PAWS 2025: MATHEMATICAL CRYPTOGRAPHY
PROBLEM SET 1

GIACOMO BORIN, JOLIJN COTTAAR, ELI ORVIS, GABRIELLE SCULLARD

The goal for Problem Set 1 is to gain some more familiarity with ciphers, public-key cryptography, and Diffie-Hellman. The exercises are organized into beginner, intermediate, and advanced levels. If this is your first time seeing these topics, we suggest focusing on the beginner and intermediate problems to start. However, you're welcome and encouraged to tackle any problems that catch your interest!

(1) (Beginner) Instead of using a shift of exactly 3 letters in Caesar's cipher, one could also use a secret shift depending on a key $k \in \{0, ..., 25\}$. Describe $\mathcal{M}$, $\mathcal{C}$, $\mathcal{K}$, **Dec** and **Enc** for this new encryption method.

(2) (Intermediate) To increase the number of keys for the Caesar cipher, one may also choose a key of the form $k = (a, b)$ with $a \in (\mathbb{Z}/26\mathbb{Z})^*$, $b \in (\mathbb{Z}/26\mathbb{Z})$, and

$$\mathbf{Enc}_k(m_1, ..., m_n) = (c_1, ..., c_n) \quad \text{with} \quad c_i = am_i + b.$$

The corresponding scheme is known as an affine cipher.

(a) Describe the corresponding decryption function $\mathbf{Dec}_k$. Why is it necessary that a is a unit in $(\mathbb{Z}/26\mathbb{Z})$?

(b) You (Eve) read a cipher text starting with BMVVK and you think that it means HELLO. Is it possible that Alice and Bob used an affine cipher in their communication? Can you recover their secret key?

(c) (sage) Alice and Bob noticed that you found their secret, and chose a new private key. This time you intercept the cipher text:

IFELTKHURFENHAFEEFSFUTSVGEDNULTKFBF

Can you find the plain text message?

(3) (Beginner/Intermediate) How many affine ciphers are possible using the 26-letter English alphabet? How many are possible if we also allow the symbols "?", ".", "," , and "!"?

(4) (sage, Intermediate) You intercepted the cipher text

JIVQOJIV LEALAVQO KGOONDTV QOAELONE OAINYNGJ SOBVQODB CLAVQOKG OONDTJIV QOJIVLEA EIBHTBLO YBLEQPIG AA

from a conversation between Alice and Bob. You know that they used a substitution cipher. Can you recover the plain text m? Note that the spacing is only used for readability and does not coincide with the spacing of the original text.

(5) (sage, Beginner) Alice and Bob want to create a shared Diffie-Hellman key. They use setups with varying security levels. In all of these, the public parameters are a prime $p = 2q + 1$, and the element $g = 4 \in \mathbb{F}_p^*$ with order q. You observe the following conversations. Can you find the shared keys?

(a) $q = 4294967681 \approx 2^{32}$,
$A = 5104411285$, $B = 7620748646$.

(b) $q = 18446744073709552109 \approx 2^{64}$,
$A = 17485644247020728566$, $B = 17485644247020728566$.

(c) $q = 340282366920938463463374607431768219863 \approx 2^{128}$,
$A = 158556695861572453782110953476057063 05$,
$B = 643791185530305885858740134964520672205$

In SageMath, you can use the *log* function to compute discrete logarithms, i.e., $a = A.log(g)$ (provided that g, A are defined as elements over $\mathbb{F}_p$). Further, you can use *%time* to time your results. How does the runtime evolve for increasing values of q?

(6) (Beginner) We now try to set up the ElGamal public key encryption scheme. We will start by doing that in the group $G = \mathbb{F}_{29}^*$.

(a) First we need to fix a generator for G, say $g = 2$. Check that g is actually a generator of G. Could we have chosen $g = 3$ or $g = 5$?

(b) Bob chooses his secret key $b = 5$, compute his public key.

(c) Bob receives the encryption $(c_1, c_2) = (7, 9)$ from Alice, decrypt it and find the message. (Bonus: do you think Bob can also find k?)

