Cryptography, winter term 16/17: Sample solution to assignment 10

Cornelius Brand, Marc Roth

Exercise 10.1 (Be nice to your tutors and TAs, 1 Bonus Point)
Write the name and matriculation number of every author as well as number, time slot and the name of the tutor of your tutorial group on the first page of your solution. Furthermore write in english and staple *all* sheets of your solution.

Exercise 10.2 (Key-Exchange, 1 + 3 Points) Consider the following key-exchange protocol:

(1) Alice chooses $k, r \in \{0,1\}^n$ uniformly at random and sends $c := k \oplus r$ to Bob.

(2) Bob recieves c, chooses $s \in \{0,1\}^n$ uniformly at random and sends $d := c \oplus s$ to Alice.

(3) Alice recieves d and sends $e := r \oplus d$ to Bob.

(4) Finally, Alice outputs k and Bob outputs $e \oplus s$.

Now it is your task to analyse correctness and security of this protocol:

(a) Show that Alice and Bob output the same key.

(b) Decide whether the protocol is secure and prove your answer.

Solution 10.2 (Key-Exchange, 1 + 3 Points) Let k, r, c, s, d and e be defined as above.

(a) It holds that

$$e \oplus s = r \oplus d \oplus s = r \oplus c \oplus s \oplus s = r \oplus k \oplus r \oplus s \oplus s = k$$

(b) The protocol is not secure: Given a transcript (c, d, e), an attacker can compute

$$c \oplus d \oplus e = c \oplus d \oplus r \oplus d = c \oplus r = k \oplus r \oplus r = k\,.$$

This attack can clearly be done in polynomial time.

Exercise 10.3 (Public-Key Encryption, 3 + (2 + 3) Points) Let $\Pi = (\mathrm{Gen}, \mathrm{Enc}, \mathrm{Dec})$ be a public-key encryption scheme for single-bit messages.

(a) Assuming Π has no decryption error, show that there is an *unbounded* adversary $\mathcal{A}$[1] that, on input pk and $c \leftarrow \mathrm{Enc}_{\mathrm{pk}}(m)$ outputs m with probability 1.

[1] *Unbounded* means that $\mathcal{A}$ has no time constraints. In particular, the running time of $\mathcal{A}$ may be exponential or even worse.

(b) Assuming the length of every ciphertext obtained by Enc is bounded logarithmically[2] in the security parameter, show that Π is *not* CPA-secure. Proceed as follows:

(i) Prove that there exists a polynomial p such that for every $n \in \mathbb{N}$:

$$|\{x \in \{0,1\}^* \mid |x| \leq \log n\}| \leq p(n)$$

(ii) Consider the following ppt adversary $\mathcal{A}$:

(1) Given pk, $\mathcal{A}$ computes $c_0 = \mathrm{Enc}_{\mathrm{pk}}(0)$ and $c_1 = \mathrm{Enc}_{\mathrm{pk}}(1)$.

(2) Then $\mathcal{A}$ outputs $m_0 = 0$ and $m_1 = 1$.

(3) When $\mathcal{A}$ recieves c it checks whether $c = c_0$ or $c = c_1$. If this is the case, $\mathcal{A}$ outputs 0 or 1, respectively. Otherwise $\mathcal{A}$ outputs a random bit.

Use (i) to prove that $\mathcal{A}$ wins the PubK-experiment with probability non-negligible greater than $\frac{1}{2}$, that is, show that there is a polynomial q such that

$$\Pr[\mathrm{PubK}^{\mathrm{eav}}_{\mathcal{A},\Pi}(n) = 1] \geq \frac{1}{2} + \frac{1}{q(n)}\,.$$

Solution 10.3 (Public-Key Encryption, 3 + (2 + 3) Points)

(a) $\mathcal{A}$ is constructed as follows: On input pk and c, $\mathcal{A}$ simulates $\mathrm{Enc}_{\mathrm{pk}}(0)$ for every possible sequence of random bits r (used by Enc). If at least one of the simulations results in c then $\mathcal{A}$ outputs 0. Otherwise it outputs 1. We claim that $\mathcal{A}$ outputs the correct bit with probability 1:
First assume that $m = 0$. In this case one of the simulations has to result in c. (Otherwise c would not be an encryption of 0). If $m = 1$, no simulation can result in c as Π has no decryption error.

