ETH Zürich, D-INFK
Spring 2018

Prof. Ueli Maurer
Christian Badertscher
Fabio Banfi

Cryptography Foundations
Solution Exercise 2

2.1 Block Ciphers in ECB and CBC Mode

a) If a block cipher is used in ECB mode, the encryption of two equal n-bit blocks (aligned at n-bit boundaries) of the plaintext will yield the same n-bit blocks in the ciphertext. This can be hazardous for security in applications: If the plaintext encodes a bitmap with a schematic image (think of Mickey Mouse ears), then most of the n-bit blocks will either be 0^n or 1^n, so the ciphertext will mostly consist of the corresponding blocks $F(0^n, k)$ and $F(1^n, k)$. In fact, if the resulting ciphertext is drawn as a bitmap, then the schematic structure will still be visible.[1]

b) The problem here is that the same message is always encrypted to the same ciphertext. Hence, message repetitions can be detected. To see that this is a problem, consider the following scenario as an example: Assume a trader encrypts either "buy" or "sell" and an attacker observes the resulting ciphertexts. If the attacker can find out whether something was bought or sold at some point, she can learn all future messages by simply comparing the ciphertext to the one that was sent before the known transaction.

c) Let F_k^{-1} be the inverse function of $F(\cdot, k)$, which exists since $F(\cdot, k)$ is a permutation. To decrypt the i-th block of a ciphertext $c = c_0|\dots|c_\ell$, compute

$$F_k^{-1}(c_i) \oplus c_{i-1} = F_k^{-1}(F(m_i \oplus c_{i-1}, k)) \oplus c_{i-1} = m_i \oplus c_{i-1} \oplus c_{i-1} = m_i.$$

d) To decrypt the i-th block of a ciphertext $c = r|c_1|\dots|c_\ell$, compute

$$c_i \oplus F(r|\langle i \rangle, k) = \big(m_i \oplus F(r|\langle i \rangle, k)\big) \oplus F(r|\langle i \rangle, k) = m_i.$$

Note that F_k^{-1} is not needed for the decryption.

Without a nonce, the same message is always encrypted to the same ciphertext, and we have the same problem as in subtask b). The important property of the nonce is not that it is uniform, but that it does not repeat (thus the name). Choosing a uniform nonce is a straightforward way to ensure that nonce repetitions only occur with small probability. If the sender and receiver can keep state, the nonce can be replaced by keeping the value of the counter between messages can continue counting from there.

2.2 Information Theoretically Secure Message Authentication

In the 1-message MAC-forgery game for f, the adversary can obtain the tags for up to one message of her choice. Hence, there are two ways to win the game: One option for the adversary is to ask for the tag for a message m and then submit a message $m' \neq m$ together with a valid tag for m' (this is called a *substitution attack*). The other option is directly submit a message m together with a valid tag for this message (this is called *impersonation attack*).

[1] Such an example can be found here:
`http://en.wikipedia.org/wiki/Block_cipher_modes_of_operation#Electronic_Codebook_.28ECB.29`.

In the MAC-forgery game, the key K is uniformly distributed over $\mathcal{K}$. This induces for each message $m \in \mathcal{M}$ a random variable $f(m, K)$, which corresponds to the tag for m. The MACs we provide in all subtasks will have the following two properties:

1. For each message $m \in \mathcal{M}$, the tag $f(m, K)$ is uniformly distributed. This guarantees that the success probability of an impersonation attack is bounded by $1/|\mathcal{T}|$.
2. For two different messages $m, m' \in \mathcal{M}$, the tags $f(m, K)$ and $f(m', K)$ are independent. This ensures that asking for a valid tag for the message m does not help the adversary to guess the tag for m'. Together with the property above, we thus have that the success probability of a substitution attack is also bounded by $1/|\mathcal{T}|$.

