A preliminary version of this paper appears in the proceedings of the *22nd International Conference on Information and Communications Security (ICICS 2020)*, DOI: 10.1007/978-3-030-61078-4_17.

Information-Theoretic Security of Cryptographic Channels

Marc Fischlin[1] Felix Günther[2] Philipp Muth[1]

[1] Department of Computer Science, TU Darmstadt, Darmstadt, Germany
[2] Department of Computer Science, ETH Zürich, Zürich, Switzerland

marc.fischlin@cryptoplexity.de
mail@felixguenther.info muth@seceng.informatik.tu-darmstadt.de

Abstract. We discuss the setting of information-theoretically secure channel protocols where confidentiality of transmitted data should hold against unbounded adversaries. We argue that there are two possible scenarios: One is that the adversary is currently bounded, but stores today's communication and tries to break confidentiality later when obtaining more computational power or time. We call channel protocols protecting against such attacks *future-secure*. The other scenario is that the adversary already has extremely strong computational powers and may try to use that power to break current executions. We call channels withstanding such stronger attacks *unconditionally-secure*.

We discuss how to instantiate both future-secure and unconditionally-secure channels. To this end we first establish according confidentiality and integrity notions, then prove the well-known composition theorem to also hold in the information-theoretic setting: Chosen-plaintext security of the channel protocol, together with ciphertext integrity, implies the stronger chosen-ciphertext notion. We discuss how to build future-secure channel protocols by combining computational message authentication schemes like HMAC with one-time pad encryption. Chosen-ciphertext security follows easily from the generalized composition theorem. We also show that using one-time pad encryption with the unconditionally-secure Carter-Wegman MACs we obtain an unconditionally-secure channel protocol.

1 Introduction

In today's information infrastructure the time intervals over which sensitive data are stored increase rapidly. Striking examples are digital tax data or electronic medical records which need to be kept for years or even decades according to legal stipulations, requiring also to uphold the involved individuals' right to privacy for such time periods. In some cases the protection time span is quasi indefinite, if one considers for example genetic data which descendants (partially) inherit from their ancestors.

The cryptographic challenge here is that the long-term protecting schemes must be able to withstand unexpected cryptanalytic advances, but also predictable advances in computational power. An adversary may store digital data and aim to break the underlying cryptographic scheme later with new methods or by pure advances in technology. Remarkably, this does not only hold for data at rest but also for data in transmission: An adversary may record encrypted communication today and try to break confidentiality tomorrow. If we talk about transmissions over unreliable networks then the adversary may also use additional means to attack schemes, such as omission, injection or modification of transmitted ciphertexts.

The above challenge is the starting point of our work. We consider security of cryptographic channels against potentially unbounded adversaries, denoted as information-theoretically secure channels.[1] The

[1] Our notion of (cryptographic) channels should not be confused with other concepts like Wyner's wire-tap channels [Wyn75] or other measures to generate information-theoretically secure keys from physical assumptions. We are interested in how to transmit data securely once the sender and the receiver already share a key.

question we address is what kind of channel security can we achieve in settings with unbounded adversaries, and how can we accomplish this.

1.1 Modeling Information-Theoretically Secure Channels

If we look at the long-term security of channel protocols, in order to completely rule out unforeseen cryptanalytic advancements, this boils down to unconditional security. In this context Shannon's famous result [Sha49] tells us that we need keying material as long as the cumulative size of transmitted messages which should be protected. Ensuring that sufficient keying material is available when required is beyond our scope; the most prominent option today would be to use quantum key distribution (QKD) [GND+19]. Clearly, this attaches a high-price tag to information-theoretic security in practical deployment. When securing high-stake data transmission, truly long-term security however is and will be called for, and hence ought to be formally understood. Focusing on the channel protocol, we make the simplifying assumption that sender and receiver readily have secure shared keys K available with each operation; our channel notions will allow to precisely quantify the amount of required keying material per operation.

For modeling unconditional security of channels we use a two-stage adversary model similar to the one introduced by Bindel et al. [BHMS17]. They consider signature-based public-key infrastructures and the question how security is affected by quantum adversaries. Among other, they distinguish adversaries which are classical when interacting with the certificate authority and gain quantum power only much later in the future, versus adversaries which have quantum capabilities even when interacting with the signer. The idea has also been adapted in subsequent works like [BBF+19].

In our setting we distinguish between adversaries which are bounded or unbounded in the first phase, during the channel protocol execution, but definitely become unbounded in the second phase, after the receiver closed the connection:

- For *future-secure* channels the first-stage adversary is bounded in computational resources when the channel protocol is running, but may store the communication data and later try to decrypt when having more computational power or more time.
- For *unconditionally-secure* channels the first-stage adversary already has extreme computational power when the channel protocol is executed, such that we need to protect against unbounded adversaries immediately.

In both cases we assume an active adversary which can tamper with the network communication, thereby capturing (and preventing) re-ordering and replay attacks. This in particular distinguishes our setting from prior works concerned with the unconditional security of *individual* messages (but without ordering requirements), e.g., aiming at everlasting privacy in e-voting [MN06].

1.2 Achieving Information-Theoretically Secure Channels

We next show how one can build future-secure and unconditionally-secure channel protocols. We follow the common paradigm to encrypt and authenticate the data in transmission. For encryption we need unconditional security for both channel types, because any break of confidentiality, during the protocol execution or afterwards, violates long-term secrecy of the data. This suggests to use the one-time pad encryption.

Authenticity, on the other hand, is a property which has to hold only during the channel's life time, in order to decide if a transmission comes from the expected sender. This is also remarked in [MSU13] where the authors combine quantum key distribution with short-term authentication methods. In our channel instantiation aiming at future security we can thus use computationally-secure authentication

methods like HMAC [BCK96]. For unconditionally-secure channels we need information-theoretically secure authentication schemes like Carter-Wegman MACs [WC81].

Before diving into the construction we first carefully adapt the classic composition theorem of Bellare and Namprempre [BN00] to the setting of information-theoretically secure channels: we show that an IND-CPA secure protocol which additionally provides INT-CTXT integrity of ciphertexts is also IND-CCA secure. As we will see, in our setting IND-CPA (even against unbounded adversaries) holds based on using one-time pad encryption; the composition result hence elegantly allows us to focus on establishing INT-CTXT (computationally or unconditionally) via appropriate authentication methods. This way, we obtain IND-CCA future-secure channels if we use computational authentication, and even IND-CCA unconditionally-secure channels if we use information-theoretically secure authentication.

We then give two concrete channel protocols, combining one-time pad encryption with computationally-secure MACs like HMAC, resp. with information-theoretically secure schemes like Carter-Wegman MACs. For the future-secure channel we use a counter to prevent repetition and out-of-order attacks, and show that the channel is IND-CPA secure and (computationally) INT-CTXT secure. Our general composition theorem therefore shows that the channel is IND-CCA future-secure. For the unconditional case it turns out that we do not need counters since we use a one-time key in each authentication step. We show, applying once more the composition theorem, that we achieve unconditional security of the channel if we apply Carter-Wegman MACs to the (plain) one-time pad encryption. Due to unforgeability of Carter-Wegman MACs linearly degrading with the number of transmitted messages, our results exhibit a noteworthy trade-off between the future- and unconditionally-secure constructions.

1.3 Further Related Work

Alternative approaches to unconditionally-secure encryption include limiting the adversary's memory capacity in the bounded-storage model [Mau92, CM97]. As such restriction may regularly not apply in practice for small-bandwidth, but highly-critical communication data, we in contrast consider fully-unbounded adversaries (and hence have to resort to the one-time pad for confidentiality).

