

Cryptography and Security — Final Exam

Solution

Serge Vaudenay

26.1.2022

- duration: 3h
- no documents allowed, except one 2-sided sheet of handwritten notes
- a pocket calculator is allowed
- communication devices are **not** allowed
- the exam invigilators will **not** answer any technical question during the exam
- readability and style of writing will be part of the grade
- answers should not be written with a pencil

The exam grade follows a linear scale in which each question has the same weight.

1 Diffie-Hellman vs ElGamal

We consider a group (with multiplicative notation) of prime order q generated by a given element g . The Diffie-Hellman protocol was specified as follows in the course:

- Alice picks $x \in \mathbf{Z}_q^*$ uniformly, computes $X = g^x$, and sends X to Bob.
- Bob checks that $X \in \langle g \rangle - \{1\}$ and aborts if it is not the case.
- Bob picks $y \in \mathbf{Z}_q^*$ uniformly, computes $Y = g^y$, and sends Y to Alice.
- Bob computes $K = \text{KDF}(X^y)$ and uses it as the output.
- Alice checks that $Y \in \langle g \rangle - \{1\}$ and aborts if it is not the case.
- Alice computes $K = \text{KDF}(Y^x)$ and uses it as the output.

The ElGamal cryptosystem was specified as follows in the course:

- The receiver generates his key pairs as follows. He picks $x \in \mathbf{Z}_q$ uniformly as a secret decryption key and computes the public encryption key $y = g^x$.
- The sender who wants to encrypt $m \in \langle g \rangle$ with key y picks $r \in \mathbf{Z}_q$ and computes $u = g^r$ and $v = my^r$. The ciphertext is (u, v) .
- The receiver who wants to decrypt (u, v) with key x computes $m = vu^{-x}$.

Q.1 Explain how the ElGamal cryptosystem is built from the Diffie-Hellman protocol. For this, explain each element of the ElGamal cryptosystem and how it comes from the Diffie-Hellman protocol.

$y = g^x$ from the receiver plays the role of $Y = g^y$ from Bob. $u = g^r$ from the sender plays the role of $X = g^x$ from Alice. In the cryptosystem, there is an implicit exchange of y (the encryption key) and u (part of the ciphertext). The derived key is g^{xr} instead of g^{xy} and it is used without a KDF to encrypt m using one-time pad in the group.

Hence, we can say that the ElGamal cryptosystem is one-time pad with the key derived from a semi-static Diffie-Hellman protocol. Here, $y = g^x$ is a static key for the receiver and $u = g^r$ is the ephemeral key for the sender.

Q.2 In the Diffie-Hellman protocol, x and y are selected in \mathbf{Z}_q^* . What is the problem if we select them in \mathbf{Z}_q ? Explain the security threat.

A man-in-the-middle could replace X and Y by 1 and the protocol would end with a correct $K = \text{KDF}(1)$ which is also known by the adversary. To defeat this trivial active attack, Alice and Bob check that the key they receive is not equal to 1. To have the protocol correct (i.e. not failing although participants are honest), x and y must be nonzero.

Q.3 In the ElGamal cryptosystem, x and r are instead selected in \mathbf{Z}_q . Why isn't it a problem?

In the security model of the cryptosystem, the adversary is not supposed to corrupt the encryption key y . The probability that $x = 0$ is negligible so the probability that $y = 1$ is also negligible. The adversary could replace u by 1. The effect would be that the receiver decrypts to v which could also be selected by the adversary. In any case, the adversary would not learn m , which is the purpose of the cryptosystem.

2 Immunity Certificate

Disclaimer: this exercise was imagined in 2021. This is a work of fiction. Any similarity to actual persons, living or dead, or actual events, is purely coincidental.

A virus is currently striking many countries. Fortunately, there exists an effective vaccine, people who caught the virus are immune, and there exists a reliable infection test. Hence, countries can assess if people are immune in three categories: *vaccinated*, *healed*, and (negatively) *tested*. In the country Weirdoland, the Badass Office for Public Health (BOPH) — the local health authority — has deployed a way to give digital immunity certificates to people and a way to show them from inspection and to verify them.

