit elements for the operation \cdot (not only for $+$).

Definition A.2.2 (commutative ring with unit element). We say that $(R, +, \cdot)$ is a commutative ring with unit element, if

- for any $x, y \in R$, $x \cdot y = y \cdot x$ (commutativity);

- there exists $1 \in R$ such that $1 \neq 0$, and for any $x \in R$, $x \cdot 1 = 1 \cdot x = x$ (unit element for \cdot).

Definition A.2.3. Assume $(\mathbf{F}, +, \cdot)$ is a commutative ring with unit element 1, where $(\mathbf{F} \setminus \{0\}, \cdot)$ is an abelian group. Then we say that $(\mathbf{F}, +, \cdot)$ is a field.

Terminological and notational conventions: $(\mathbf{F}, +)$ is called the additive group, $(\mathbf{F}^\times, \cdot)$ is called the multiplicative group, where \mathbf{F}^\times stands for $\mathbf{F} \setminus \{0\}$.

Example A.2.4. The following structures are all commutative rings with unit element: $(\mathbf{Z}, +, \cdot)$, $(\mathbf{Q}, +, \cdot)$, $(\mathbf{R}, +, \cdot)$, $(\mathbf{C}, +, \cdot)$. The following structures are all fields: $(\mathbf{Q}, +, \cdot)$, $(\mathbf{R}, +, \cdot)$, $(\mathbf{C}, +, \cdot)$. Given any field \mathbf{F} and $n \in \mathbf{N}$, the $n \times n$ matrices over \mathbf{F} form the ring $\text{Mat}(n, \mathbf{F})$, which has a unit element but is not commutative, if $n \geq 2$.

A.3 Vector spaces and projective spaces

Definition A.3.1 (vector space). Let \mathbf{F} be a field. Assume V is an abelian group such that \mathbf{F} acts on V via multiplication (i.e. for any $\lambda \in \mathbf{F}, v \in V$, there is an element $\lambda v \in V$). If this action satisfies the properties

- for any $\lambda, \mu \in \mathbf{F}, v \in V$, $(\lambda + \mu)v = \lambda v + \mu v$;
- for any $\lambda \in \mathbf{F}, u, v \in V$, $\lambda(u + v) = \lambda u + \lambda v$;
- for any $\lambda, \mu \in \mathbf{F}, v \in V$, $(\lambda\mu)v = \lambda(\mu v)$;
- for any $v \in V$, $1v = v$ (where 1 stands for the unit element of \mathbf{F}),

then we say that V is a vector space over \mathbf{F} .

Remark A.3.2. The elements of V are called vectors, while those of \mathbf{F} are called scalars.

Proposition A.3.3 (d -dimensional vector space over \mathbf{F}). Let \mathbf{F} be a field, and $d \in \mathbf{N}$. Then, as a set:

$$\mathbf{F}^d = \left\{ \begin{pmatrix} v_1 \\ \dots \\ v_d \end{pmatrix} : v_1, \dots, v_d \in \mathbf{F} \right\}.$$

On \mathbf{F}^d , we define the coordinatewise addition:

$$\begin{pmatrix} u_1 \\ \dots \\ u_d \end{pmatrix} + \begin{pmatrix} v_1 \\ \dots \\ v_d \end{pmatrix} = \begin{pmatrix} u_1 + v_1 \\ \dots \\ u_d + v_d \end{pmatrix}.$$

Finally, \mathbf{F} acts on \mathbf{F}^d via multiplication in coordinates:

$$\lambda \begin{pmatrix} v_1 \\ \dots \\ v_d \end{pmatrix} = \begin{pmatrix} \lambda v_1 \\ \dots \\ \lambda v_d \end{pmatrix} \text{ for } \lambda \in \mathbf{F}.$$

Then \mathbf{F}^d is a vector space (often called the d -dimensional vector space) over \mathbf{F} .

Proof. All are straight-forward calculations. □

Remark A.3.4. Often \mathbf{F}^0 is meant to be the trivial group with the trivial \mathbf{F} -action.

Definition A.3.5 (linear function). Assume U, V are vector spaces over \mathbf{F} . We say that $\phi : U \rightarrow V$ is a linear function, if

- for any $u_1, u_2 \in U$, $\phi(u_1 + u_2) = \phi(u_1) + \phi(u_2)$;
- for any $\lambda \in \mathbf{F}, u \in U$, $\phi(\lambda u) = \lambda\phi(u)$.

