Problem Set 6
Course **Secure Channels**
(Summer Term 2018)

Bauhaus-Universität Weimar, Chair of Media Security
Prof. Dr. Stefan Lucks, Eik List
URL: `http://www.uni-weimar.de/de/medien/professuren/mediensicherheit/teaching/`
Due Date: 28 June 2018, 11:00 AM, via email to `eik.list(at)uni-weimar.de`.

Task 1 – Online Encryption Schemes (8 Credits)
A tweakable block cipher $\widetilde{E} : \mathcal{K} \times \mathcal{T} \times \{0,1\}^n \to \{0,1\}^n$ is a family of keyed permutations s.t. $\widetilde{E}_K^T(\cdot)$ is a independent permutation over $\{0,1\}^n$ for every key-tweak pair (K,T). So, it adds a public tweak T to a classical block cipher.

Consider the following simplified version of McOE that takes a nonce N and encrypts a message $M = (M_1, \ldots, M_m)$ to a ciphertext $C = (C_1, \ldots, C_m)$. Z is not returned and kept secret. The chaining values $U_1 = N \oplus Z$ and later $U_i = M_{i-1} \oplus C_{i-1}$ are used as tweak for $\widetilde{E}$ as shown below.

a) Describe *either* an attack on in the RoR-OPRP-CPA model *or* explain why it is secure.

b) Now consider that the first call to $\widetilde{E}$ uses 0^n as message input and N as tweak. Again, describe *either* an attack on in the RoR-OPRP-CPA model *or* explain why this variant is secure.

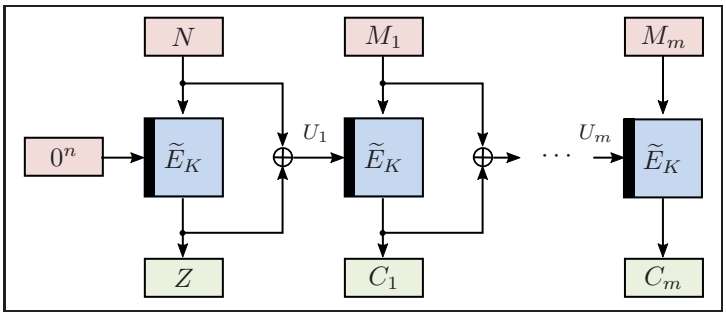

Task 2 – OAE2 (4 Credits)
Hoang et al. [3][1] proposed OAE2 as a better on-line authenticated encryption. Actually, their notion defines a segmented AE scheme that wraps a normal AE scheme internally. According to them, a segmented AE scheme define encryption and decryption as 3-tuples of algorithms each:

$$\begin{array}{lll} \mathcal{E}.\text{INIT}: & \mathcal{K} \times \mathcal{N} \to \mathcal{S} & \mathcal{D}.\text{INIT}: \quad \mathcal{K} \times \mathcal{N} \to \mathcal{S} \\ \mathcal{E}.\text{NEXT}: & \mathcal{S} \times \mathcal{A} \times \mathcal{M} \to \mathcal{C} \times \mathcal{S} & \mathcal{D}.\text{NEXT}: \quad \mathcal{S} \times \mathcal{A} \times \mathcal{C} \to (\mathcal{M} \times \mathcal{S}) \times \{\perp\} \\ \mathcal{E}.\text{LAST}: & \mathcal{S} \times \mathcal{A} \times \mathcal{M} \to \mathcal{C} & \mathcal{D}.\text{LAST}: \quad \mathcal{S} \times \mathcal{A} \times \mathcal{M} \to \mathcal{M} \times \{\perp\}. \end{array}$$

[1] Please mind the version, the paper might get updated.