(b) We first give the intuitive reason why the claim should hold: If the length of the ciphertext is bounded logarithmically in the security parameter, then the size of the set of all ciphertexts of 0 and 1, respectively, is bounded by a polynomial. It follows that the probability of 0 (or 1 respectively) being mapped to the *most likely* ciphertext by Enc is non-negligible.

(i) It holds that

$$|\{x \in \{0,1\}^* \mid |x| \leq \log n\}| = \sum_{i=0}^{\log n} \left|\{0,1\}^i\right| = \sum_{i=0}^{\log n} 2^i = 2^{\log n+1} - 1 = 2n - 1 =: p(n)$$

where the third step follows from the closed form of the geometric series.

(ii) We claim that $\mathcal{A}$ wins with non-negligible probability. For the proof, let $C(b)$ denote the set of all possible ciphertexts of the message b. By assumption there exists a polynomial p such that $|C(b)| \leq p(n)$ where n is the security parameter.

[2] That is, the length of every ciphertext is $\leq \log n$.

It follows that there exists $c_b^* \in |C(b)|$ such that $\Pr[\mathrm{Enc}_{\mathrm{pk}}(b) = c_b^*] \geq \frac{1}{p(n)}$, for both $b = 0$ and $b = 1$ (c_b could be the most likely ciphertext as explained above). By law of total probability it holds that:

$$\begin{aligned}
\Pr[\mathcal{A} \text{ outputs } 0 | c \leftarrow \mathrm{Enc}_{\mathrm{pk}}(0)] &\geq 1 \cdot \Pr[c = c_0] + \frac{1}{2} \cdot \Pr[c \neq c_0] \\
&= \frac{1}{2} \cdot \Pr[c = c_0] + \frac{1}{2} \cdot \Pr[c = c_0] + \frac{1}{2} \cdot \Pr[c \neq c_0] \\
&= \frac{1}{2} \cdot \Pr[c = c_0] + \frac{1}{2} \\
&\geq \frac{1}{2} \cdot \Pr[c = c_0 = c_0^*] + \frac{1}{2} \\
&\geq \frac{1}{2} \cdot \frac{1}{p(n)^2} + \frac{1}{2}
\end{aligned}$$

Similary it holds that

$$\Pr[\mathcal{A} \text{ outputs } 1 | c \leftarrow \mathrm{Enc}_{\mathrm{pk}}(1)] \geq \frac{1}{2} \cdot \frac{1}{p(n)^2} + \frac{1}{2}$$

We conclude that

$$\begin{aligned}
\Pr[\mathrm{PubK}^{\mathrm{eav}}_{\mathcal{A},\Pi}(n) = 1] &= \frac{1}{2}\Pr[\mathcal{A} \text{ outputs } 1 | c \leftarrow \mathrm{Enc}_{\mathrm{pk}}(1)] + \frac{1}{2}\Pr[\mathcal{A} \text{ outputs } 0 | c \leftarrow \mathrm{Enc}_{\mathrm{pk}}(0)] \\
&\geq \frac{1}{2} \cdot 2 \cdot \left(\frac{1}{2} \cdot \frac{1}{p(n)^2} + \frac{1}{2} \right) \\
&= \frac{1}{2} + \frac{1}{2p(n)^2}
\end{aligned}$$

Exercise 10.4 (Number theory, 4 Points + 4 Bonus Points)

(a) Recall the definition of the Euler totient function $\phi(N) := |\mathbb{Z}_N^*|$, i.e. the number of integers between 1 and N that are coprime to N. In this exercise, you will prove an explicit formula: If N has the prime decomposition $p_1^{e_1} \dots p_t^{e_t}$, then

$$\phi(N) = N \cdot \prod_{i=1}^{t} (1 - \frac{1}{p_i}) \,.$$

You can do this as follows: First, prove the formula for ϕ on prime powers, i.e., if p is a prime and $k \geq 1$, show that $\phi(p^k) = p^k(1 - \frac{1}{p})$. Then, use the Chinese Remainder Theorem to argue about the value of $\phi(N)$ for arbitrary N (recall that you can decompose N into powers of distinct primes).