Künzler et al. [KMR09] consider which functions are securely computable in the long-term scenario when one assumes short-term authenticated channels, i.e., channels which are only computationally secure during the computation. In a similar vein, Müller-Quade and Unruh [MU10] define a statistical version of the universal composition framework, enabling long-term security considerations. The work shows how to build commitments and zero-knowledge protocols in this setting, again assuming that secure channels are available.

2 Security of Information-Theoretically Secure Channels

2.1 Channels

We first define the notion of a channel protocol. It consists of an intialization step in which some shared key material K_I is generated, usually for authentication purposes, and the sender's and receiver's states are initialized. The OTKey algorithm lets the sender and receiver generate fresh key material, e.g., through authenticated quantum key distribution, to be used only once and in a pre-determined sequence (e.g., the order they are established in QKD). We do not specify in our abstract model how this is accomplished. Finally, the Send and Recv algorithms allow to process data for the communication.

Definition 2.1. *[Syntax of channels] A* channel $\mathsf{Ch} = (\mathsf{Init}, \mathsf{OTKey}, \mathsf{Send}, \mathsf{Recv})$ *with associated sending and receiving state space $\mathcal{S}_S$, resp. $\mathcal{S}_R$, message space $\mathcal{M} \subseteq \{0,1\}^{\leq M}$ for some maximum message length $M \in \mathbb{N}$, initialization key space $\mathcal{K}_{init} = \{0,1\}^{N_{init}}$ and per-message key space $\mathcal{K}_{msg} = \{0,1\}^N$ for some key*

lengths $N_{init}, N \in \mathbb{N}$, error space $\mathcal{E}$ with $\mathcal{E} \cap \{0,1\}^* = \emptyset$, consists of four efficient algorithms defined as follows.

- $\mathsf{Init}() \xrightarrow{\$} (K_I, \mathsf{st}_S, \mathsf{st}_R)$. *This probabilistic algorithm outputs an initial key $K_I \in \mathcal{K}_{init}$ and initial sending and receiving states $\mathsf{st}_S \in \mathcal{S}_S$, resp. $\mathsf{st}_R \in \mathcal{S}_R$.*
- $\mathsf{OTKey}() \xrightarrow{\$} K \in \{0,1\}^N$. *This algorithm generates the next per-message key K for both parties, to be used only once.*
- $\mathsf{Send}(\mathsf{st}_S, K_I, K, m) \xrightarrow{\$} (\mathsf{st}_S, c)$. *On input a sending state $\mathsf{st}_S \in \mathcal{S}_S$, an initial key $K_I \in \mathcal{K}_{init}$, a per-message key $K \in \mathcal{K}_{msg}$, and a message $m \in \mathcal{M}$, this (possibly) probabilistic algorithm outputs an updated state $\mathsf{st}_S \in \mathcal{S}_S$ and a ciphertext (or error symbol) $c \in \{0,1\}^* \cup \mathcal{E}$.*
- $\mathsf{Recv}(\mathsf{st}_R, K_I, K, c) \to (\mathsf{st}_R, m)$. *On input a receiving state $\mathsf{st}_R \in \mathcal{S}_R$, an initial key $K_I \in \mathcal{K}_{init}$, a per-message key $K \in \mathcal{K}_{msg}$, and a ciphertext $c \in \{0,1\}^*$, this deterministic algorithm outputs an updated state $\mathsf{st}_R \in \mathcal{S}_R$ and a message (or error symbol) $m \in \mathcal{M} \cup \mathcal{E}$.*

We say that a channel is *correct* if for any $i \in \mathbb{N}$, any $(K_I, \mathsf{st}_{S,0}, \mathsf{st}_{R,0}) \leftarrow\!\!\$\, \mathsf{Init}()$, any $(K_1, \ldots, K_i) \in (\mathcal{K}_{msg})^i$ with $K_j \leftarrow\!\!\$\, \mathsf{OTKey}()$ in sequence for $j = 1$ to $j = i$, any $(m_1, \ldots, m_i) \in \mathcal{M}^i$, any sequence $(\mathsf{st}_{S,1}, c_1) \leftarrow\!\!\$\, \mathsf{Send}(\mathsf{st}_{S,0}, K_I, K_1, m_1), \ldots, (\mathsf{st}_{S,i}, c_i) \leftarrow\!\!\$\, \mathsf{Send}(\mathsf{st}_{S,i-1}, K_I, K_i, m_i)$, and $(\mathsf{st}_{R,1}, m'_1) \leftarrow \mathsf{Recv}(\mathsf{st}_{R,0}, K_I, K_1, c_1), \ldots, (\mathsf{st}_{R,i}, m'_i) \leftarrow \mathsf{Recv}(\mathsf{st}_{R,i-1}, K_I, K_i, c_i)$, it holds that $(m_1, \ldots, m_i) = (m'_1, \ldots, m'_i)$.

2.2 Channel Security

Our core security notion follows the common ones for channels (or stateful authenticated encryption) by Bellare, Kohno, and Namprempre [BKN04], but combines confidentiality and integrity in a single game, following what is sometimes referred to as CCA3 security [Shr04]. The adversary $\mathcal{A}$ can repeatedly ask the sender (oracle) to encrypt one of two messages. The choice of which message to encrypt is based on a secret bit b which the adversary tries to predict eventually. On the receiver's side the adversary may submit arbitrary ciphertexts C in order to learn something about the bit b. Indeed, if the adversary manages to forge a ciphertext (decrypting to a non-error) on the receiver's side, either by creating a fresh valid ciphertext or by changing the order of the sender's ciphertexts, then we give the adversary enough information to predict b. The latter is achieved for a ciphertext forgery by returning the encapsulated message m if $b = 0$, and $\perp$ otherwise.

In more detail, the corresponding security experiment (in Figure 1) works as follows: The adversary can call the sending oracle O_{Send} about two equal-length messages m_0, m_1, then the sender encapsulates m_b (and updates its state st_S) and returns the ciphertext. We keep track of the order of ciphertexts by a counter i. The receiver's oracle O_{Recv} is more involved. When called with a ciphertext C it first increments its counter j and then decapsulates the message and updates its state st_R. There are now various cases to distinguish, relating to the question whether the ciphertext C is a forgery or not:

- If $j > i$ or $C \neq C_j$, i.e., if this is a new ciphertext or one which has not been produced by the sender as the i-th ciphertext before, then we say that the ciphertext sequences are not in-sync anymore. This is captured by a flag OUT-OF-SYNC.
- If we have reached an OUT-OF-SYNC situation, either in this call to O_{Recv} or an earlier one, then we provide the adversary with the received message in case $b = 0$. This enforces that, for a scheme to be secure, whenever the received ciphertext sequences goes out of sync, the output of Recv must be $\perp$, as otherwise it would be easily distinguishable from the case $b = 1$ always outputting $\perp$.