Definition A.3.6 (endomorphism, automorphism of vector spaces). Assume V is a vector space. If $\phi : V \rightarrow V$ is a linear function, then we say it is an endomorphism (or a linear transformation) of V . If ϕ is a bijective endomorphism of V , then we say it is an automorphism of V .

Proposition A.3.7. *The endomorphisms of \mathbf{F}^d are the elements of $\text{Mat}(d, \mathbf{F})$, when matrices act on vectors via left matrix-multiplication. The automorphisms of \mathbf{F}^d are the elements of $\text{GL}(d, \mathbf{F})$, when matrices act on vectors via left matrix-multiplication.*

Proof. It is a straight-forward calculation that matrices act linearly. Also, bijectivity is equivalent to the invertibility of the corresponding matrix, which is nothing else but being in $\text{GL}(d, \mathbf{F})$. Therefore, it suffices to see that every endomorphism can be written in matrix form. Given a linear transformation ϕ of \mathbf{F}^d , introduce

$$v_1 = \phi \left(\begin{pmatrix} 1 \\ 0 \\ \dots \\ 0 \end{pmatrix} \right), \dots, v_d = \phi \left(\begin{pmatrix} 0 \\ \dots \\ 0 \\ 1 \end{pmatrix} \right).$$

Now one can verify via calculation that the matrix whose i th column is v_i for $1 \leq i \leq d$ stands for ϕ . \square

Remark A.3.8. An alternative – although less widespread – notation is to use row vectors, i.e. $\mathbf{F}^d = \{(v_1 \ \dots \ v_d) : v_1, \dots, v_d \in \mathbf{F}^d\}$. In this notation, matrices (corresponding to endomorphisms) act on the right.

Definition A.3.9 (projective space). Assume V is a vector space over the field \mathbf{F} . Then the corresponding projective space $\mathbf{P}V$ is defined as follows. Consider the set of nonzero vectors, and identify two of them, if one is a scalar multiple of the other. The elements of $\mathbf{P}V$ are the elements obtained after this identification. In more mathematical terms:

$$\mathbf{P}V = V \setminus \{0\} / \sim, \quad \text{where } u \sim v, \text{ if } u = \lambda v \text{ for some } \lambda \in \mathbf{F}^\times.$$

Definition A.3.10 (d -dimensional projective space). The d -dimensional projective space over the field \mathbf{F} is the above, when the vector space is \mathbf{F}^{d+1} . We usually denote it by $\mathbf{P}^d(\mathbf{F})$.

Example A.3.11 (projective plane and its affine coverings). Let \mathbf{F} be any field, and consider the projective plane $\mathbf{P}^2(\mathbf{F})$. Any element of $\mathbf{P}^2(\mathbf{F})$ can be represented by a vector $(X, Y, Z) \in \mathbf{F}^3$, where not all of X, Y, Z are zero. However, this representation is not unique (unless \mathbf{F} is the exceptional 2-element field): we can multiply each of X, Y, Z by the same $\lambda \in \mathbf{F}^\times$. The resulting vector will be $(\lambda X, \lambda Y, \lambda Z) \in \mathbf{F}^3$, it is admissible in the sense that not all of $\lambda X, \lambda Y, \lambda Z$ equals zero, and it stands for the same element of $\mathbf{P}^2(\mathbf{F})$ as (X, Y, Z) by definition. In notation, we write this element $[X, Y, Z]$ (or equivalently, $[\lambda X, \lambda Y, \lambda Z]$), i.e. we change $(,)$ around the coordinates to $[,]$ in order to indicate that we are speaking about the equivalence class, not the vector.

The space $\mathbf{P}^2(\mathbf{F})$ can be covered with three *affine* planes, i.e. three copies of \mathbf{F}^2 , which underlines the fact that $\mathbf{P}^2(\mathbf{F})$ is indeed a two-dimensional structure (which explains the 2 in the exponent, and the name 'plane'). This goes as follows. Consider the affine xy -plane, i.e.

$$\mathbf{F}_{xy}^2 = \{(x, y) : x, y \in \mathbf{F}^2\},$$

this is nothing else but the two-dimensional vector space over \mathbf{F} (we only introduce x, y for notational convenience). Now consider the mapping

$$\mathbf{F}_{xy}^2 \ni (x, y) \mapsto [x, y, 1] \in \mathbf{P}^2(\mathbf{F}).$$

We see that this covers any $[X, Y, Z]$ which satisfies $Z \neq 0$.