The immunity certificate consists of a message encoding several pieces of information together with a digital signature using ECDSA. The information includes:

- the signing authority and how to verify the immunity certificate,
- the identity of the holder (full name and date of birth),
- the type of immunity (*vaccinated*, *healed*, or *tested*),
- and extra information about immunity which contains a reference date.

If *vaccinated*, this includes the date of the vaccine. If *healed*, this includes the date of the positive diagnosis. If *tested*, this includes the date of the test. An immunity certificate can be issued through a specific online platform with appropriate access control by vaccine/test centers and medical doctors. Signing keys are often changed and a PKI certificate for the verification key is provided online by BOPH. The validity period of a certificate is announced by BOPH following a policy. It depends on the type of immunity certificate and the reference date.

BOPH provides two apps which people install on their smartphones. One app is the Certshow and the other is Certscan. Both frequently connect to the BOPH server to retrieve existing PKI certificates and possible lists of revoked certificates. Indeed, when it appears that a signing key leaked or was misused, the PKI certificate of the corresponding verification key can be revoked and the immunity certificates which have been signed by this key are not valid any more. Immunity certificates which have a correct signature, a correct PKI certificate which is also non-revoked, and which are still in their validity period are called *valid*. Certshow allows people to store their immunity certificate, to show it in the form of a QR code, and to verify its validity. The QR code encodes all information. Certscan allows to scan the QR code of an immunity certificate, to decode it, to display the data it contains, and to check its validity.

Q.1 The Weirdoland government decided that restaurants should only give access to customers holding a valid immunity certificate. Why are restaurants asking the ID of customers in addition to scanning their QR code?

Why should restaurants use their Certscan app instead of relying on the Certshow app of their customers to verify the immunity certificate?

The certificate only assesses that a person with a given identity holds a valid certificate. To verify that it is the same as the person entering in the restaurant, the identity of this person must be checked.

If the restaurant relies on the smartphone of the customer, the customer can install an app which mimics Certshow but always say that the certificate is valid.

Q.2 Mr Antigov wants to enter a restaurant without following the correct procedure (i.e. he had no vaccine, no recent test, and was not sick before). Suggest 3 ways to do so and sort them in increasing difficulty. (Note: the difficulty of some methods may be subjective, it is recommended to justify the order.)

- *The easiest way is certainly to follow the rules, but this may be a no-go for Mr Antigov.*
- *Otherwise, he can try social engineering to enter the restaurant, e.g. claim that he forgot his certificate, show a smartphone out of battery, etc.*
- *He can try to use the certificate of someone else and hope that ID verification will be poor (or skipped).*
- *He can try to jam the mobile communication so that the Certscan of the restaurant will not be able to connect to BOPH, or more generally some DoS attack to stop the communication between Certscan and BOPH.*
- *He can use an outdated certificate and corrupt the clock of the Certscan of the restaurant (clock adjustment is often insecure).*
- *He can also corrupt someone who has access to the online platform (someone from a vaccine/test center or a physician) with the risk that this person may be caught and all the certificates issued by the same key being revoked.*
- *If foreign certificates are also accepted, he can try to find a country with lower security standards to try the previous attack and get a certificate.*
- *He can use bruteforce to enter the restaurant.*
- *He can do revolution to abolish the rules.*
- *Finally, he can try to solve the discrete logarithm problem to forge a certificate.*

Q.3 The Badass Officer for Privacy Protection (BOPP) complains that some medical information about the holder leaks from the certificate. For that, BOPH proposes a “blind certificate” which only says “valid” or not, together with the identity. The way to get this blind certificate consists of having Certshow to send the full certificate to BOPH and to get in return an immunity certificate of a new type (the type *blind*) which is valid for only one day. This question is split into 3 subquestions.

Compare the privacy issues with the regular immunity certificate and the blind certificate.

In parallel, the Weirdoland government decided that immunity certificate of type *tested* should not be allowed any more in restaurants. **Why** doesn't the blind certificate work in this case?

Next, the Weirdoland government requires that people going to bars must have a valid immunity certificate of type other than *tested* and be tested. **How** could we update the immunity certificate and the blind certificate to enforce that?