$\mathsf{Exp}^{\mathsf{ind\text{-}sfcca}}_{\mathsf{Ch}}(\mathcal{A})$

1 : $b \leftarrow\!\!\$\ \{0,1\}$
2 : $(K_I, \mathsf{st}_S, \mathsf{st}_R) \leftarrow\!\!\$\ \mathsf{Init}()$
3 : $K_1, K_2, K_3, \ldots \leftarrow\!\!\$\ \mathsf{OTKey}()$
4 : OUT-OF-SYNC $\leftarrow \mathsf{false}$
5 : $i, j \leftarrow 0$
6 : $\mathsf{st}_{\mathcal{A}} \leftarrow\!\!\$\ \mathcal{A}_1^{O_{\mathrm{Send}}(\mathsf{st}_S, K_I, \cdot, \cdot), O_{\mathrm{Recv}}(\mathsf{st}_S, K_I, \cdot)}()$
7 : $b' \leftarrow\!\!\$\ \mathcal{A}_2^{O_{\mathrm{Send}}(\mathsf{st}_S, K_I, \cdot, \cdot)}(\mathsf{st}_{\mathcal{A}})$
8 : **return** $b == b'$

$O_{\mathrm{Send}}(\mathsf{st}_S, K_I, m_0, m_1)$

1 : **assert** $|m_0| = |m_1|$
2 : $i \leftarrow i + 1$
3 : $(C_i, \mathsf{st}_S) \leftarrow \mathsf{Send}(\mathsf{st}_S, K_I, K_i, m_b)$
4 : **return** C_i

$O_{\mathrm{Recv}}(\mathsf{st}_R, K_I, C)$

1 : $j \leftarrow j + 1$
2 : $(m, \mathsf{st}_R) \leftarrow \mathsf{Recv}(\mathsf{st}_R, K_I, K_j, C)$
3 : **if** $(j > i$ **or** $C \neq C_j)$ **then**
4 : OUT-OF-SYNC $\leftarrow \mathsf{true}$
5 : **endif**
6 : **if** (OUT-OF-SYNC **and** $b == 0$) **then**
7 : **return** m
8 : **endif**
9 : **return** $\perp$

Figure 1: Experiment $\mathsf{Exp}^{\mathsf{ind\text{-}sfcca}}_{\mathsf{Ch}}(\mathcal{A})$

The overall goal of the adversary is to predict b, either by distinguishing the messages encapsulated by the sender, or by breaking integrity and learning about b through a receiver's reply.

To capture unconditionally-secure channels and future-secure ones in a single game we divide the adversary $\mathcal{A}$ in two phases, $\mathcal{A}_1$ and $\mathcal{A}_2$. In the first phase the adversary has access to both the sender and receiver oracle. In this first stage the adversary may still be bounded in running time (for future-secure channels), resp. already be unbounded (for unconditionally-secure channels). In the second stage the adversary is in both cases unbounded but can no longer access the receiver oracle. This allows us to model future-secure channels where $\mathcal{A}_1$ is restricted and the authentication only needs to be temporarily secure, and in the second phase of the unbounded $\mathcal{A}_2$ past and future sender messages remain confidential (but computational authentication may now be broken). For unconditionally-secure channels we allow already $\mathcal{A}_1$ to be unbounded such that $\mathcal{A}_2$ merely acts as a dummy.

We stress, however, that we do not formalize the notion of being bounded or unbounded in our concrete security analysis. Instead, we give reductions to underlying problems, e.g., if $\mathcal{A}_1$ breaks integrity of the channel then we break some underlying primitive with (roughly) the same running time. By this we get a reasonable security guarantee from computationally secure authentication schemes such as HMAC, as well as from unconditionally secure ones such as Carter-Wegman MACs.

Definition 2.2 (Chosen-Ciphertext Security)**.** *For an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ define its advantage in Experiment $\mathsf{Exp}^{\mathsf{ind\text{-}sfcca}}_{\mathsf{Ch}}(\mathcal{A})$ (Figure 1) as*

$$\mathsf{Adv}^{\mathsf{ind\text{-}sfcca}}_{\mathsf{Ch}}(\mathcal{A}) = \Pr\left[\mathsf{Exp}^{\mathsf{ind\text{-}sfcca}}_{\mathsf{Ch}}(\mathcal{A}) == \mathsf{true}\right]. \tag{1}$$

Note that for a secure channel we expect the advantage to be close to the pure guessing probability $\frac{1}{2}$.

We argue below that one can achieve the CCA notion by considering a weaker CPA requirement on confidentiality, and combining it with an integrity notion. The CPA indistinguishability game is identical to the CCA game but does not give the adversary access to the receiver oracle O_{Recv}, merging the two-stage adversary into a single one. The integrity experiment allows the adversary to see ciphertexts of chosen

$\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\mathsf{Ch}}(\mathcal{B})$

1 : $c \leftarrow\!\!\$\ \{0,1\}$
2 : $(K_I, \mathsf{st}_S, \mathsf{st}_R) \leftarrow\!\!\$\ \mathsf{Init}()$
3 : $K_1, K_2, K_3, \ldots \leftarrow\!\!\$\ \mathsf{OTKey}()$
4 : $i \leftarrow 0$
5 : $c' \leftarrow\!\!\$\ \mathcal{B}^{O_{\mathrm{Send}}(\mathsf{st}_S, K_I, \cdot, \cdot)}$
6 : **return** $c == c'$

$O_{\mathrm{Send}}(\mathsf{st}_S, K_I, K_i, m_0, m_1)$

1 : **assert** $|m_0| = |m_1|$
2 : $i \leftarrow i + 1$
3 : $(C_i, \mathsf{st}_S) \leftarrow \mathsf{Send}(\mathsf{st}_S, K_I, K_i, m_c)$
4 : **return** C_i

Figure 2: Experiment $\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\mathsf{Ch}}(\mathcal{B})$

$\mathsf{Exp}^{\mathsf{int\text{-}sfctxt}}_{\mathsf{Ch}}(\mathcal{I})$

1 : $(K_I, \mathsf{st}_S, \mathsf{st}_R) \leftarrow\!\!\$\ \mathsf{Init}()$
2 : $K_1, K_2, K_3, \ldots \leftarrow\!\!\$\ \mathsf{OTKey}()$
3 : OUT-OF-SYNC $\leftarrow$ false
4 : INT-BROKEN $\leftarrow$ false
5 : $i, j \leftarrow 0$
6 : $\mathcal{I}^{O_{\mathrm{Send}}(\mathsf{st}_S, K_I, \cdot), O_{\mathrm{Recv}}(\mathsf{st}_R, K_I, \cdot)}$
7 : **return** INT-BROKEN

$O_{\mathrm{Send}}(\mathsf{st}_S, K_I, m)$

1 : $i \leftarrow i + 1$
2 : $(C_i, \mathsf{st}_S) \leftarrow \mathsf{Send}(\mathsf{st}_S, K_I, K_i, m)$
3 : **return** C_i

$O_{\mathrm{Recv}}(\mathsf{st}_R, K_I, C)$

1 : $j \leftarrow j + 1$
2 : $(m, \mathsf{st}_R) \leftarrow \mathsf{Recv}(\mathsf{st}_R, K_I, K_j, C)$
3 : **if** $(j > i$ **or** $C \neq C_j)$ **then**
4 : OUT-OF-SYNC $\leftarrow$ true
5 : **endif**
6 : **if** $(m \neq \perp$ **and** OUT-OF-SYNC$)$
7 : INT-BROKEN $\leftarrow$ true
8 : **endif**
9 : **return** $\perp$

Figure 3: Experiment $\mathsf{Exp}^{\mathsf{int\text{-}sfctxt}}_{\mathsf{Ch}}(\mathcal{I})$

messages via oracle O_{Send}, and merely checks if the adversary manages to send a new or out-of-order ciphertext which decrypts correctly.

