Cryptography and Security — Midterm Exam
Solution

Serge Vaudenay

30.10.2024

- duration: 1h45
- 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 Perfect Secrecy Except Message Length

We consider the set of finite bitstrings $\{0,1\}^*$. Given a string s, we denote by $|s|$ the length of s (i.e. the number of bits). We denote by $\perp$ a special symbol which is not an element of $\{0,1\}^*$ and which represents an exception in computation. A *cipher* $C = (X, K, \mathsf{Enc}, \mathsf{Dec})$ is defined by random variables X and K in their respective domains $\mathcal{X} \subseteq \{0,1\}^*$ and $\mathcal{K} \subseteq \{0,1\}^*$, a function $\mathsf{Enc} : \mathcal{K} \times \mathcal{X} \to \{0,1\}^*$, and a function $\mathsf{Dec} : \mathcal{K} \times \{0,1\}^* \to \mathcal{X} \cup \{\perp\}$, such that for any $x \in \mathcal{X}$ and any $k \in \mathcal{K}$, we have $\mathsf{Dec}(k, \mathsf{Enc}(k, x)) = x$. We denote $Y = \mathsf{Enc}(K, X)$. In the Shannon model, X and K are independent.

Q.1 Recall what it means for C to provide perfect secrecy for X of support $\mathcal{X}$.

For any $x \in \mathcal{X}$ and $y \in \{0,1\}^*$ such that $\Pr[Y = y] \neq 0$, we have $\Pr[X = x | Y = y] = \Pr[X = x]$.

Q.2 Can C provide perfect secrecy in the Shannon model for X of support $\mathcal{X} = \{0,1\}^*$? Why?

This is not possible. It is another aspect of the Shannon theorem. Here is the proof which was seen in class.
When we have perfect secrecy, if x and y belong to the supports of X and Y, we have

$$\Pr[X = x] \Pr[Y = y] = \Pr[X = x, Y = y] = \Pr[X = x, \mathsf{Enc}(K, x) = y]$$

By using the independence of X and K, this is equal to $\Pr[X = x] \Pr[\mathsf{Enc}(K, x) = y]$. So, if we divide by $\Pr[X = x]$ (which is nonzero), we obtain $\Pr[\mathsf{Enc}(K, x) = y] = \Pr[Y = y]$, which does not depend on x. Due to correctness, we obtain

$$\Pr[\mathsf{Dec}(K, y) = x] \geq \Pr[\mathsf{Enc}(K, x) = y] = \Pr[Y = y]$$

This means that given a fixed y, for any x in the support of X, the probability that $\mathsf{Dec}(K, y)$ is x is at least a fixed $\Pr[Y = y]$. Therefore, the support of X cannot be infinite. So it cannot be $\{0,1\}^*$.

Q.3 In practice, we want to be able to encrypt a message X of arbitrary length but we want to have a length-preserving cipher, i.e. such that $|\mathsf{Enc}(k,x)| = |x|$ for any $x \in \mathcal{X}$ and $k \in \mathcal{K}$. Justify why we want to encrypt messages of arbitrary length and to have a length-preserving cipher. (There are many good answers for this.)

We want the cipher to be able to encrypt a short message (for instance, an email of a few kilobytes) or a full hard drive (or terabytes) to be usable in many applications. However, we would not like the encryption of a short message to take the space of a full hard drive, because we have to pay for the transmission. So the encryption should somehow respect the length. Ideally, it would exactly preserve it to save spave in an optimal manner but it also comes with privacy risks.

Q.4 We consider a random variable L which models what leaks about X. We assume that L can be easily deduced from Y. Formally define the notion of "perfect secrecy except for the leakage of L".

We can define the notion as follows. For any x, y, ℓ such that $\Pr[Y = y, L = \ell] \neq 0$, we have $\Pr[X = x | Y = y, L = \ell] = \Pr[X = x | L = \ell]$. This means that Y brings no more information about X than what L already leaks. If L is a deterministic function $L = f(Y)$ of Y, the notion boils down to

$$\forall x, y \quad \Pr[Y = y] \neq 0 \Longrightarrow \Pr[X = x | Y = y] = \Pr[X = x | L = f(y)]$$

Q.5 Construct a cipher providing perfect secrecy except for the leakage of $L = |X|$ for any X of support included in $\{0,1\}^*$. Prove that it provides perfect secrecy except for the leakage of $L = |X|$.