Without the blind certificate, the content of the immunity certificate is given to every restaurant. They are not supposed to keep records but there is no technical measure to prevent them from doing so. With the blind certificate, the content of the immunity certificate is given to BOPH every time the holder needs one blind certificate. Hence, BOPH learns the content and when the person needs it. BOPH is not supposed to keep records but there is no technical measure to prevent BOPH from doing so. Finally, both methods leak and it is a matter of who the holder would trust the most. Since the blind certificate is issued even to tested people and it does not say if the person was vaccinated or tested, the blind certificate cannot work to exclude one type of immunity certificate.

The problem is now to show a blind certificate assessing a specific predicate. Before, the default predicate was whether or not the immunity certificate was valid. Now, we need more elaborate predicates such as “valid immunity certificate and (vaccinated or healed) and tested”. Obviously, we need to store several types in the same certificate or to have several certificates. One way to have a blind certificate for the predicate is to send the immunity certificate(s) to BOPH together with the required predicate. BOPH would then verify the predicate and sign that this specific predicate was verified inside a blind certificate.

- Q.4** BOPH claims that they store no private information. They claim that immunity certificates are decentralized and only stored on the holder’s smartphone. However, when Mrs Jab had her last dose, she received a certificate containing her previous data and the new injection. Do you believe that BOPH really stores no information? Explain why.

It may be correct that BOPH stores no information but the vaccination center does. However, it is misleading to say that data is only on the smartphone as it is clearly elsewhere as well. In practice, BOPH keeps immunity certificate serial numbers and who has issued them. This can be useful to revoke them or to investigate about possible corruption. This can obviously be used to identify the vaccine center and to contact the center to get the stored data.

3 Blind Schnorr Signatures

The following exercise is inspired from On the (in)security of ROS by Benhamouda, Lepoint, Loss, Orrù, and Raykova, published in proceedings of EUROCRYPT 2021.

We recall the Schnorr signature scheme (with a slight difference). We assume a group (with multiplicative notations) of prime order p generated by an element G . For key generation, the key holder selects $x \in \mathbf{Z}_p$ as a signing key and publishes $X = G^x$ as a verifying key. A valid signature for a message m is a pair (R, s) made of a group element R and an integer s which satisfies $G^s = X^{H(R,m)}R$. To sign m , the signer picks uniformly a random $r \in \mathbf{Z}_p$, computes $R = G^r$ and $s = (x \cdot H(R, m) + r) \bmod p$. We adapt the scheme into a *blind* signature scheme where the signer signs in a blind way, i.e. he does not see what he is signing but he signs one and only one document. In the blind signature protocol, the enquirer has (X, m) as input and the signer has x as input (in addition to the public group parameters such as p and G). Then, the protocol works as follows:

- The enquirer sends a signature request to the signer.
- The signer picks $r \in \mathbf{Z}_p$ uniformly at random and sends back $R = G^r$.
- The enquirer picks $\alpha, \beta \in \mathbf{Z}_p$ uniformly at random, computes $R' = G^\alpha X^\beta R$ and $c = (H(R', m) + \beta) \bmod p$, and sends c to the signer.
- The signer responds with $s = (c \cdot x + r) \bmod p$.
- The enquirer sets $s' = (s + \alpha) \bmod p$ and checks that (R', s') is a valid signature for m .

The security is such that by querying the signer ℓ times, the enquirer should not be able to produce $\ell + 1$ different triplets (m_i, R'_i, s'_i) such that (R'_i, s'_i) is a valid signature of m_i . We model the signer with secret x by the following stateful oracle which takes as input a *session identifier* sid :

$\text{OSig}(\text{sid}, c)$:

- 1: **if** $\text{state}\{\text{sid}\}$ does not exist **then**
- 2: pick $r \in \mathbf{Z}_p$ uniformly
- 3: set $R \leftarrow G^r$
- 4: set $\text{state}\{\text{sid}\} \leftarrow r$
- 5: **return** R
- 6: **end if**
- 7: **if** $\text{state}\{\text{sid}\} = \perp$ **then** abort
- 8: set $r \leftarrow \text{state}\{\text{sid}\}$
- 9: set $s \leftarrow (c \cdot x + r) \bmod p$
- 10: set $\text{state}\{\text{sid}\} \leftarrow \perp$
- 11: **return** s

Here, state is an associative array (a.k.a. a dictionary). $\text{OSig}(\text{sid}, c)$ can be queried the first time with a dummy input $c = \perp$ (which is unused) to get R , then the second time with $c \in \mathbf{Z}_p$ to get s . Both oracle calls with the same sid count for *one* signature query.