Definition 2.3 (Chosen-Plaintext Security)**.** *For an adversary $\mathcal{B}$ define the advantage in Experiment $\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\mathsf{Ch}}(\mathcal{B})$ (Figure 2) as:*

$$\mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{Ch}}(\mathcal{B}) = \Pr\left[\mathsf{Exp}^{\mathsf{ind\text{-}cpa}}_{\mathsf{Ch}}(\mathcal{B}) == \mathsf{true}\right]. \tag{2}$$

Finally we define integrity by demanding that the adversary is able to forge a valid ciphertext with negligible probability only:

Definition 2.4 (Ciphertext Integrity)**.** *For an adversary $\mathcal{I}$ define the advantage in Experiment $\mathsf{Exp}^{\mathsf{int\text{-}sfctxt}}_{\mathsf{Ch}}(\mathcal{I})$ (Figure 3) as:*

$$\mathsf{Adv}^{\mathsf{int\text{-}sfctxt}}_{\mathsf{Ch}}(\mathcal{I}) = \Pr\left[\mathsf{Exp}^{\mathsf{int\text{-}sfctxt}}_{\mathsf{Ch}}(\mathcal{I}) == \mathsf{true}\right]. \tag{3}$$

3 Composition Theorem

We next show that for any channel protocol Ch chosen-ciphertext security follows from chosen-plaintext security and integrity, similar to the composition result for classical channels [BKN04]. The security

$\mathsf{Exp}^{\mathsf{ind\text{-}sfcca}}_{\mathsf{Ch}}(\mathcal{A})$

$b \leftarrow\!\!\$\, \{0,1\}$
$(K_I, \mathsf{st}_S, \mathsf{st}_R) \leftarrow\!\!\$\, \mathsf{Init}()$
$K_1, K_2, K_3, \ldots \leftarrow\!\!\$\, \mathsf{OTKey}()$
OUT-OF-SYNC $\leftarrow \mathsf{false}$
$i, j \leftarrow 0$
$\mathsf{st}_{\mathcal{A}} \leftarrow\!\!\$\, \mathcal{A}_1^{O_{\mathrm{Send}}(\mathsf{st}_S, K_I, \cdot, \cdot), O_{\mathrm{Recv}}(\mathsf{st}_S, K_I, \cdot)}()$
$b' \leftarrow\!\!\$\, \mathcal{A}_2^{O_{\mathrm{Send}}(\mathsf{st}_S, K_I, \cdot, \cdot)}(\mathsf{st}_{\mathcal{A}})$
return $b == b'$

$O_{\mathrm{Send}}(\mathsf{st}_S, K_I, m_0, m_1)$

assert $|m_0| = |m_1|$
$i \leftarrow i + 1$
$(C_i, \mathsf{st}_S) \leftarrow \mathsf{Send}(\mathsf{st}_S, K_I, K_i, m_b)$
return C_i

$O_{\mathrm{Recv}}(\mathsf{st}_R, K_I, C)$

$j \leftarrow j + 1$
$(m, \mathsf{st}_R) \leftarrow \mathsf{Recv}(\mathsf{st}_R, K_I, K_j, C)$
if $(j > i$ **or** $C \neq C_j)$ **then**
  OUT-OF-SYNC $\leftarrow \mathsf{true}$
endif
if (OUT-OF-SYNC **and** $b == 0$) **then**
  return $\perp$ // instead of m
endif
return $\perp$

Figure 4: Modified receiver oracle experiment $\mathsf{Exp}^{\mathsf{ind\text{-}sfcca}}_{\mathsf{Ch}}(\mathcal{A})$ for $\textsc{Game}_1$ in the proof of Theorem 3.1.

reduction shows that the derived attackers $\mathcal{B}$ against ind-cpa and $\mathcal{I}$ against int-sfctxt have roughly the same running time characteristics as the adversary against ind-sfcca. In particular, if the first-stage adversary $\mathcal{A}_1$ against ind-sfcca is bounded (or unbounded) then so is the adversary $\mathcal{B}$ against ind-cpa and also $\mathcal{I}$.

Theorem 3.1 (ind-cpa $\wedge$ int-sfctxt $\Rightarrow$ ind-sfcca)**.** *For any channel protocol* Ch *and any* ind-sfcca *adversary* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$*, we can construct an* int-sfctxt *adversary* $\mathcal{I}$ *and an* ind-cpa *adversary adversary* $\mathcal{B}$ *such that*

$$\mathsf{Adv}^{\mathsf{ind\text{-}sfcca}}_{\mathsf{Ch}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{int\text{-}sfctxt}}_{\mathsf{Ch}}(\mathcal{I}) + \mathsf{Adv}^{\mathsf{ind\text{-}cpa}}_{\mathsf{Ch}}(\mathcal{B})\,. \tag{4}$$

Here, $\mathcal{B}$ *and* $\mathcal{I}$ *use approximately the same resources as* $\mathcal{A}_1$*.*

Proof. The proof follows the common game-hopping technique, where $\textsc{Game}_0$ denotes $\mathcal{A}$'s attack in experiment $\mathsf{Exp}^{\mathsf{ind\text{-}sfcca}}_{\mathsf{Ch}}$. In $\textsc{Game}_1$ we modify the receiver oracle O_{Recv} by letting it return $\perp$ instead of m for an out-of-sync query (for which in addition $b == 0$). This is depicted in Figure 4. The other steps of the experiment remain unchanged.

We argue that the difference of $\mathcal{A}$'s advantage between the two games lies in a potential first-stage query of $\mathcal{A}_1$ to the receiver oracle which returns a message $m \neq \perp$ in $\textsc{Game}_0$ but not in $\textsc{Game}_1$. We show that the probability of this happening is bounded by the integrity guarantees of the channel. To this end we build a reduction $\mathcal{I}$ mounting an attack according to experiment $\mathsf{Exp}^{\mathsf{int\text{-}sfctxt}}_{\mathsf{Ch}}(\mathcal{I})$. This algorithm $\mathcal{I}$ runs a black-box simulation of $\mathcal{A}_1$ (in $\textsc{Game}_0$). Any oracle call O_{Recv} of $\mathcal{A}_1$ is forwarded directly to the corresponding oracles of $\mathcal{I}$. Algorithm $\mathcal{I}$ initially also picks a random bit $b \leftarrow\!\!\$\, \{0,1\}$ and whenever $\mathcal{A}_1$ makes an oracle call m_0, m_1 to O_{Send}, then $\mathcal{I}$ first checks that $|m_0| = |m_1|$ and returns $\perp$ if not; else it forwards m_b to its own oracle O_{Send} to receive a ciphertext C_i. Algorithm $\mathcal{I}$ returns C_i in the simulation of $\mathcal{A}_1$. Algorithm $\mathcal{I}$ stops if $\mathcal{A}_1$ stops.

Note that the only difference between the two games from $\mathcal{A}$'s perspective is that $\textsc{Game}_0$, in case $b = 0$, returns an actual message m in a call to O_{Recv} if (a) $m \neq \perp$, and (b) OUT-OF-SYNC has been set to true (in this call or a previous call). This, however, means that all prerequisites in the O_{Recv} oracle of

the integrity experiment are satisfied, causing INT-BROKEN to become true and to make $\mathcal{I}$ win the game. Hence, any difference between the games can be bounded by the advantage against integrity.

A careful inspection of the modified O_{Recv} oracle now shows that this oracle always returns $\perp$ and only changes the state of the OUT-OF-SYNC variable. The latter only affects the O_{Recv} oracle itself. It follows that we can simulate this oracle by returning $\perp$ immediately for any query to O_{Recv}. Formally, this is a black-box simulation $\mathcal{B}$ of $\mathcal{A}$, where $\mathcal{B}$ relays all communication of $\mathcal{A}_1$ and $\mathcal{A}_2$ with oracle O_{Send} without modification, but returns $\perp$ to $\mathcal{A}_1$ for any call of $\mathcal{A}_1$ to O_{Recv}. Hence, in the next game hop we can eliminate the O_{Recv} oracle altogether, obtaining the CPA-game GAME_2. For this game we can bound the advantage by the CPA-security of the channel. □

4 Instantiations

In this section we discuss that instantiations combining the one-time pad encryption scheme with a computationally-secure MAC like HMAC, and with an unconditionally-secure one like Carter-Wegman MACs, provide future security, resp. unconditional security for the channel protocol. This of course requires additional steps to prevent replay attacks or protection against omission of ciphertext. For the computational case we choose here for the sake of concreteness a sequence number on the sender's and receiver's side. For the unconditional MAC we can omit the sequence number because we use a fresh key portion with each message anyway.