HINT: we can leave the Shannon model.

Let X be a random variable with support included in $\{0,1\}^*$. We define $L = |X|$. We define K depending on X. Once $X = x$ is given, we sample $K \in \{0,1\}^{|x|}$ with uniform distribution. Then, we define $\mathsf{Enc}(K, X) = K \oplus X$, and $\mathsf{Dec}(K, Y) = K \oplus Y$. This is not in the Shannon model since X and K are now dependent.
We can verify the new notion of perfect secrecy. Let x, y, ℓ be such that $\Pr[Y = y, L = \ell] \neq 0$. We have

$$\begin{aligned}\Pr[X = x | Y = y, L = \ell] &= \frac{\Pr[X = x, Y = y | L = \ell]}{\Pr[Y = y | L = \ell]} \\ &= \frac{\Pr[X = x, Y = y | L = \ell]}{\sum_{x' \in \{0,1\}^\ell} \Pr[X = x', Y = y | L = \ell]} \\ &= \frac{\Pr[X = x, K = x \oplus y | L = \ell]}{\sum_{x' \in \{0,1\}^\ell} \Pr[X = x', K = x' \oplus y | L = \ell]}\end{aligned}$$

When $L = \ell$, we know that K is independent from X and uniform over a domain of 2^ℓ. So, the fraction boils down to

$$\frac{2^{-\ell} \Pr[X = x | L = \ell]}{\sum_{x' \in \{0,1\}^\ell} 2^{-\ell} \Pr[X = x' | L = \ell]} = \Pr[X = x | L = \ell]$$

and we obtain $\Pr[X = x | Y = y, L = \ell] = \Pr[X = x | L = \ell]$.

2 Diffie-Hellman as a Group Action

A group action by a group G on a set E is a function α with input $(a, u) \in G \times E$ returning an output $\alpha(a, u) \in E$. (We assume that G is multiplicatively denoted.) It must satisfy $\alpha(1, u) = u$ and $\alpha(ab, u) = \alpha(a, \alpha(b, u))$ for any $a, b \in G$ and $u \in E$. For simplicity, we denote $\alpha(a, u) = a * u$. Given a, the function $\alpha_a : u \mapsto a * u$ is a permutation of E. Actually, we can see the group action as a group homomorphism from G to the the group of permutations over E (i.e., the symmetric group of E). We say that the action is transitive if for any pair $u, v \in E$, there exists a such that $a * u = v$.

We define an algorithm $\mathsf{Setup}(1^\lambda)$ which essentially defines a transitive group action by a group of order n, with n of length depending on λ, and which returns some group action parameters, n, and a fixed element $w \in E$.

Q.1 Assume that E is a multiplicative group of prime order q in which we removed the neutral element. Show that $a * u = u^a$ defines a group action from a group G to E and that there is a set $Z \subseteq \mathbf{Z}$ and a surjective function $\mathsf{rep} : Z \to G$ such that $\mathsf{rep}(xy) = \mathsf{rep}(x)\mathsf{rep}(y)$ for all $x, y \in Z$. (For $\mathsf{rep}(xy)$, the multiplication is the one in $\mathbf{Z}$. For $\mathsf{rep}(x)\mathsf{rep}(y)$, the multiplication is the one in G.) Precisely define Z, rep, and G, and give its order n.

When we compute u^a, a is only taken modulo q. The group action is obtained with a multiplicative group. Hence, we take $G = Z_q^*$ which has order $n = q - 1$.
We can see that $u^1 = u$ and $u^{ab} = (u^b)^a$. Therefore, we have a group action.
We let $Z = \{x \in \mathbf{Z}; \mathsf{gcd}(x, q) = 1\}$ and $\mathsf{rep}(x) = x \bmod q$. Clearly, the image set of rep is G and rep has the homomorphic property.

Q.2 Prove that it is transitive.

In E, all elements have orders factor of q which is not 1. Since q is prime, all elements of E have order q. Thus, the mapping $a \mapsto u^a$ is injective on $\{1, \ldots, q-1\}$. Since the mapping is from $\mathbf{Z}_q^*$ of cardinality $q - 1$ to E of cardinality $q - 1$, it is surjective too. Hence, for any pair $u, v \in E$, there exists a such that $a * u = v$: the action is transitive.