If we similarly add further

$$\mathbf{F}_{xz}^2 = \{(x, z) : x, z \in \mathbf{F}^2\}, \quad \mathbf{F}_{yz}^2 = \{(y, z) : y, z \in \mathbf{F}^2\},$$

and take the mappings

$$\mathbf{F}_{xz}^2 \ni (x, z) \mapsto [x, 1, z] \in \mathbf{P}^2(\mathbf{F}), \quad \mathbf{F}_{yz}^2 \ni (y, z) \mapsto [1, y, z] \in \mathbf{P}^2(\mathbf{F}),$$

then we see that we obtain each point of $\mathbf{P}^2(\mathbf{F})$, since for any element of $\mathbf{P}^2(\mathbf{F})$, there is a non-vanishing coordinate of any representative (on this point, it may be worth to note that although coordinates of representatives are usually not well-defined, vanishing or non-vanishing of coordinates of representatives are: these properties are invariant under scalar multiplications).

These mappings can be easily inverted, e.g.

$$\mathbf{P}^2(\mathbf{F}) \setminus \{Z = 0\} \ni [X, Y, Z] \text{ with } Z \neq 0 \mapsto (XZ^{-1}, YZ^{-1}) \in \mathbf{F}^2,$$

(and similarly for the other non-vanishing coordinates).

Remark A.3.12. In general, $\mathbf{P}^d(\mathbf{F})$ can be covered with $d + 1$ copies of \mathbf{F}^d .

Chapter B

Number theory

B.1 The fundamental theorem of arithmetic

In this section, we state and prove the fundamental theorem of arithmetic: the fact that any nonzero integer can be written – essentially uniquely – as the product of prime (irreducible) numbers. The heart of the matter is in fact that primes and irreducibles coincide among rational integers.

Proposition B.1.1 (euclidean division). *Given integers a, b , $b \neq 0$. Then there exist integers c, d satisfying $a = bc + d$ and $|d| < |b|$.*

Proof. Let $a \geq 0$, $b > 0$, the remaining cases are similar. Induct on a . For $a = 1$, the statement is trivial ($c = 1, d = 0$ if $b = 1$ and $c = 0, d = 1$ if $b > 1$). Now assume that the statement holds for any $0 \leq a' < a$. If $a < b$, then $c = 0, d = a$. If $a \geq b$, then $a - b = bc' + d'$ with $|d'| < |b|$ by induction, so $a = b(c' + 1) + d'$. \square

Proposition B.1.2. *Assume $a, b \in \mathbf{Z}$. Then there exists an integer $\gcd(a, b)$ satisfying $\gcd(a, b) \mid a, b$ and also that whenever $d \mid a, b$, $d \mid \gcd(a, b)$.*

Proof. If $b = 0$, then $\gcd(a, b) = a$ does the job. Otherwise, consider the sequence $(a, b, d_1, \dots, d_n, 0)$, where each d_i is defined via the euclidean division $d_{i-2} = c_{i-1}d_{i-1} + d_i$ (with $d_{-1} = a, d_0 = b, d_{n+1} = 0$). It is clear that such a sequence of euclidean divisions terminates, since the absolute value decreases in each step. Set $\gcd(a, b) = d_n$. It is clear that $d_n \mid d_{n+1}, d_n$, and then by induction, $d_n \mid d_i, d_{i-1}$ implies $d_n \mid d_{i-2}$. Also, if $d \mid d_{i-2}, d_{i-1}$ (which holds for $i = 1$), then $d \mid d_i$, yielding $d \mid d_n = \gcd(a, b)$. \square

Observe that $\gcd(a, b)$ is well-defined only up to sign.

Definition B.1.3 (greatest common divisor). The greatest common divisor of $a, b \in \mathbf{Z}$ is the nonnegative number which satisfies the conditions imposed on $\gcd(a, b)$ in Proposition B.1.2.

Definition B.1.4 (euclidean algorithm). The sequence of euclidean divisions in the proof of Proposition B.1.2 is called the euclidean algorithm.