In both cases we use our composition result from Theorem 3.1 to argue security. $\mathsf{ind\text{-}cpa}$ security of the encryption scheme follows by the perfect secrecy of the one-time pad encryption and the fact that we use a fresh key for each ciphertext. This holds even against unbounded adversaries. It hence suffices to argue $\mathsf{int\text{-}sfctxt}$ security to conclude $\mathsf{ind\text{-}sfcca}$ security of the channel protocol. For this we need the strong unforgeability of the authentication algorithm.

4.1 Message Authentication

We first define message authentication codes and their security:

Definition 4.1 (Message Authentication Codes)**.** *A MAC scheme* $\mathsf{M} = (\mathsf{MKGen}, \mathsf{MAC}, \mathsf{Verify})$ *with associated message space* $\mathcal{M}$ *consists of three algorithms such that*

- $\mathsf{MKGen}() \xrightarrow{\$} K_{\mathsf{MAC}}$*. The key generation algorithm outputs a key* K_{MAC}*.*
- $\mathsf{MAC}(K_{\mathsf{MAC}}, m) \xrightarrow{\$} t$*. The (possibly probabilistic) MAC algorithm maps the key* K_{MAC} *and a message* $m \in \mathcal{M}$ *to a tag* t*.*
- $\mathsf{Verify}(K_{\mathsf{MAC}}, m, t) \to \{\mathsf{true}, \mathsf{false}\}$*. The verification algorithm takes a key, a message, and a tag as input, and outputs a decision.*

Correctness says that for all keys $K_{\mathsf{MAC}} \leftarrow \mathsf{MKGen}()$*, any message* $m \in \mathcal{M}$*, any tag* $t \leftarrow\!\!\$\, \mathsf{MAC}(K_{\mathsf{MAC}}, m)$ *we always have* $\mathsf{Verify}(K_{\mathsf{MAC}}, m, t) == \mathsf{true}$*.*

As mentioned earlier we require strong unforgeability of the MAC, demanding that is not only infeasible to find a valid tag for a previously untagged message, but that one also cannot find a different valid tag to a previously tagged message. Strong unforgeability follows for example for unforgeable MACs where authentication is deterministic and verification is done by recomputing the tag and checking the result against the given tag [BGM04].

$\underline{\mathsf{Exp}^{\mathsf{seuf\text{-}cma}}_{\mathsf{M}}(\mathcal{F})}$

1 : $K_{\mathsf{MAC}} \leftarrow\!\!\$\ \mathsf{MKGen}()$
2 : $Q \leftarrow \emptyset$
3 : $(m^*, t^*) \leftarrow\!\!\$\ \mathcal{F}^{O_{\mathrm{MAC}}(K_{\mathsf{MAC}},\cdot)}$
4 : **if** $\mathsf{Verify}(K_{\mathsf{MAC}}, m^*, t^*) == \mathsf{true}$
5 : **and** $(m^*, t^*) \notin Q$ **then**
6 : **return** true
7 : **else**
8 : **return** false

$\underline{O_{\mathrm{MAC}}(K_{\mathsf{MAC}}, m)}$

1 : $t \leftarrow\!\!\$\ \mathsf{MAC}(K_{\mathsf{MAC}}, m)$
2 : $Q \leftarrow Q \cup \{(m, t)\}$
3 : **return** t

Figure 5: Experiment $\mathsf{Exp}^{\mathsf{seuf\text{-}cma}}_{\mathsf{M}}(\mathcal{F})$

Definition 4.2 (Strong Unforgeability)**.** *For an adversary $\mathcal{F}$ define the advantage in Experiment $\mathsf{Exp}^{\mathsf{seuf\text{-}cma}}_{\mathsf{M}}(\mathcal{F})$ (Figure 5) as:*

$$\mathsf{Adv}^{\mathsf{seuf\text{-}cma}}_{\mathsf{M}}(\mathcal{F}) = \Pr\left[\mathsf{Exp}^{\mathsf{seuf\text{-}cma}}_{\mathsf{M}}(\mathcal{F}) == \mathsf{true}\right]. \tag{5}$$

We say that $\mathcal{F}$ is q-query bounded if $|Q| \leq q$ in the experiment.

Note that here adversary $\mathcal{F}$ may be bounded or unbounded in computation time. For unbounded $\mathcal{F}$ we usually assume that the adversary can only make a single query to oracle during the attack O_{MAC} and is thus 1-query bounded.

Two possible instantiations which are relevant for us here are the HMAC algorithm which provides strong unforgeability under reasonable assumptions about the compression function in the underlying hash function [BCK96, Bel15], and Carter-Wegman MACs which are unconditionally secure for 1-bounded adversaries [WC81] and also follow the verification-through-recomputation paradigm.

4.2 Future-Secure Channels

For a future-secure channel we define the sender and receiver algorithms as follows. We initialize counters for the sender and the receiver, respectively, both as zero. Algorithm Send first generates a ciphertext c via one-time pad encryption $\mathsf{OTP.Enc}(K, m) = m \oplus K$ under the fresh per-message key K. It then authenticates the ciphertext c, prepended with a fixed-length encoding of the counter value in st_S, under a computationally-secure MAC, using the steady key K_I.[2] The sender then increments its counter to be stored in the updated state st_S, and outputs the full ciphertext consisting of the OTP ciphertext and MAC tag.

The receiver algorithm Recv, when receiving a ciphertext $C = (c, t)$, first checks if the state st_R indicates a previous failed decryption or if the MAC is invalid. If so, Recv returns the error symbol $\perp$ and keeps this information in its state. Otherwise Recv decrypts the ciphertext part c with the per-message key, $\mathsf{OTP.Dec}(K, c) = c \oplus K$, increments the counter, and stores the updated value in its state st_R.

Construction 4.3 (Future-Secure Channel)**.** *Define the channelprotocol $\mathsf{FSCh} = (\mathsf{Init}, \mathsf{OTKey}, \mathsf{Send}, \mathsf{Recv})$ for message space $\mathcal{M} = \{0,1\}^{\leq M}$ and key space $\mathcal{K} = \{0,1\}^M$ by the algorithms in Figure 6.*

We next argue int-sfctxt security of the channel protocol, assuming that the underlying MAC scheme M is strongly unforgeable:

[2]Technically, the encoded counter restricts the number of messages that can be sent. If there are n bits reserved for the counter value then one can transmit at most 2^n messages. In practice this is not an issue and deployed channel protocools today commonly have such restrictions as well (e.g., TLS 1.3 [Res18] uses an $n = 64$ bit sequence number).

$\mathsf{Init}()$

1 : $K_I \leftarrow\!\!\$\ \mathsf{MKGen}()$
2 : $\mathsf{st}_S \leftarrow 0$
3 : $\mathsf{st}_R \leftarrow (\top, 0)$
4 : **return** $(K_I, \mathsf{st}_S, \mathsf{st}_R)$

$\mathsf{OTKey}()$

1 : $K \leftarrow\!\!\$\ \mathcal{K}$
2 : **return** K

$\mathsf{Send}\,(\mathsf{st}_S, K_I, K, m)$

1 : $c \leftarrow \mathsf{OTP.Enc}\,(K, m)$
2 : $t \leftarrow \mathsf{MAC}\,(K_I, \mathsf{st}_S || c)$
3 : $C \leftarrow (c, t)$
4 : $\mathsf{st}_S \leftarrow \mathsf{st}_S + 1$
5 : **return** (C, st_S)