If both properties hold, the winning probability of any adversary in the 1-message MAC-forgery game is hence bounded by $1/|\mathcal{T}|$.

a) Consider the MAC $f\colon \{0,1\} \times \{0,1\}^n \to \{0,1\}^{\frac{n}{2}}$ given by

$$f\big(m, (k_1, \ldots, k_n)\big) = \begin{cases} (k_1, \ldots, k_{\frac{n}{2}}), & m = 0 \\ (k_{\frac{n}{2}+1}, \ldots, k_n), & m = 1. \end{cases}$$

Since the key is uniformly distributed, we obviously have that the tags for both possible messages are uniform, too. Furthermore, the tags for the messages 0 and 1 are independent since the bits in the key are independent. Therefore, the winning probability of any adversary is upper bounded by $1/|\mathcal{T}| = 2^{-\frac{n}{2}}$.

b) We define the MAC function as follows:

$$f\big(m, (k_1, \ldots k_n)\big) = \begin{cases} (k_1, \ldots, k_{\frac{n}{2}}), & m = 0 \\ (k_{\frac{n}{2}+1}, \ldots, k_n), & m = 1 \\ (k_1 \oplus k_{\frac{n}{2}+1}, \ldots, k_{\frac{n}{2}} \oplus k_n), & m = 2. \end{cases}$$

As in subtask a), $f(0, K)$ and $f(1, K)$ are uniformly distributed and independent. By Exercise 1.2 we also have that $f(2, K)$ is uniformly distributed. Using Proposition 2.1 from the lecture notes, we can further deduce that $f(0, K)$ and $f(2, K)$ are independent: We have $f(2, K) = f(0, K) \oplus (k_{\frac{n}{2}+1}, \ldots, k_n)$. Since $(k_{\frac{n}{2}+1}, \ldots, k_n)$ is uniformly distributed and independent from $(k_1, \ldots, k_{\frac{n}{2}}) = f(0, K)$, we have that $f(0, K)$ and $f(2, K)$ are independent. Similarly, we obtain that $f(1, K)$ and $f(2, K)$ are independent. Hence, the winning probability of any adversary is again upper bounded by $2^{-\frac{n}{2}}$.

c) The idea from the previous subtask can be generalized to the message space $\{0,1\}^{\frac{n}{2}}$ in the following way. Let $\varphi : \{0,1\}^{\frac{n}{2}} \to F$ be a bijection from the bitstrings of length $\frac{n}{2}$ to the field elements $\mathbb{F} := \mathrm{GF}(2^{\frac{n}{2}})$ of $2^{\frac{n}{2}}$ elements. In fact, as you might remember from your discrete mathematics course, the standard example of $\mathrm{GF}(2^{\frac{n}{2}})$ consists of the polynomials of degree at most $\frac{n}{2} - 1$ over $\mathbb{Z}_2$. Thus we have the canonical bijection

$$\varphi\colon (b_1, \ldots, b_{\frac{n}{2}}) \mapsto b_1 X^{\frac{n}{2}-1} + \cdots + b_{\frac{n}{2}-1} X + b_{\frac{n}{2}}.$$

With the help of φ we can identify the message space with the set $\mathbb{F}$. The key space is identified with the set $\mathbb{F}^2$ via the bijection

$$\varphi \times \varphi\colon (k_1, \ldots, k_n) \mapsto \big(\varphi(k_1, \ldots, k_{\frac{n}{2}}), \varphi(k_{\frac{n}{2}+1}, \ldots, k_n)\big).$$

Now let the MAC $f\colon \mathbb{F} \times \mathbb{F}^2 \to \mathbb{F}$ be given by

$$f\big(m, (k_1, k_2)\big) = k_1 m + k_2.$$

Since the second half K_2 of the key $K = K_1|K_2$ is uniformly distributed, it follows from Exercise 1.2 that, for each message $m \in \mathbb{F}$, the tag $f(m, K)$ is uniformly distributed. Moreover, we have for $m, m' \in \mathcal{M}$, $m \neq m'$,

$$\begin{aligned}\Pr[f(m,K) = t \ \wedge \ f(m',K) = t'] &= \Pr[K_1 m + K_2 = t \ \wedge \ K_1 m' + K_2 = t'] \\ &= \Pr\left[\begin{pmatrix} m & 1 \\ m' & 1 \end{pmatrix} \cdot \begin{pmatrix} K_1 \\ K_2 \end{pmatrix} = \begin{pmatrix} t \\ t' \end{pmatrix}\right] \\ &= 2^{-n} \\ &= \Pr[f(m,K) = t] \cdot \Pr[f(m',K) = t'],\end{aligned}$$