We state the exact version of the theorem you can use to solve this exercise:

Chinese Remainer Theorem. *Let $m_1, \dots, m_t$ be pairwise coprime (i.e., not necessarily prime) integers, and let $m = m_1 \cdots m_t$ be their product. Then, there is an isomorphism between the (multiplicative) groups*

$$\mathbb{Z}_m^* \cong \mathbb{Z}_{m_1}^* \times \dots \times \mathbb{Z}_{m_t}^* \,.$$

In particular, they have the same number of elements.

(b) **Bonus:** Consider the elliptic curve

$$E \ : \ y^2 = x^3 + 3x + 3$$

over $\mathbb{Z}_{11}$. How many points does $E(\mathbb{Z}_{11})$ have? Also, compute $(0,5) + (7,2)$ and $(0,5) + (0,6)$ in the group defined by E.

Solution 10.4 (Number theory, 4 Points + 4 Bonus Points) (a) Let $N = p^k$ for some prime p. Now, the numbers that have a common divisor with p^k are exactly the multiples of p less than p^k, i.e. $1, p, 2p, 3p, \ldots, p^{k-1}p$, of which there are p^{k-1} many. Therefore, the numbers *not* having a common divisor with p^k, i.e. that are coprime to p^k, are $p^k - p^{k-1} = p^k(1 - \frac{1}{p})$ many. This shows $\phi(p^k) = p^k(1 - \frac{1}{p})$.

Now, let N have the prime decomposition $p_1^{e_1} \ldots p_t^{e_t}$. Hence, $\mathbb{Z}_N = \mathbb{Z}_{p_1^{e_1} \ldots p_t^{e_t}}$. Because any two primes are coprime, we can use the Chinese Remainder Theorem, which gives an isomorphism of multiplicative groups:

$$\mathbb{Z}_N^* \cong \mathbb{Z}_{p_1^{e_1}}^* \times \mathbb{Z}_{p_2^{e_2}}^* \times \ldots \times \mathbb{Z}_{p_t^{e_t}}^*$$

As usual, the order of the product of finitely many groups is the product of the orders of the groups. Together with the formula for prime powers that we proved before, this yields

$$\phi(N) = |\mathbb{Z}_N^*| = \prod_{i=1}^{t} |\mathbb{Z}_{p_i^{e_i}}^*| = \prod_{i=1}^{t} p_i^{e_i}(1 - \frac{1}{p_i}) = \left(\prod_{i=1}^{t} p_i^{e_i}\right)\left(\prod_{i=1}^{t}(1 - \frac{1}{p_i})\right) = N \cdot \prod_{i=1}^{t}(1 - \frac{1}{p_i})$$

(b) The quadratic residues modulo 11 are $1, 3, 4, 5, 9$.

- $f(0) = 3$, a quadratic residue modulo 11, with square roots 6 and 5. Hence $(0,6), (0,5) \in E(\mathbb{Z}_{11})$.
- $f(1) = 7$, a quadratic non-residue modulo 11.
- $f(2) = 6$, a quadratic non-residue modulo 11.
- $f(3) = 6$, a quadratic non-residue modulo 11.
- $f(4) = 2$, a quadratic non-residue modulo 11.
- $f(5) = 0$, so we obtain the single point $(5,0) \in E(\mathbb{Z}_{11})$.
- $f(6) = 6$, a quadratic non-residue modulo 11.
- $f(7) = 4$, a quadratic residue modulo 11, with square roots 2 and 9. Hence $(7,2), (7,9) \in E(\mathbb{Z}_{11})$.
- $f(8) = 0$, so we obtain the single point $(8,0) \in E(\mathbb{Z}_{11})$.
- $f(9) = 0$, so we obtain the single point $(9,0) \in E(\mathbb{Z}_{11})$.
- $f(10) = 10$, a quadratic non-residue modulo 11.

Thus, $E(\mathbb{Z}_{11}) = \{(0,6),(0,5),(5,0),(7,2),(7,9),(8,0),(9,0)\}$, which are 7 points.

To compute $P = P_1 + P_2 = (0,5) + (7,2)$, we calculate $m = 8 \cdot 8 = -2$ modulo 11. By definition, if $P = (x,y)$, then $x = m^2 - 7 = 4 - 7 = 8$ modulo 11, and since the only point that comes into question (from the list above) with this x-coordinate is $(8,0)$, we conclude that $P = (8,0)$. For $(0,5)+(0,6)$, observe that they have the same x-coordinate but distinct y-coordinates, hence their sum vanishes and $(0,5) + (0,6) = \mathcal{O}$, the point at infinity.