(d) Suppose now that Alice receives the public key $B = 14$ from Bob. Encrypt the message $m = 23$.

(e) Suppose that Bob sets his public key to $B = 28$. Do you think this is secure? Suppose Eve sees the encryption $(c_1, c_2) = (14, 22)$. Can she say something about the message?

(7) (sage, Intermediate) Try now to implement the ElGamal public key encryption scheme in sagemath using prime fields. You need to implement the following functionalities:

- **KeyGen**, taking as input a prime p and a generator g, returning a key pair $(\mathsf{sk}, \mathsf{pk})$.
- **Enc**, taking as input a prime p, a generator g, a public key pk and a message m, returning an encryption (c_1, c_2).
- **Dec** taking as input a prime p, a generator g, the secret key sk and an encryption (c_1, c_2), returning a message m.

For the prime field and the generator use the parameter set ffdhe3072, in which

$$p = 2^{3072} - 2^{3008} + (\lfloor 2^{2942} \cdot e \rfloor + 2625351) \cdot 2^{64} - 1$$

and $g = 2$. You can find more information and the hexadecimal representation for the prime in the Appendix A and A.2 of the IETF standard Negotiated Finite Field Diffie-Hellman Ephemeral Parameters for Transport Layer Security (TLS). Try to time your implementations of **KeyGen**, **Enc** and **Dec**.

Some useful function you can use in SageMath are the time library (load it with `import time`, then use `time.time()` to get the Unix time), the function `pow` (try to run `pow?` to see how to use it) or the constructor `GF(p)` that create the finite field $\mathbb{F}_p$.

(8) (Intermediate) Consider the following setup:

$$p = 8589935363, \quad g = 4 \in \mathbb{F}_p^*,$$

and assume that Bob's public key is

$$B = 1865230978.$$

(For readability, we chose a small prime for which the dlog can still be computed efficiently. For the sake of this exercise, assume however that you cannot compute $b = \mathrm{dlog}_g(B)$.)

Bob is asking some yes/no questions to Alice. Alice encrypts her answers ($Y = 25 \in \mathbb{F}_p^*$ for yes and $N = 14 \in \mathbb{F}_p^*$ for no) using Bob's public key and the ElGamal encryption scheme.

(a) Eve intercepts Alice's answers:

$$\text{Answer 1:}(2456530342, 8487632028),$$
$$\text{Answer 2:}(2456530342, 1660697205),$$
$$\text{Answer 3:}(2456530342, 1660697205),$$

and immediately sees that Alice is reusing the random integer $k \in \mathbb{Z}$.

(i) Without doing any computations: What are the possible answers that Alice could have sent?

(ii) With some (computationally easy) computation: What are Alice's answers to Questions 1,2,3?

(b) Alice notices her mistake and uses different random exponents for the next answers. However she decides that it is easier to encode $Y = 1 \in \mathbb{F}_p^*$ and $N = -1 \in \mathbb{F}_p^*$. Now Eve intercepts the following messages:

$$\text{Answer 4:}(6324669601, 8569725934),$$
$$\text{Answer 5:}(5864877653, 1038689194),$$
$$\text{Answer 6:}(1841857395, 573429127).$$

Can you recover Alice's answers to Question 4,5,6 as well?

(9) (Intermediate) The most widely used cryptosystem is RSA. The RSA algorithm works as follows: Bob chooses two secret large primes p and q; he computes $N = pq$, e such that $\gcd(e, (p-1)(q-1)) = 1$, and $d = e^{-1} \pmod{(p-1)(q-1)}$. Bob's public key is (N, e). Alice encrypts the message $m < N$ by computing $c = m^e \pmod N$. To decrypt, Bob computes $c^d \pmod N$.

(a) Suppose Bob's public key is $(55, 3)$. Decrypt the ciphertext $c = 12$.

(b) Why would an efficient algorithm for factoring make RSA insecure? (A fun fact: Shor's algorithm is an algorithm which enables *quantum* computers to factor efficiently; finding alternatives which are believed to be secure against quantum computers is an active area of research!)