Q.3 Rewrite the Diffie-Hellman protocol with that group action.

We assume that a fixed element w of E is pre-determined in setup. Alice picks $a \in G$ and sends $A = a * w$ to Bob. Bob verifies that $A \in E$. Bob picks $b \in G$ and sends $B = b * w$ to Alice. Alice verifies that $B \in E$. Alice computes $K = a * B$ while Bob computes $K = b * A$. We indeed have $a * B = a * (b * w) = (ab) * w = (ba) * w = b * (a * w) = b * A$.

Alice		***Bob***
pick $a \in G$ at random		
$A \leftarrow a * w$	$\xrightarrow{A}$	*verify $A \in E$*
		pick $b \in G$ at random
verify $B \in E$	$\xleftarrow{B}$	$B \leftarrow b * w$
$K \leftarrow a * B$		$K \leftarrow b * A$
$(K = (ab) * w)$		

Q.4 Reformulate the discrete logarithm problem in terms of group action.

1: $\mathsf{Setup}(1^\lambda) \to (\mathsf{params}, n, w)$
2: *pick* $a \in G$
3: $A \leftarrow a * w$
4: $\mathcal{A}(\mathsf{params}, n, w, A) \to a'$
5: ***return*** $1_{A=a'*w}$

3 Square Root Modulo a Prime p s.t. $p \bmod 8 = 5$

Let p be a prime number.

Q.1 When $p \bmod 4 = 3$, recall a method to compute the square root of a quadratic residue modulo p.

We just raise to the power $e = \frac{p+1}{4}$, which is an integer in that case. Note that $e \equiv \frac{1}{2} \pmod{p-1}$.

Q.2 When $p \bmod 4 \neq 3$, what values can $p \bmod 8$ be?

We can have $p = 2$ in which case $p \bmod 8 = 2$. Otherwise, p is odd, so $p \bmod 4$ is either 1 or 3. If it is not 3, it is 1. This means that $p = 1 + 4k$ for some integer k. If k is even, we have $p \bmod 8 = 1$. Otherwise, $p \bmod 8 = 5$. So, $p \bmod 8 \in \{1, 5\}$ except when $p = 2$.

Q.3 In the $p \bmod 8 = 5$ case, prove that if x is a quadratic residue in $\mathbf{Z}_p^*$, then $x^{\frac{p+3}{8}}\theta^{\frac{p-1}{4}}$ is a square root of either x or $-x$ modulo p, for any $\theta \in \mathbf{Z}_p^*$.

Say $x = y^2 \bmod p$. We have

$$\left(x^{\frac{p+3}{8}}\theta^{\frac{p-1}{4}}\right)^4 \equiv y^{p+3}\theta^{p-1} \equiv y^4(y\theta)^{p-1} \equiv y^4 \equiv x^2 \pmod{p}$$

In $\mathbf{Z}_p$, x and $-x$ are the only square roots of x^2. Hence,

$$\left(x^{\frac{p+3}{8}}\theta^{\frac{p-1}{4}}\right)^2 \equiv \pm x \pmod{p}$$

$x^{\frac{p+3}{8}}\theta^{\frac{p-1}{4}}$ is a square root of either x or $-x$ modulo p.

Q.4 In the $p \bmod 8 = 5$ case, let θ be a non-quadratic residue modulo p. Prove that if x is a quadratic residue in $\mathbf{Z}_p^*$, then either $x^{\frac{p+3}{8}}$ or $x^{\frac{p+3}{8}}\theta^{\frac{p-1}{4}}$ is a square root or x.

We have

$$\frac{\left(x^{\frac{p+3}{8}}\theta^{\frac{p-1}{4}}\right)^2}{\left(x^{\frac{p+3}{8}}\right)^2} \equiv \theta^{\frac{p-1}{2}} \equiv -1 \pmod{p}$$

since θ is a non-quadratic residue. Hence, $x^{\frac{p+3}{8}}$ and $x^{\frac{p+3}{8}}\theta^{\frac{p-1}{4}}$ are square roots of different values. Since they both are square roots of either x or $-x$, exactly one of them is a square root of x.