$\mathsf{Recv}\,(\mathsf{st}_R, K_I, K, C)$

1 : parse $\mathsf{st}_R = (b, j)$ and $C = (c, t)$
2 : $d \leftarrow \mathsf{Verify}(K_I, c, t)$
3 : **if** $b == \perp$ **or** $d == \mathsf{false}$ **then**
4 : $\quad m \leftarrow \perp$
5 : $\quad \mathsf{st}_R \leftarrow (\perp, 0)$
6 : **else**
7 : $\quad m \leftarrow \mathsf{OTP.Dec}\,(K, c)$
8 : $\quad \mathsf{st}_R \leftarrow (\top, j + 1)$
9 : **fi**
10 : **return** (m, st_R)

Figure 6: Future-Secure Channel Protocol FSCh

Lemma 4.4. *For any* int-sfctxt *adversary* $\mathcal{I}$ *there exists an adversary* $\mathcal{F}$ *such that*

$$\mathsf{Adv}^{\mathsf{int\text{-}sfctxt}}_{\mathsf{FSCh}}(\mathcal{I}) \leq \mathsf{Adv}^{\mathsf{seuf\text{-}cma}}_{\mathsf{M}}(\mathcal{F})\,. \tag{6}$$

Furthermore, $\mathcal{F}$ *uses approximately the same resources as* $\mathcal{I}$*.*

Proof. We show that if $\mathcal{I}$ at some point during the integrity experiment sets INT-BROKEN to true, then we can break (strong) unforgeability of the MAC scheme. To this end we let a forger $\mathcal{F}$ run a black-box simulation of $\mathcal{I}$, simulating the other steps of the channel protocol FSCh like encryption locally, and only using the oracle access to $O_{\mathrm{MAC}}(K_I, \cdot)$ to compute MACs when required. For the simulated receiver oracle $\mathcal{F}$ always answers $\perp$. Algorithm $\mathcal{F}$ also keeps track of sent and received ciphertexts in the simulation, including the values i and j. When $\mathcal{I}$ sends the first ciphertext $C^* = (c^*, t^*)$ to the receiver oracle such that C^* has not been the next ciphertext prepared by the sender (i.e., C^* is entirely new or a modification of the j-th sent ciphertext $C_j = (c_j, t_j)$), then $\mathcal{F}$ outputs $(j||c^*, t^*)$ as its forgery attempt.

Note that the simulation is perfect, as the receiver oracle always returns $\perp$. Furthermore, $\mathcal{F}$ outputs a forgery as soon as INT-BROKEN is set to true. This can only happen if OUT-OF-SYNC has become true (according to the model) but the MAC verification has returned true (according to the protocol). The former implies that the ciphertext C^* must have been new or reordered ($j > i$ or $C^* \neq C_j$). And since the channel starts returning error symbols $\perp$ whenever it has encountered an invalid MAC, it must be the first such out-of-sync ciphertext C^* which, too, carries a valid MAC, to get some output $m \neq \perp$ from the receiver oracle.

Assume that $j > i$ for the first valid out-of-sync ciphertext $C^* = (c^*, t^*)$. In this case, since the receiver in the protocol holds the same counter value j in st_R up to this point, the receiver verifies t^* with regard to $j||c^*$. Since $j > i$ the sender oracle (and thus the MAC oracle in the simulation) has not issued any MAC for this counter value yet, such that the "message" $j||c^*$ for valid tag t^* in $\mathcal{F}$'s output constitutes a fresh forgery. Analogously, if $j \leq i$ and $C^* = (c^*, t^*)$ is different from $C_j = (c_j, t_j)$, then the pair $(j||c_j, t_j)$ is a successful strong forgery for $\mathcal{F}$ because the sender oracle (and thus MAC oracle) has only issued one tag for value j, with a different result $(j||c_j, t_j) \neq (j||c^*, t^*)$.

It follows that whenever $\mathcal{I}$ breaks integrity of the channel protocol we have a forgery for the underlying MAC scheme. For efficient $\mathcal{I}$ the resulting forger $\mathcal{F}$ is also efficient. □

We can now apply the composition theorem (Theorem 3.1), noting that the one-time-pad encryption

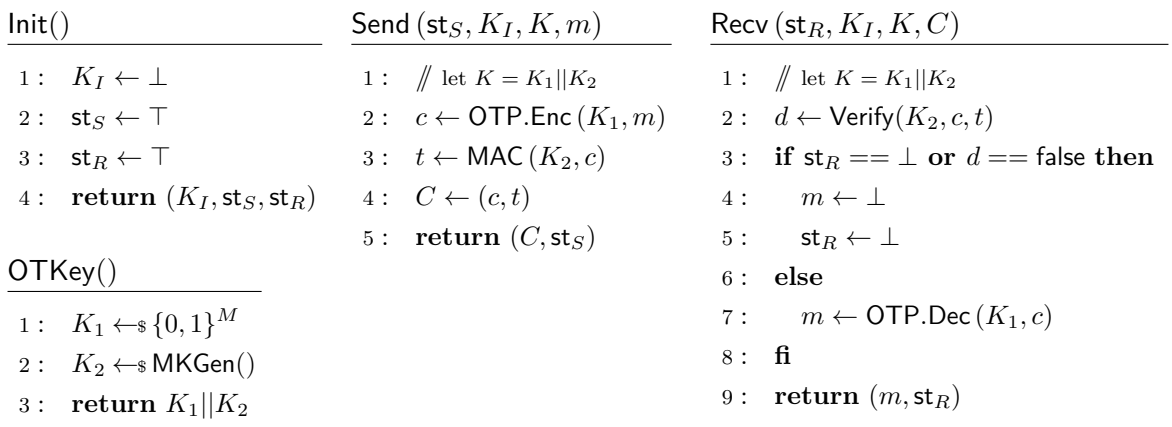

Figure 7: Unconditionally-Secure Channel Protocol USCh

ensures perfect ind-cpa security (such that independently of the adversarial resources the advantage is $\frac{1}{2}$), and using that integrity is bounded by the security of the strong unforgeability of the MAC scheme:

Theorem 4.5 (Future-Secure Channel)**.** *For the channel protocol* FSCh *in Construction 4.3 and any* ind-sfcca *adversary* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$*, we can construct and* seuf-cma *adversary* $\mathcal{F}$ *such that*

$$\mathsf{Adv}^{\mathsf{ind\text{-}sfcca}}_{\mathsf{FSCh}}(\mathcal{A}) \leq \tfrac{1}{2} + \mathsf{Adv}^{\mathsf{seuf\text{-}cma}}_{\mathsf{M}}(\mathcal{F}). \tag{7}$$

Here, $\mathcal{F}$ *uses approximately the same resources as* $\mathcal{A}_1$*.*

For an unbounded $\mathcal{A}_1$ —and hence an unbounded $\mathcal{I}$ in the proof— however, equation (7) may become void, since $\mathcal{I}$ may win Experiment $\mathsf{Exp}^{\mathsf{seuf\text{-}cma}}_{\mathsf{M}}(\mathcal{F})$ with significant probability.