(10) (Advanced) Let p be a prime and $g \in \mathbb{F}_P^*$ a primitive element. We denote $p - 1 = p_1^{e_1} \cdot \ldots \cdot p_n^{e_n}$ for the prime factorization of $p - 1$. The goal of this exercise is to show that solving the DLP in $\mathbb{F}_p^*$ is essentially as hard as solving the DLP in a subgroup $G \subset \mathbb{F}_p^*$ of prime order $\#G = \max\{p_i \mid i \in \{1, .., n\}\}$. To make this more formal, let us say that the DLP in a subgroup $G_i \subset \mathbb{F}_p$ of order p_i can be solved in time $O(S_i)$.

Let $g \in \mathbb{F}_p^*$ be a primitive element and $A \in \mathbb{F}_p^*$ the challenge for which we want to solve the DLP, i.e., we want to find $a \in \mathbb{Z}$ with $g^a = A$.

(a) Use the Chinese Remainder Theorem to translate the problem into solving n smaller instances of the DLP in subgroups of order $p_i^{e_i}$ with $i \in \{1, ..., n\}$, respectively.

(b) Fix $i \in \{1, ..., n\}$ and say you want to solve one of the small DLP instances on input $A_i, g_i \in \mathbb{F}_p^*$, where $\operatorname{ord}(g_i) = p_i^{e_i}$, i.e., you want to find $a_i \in \mathbb{Z}$ with

$$g_i^{a_i} \equiv A_i \mod p_i^{e_i}.$$

Show that this can be done in time $O(e_i S_i)$. Hint: Use a p_i-ary representation of $a_i = \alpha_o + \alpha_1 p_i + \ldots + \alpha_{e_1 - 1} p_i^{e_i - 1}$.

(c) Combine the results of (a) and (b) to show that the DLP can be solved in time

$$O(\text{polylog}(p) \max\{S_i\}).$$

(11) (Intermediate) To avoid the attack described above in implementations we use *safe primes*, i.e. primes p so that $p - 1 = 2q$ with q also prime (primes as q are also know as Sophie Germain primes). Try to understand if the following are safe primes or not:

(a) $p_1 = 140970312865529183$,

(b) $p_2 = 282481863544496003$,

(c) $p_3 = 289942627069958089$,

(d) $p_4 =$ 0x57f6fbca677315519342e1adf372f2402eb9ce3f77dbe8e4fcce5052bee98efb,

(e) $p_4 =$ 0x7404cc9709ed7da6a4e7551e85465df1c5bd4855274bff5d392da63732baa65,

(f) the prime from the parameter set ffdhe4096 (see the IETF standard Negotiated Finite Field Diffie-Hellman Ephemeral Parameters for Transport Layer Security (TLS)).

(12) (Advanced) Alice and Bob create a symmetric cipher in the following way: Their private key k is a large integer and their messages are d-digit integers, so

$$\mathcal{M} = \{m \in \mathbb{Z} : 0 \leq m < 10^d\}.$$

To encrypt a message, Alice computes $\sqrt{k}$ to d decimal places and lets α be the d-digit number to the right of the decimal place. (For example, if $k = 87$ and $d = 6$ then $\sqrt{87} = 9.32737905...$ and $\alpha = 327379$.)

Alice encrypts m as $c \equiv m + \alpha \pmod{10^d}$. Bob decrypts c by computing $m \equiv c - \alpha \pmod{10^d}$.

(a) Alice and Bob choose the secret key $k = 11$ and use $d = 6$. Bob wants to send Alice the message $m = 328973$. What is the ciphertext that he sends?

(b) Alice and Bob choose the secret key $k = 23$ and use it to encrypt 8-digit integers. Alice receives $c = 78183903$. What is the plaintext m?

(c) Show that

$$\alpha = \lfloor 10^d(\sqrt{k} - \lfloor \sqrt{k} \rfloor) \rfloor.$$

(d) (Challenge!) If Eve steals a plaintext/ciphertext pair (m, c), then she can easily compute $\alpha \equiv c - m \pmod{10^d}$. If 10^d is large compared to k, can she also recover k? (This would be useful if k is reused as a secret key, for greater values of d.)