11: **procedure** $\mathcal{E}.\text{INIT}(K, N)$
12: **return** $(K, N, 0)$

21: **procedure** $\mathcal{E}.\text{LAST}(S, A, M)$
22: $(K, V, d) \leftarrow S$
23: **if** $d = 0$ **then** $d \leftarrow 4$
24: **else** $d \leftarrow 2$
25: **return** $\mathbf{E}_K(V, \langle A, d \rangle, M)$

31: **procedure** $\mathcal{E}.\text{NEXT}(S, A, M)$
32: $(K, V, d) \leftarrow S$
33: $C \leftarrow \mathbf{E}_K(V, \langle A, d \rangle, M)$
34: **if** $M \geq n$ **then** $V \leftarrow (C[1..n] \oplus M[1..n])$
35: **else** $V \leftarrow \mathbf{E}_K(V, \langle A, 3 \rangle, M \,\|\, 0^n)[1..n]$
36: **return** $(C, (K, V, 1))$

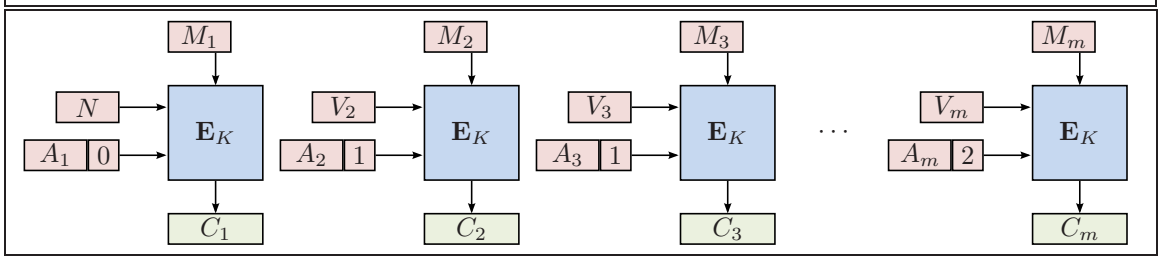

The **CHAIN** construction.

$\mathcal{K}$, $\mathcal{N}$, $\mathcal{S}$, $\mathcal{M}$, $\mathcal{C}$ denote key, nonce, state, message and ciphertext spaces, respectively. A message vector $\mathbf{M} = (M_1, \ldots, M_m)$ can be segmented arbitrarily. Each segment is authenticated-encrypted separately by a secure nonce-based AE scheme $(\mathbf{E}, \mathbf{D})$ to $\mathbf{C} = (C_1, \ldots, C_m)$. Each ciphertext segment contains its tag, so the $|C_i| = |M_i| + \tau$, i.e., they are τ bits longer than their respective messages. The decryption returns $\perp$ for invalid ciphertexts and, if C_i was invalid, all subsequent ciphertext segments C_j for $j > i$ for that message vector will also be decrypted to $\perp$.

They went on to propose the **CHAIN** construction. The encryption takes M_i, A_i, an IV V_i, and a domain parameter d. For the first segment, the encryption takes a nonce $N = V_1$ that must not repeat, and $d = 0$. The next segments $M_2, \ldots, M_{m-1}$ derive their IVs V_i from the previous segment and use domain $d = 1$. The last segment uses domain $d = 2$. If a message consists of a single segment, it uses domain $d = 4$. V_i is either the XOR of the first n bits of $M_{i-1} \oplus C_{i-1}$, or, if M_{i-1} is too short, it is derived from the encryption under domain $d = 3$.

This is similar to McOE; but take a look at the difference in the computation of chaining values, here in Lines 34 and 35. *Either* describe an efficient attack on the privacy of the **CHAIN** construction *or* explain in your own words why it is secure.

Task 3 – Stateful Authenticated Encryption (6 Credits)
Read autonomously into the levels of stateful AE in [2], Section 1.1 and **(a)** fill the table below.