Proposition B.1.5. *Assume $a, b \in \mathbf{Z}$. Then $\gcd(a, b) = au + bv$ for some $u, v \in \mathbf{Z}$.*

Proof. If $b = 0$, the statement is trivial. Otherwise, we can create the same sequence $(a, b, d_1, \dots, d_n, 0)$ as in the proof of Proposition B.1.2. Clearly $d_{-1} = a, d_0 = b$ are integer combinations of a and b . Also, if d_{i-2}, d_{i-1} are integer combinations, then so is d_i . \square

Definition B.1.6 (prime numbers). A nonzero integer p is said to be prime, if $p \nmid 1$, and whenever $p \mid ab$, $p \mid a$ or $p \mid b$.

Definition B.1.7 (irreducible numbers). A nonzero integer p is said to be irreducible, if $p \nmid 1$, and whenever $p = ab$, $a \mid 1$ or $b \mid 1$.

Proposition B.1.8. *An integer p is prime if and only if it is irreducible.*

Proof. Assume p is prime, and let $p = ab$. Then $a, b \neq 0$. If $a \nmid 1$ and $b \nmid 1$, then $1 < |a|, |b| < p$. Therefore $p \nmid a, b$, which is a contradiction.

Assume p is irreducible, and $p \mid ab$. If $p \mid a$, we are done. If $p \nmid a$, then $\gcd(a, p) = 1$, since p is irreducible. Then there exist integers u, v satisfying $au + pv = 1$. Multiplying by b , we obtain $abu + pbv = b$, the left-hand side is divisible by p , so is the right-hand side. \square

Theorem B.1.9 (fundamental theorem of arithmetic). *Every nonzero integer can be written as a product of prime (irreducible) numbers. The decomposition is unique, apart from factors dividing 1.*

Proof. First we prove the existence by induction on $|n|$. For $|n| = 1$, it is trivial. Assume that the statement holds for any n' with $|n'| < |n|$. If n is irreducible, we are done. If not, we can write it as a product $n = ab$ with $|a|, |b| < |n|$. We are done by induction.

Now we prove the uniqueness. Assume n has two decompositions $p_1 \cdot \dots \cdot p_k = q_1 \cdot \dots \cdot q_l$. Here, p_1 divides the left-hand side, so it divides the right-hand side as well. Then, since it is a prime, it divides a factor of the right-hand side, say, q_1 . Then $p_1 \mid q_1$, and also $q_1 \mid p_1$, since q_1 is irreducible. Dividing by them, we can complete the proof by induction. \square

B.2 Residue classes

Definition B.2.1 (congruence). Given $m \in \mathbf{Z}$, we say that $a \equiv b \pmod{m}$ (in words: a is congruent to b modulo m) if $m \mid (a - b)$.

Proposition B.2.2 (remainders). *Given $m \in \mathbf{Z} \setminus \{0\}$ and $a \in \mathbf{Z}$, there exist $0 \leq b < |m|$ and $-|m|/2 < c \leq |m|/2$ satisfying $a \equiv b \equiv c \pmod{m}$.* \square

Proposition B.2.3. *Being congruent modulo m is an equivalence relation, by which we mean that $a \equiv a \pmod{m}$ (reflexivity), $a \equiv b \pmod{m}$ implies $b \equiv a \pmod{m}$ (symmetry), $a \equiv b \pmod{m}$ and $b \equiv c \pmod{m}$ imply $a \equiv c \pmod{m}$ (transitivity).* \square

Definition B.2.4 (residue classes). Let $m \in \mathbf{Z}$. Then the equivalence classes defined via modulo m congruency are said to be modulo m residue classes.

Proposition B.2.5. *Let $m \in \mathbf{Z} \setminus \{0\}$. The number of residue classes modulo m is $|m|$.* \square

Proposition B.2.6. *Let $m \in \mathbf{Z} \setminus \{0\}$. The modulo m residue classes form a commutative ring, where addition, multiplication and taking additive inverse are the usual addition, multiplication and taking additive inverse, all reduced modulo m . The residue class of 1 is the multiplicative unit.*

Proof. Let a, b be arbitrary representatives of two residue classes. Then for any $k, l \in \mathbf{Z}$, we have

$$\begin{aligned}(a + km) + (b + lm) &= (a + b) + (k + l)m \equiv a + b \pmod{m}, \\(a + km) \cdot (b + lm) &= ab + (kb + la)m + klm^2 \equiv ab \pmod{m}, \\-(a + km) &= -a - km \equiv -a \pmod{m},\end{aligned}$$

showing that the operations can be performed via any representatives. Then obviously the residue class of 1 is a multiplicative unit. \square

Theorem B.2.7 (Chinese remainder theorem). *Assume $m, n \in \mathbf{N}$ satisfy $\gcd(m, n) = 1$. Then for any $a, b \in \mathbf{Z}$, there exists a unique residue class c modulo mn satisfying $c \equiv a \pmod{m}$, $c \equiv b \pmod{n}$.*

Proof. Define the function $f : x \pmod{mn} \mapsto (x \pmod{m}, x \pmod{n})$. There are mn residue classes modulo mn , and the number of possible values of this function is also mn . It suffices to prove that f is a bijection, which holds if and only if it is a surjection.