4.3 Unconditionally-Secure Channels

For an unconditionally-secure channel we assume that both adversarial stages $\mathcal{A}_1$ and $\mathcal{A}_2$ in Experiment $\mathsf{Exp}^{\mathsf{ind\text{-}sfcca}}_{\mathsf{Ch}}(\mathcal{A})$ are unbounded, that is, we consider an unbounded adversary throughout the entire Experiment $\mathsf{Exp}^{\mathsf{ind\text{-}sfcca}}_{\mathsf{Ch}}(\mathcal{A})$. Our construction therefore asks for a fresh authentication key (part) with each send operation: we first split the per-message key K into two parts, K_1 and K_2. The former, K_1, is used for encryption via OTP, the latter, K_2, is used for authentication via an unconditionally-secure Carter-Wegman-MAC. For messages of length M bits we typically need M bits for the one-time pad and $2M$ bits for the Carter-Wegman MAC. More abstractly we consider a 1-query bounded MAC M in the construction below:

Construction 4.6 (Unconditionally-Secure Channel)**.** *Define the channel protocol* $\mathsf{USCh} = (\mathsf{Init}, \mathsf{OTKey}, \mathsf{Send}, \mathsf{Recv})$ *for message space* $\mathcal{M} = \{0,1\}^{\leq M}$ *by the algorithms in Figure 7.*

Once more we first argue int-sfctxt security of the channel protocol, assuming that the underlying MAC scheme M is strongly unforgeable against unbounded adversaries. The noteworthy fact here is that we lose a factor of $q_{\mathsf{Send}} + 1$ of sender queries in the security bound:

Lemma 4.7. *For any* int-sfctxt *adversary* $\mathcal{I}$ *making at most* q_{Send} *sender oracle queries there exists a* 1*-query bounded adversary* $\mathcal{F}$ *such that*

$$\mathsf{Adv}^{\mathsf{int\text{-}sfctxt}}_{\mathsf{USCh}}(\mathcal{I}) \leq (q_{\mathsf{Send}} + 1) \cdot \mathsf{Adv}^{\mathsf{seuf\text{-}cma}}_{\mathsf{M}}(\mathcal{F}). \tag{8}$$

Furthermore, $\mathcal{F}$ *uses the same resources as* $\mathcal{I}$*.*

Note that a Carter-Wegman MAC satisfies $\mathsf{Adv}^{\mathsf{seuf\text{-}cma}}_{\mathsf{M}}(\mathcal{F}) \leq 2^{-M}$ if we authenticate messages of at most M bits with $2M$ key bits [WC81]. This means that, as long as the number q_{Send} of sent ciphertexts is limited, the bound in the lemma is still reasonably small. Interestingly, for small message sizes M though and with a focus on "only" future-secure protection, an HMAC-based instantiation of Construction 4.3 can provide better concrete security.

Proof. The proof follows the one for the computational case closely. Only this time $\mathcal{F}$ guesses in advance, with probability $\frac{1}{q_{\mathsf{Send}}+1}$, the number i of the sender query for which $\mathcal{I}$ sends the first modified ciphertext $C^* \neq C_i$ to the receiver oracle, where we account for the possibility that $j > i$ with the additional choice $i = q_{\mathsf{Send}} + 1$. Algorithm $\mathcal{F}$ simulates an execution of $\mathcal{I}$ by doing all steps locally, and answering each receiver request with $\perp$. Only in the i-th sender oracle query $\mathcal{F}$ uses the external MAC oracle to compute the tag (still using a self-chosen, independent key part K_1 to encrypt the message before). When the integrity adversary $\mathcal{I}$ outputs the first modified ciphertext $C^* = (c^*, t^*)$ to the receiver oracle then $\mathcal{F}$ returns the pair (c^*, t^*) as its forgery attempt.

Given that the guess i is correct it follows as in the computational case that $\mathcal{F}$ wins the 1-query bounded unforgeability game if $\mathcal{I}$ wins the integrity game. Here we use that $\mathcal{F}$ at most makes a single external MAC query—or none if $i = q_{\mathsf{Send}} + 1$—and creates a (strong) forgery against the MAC scheme, because the pair (c^*, t^*) must be distinct from the MAC query (for $i \leq q_{\mathsf{Send}}$) or even new (for $i = q_{\mathsf{Send}} + 1$). □

It follows as in the computational case that Theorem 3.1 yields overall security of the unconditionally-secure channel protocol.

Theorem 4.8 (Unconditionally-Secure Channel)**.** *For the channel protocol* USCh *in Construction 4.6 and any* ind-sfcca *adversary* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ *where* $\mathcal{A}_1$ *makes at most* q_{Send} *sender oracle queries, we can construct an* int-sfctxt *adversary* $\mathcal{F}$ *such that*

$$\mathsf{Adv}^{\mathsf{ind\text{-}sfcca}}_{\mathsf{FSCh}}(\mathcal{A}) \leq \tfrac{1}{2} + (q_{\mathsf{Send}} + 1) \cdot \mathsf{Adv}^{\mathsf{seuf\text{-}cma}}_{\mathsf{M}}(\mathcal{F}) \,. \tag{9}$$

Here, $\mathcal{F}$ *uses the same resources as* $\mathcal{A}_1$ *and is* 1*-query bounded.*

5 Conclusion

We have shown how to achieve long-term confidentiality for channels, modeling security along the common notions for the computational setting like [BN00, BKN04] and adopting the two-stage adversaries of [BHMS17] to account for unbounded adversarial resources. We have shown how one-time pad encryption with authentication can be used to achieve the notion, where the proof is simplified through our translated general composition theorem that chosen-plaintext confidentiality and integrity gives chosen-ciphertext confidentiality in this setting. This provides fundamental security guarantees for such channels from which one can extend the result in several directions, as we discuss next.

We considered atomic channel protocols in which it is assumed that a transmitted ciphertext is fully received on the other side. Depending on the network, however, ciphertexts may be fragmented. It has been shown in attacks on actual channel protocols like SSH and IPSec [APW09, DP10] that this fragmentation behavior could potentially be exploited. A more formal treatment of ciphertext fragmentation can be found in [BDPS12, ADHP16]. One can also consider, on top, the possibility that the channel protocol itself may distribute input messages arbitrarily over ciphertexts, leading to the notion of stream-based channels [FGMP15]. It would be interesting to see how the requirement of unconditional security affects such models.

A possible extension in regard of security may be to allow exposure of some per-message keys, in which case these messages would not be confidential anymore. Still, the "fresh" keys should uphold security for

the other messages. This is similar to key updates in (computationally-secure) channel protocols where leakage of keys should not affect other keys and phases [GM17]. It would be interesting to augment the model here by similar considerations.

We followed earlier work and used a game-based definition for the security of channels, where keying material is provided by external means. If one now uses, say, a secure QKD protocol to generate the keys, then it remains yet to prove formally that the combined protocol is secure (albeit no attack on the joint execution is obvious). This is called compositional security. In stronger, simulation-based notions for key exchange and channels such as [CK02, DF18] compositional guarantees usually follow immediately. Compositional security for game-based notions of key exchange, as here, have been discussed in [BFWW11]. Again, both types, simulation-based and game-based models, usually only consider computationally bounded adversaries, leaving open the question if they still hold in the information-theoretic setting.

Acknowledgments

We thank Matthias Geihs and Lucas Schabhüser for discussions about long-term security, and the anonymous reviewers for valuable comments. Marc Fischlin and Philipp Muth have been (partially) funded by the Deutsche Forschungsgemeinschaft (DFG) – SFB 1119 – 236615297. Felix Günther is supported by the research fellowship grant GU 1859/1-1 of the DFG.

References

[ADHP16] Martin R. Albrecht, Jean Paul Degabriele, Torben Brandt Hansen, and Kenneth G. Paterson. A surfeit of SSH cipher suites. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016: 23rd Conference on Computer and Communications Security*, pages 1480–1491, Vienna, Austria, October 24–28, 2016. ACM Press. (Cited on page 12.)

[APW09] Martin R. Albrecht, Kenneth G. Paterson, and Gaven J. Watson. Plaintext recovery attacks against SSH. In *2009 IEEE Symposium on Security and Privacy*, pages 16–26, Oakland, CA, USA, May 17–20, 2009. IEEE Computer Society Press. (Cited on page 12.)

[BBF+19] Nina Bindel, Jacqueline Brendel, Marc Fischlin, Brian Goncalves, and Douglas Stebila. Hybrid key encapsulation mechanisms and authenticated key exchange. In Jintai Ding and Rainer Steinwandt, editors, *PQCrypto 2019*, volume 11505 of *LNCS*, pages 206–226. Springer, 2019. (Cited on page 2.)