Q.1 Define correctness for blind signatures and prove that the above scheme is correct.

Correctness means that if participants follow the protocol, then the outcome is as expected. Namely, after generating a key pair, for every m , following the above scheme gives an (R', s') signature for m which is valid. I.e., it satisfies $G^{s'} = X^{H(R', m)} R'$. We can indeed check that

$$G^{s'} = G^{c \cdot x + r + \alpha} = X^{H(R', m) + \beta} G^r G^\alpha = X^{H(R', m)} R'$$

Q.2 Define the blind signature security with ℓ sessions by using a game Γ_ℓ and properly define the advantage of an adversary.

Game Γ_ℓ :

- 1: set up p and G
- 2: run the key generation to define x and X
- 3: $\mathcal{A}^{\text{OSig}}(p, G, X) \rightarrow (m_0, R'_0, s'_0, \dots, m_\ell, R'_\ell, s'_\ell)$
- 4: **if** state has more than ℓ entries **then** abort
- 5: **if** the m_i 's are not pairwise different **then** abort
- 6: **if** there exists i such that (R'_i, s'_i) is not a valid signature of m_i **then** abort
- 7: **return** 1

The advantage is the probability that the game outputs 1. It happens if the adversary does not make more than ℓ signing queries and produces $\ell + 1$ pairwise different messages together with their valid signatures.

Q.3 We consider an adversary making ℓ oracle calls normally as follows:

$\mathcal{A}(p, G, X)$:

- 1: **for** $i = 0$ to $\ell - 1$ **do**
- 2: $R_i \leftarrow \text{OSig}(i, \perp)$
- 3: pick $\alpha_i, \beta_i \in \mathbf{Z}_p$ uniformly
- 4: set $R'_i \leftarrow G^{\alpha_i} X^{\beta_i} R_i$
- 5: **end for**
- 6: [to be continued...]

After that, the adversary selects 2ℓ pairwise different random messages m_i^b for $i = 0, \dots, \ell - 1$ and $b = 0, 1$ such that $H(R'_i, m_i^0) \bmod p \neq H(R'_i, m_i^1) \bmod p$ for every i . Then, the adversary defines $c_i^b = (H(R'_i, m_i^b) + \beta_i) \bmod p$, $z_i = \frac{2^i}{c_i^1 - c_i^0} \bmod p$, and

$$z_\ell = - \sum_{i=0}^{\ell-1} z_i \cdot c_i^0$$

Note that in notations m_i^b and c_i^b , the superscript b represents an index and not an exponent.

Prove that for any set of ℓ bits $b_0, \dots, b_{\ell-1}$, we have

$$z_0 \cdot c_0^{b_0} + \dots + z_{\ell-1} \cdot c_{\ell-1}^{b_{\ell-1}} + z_\ell = b_0 + 2b_1 + \dots + 2^{\ell-1} \cdot b_{\ell-1}$$

We have the following equation in \mathbf{Z}_p :

$$\sum_{i=0}^{\ell-1} z_i c_i^{b_i} + z_\ell = \sum_{i=0}^{\ell-1} z_i (c_i^{b_i} - c_i^0) = \sum_{i=0}^{\ell-1} 2^i \frac{c_i^{b_i} - c_i^0}{c_i^1 - c_i^0}$$

In each term, the fraction is always equal to b_i . So we can conclude.

Q.4 We first assume that the α_i and β_i that the adversary selected are all equal to 0 for simplicity. Next, the adversary selects a random message m_ℓ which is different from the 2ℓ previous ones. The adversary defines $R'_\ell = R_0^{z_0} \dots R_{\ell-1}^{z_{\ell-1}} X^{-z_\ell}$ and $c_\ell = H(R'_\ell, m_\ell) \bmod p$. We assume that $\ell > \log_2 p$ so that the adversary can make the binary decomposition $c_\ell = \sum_{i=0}^{\ell-1} 2^i b_i$ in ℓ bits b_i . The adversary sets $m_i = m_i^{b_i}$ and $c_i = c_i^{b_i}$. Finally, the adversary runs $s'_i \leftarrow \text{OSig}(i, c_i)$ for $i = 0, \dots, \ell - 1$.