To see this, take $u, v \in \mathbf{Z}$ satisfying $mu + nv = 1$. Then take the number $c = mub + nva$. Then

$$c \equiv nva \equiv mua + nva \equiv a \pmod{m}, \quad c \equiv mub \equiv mub + nvb \equiv b \pmod{n}.$$

The proof is complete. \square

Corollary B.2.8. *Assume m_1, \dots, m_n satisfy $\gcd(m_i, m_j) = 1$ for all $1 \leq i < j \leq n$. Then for any $a_1, \dots, a_n \in \mathbf{Z}$, there exists a unique residue class c modulo $m_1 \cdot \dots \cdot m_n$ satisfying $c \equiv a_i \pmod{m_i}$ for each $1 \leq i \leq n$.* \square

B.3 The multiplicative group of the residue class rings

Given $m \in \mathbf{N}$, we denote by \mathbf{Z}_m the ring of residue classes modulo m . Assume $\gcd(a, m) = 1$.

Proposition B.3.1. *The function $x \mapsto xa : \mathbf{Z}_m \rightarrow \mathbf{Z}_m$ is a bijection.*

Proof. Since \mathbf{Z}_m is finite, it suffices to show that the function is surjective. Since $\gcd(a, m) = 1$, for some $u, v \in \mathbf{Z}$, $au + mv = 1$. Then if the residue class c modulo m is given, $x \equiv uc \pmod{m}$ is mapped to c via $x \mapsto xa \pmod{m}$, since $c = auc + mvc \equiv auc \pmod{m}$. \square

Corollary B.3.2. *The coprime residue classes form a group. (This is the unit group of \mathbf{Z}_m and will be denoted by \mathbf{Z}_m^\times from now on.)* \square

Corollary B.3.3. *If p is a prime, the residue classes modulo p form a field (denoted by \mathbf{F}_p from now on).* \square

The following function is of extreme importance in number theory.

Definition B.3.4 (Euler's number of coprime residue classes function). Define

$$\varphi(n) = \sum_{\substack{d \leq n \\ \gcd(d, n) = 1}} 1.$$

With this notation, we see that the group \mathbf{Z}_m^\times has $\varphi(m)$ elements.

Proposition B.3.5. *Assume $n = p_1^{\alpha_1} \cdots p_r^{\alpha_r}$. Then*

$$\varphi(n) = n \prod_{j=1}^r \left(1 - \frac{1}{p_j}\right).$$

Proof. Assume n has prime factors p_1, \dots, p_r . From the set $\{1, \dots, n\}$, sift out the numbers that are divisible by some of p_1, \dots, p_r . That is, by the inclusion-exclusion principle,

$$\varphi(n) = n + \sum_{j=1}^r (-1)^j \sum_{1 \leq i_1 < \dots < i_j \leq r} \frac{n}{p_{i_1} \cdots p_{i_j}} = n \prod_{j=1}^r \left(1 - \frac{1}{p_j}\right).$$

The proof is complete. \square

Now Corollary A.1.18 has the following consequences.

Corollary B.3.6. *Assume $m \in \mathbf{N}$ and a is coprime to m . Then the order of a (in the multiplicative group modulo m) divides $\varphi(m)$.* \square

Corollary B.3.7 (Euler-Fermat). *Assume $m \in \mathbf{N}$. Then*

$$a^{\varphi(m)} \equiv 1 \pmod{m}$$

for any a coprime to m . \square

Corollary B.3.8 (Fermat). *Assume p is a prime. Then*

$$a^p \equiv a \pmod{p}$$

for any $a \in \mathbf{Z}$. \square

Finally, we describe the group structure of \mathbf{Z}_m^\times . By the Chinese remainder theorem (Corollary B.2.8),

$$\mathbf{Z}_m^\times \cong \mathbf{Z}_{p_1^{\alpha_1}}^\times \times \cdots \times \mathbf{Z}_{p_r^{\alpha_r}}^\times,$$

if $m = p_1^{\alpha_1} \cdots p_r^{\alpha_r}$ is the canonical form. Therefore, it suffices to describe $\mathbf{Z}_{p^\alpha}^\times$ for primes p and positive integers α . This group has $p^\alpha - p^{\alpha-1} = p^{\alpha-1}(p-1)$ elements.