	Protection against			
Level	Forgeries	Replay	Reordering	Dropping
Level 1				
Level 2				
Level 3				
Level 4				

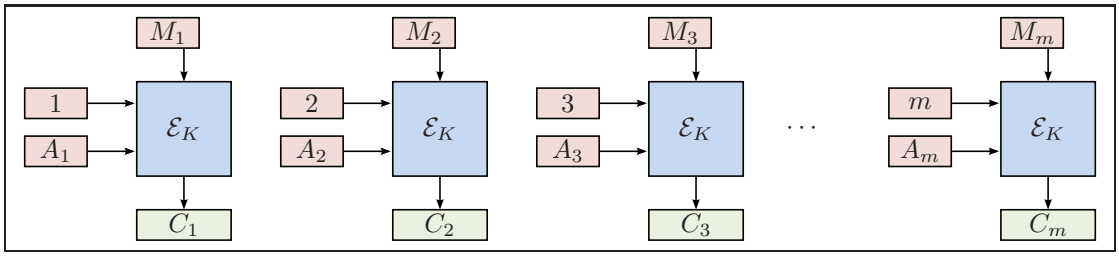

The STREAM′ scheme.

Then, consider the STREAM′ that encrypts a sequence of m messages. Here, the M_i's are messages each (not segments) and $(\mathcal{E}, \mathcal{D})$ is a secure nonce-based AE scheme. The nonce is a sequence number (= a counter) that is maintained by sender and receiver, and not by the user. Again, if the decryption of a message fails, all subsequent messages will also decrypt to $\perp$. Analyze which level(s) of stateful AE are fulfilled by this scheme.

References

[1] Mihir Bellare, Alexandra Boldyreva, Lars R. Knudsen, Chanathip Namprempre: Online Ciphers and the Hash-CBC Construction. CRYPTO 2001: 292-309. Full version (On-Line Ciphers and the Hash-CBC Constructions) at `https://eprint.iacr.org/2007/197.pdf`.

[2] Colin Boyd, Britta Hale, Stig Frode Mjølsnes, Douglas Stebila: From Stateless to Stateful: Generic Authentication and Authenticated Encryption Constructions with Application to TLS. CT-RSA 2016: 55-71. `https://eprint.iacr.org/2015/1150`.

[3] Viet Tung Hoang, Reza Reyhanitabar, Phillip Rogaway, Damian Vizár: Online Authenticated-Encryption and its Nonce-Reuse Misuse-Resistance. CRYPTO (1) 2015. Full version available at `https://eprint.iacr.org/2015/189`, version: 20180522:140429.

Definitions for On-line Security

Let $\mathcal{B} = \{0,1\}^n$ for positive integer n and let $x \twoheadleftarrow \mathcal{X}$ denote uniform independent random sampling of an element x from a given set $\mathcal{X}$. Let $\mathcal{B}^{\geq m} \stackrel{\text{def}}{=} \bigcup_{i=1}^{m} \mathcal{B}^i$.

Definition 1 (Length of Longest Common Prefix)**.** Given two arbitrary $M, M' \in \mathcal{B}^*$, let M_i denote the i-th block of M and M'_i the i-th block of M', for all $i \geq 1$. The length of the longest common $\mathcal{B}$-block prefix of M and M' is then given by

$$\mathsf{LLCP}_{\mathcal{B}}(M, M') \stackrel{\text{def}}{=} \max_i \left\{ \forall j \in 1, \ldots, i : M_i = M'_i \right\}.$$

For any two distinct m-block messages M and M' that share an exactly p-block common prefix $M_1 \parallel \ldots \parallel M_p = M'_1 \parallel \ldots \parallel M'_p$, the corresponding outputs $C = \pi(M)$ and $C' = \pi(M')$ satisfy $C_i = C'_i$ for all $i \in [1, p]$ and $p \leq p$, where π denotes an on-line permutation. However, it applies that $C_{p+1} \neq C'_{p+1}$ and all further blocks C_i and C'_i, with $i \in \{p+2, \ldots, m\}$, are *independent*. This behavior is defined by on-line permutations; we recall their definition briefly in the following.