Show that the adversary can find s'_ℓ and make $\ell + 1$ signed messages.

Due to the construction, m_0, \dots, m_ℓ are pairwise different. The adversary gets a signature for $m_0, \dots, m_{\ell-1}$ normally. What is missing is the signature for m_ℓ . For that, we already have a suggested R'_ℓ and we need s'_ℓ such that $G^{s'_\ell} = X^{H(R'_\ell, m_\ell)} R'_\ell$. We have

$$\begin{aligned} & X^{H(R'_\ell, m_\ell)} R'_\ell \\ &= X^{c_\ell} R'_\ell \\ &= X^{2^0 b_0 + \dots + 2^{\ell-1} b_{\ell-1}} R'_\ell \\ &= X^{z_0 c_0^{b_0} + \dots + z_{\ell-1} c_{\ell-1}^{b_{\ell-1}} + z_\ell} R'_\ell \\ &= X^{z_0 c_0 + \dots + z_{\ell-1} c_{\ell-1} + z_\ell} R_0^{z_0} \dots R_{\ell-1}^{z_{\ell-1}} X^{-z_\ell} \\ &= (X^{c_0} R_0)^{z_0} \dots (X^{c_{\ell-1}} R_{\ell-1})^{z_{\ell-1}} \\ &= (G^{s'_0})^{z_0} \dots (G^{s'_{\ell-1}})^{z_{\ell-1}} \\ &= G^{s'_0 z_0 + \dots + s'_{\ell-1} z_{\ell-1}} \end{aligned}$$

So, we can just define $s'_\ell = s'_0 z_0 + \dots + s'_{\ell-1} z_{\ell-1}$.

Q.5 In order to make the attack undetectable, we now assume that the α_i and β_i are totally random. The adversary further selects $\alpha_\ell, \beta_\ell \in \mathbf{Z}_p$ at random. The adversary now defines

$$R'_\ell = G^{\alpha_\ell} X^{\beta_\ell} R_0^{z_0} \dots R_{\ell-1}^{z_{\ell-1}} X^{-z_\ell - \beta_0 z_0 - \dots - \beta_{\ell-1} z_{\ell-1}}$$

Show that the previous attack adapts with a new s'_ℓ .

Like in the previous attack, the adversary gets a signature for $m_0, \dots, m_{\ell-1}$ by normally following the protocol (it is just that the messages are chosen a bit later than normal). We redo the previous computation.

$$\begin{aligned}
& X^{H(R'_\ell, m_\ell)} R'_\ell \\
&= X^{c_\ell - \beta_\ell} R'_\ell \\
&= X^{2^0 b_0 + \dots + 2^{\ell-1} b_{\ell-1} - \beta_\ell} R'_\ell \\
&= X^{z_0 c_0^{b_0} + \dots + z_{\ell-1} c_{\ell-1}^{b_{\ell-1}} + z_\ell - \beta_\ell} R'_\ell \\
&= X^{z_0 c_0 + \dots + z_{\ell-1} c_{\ell-1} + z_\ell - \beta_\ell} R_0^{z_0} \dots R_{\ell-1}^{z_{\ell-1}} X^{-z_\ell + \beta_\ell - \beta_0 z_0 - \dots - \beta_{\ell-1} z_{\ell-1}} G^{\alpha_\ell} \\
&= (X^{c_0} R_0)^{z_0} \dots (X^{c_{\ell-1}} R_{\ell-1})^{z_{\ell-1}} X^{-\beta_0 z_0 - \dots - \beta_{\ell-1} z_{\ell-1}} G^{\alpha_\ell} \\
&= (G^{s'_0} X^{\beta_0})^{z_0} \dots (G^{s'_{\ell-1}} X^{\beta_{\ell-1}})^{z_{\ell-1}} X^{-\beta_0 z_0 - \dots - \beta_{\ell-1} z_{\ell-1}} G^{\alpha_\ell} \\
&= G^{\alpha_\ell + s'_0 z_0 + \dots + s'_{\ell-1} z_{\ell-1}}
\end{aligned}$$

So, we can just define $s'_\ell = \alpha_\ell + s'_0 z_0 + \dots + s'_{\ell-1} z_{\ell-1}$.