Proposition B.3.9. *If p is an odd prime, then $\mathbf{Z}_{p^\alpha}^\times$ is cyclic for any $\alpha \in \mathbf{N}$. For $p = 2$, $\mathbf{Z}_{2^\alpha}^\times$ is cyclic for $\alpha = 1, 2$, and is not cyclic but generated by the elements -1 (of order 2) and 5 (of order $2^{\alpha-2}$) for $\alpha \geq 3$.*

Proof omitted.

Chapter C

Probability theory

C.1 Discrete probability spaces

Definition C.1.1 (probability space). By a (finite, discrete) probability space, we mean a pair (Ω, \Pr) , where Ω is a finite set, and \Pr is a function from 2^Ω to $\mathbf{R}_{\geq 0}$ satisfying

- $\Pr(\Omega) = 1, \Pr(\emptyset) = 0$;
- if $A, B \subseteq \Omega$ are disjoint (i.e. $A \cap B = \emptyset$), then $\Pr(A \cup B) = \Pr(A) + \Pr(B)$.

Example C.1.2 (uniform distribution). On any finite set Ω , we can define the uniform distribution by assigning the same probability $(\#\Omega)^{-1}$ to each one-element set, i.e. for any $A \subseteq \Omega$,

$$\Pr(A) = \frac{\#A}{\#\Omega}.$$

The subsets of Ω are the *events*, and the corresponding \Pr values are their *probabilities*. Intuitively we really think of some randomness and the probability of a given event is the chance that the given event occurs.

For any $A \subseteq \Omega$, we define A^c to be the complement of A , obviously

$$\Pr(A^c) = 1 - \Pr(A).$$

Definition C.1.3 (conditional probability). Given a probability space Ω , for any $A, B \subseteq \Omega$, if $\Pr(B) > 0$, we define the conditional probability

$$\Pr(A \mid B) = \frac{\Pr(A \cap B)}{\Pr(B)}.$$

The meaning of this notion is the following: the probability of A given that B occurs.

Proposition C.1.4. Assume $A, B \subseteq \Omega$ satisfy $\Pr(B), \Pr(B^c) > 0$. Then

$$\Pr(A) = \Pr(A \mid B)\Pr(B) + \Pr(A \mid B^c)\Pr(B^c).$$

Proof. The statement follows by a simple calculation from the definitions. Indeed,

$$\Pr(A) = \Pr((A \cap B) \cup (A \cap B^c)) = \Pr(A \cap B) + \Pr(A \cap B^c) = \Pr(A \mid B)\Pr(B) + \Pr(A \mid B^c)\Pr(B^c),$$

and the proof is complete. \square

Proposition C.1.5 (Bayes's theorem). Assume $A, B \subseteq \Omega$ satisfy $\Pr(A), \Pr(B) > 0$. Then

$$\Pr(A \mid B) = \frac{\Pr(B \mid A)\Pr(A)}{\Pr(B)}.$$

Proof. Again, by a simple calculation,

$$\Pr(A | B) = \frac{\Pr(A \cap B)}{\Pr(B)} = \frac{\Pr(B \cap A)}{\Pr(B)} = \frac{\Pr(B | A)\Pr(A)}{\Pr(B)},$$

and the proof is complete. \square

Definition C.1.6. We say that the events A_1, \dots, A_n are independent, if for any $S \subseteq \{1, \dots, n\}$,

$$\Pr\left(\bigcap_{i \in S} A_i\right) = \prod_{i \in S} \Pr(A_i),$$

(we can freely exclude $S = \emptyset$; or use the conventions: the empty product is 1, and the empty intersection is Ω).

Bibliography

- [1] J. Hoffstein, J. Pipher, and J. H. Silverman. *An Introduction to Mathematical Cryptography*. Undergraduate Texts in Mathematics. Springer, 2008.
- [2] L. Lovász. *Algoritmusok bonyolultsága*. University text, in Hungarian.
- [3] W. Rudin. *Principles of mathematical analysis*. International Series in Pure and Applied Mathematics. McGraw-Hill, 1976 (3rd edition).