where we have used in the third step that there is exactly one pair (k_1, k_2) for which the equation holds, which follows from the fact that the determinant of the matrix is $m' - m \neq 0$. Hence, $f(m, K)$ and $f(m', K)$ are independent, and we again obtain that the winning probability of any adversary is upper bounded by $2^{-\frac{n}{2}}$.

2.3 MAC Construction

a) The real and the ideal settings are depicted below.

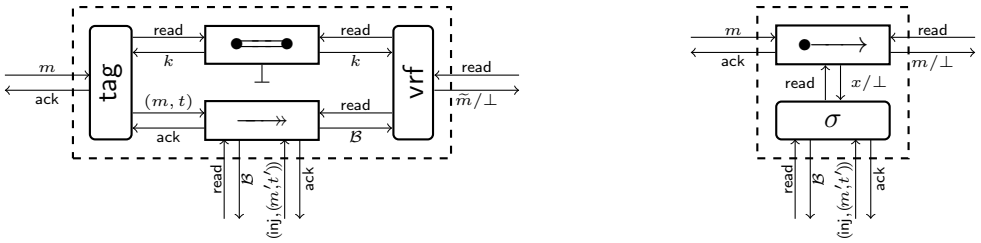

b) The converters tag, vrf, and σ are described below.

The converter tag
(for key space $\mathcal{M}$)

On the first input $x \in \mathcal{M}$ at outside interface, output read at inside sub-interface attached to •═══•, obtain key k at this inside sub-interface, and then set $t := f(x, k)$ and output (x, t) at inside sub-interface attached to ——». Ignore all subsequent inputs.

The converter vrf
(for key space $\mathcal{M}$)

On input read at outside interface, output read at inside sub-interface attached to •═══•, and obtain key k at this inside sub-interface. Then output read at inside sub-interface attached to ——», obtain the multiset $\mathcal{B}$ at the same inside sub-interface, and then:

- If $\mathcal{B} = \emptyset$, output $\perp$ at outside interface.
- If otherwise $\mathcal{B} \neq \emptyset$, then go through each element $(x, t) \in \mathcal{B}$ (according to some specified ordering), and do the following:

  (1) If $f(x, k) \neq t$, continue.

  (2) If $f(x, k) = t$, stop the loop and output x at outside interface.

  If no element satisfying the condition in (2) was found, then output $\perp$ at outside interface.

The converter σ
(for key space $\mathcal{M}$)

Variable $k \in \mathcal{K}$: Initially chosen uniformly at random from $\mathcal{K}$.
Multiset $\mathcal{B}$ on $\mathcal{M}'$: Initialized to $\emptyset$.

- On input read at outside interface, output read at inside interface (attached to $\bullet\longrightarrow$), and obtain value $x \in \mathcal{M} \cup \{\perp\}$. Then:
  - If $x = \perp$, output $\mathcal{B}$ at outside interface.
  - If otherwise $x \neq \perp$ ($x \in \mathcal{M}$), then compute $t := f(x, k)$, and if $(x, t) \notin \mathcal{B}$, then set $\mathcal{B} := \mathcal{B} \uplus \{(x, t)\}$. Finally output $\mathcal{B}$ at outside interface.
- On input $(\mathsf{inj}, (x', t'))$, just set $\mathcal{B} := \mathcal{B} \uplus \{(x, t)\}$ (in particular, do not interact with $\bullet\longrightarrow$).

c) The solution to this tasks consists of three steps. First, we have to be convinced that the simulator is perfect until a distinguisher inputs a valid tag t' for a message m' that was never input at interface A. A distinguisher can submit $q \leq q_E$ queries to provoke such an event. Second, we need to bound the probability that this event occurs. Finally, we need to be convinced that the distinguishing advantage is indeed upper bounded by this error probability given the previous two steps.