[BCK96] Mihir Bellare, Ran Canetti, and Hugo Krawczyk. Keying hash functions for message authentication. In Neal Koblitz, editor, *Advances in Cryptology – CRYPTO'96*, volume 1109 of *Lecture Notes in Computer Science*, pages 1–15, Santa Barbara, CA, USA, August 18–22, 1996. Springer, Heidelberg, Germany. (Cited on pages 3 and 9.)

[BDPS12] Alexandra Boldyreva, Jean Paul Degabriele, Kenneth G. Paterson, and Martijn Stam. Security of symmetric encryption in the presence of ciphertext fragmentation. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 682–699, Cambridge, UK, April 15–19, 2012. Springer, Heidelberg, Germany. (Cited on page 12.)

[Bel15] Mihir Bellare. New proofs for NMAC and HMAC: Security without collision resistance. *Journal of Cryptology*, 28(4):844–878, October 2015. (Cited on page 9.)

[BFWW11] Christina Brzuska, Marc Fischlin, Bogdan Warinschi, and Stephen C. Williams. Composability of Bellare-Rogaway key exchange protocols. In Yan Chen, George Danezis, and Vitaly Shmatikov, editors, *ACM CCS 2011: 18th Conference on Computer and Communications Security*, pages 51–62, Chicago, Illinois, USA, October 17–21, 2011. ACM Press. (Cited on page 13.)

[BGM04] Mihir Bellare, Oded Goldreich, and Anton Mityagin. The power of verification queries in message authentication and authenticated encryption. Cryptology ePrint Archive, Report 2004/309, 2004. `http://eprint.iacr.org/2004/309`. (Cited on page 8.)

[BHMS17] Nina Bindel, Udyani Herath, Matthew McKague, and Douglas Stebila. Transitioning to a quantum-resistant public key infrastructure. In Tanja Lange and Tsuyoshi Takagi, editors, *PQCrypto 2017*, volume 10346 of *LNCS*, pages 384–405. Springer, 2017. (Cited on pages 2 and 12.)

[BKN04] Mihir Bellare, Tadayoshi Kohno, and Chanathip Namprempre. Breaking and provably repairing the SSH authenticated encryption scheme: A case study of the encode-then-encrypt-and-MAC paradigm. *ACM Transactions on Information and System Security*, 7(2):206–241, 2004. (Cited on pages 4, 6, and 12.)

[BN00] Mihir Bellare and Chanathip Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. In Tatsuaki Okamoto, editor, *Advances in Cryptology – ASIACRYPT 2000*, volume 1976 of *Lecture Notes in Computer Science*, pages 531–545, Kyoto, Japan, December 3–7, 2000. Springer, Heidelberg, Germany. (Cited on pages 3 and 12.)

[CK02] Ran Canetti and Hugo Krawczyk. Universally composable notions of key exchange and secure channels. In Lars R. Knudsen, editor, *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 337–351, Amsterdam, The Netherlands, April 28 – May 2, 2002. Springer, Heidelberg, Germany. (Cited on page 13.)

[CM97] Christian Cachin and Ueli M. Maurer. Unconditional security against memory-bounded adversaries. In Burton S. Kaliski Jr., editor, *Advances in Cryptology – CRYPTO'97*, volume 1294 of *Lecture Notes in Computer Science*, pages 292–306, Santa Barbara, CA, USA, August 17–21, 1997. Springer, Heidelberg, Germany. (Cited on page 3.)

[DF18] Jean Paul Degabriele and Marc Fischlin. Simulatable channels: Extended security that is universally composable and easier to prove. In Thomas Peyrin and Steven Galbraith, editors, *Advances in Cryptology – ASIACRYPT 2018, Part III*, volume 11274 of *Lecture Notes in Computer Science*, pages 519–550, Brisbane, Queensland, Australia, December 2–6, 2018. Springer, Heidelberg, Germany. (Cited on page 13.)

[DP10] Jean Paul Degabriele and Kenneth G. Paterson. On the (in)security of IPsec in MAC-then-encrypt configurations. In Ehab Al-Shaer, Angelos D. Keromytis, and Vitaly Shmatikov, editors, *ACM CCS 2010: 17th Conference on Computer and Communications Security*, pages 493–504, Chicago, Illinois, USA, October 4–8, 2010. ACM Press. (Cited on page 12.)

[FGMP15] Marc Fischlin, Felix Günther, Giorgia Azzurra Marson, and Kenneth G. Paterson. Data is a stream: Security of stream-based channels. In Rosario Gennaro and Matthew J. B. Robshaw,

editors, *Advances in Cryptology – CRYPTO 2015, Part II*, volume 9216 of *Lecture Notes in Computer Science*, pages 545–564, Santa Barbara, CA, USA, August 16–20, 2015. Springer, Heidelberg, Germany. (Cited on page 12.)

[GM17] Felix Günther and Sogol Mazaheri. A formal treatment of multi-key channels. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017, Part III*, volume 10403 of *Lecture Notes in Computer Science*, pages 587–618, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Heidelberg, Germany. (Cited on page 13.)

[GND+19] M. Geihs, O. Nikiforov, D. Demirel, A. Sauer, D. Butin, F. Günther, G. Alber, T. Walther, and J. Buchmann. The status of quantum-key-distribution-based long-term secure internet communication. *IEEE Transactions on Sustainable Computing*, 2019. (Cited on page 2.)

[KMR09] Robin Künzler, Jörn Müller-Quade, and Dominik Raub. Secure computability of functions in the IT setting with dishonest majority and applications to long-term security. In Omer Reingold, editor, *TCC 2009: 6th Theory of Cryptography Conference*, volume 5444 of *Lecture Notes in Computer Science*, pages 238–255. Springer, Heidelberg, Germany, March 15–17, 2009. (Cited on page 3.)

[Mau92] Ueli M. Maurer. Conditionally-perfect secrecy and a provably-secure randomized cipher. *Journal of Cryptology*, 5(1):53–66, January 1992. (Cited on page 3.)

[MN06] Tal Moran and Moni Naor. Receipt-free universally-verifiable voting with everlasting privacy. In Cynthia Dwork, editor, *Advances in Cryptology – CRYPTO 2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 373–392, Santa Barbara, CA, USA, August 20–24, 2006. Springer, Heidelberg, Germany. (Cited on page 2.)

[MSU13] Michele Mosca, Douglas Stebila, and Berkant Ustaoglu. Quantum key distribution in the classical authenticated key exchange framework. In Philippe Gaborit, editor, *PQCrypto 2013*, volume 7932 of *LNCS*, pages 136–154. Springer, 2013. (Cited on page 2.)

[MU10] Jörn Müller-Quade and Dominique Unruh. Long-term security and universal composability. *Journal of Cryptology*, 23(4):594–671, October 2010. (Cited on page 3.)

[Res18] E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446 (Proposed Standard), August 2018. (Cited on page 9.)

[Sha49] Claude E. Shannon. Communication theory of secrecy systems. *Bell Syst. Tech. J.*, 28(4):656–715, 1949. (Cited on page 2.)

[Shr04] Tom Shrimpton. A characterization of authenticated-encryption as a form of chosen-ciphertext security. Cryptology ePrint Archive, Report 2004/272, 2004. `http://eprint.iacr.org/2004/272`. (Cited on page 4.)

[WC81] Mark N. Wegman and Larry Carter. New hash functions and their use in authentication and set equality. *J. Comput. Syst. Sci.*, 22(3):265–279, 1981. (Cited on pages 3, 9, and 12.)

[Wyn75] A. D. Wyner. The wire-tap channel. *The Bell System Technical Journal*, 54(8):1355–1387, Oct 1975. (Cited on page 1.)