Definition 2 (On-Line Permutation [1])**.** A mapping $\Pi_i : \mathcal{B}^{i-1} \times \mathcal{B} \to \mathcal{B}$ is called a family of indexed permutations iff $\Pi_i(I, \cdot)$ is a permutation over $\mathcal{B}$. for all $I \in \mathcal{B}^{i-1}$, and all $i \geq 1$. An on-line permutation $\Pi : \mathcal{B}^{\leq m} \to \mathcal{B}^{\leq m}$ is a composition of m permutations $\Pi_1 \cup \ldots \cup \Pi_m$, s. t. any ℓ-block input $M \in \mathcal{B}^i$ is mapped to an ℓ-block output $C \in \mathcal{B}^i$ by

$$C_i = \Pi_i(M_1, \ldots, M_{i-1}, M_i), \qquad \text{for all } i \in [1, \ldots, \ell].$$

We denote by $\mathsf{OPerm}(\mathcal{B})$ the set of all on-line permutations with a block space $\mathcal{B}$. Let $\mathcal{K}$, $\mathcal{N}$, $\mathcal{M}$, $\mathcal{C}$ denote non-empty sets of keys, nonces, messages, and ciphertexts, respectively. We restrict our interest to bit strings, i.e., $\mathcal{N}, \mathcal{M}, \mathcal{C} \subseteq \{0,1\}^*$. Moreover, we define $\mathcal{M} = \mathcal{C} = \mathcal{B}^*$ for some fixed integer n. A nonce-based on-line encryption scheme $\Pi = (\mathcal{E}, \mathcal{D})$ is a tuple of nonce-based encryption algorithm $\mathcal{E} : \mathcal{K} \times \mathcal{N} \times \mathcal{M} \to \mathcal{C}$ and deterministic decryption algorithm $\mathcal{D} : \mathcal{K} \times \mathcal{N} \times \mathcal{C} \to \mathcal{M}$. W.l.o.g., we assume correctness, i.e., for all (K, N, M), it holds that $\mathcal{D}_K^N(\mathcal{E}_K^N(M)) = M$, tidiness, i.e., for all (K, N, C), it holds that $\mathcal{E}_K^N(\mathcal{D}_K^N(C)) = C$.
A secure on-line encryption scheme is one that is indistinguishable from an on-line permutation. We recall briefly the RoR-OPRP-CPA and RoR-OPRP-CCA notions by Bellare et al. [1]

Definition 3 (RoR-OPRP-CPA and RoR-OPRP-CCA Security [1])**.** Let $\Pi = (\mathcal{E}, \mathcal{D})$ be an on-line encryption scheme with associated key space $\mathcal{K}$ and block space $\mathcal{B}$. Let $K \twoheadleftarrow \mathcal{K}$, and let $\mathbf{A}$ be an adversary on Π that has access to an oracle, and let further $\mathbf{A}'$ be an adversary on Π that has access to two oracles. Let $\pi \twoheadleftarrow \mathsf{OPerm}(\mathcal{B})$. Then, the RoR-OPRP-CPA advantage of $\mathbf{A}$ and the RoR-OPRP-CCA advantage of $\mathbf{A}'$ w.r.t. Π are defined as

$$\mathbf{Adv}_{\Pi}^{\text{RoR-OPRP-CPA}}(\mathbf{A}) \stackrel{\text{def}}{=} \underset{\mathbf{A}}{\Delta}\left(\mathcal{E}_K; \pi\right)(\mathbf{A}), \text{ and}$$
$$\mathbf{Adv}_{\Pi}^{\text{RoR-OPRP-CCA}}(\mathbf{A}) \stackrel{\text{def}}{=} \underset{\mathbf{A}}{\Delta}\left(\mathcal{E}_K, \mathcal{D}_K; \pi, \pi^{-1}\right)(\mathbf{A}'),$$

respectively.