- We convince ourselves by observing that, in both the real and ideal system (with simulator), the message-tag pair (m, t) output at Eve's interface E after input m at interface A is computed using the same function f and using a uniformly random chosen key. Also, as long as Eve fails to forge a MAC for a new message, it holds that a message m is output at Bob's interface if and only if it was input at Alice's interface before.[2]
- Let F_i be the following event (defined in both the real world and the ideal world with simulator): D inputs a pair (m', t) at interface E as its i-th input message and the following conditions hold:
  (1) m' is not equal to the input message at interface A
  (2) This is the first input pair (m', t') by D such that $f(m', k) = t'$ (where k is the key defined in the respective random experiment).

  We first consider the probability that a distinguisher fails to provoke any forgery up to (and including) his $(i-1)$-th input conditioned on the knowledge of one correct message-tag pair that we denote by (m_A, t_A) (corresponding to a message m_A input by Alice). This probability is the same for both worlds, since they are identical random experiments up to the first successful forgery. We have

$$\begin{aligned} p_{i-1} &:= \Pr[f(m_1, K) \neq t_1 \wedge \cdots \wedge f(m_{i-1}, K) \neq t_{i-1} \mid f(m_A, K) = t_A] \\ &= \Pr[K_1 m_1 + K_2 \neq t_1 \wedge \cdots \wedge K_1 m_{i-1} + K_2 \neq t_{i-1} \mid K_1 m_A + K_2 = t_A] \\ &= \Pr\left[K_1 \in \mathbb{F} \setminus \left\{\frac{t_1 - t_A}{m_1 - m_A}, \ldots, \frac{t_{i-1} - t_A}{m_i - m_A}\right\}\right] = \frac{|S^*|}{|\mathbb{F}|}, \end{aligned}$$

  where $S^* := \mathbb{F} \setminus \{\frac{t_1 - t_A}{m_1 - m_A}, \ldots, \frac{t_{i-1} - t_A}{m_i - m_A}\}$ is used as shorthand notation to denote the restricted key space after having issued a sequence of invalid message-tag pairs.

  This allows us to compute the probability that the distinguisher provides a successful

[2] Later in the lecture, we will call this *game-equivalence* of two systems.

forgery for the first time upon query i as follows.

$$\begin{aligned}
&\Pr[F_i \mid f(m_A, K) = t_A] \\
&= p_{i-1} \cdot \Pr[f(m', K) = t' \mid f(m_1, K) \neq t_1, \ldots, f(m_{i-1}, K) \neq t_{i-1}, f(m_A, K) = t_A] \\
&= p_{i-1} \Pr[K_1 m' + K_2 = t' \mid K_1 m_1 + K_2 \neq t_1, \ldots, K_1 m_{i-1} + K_2 \neq t_{i-1}, K_1 m_A + K_2 = t_A] \\
&\leq \frac{|S^*|}{|\mathbb{F}|} \cdot \frac{1}{|S^*|} = \frac{1}{|\mathbb{F}|},
\end{aligned}$$

where the last inequality follows from inserting p_{i-1} and observing that providing a new (possible) forgery (m', t') given the valid pair (m_A, t_A) is equivalent to the problem of guessing K_1. Since K_1 is uniformly distributed, K_1 conditioned on the past invalid message-tag pairs is uniformly distributed over the restricted set S^* and hence the upper bound on the guessing probability of K_1.

A similar computation can be made not conditioned on Alice's input. In this case we have

$$\begin{aligned}
p'_{i-1} &:= \Pr[K_1 m_1 + K_2 \neq t_1 \wedge \cdots \wedge K_1 m_{i-1} + K_2 \neq t_{i-1}] \\
&= \Pr[(K_1, K_2) \in \mathbb{F}^2 \setminus \{(k_1, k_2) \in \mathbb{F}^2 \mid k_1 m_1 + k_2 = t_1, \ldots k_1 m_{i-1} + k_2 = t_{i-1}\}] = \frac{|S'^*|}{|\mathbb{F}|^2},
\end{aligned}$$

but for $S'^* := \mathbb{F}^2 \setminus \{(k_1, k_2) \in F^2 \mid k_1 m_1 + k_2 = t_1, \ldots k_1 m_{i-1} + k_2 = t_{i-1}\}$. We conclude

$$\Pr[F_i] \leq \frac{|S'^*|}{|\mathbb{F}|^2} \cdot \frac{|\mathbb{F}|}{|S'^*|} = \frac{1}{|\mathbb{F}|},$$

where $|\mathbb{F}|/|S'^*|$ is the probability that the provided pair (m', t') is a valid forgery, which is equivalent to guessing one of the $|\mathbb{F}|$ pairs $(k_1, k_2) \in S^*$ that satisfies the relation $k_1 m' + k_2 = t'$. Since (K_1, K_2) is uniformly distributed in $\mathbb{F}^2$, these variables, conditioned on the sequence of the previous invalid message-tag pairs, are distributed uniformly over the restricted subset S'^*.

- The last step of the argumentation is to argue that the probability of at least one valid forgery is an upper bound on the distinguishing advantage. Note that thanks to the analysis in step two of this proof, we see that we do not have to distinguish the cases whether we condition on an input message by Alice or not (the maximum is just $1/|\mathbb{F}|$). So we omit the explicit condition on an input message by Alice in the following argument (i.e., both cases are identical to analyze).
  Define Forge $:= F_1 \cup \cdots \cup F_q$. By the union bound we have that

$$\Pr[\text{Forge}] = \Pr[F_1 \cup \cdots \cup F_q] \leq \sum_{i=1}^{q} \Pr[F_i] \leq \frac{q}{|\mathbb{F}|} \leq \frac{q_E}{|\mathbb{F}|}. \tag{1}$$

In the following, we denote by ¬Forge the complementary event that no forgery occurs.
Let us define $\mathbf{R} := \mathsf{tag}^A\, \mathsf{vrf}^B[\bullet\!\!=\!\!=\!\!\bullet, \longrightarrow\!\!\!\!\!\!\rightarrow]$ and $\mathbf{S} := \sigma^E \bullet\!\!\longrightarrow$. We have

$$\begin{aligned}
\Delta^D(\mathbf{R}, \mathbf{S}) &= \Pr^{D\mathbf{S}}[Z = 1] - \Pr^{D\mathbf{R}}[Z = 1] \\
&= \Pr^{D\mathbf{S}}[Z = 1 \ \wedge \ \text{Forge}] + \Pr^{D\mathbf{S}}[Z = 1 \ \wedge \ \neg\text{Forge}] \\
&\quad - \Pr^{D\mathbf{R}}[Z = 1 \ \wedge \ \text{Forge}] - \Pr^{D\mathbf{R}}[Z = 1 \ \wedge \ \neg\text{Forge}] \\
&\overset{(1)}{=} \Pr^{D\mathbf{S}}[Z = 1 \ \wedge \ \text{Forge}] - \Pr^{D\mathbf{R}}[Z = 1 \ \wedge \ \text{Forge}] \\
&\overset{(2)}{\leq} \Pr^{D\mathbf{S}}[Z = 1 \ \wedge \ \text{Forge}] \overset{(3)}{\leq} \Pr^{D\mathbf{S}}[\text{Forge}] \leq \frac{q_E}{|\mathbb{F}|}.
\end{aligned}$$

Equality (1) follows from the fact the as long as no forgery was input, the two systems cannot be distinguished, thus $\Pr^{D\mathbf{S}}[Z = 1 \ \wedge \ \neg\text{Forge}] = \Pr^{D\mathbf{R}}[Z = 1 \ \wedge \ \neg\text{Forge}]$.

Inequality (2) follows by omitting a negative term and (3) since the event Forge is a superset of the event Forge $\wedge$ $(Z = 1)$. Finally, the last inequality holds due to Equation 1 and the fact that our analysis holds for both, the real and the ideal system (they are equivalent until the first successful forgery is provided by distinguisher D as argued in the first step of this proof). To conclude the statement, we just need $|\mathbb{F}| = 2^{\frac{n}{2}}$ which we have from Task 2.2.