

Secure Channels and Termination: The Last Word on TLS

Colin Boyd Britta Hale

NTNU, Norwegian University of Science and Technology, Trondheim, Norway
{colin.boyd,britta.hale}@item.ntnu.no

Abstract. Secure channels are one of the most pivotal building blocks of cryptography today. Internet connections, secure messaging, protected IoT data, etc., all rely upon the security of the underlying channel. In this work we define channel protocols, as well as security for channels constructed from stateful length-hiding authenticated encryption (stLHAE) schemes. Furthermore, we initiate the concept of *secure termination* where, upon receipt of a signifying message, a receiver is guaranteed to have received every message that has been sent, and will ever be sent, on the channel. We apply our results to real-world protocols, linking the channel environment to previous analyses of TLS 1.2, and demonstrating that TLS 1.2 achieves secure termination via fatal alerts and `close_notify` messages, per the specification of the Alert Protocol.

Keywords: Secure channels, stateful length-hiding authenticated encryption (stLHAE), authenticated encryption with associated data (AEAD), secure termination, controllable channel protocol, Transport Layer Security (TLS)

1 Introduction

Communication security is built on a fundamental cornerstone commonly referred to as a *secure channel*. Creation of secure channels is the essential goal of secure email, end-to-end encrypted messaging applications, end-to-end encrypted VOIP, HTTPS internet connections and TLS in general, WPA2 WiFi protection, SSH, IPSec, Bluetooth, etc. Examples are innumerable. Additionally, many constructs rely on the existence of secure channels once established, e.g. key transport. Despite this, a general understanding of what *secure channels* are and how they are constructed is lacking. Research relating to secure channels has spiraled concentrically around the topic with frequently contradicting goals, particularly in the analysis of real-world protocols.

Authenticated encryption with associated data (AEAD), or even simply authenticated encryption (AE), has been argued as the foundational secure channel building block. Extensive work has been done on both AEAD [32], and AE in their various forms [3, 17, 21, 34, 16, 33]. Stateful length-hiding AE (stLHAE) is often the apparent goal of real-world protocols and has consequently been used frequently in their analysis [20, 23, 29]. Work has also been undertaken for building secure channels explicitly from AE schemes [28]. However, the view of secure channels as simply AEAD or AE is incomplete. In real-world protocols, multiple instances of a protocol may be run, session resumption/renegotiation may be

performed, and message authentication (MAC), encryption, and even exporter keys may be derived from a single master session key output of a key exchange protocol. Essentially, the real-world is not simple. Cross protocol attacks, renegotiation attacks [13], and triple handshake attacks [5] are just some attack examples that cannot be captured when considering secure channels as AEAD or AE under a single communication flow.

Expanding on the modeling of channels, authenticated and secure channel establishment (ACCE) was proposed, which considers both key exchange and channels established under the derived key in the context of stLHAE [20]. Furthermore, ACCE considers parallel sessions – an improvement over the basic, low-level view of secure channels as the stLHAE primitive. Work analyzing the TLS protocol [20, 23] has employed ACCE as the foundational secure channel building block, as TLS does not achieve key indistinguishability. However, not all protocols suffer from a lack of key indistinguishability (e.g. the current TLS 1.3 draft [30]), thus rendering the ACCE pre-accept/post-accept phase combination unnecessary. Similarly to ACCE, Augmented Secure Channels (ASC) have been proposed as a means of capturing more of the channel context than an AEAD primitive allows [2]. ASC is developed in the vein of constructive cryptography, but still suffers from the same underlying drawbacks as ACCE; namely the inability to model communication flows, under potentially different security demands, which are protected by keying material derived from the same master secret. Neither ACCE nor ASC model session resumption – despite its importance in TLS 1.2 (analyzed in the ACCE model) and TLS 1.3 (analyzed in the ASC model for draft 8). Work has also been done on multi-key channels [14], focusing on the evolution of a master key over time, and requirements on it (e.g. forward secrecy), but lacking formal definitions for the channel context. It also uses a single, fixed AEAD construction, lacking flexibility in scheme selection, and does not address termination in the channel environment.

In response to these issues, we define keyed two-party controllable channel protocols which capture parallel and consecutive sessions, with security controlled via the underlying primitives. Each session at a principal is initiated via the generation of a master session key – as would be the case at the conclusion of a key exchange protocol. In turn, the session is modeled by a collection of read and write connections, with connection keys derived from the master session key. Channels are defined by a shared key between a pair of read and write connections. This captures the behavior of real-world protocols: for example, a master session key may be used to derive a MAC key as well as an encryption key, which would result in parallel connections under the different channel keys. Simultaneously, this higher-level view of the channel environment provides a framework for analysis of session resumption; connections can be closed, and channel keys destroyed, with new connections instantiated and channel keys derived from the original session keying material. Compared to the ACCE model, our model permits separate consideration of session and connection keys and even consideration of connection closure.

Security for keyed two-party controllable channel protocols is realized via the cryptographic schemes and key derivation functions (KDF) used for connections. In real-world correlations, TLS 1.2 uses cipher suites – for example, the partial suite `AES_128_GCM_SHA256`, which defines both a scheme for securing the channel (`AES_128_GCM`) and pseudo-random function for key derivation of the channel key (HMAC using SHA256). Our model captures such real-world protocols by considering channel security under both a scheme (e.g. stLHAE) goal and a KDF goal. Considering secure channels in the context of stLHAE allows certain analysis benefits; statefulness gives assurance that the i -th packet processed on the receiving end of a channel is the same as i -th packet output on the sending end. However, the stLHAE framework is again incomplete. How does a receiver know that all sent packets have been received? If an adversary drops the last n packets on a channel, the receiver could be convinced that the transmission is shorter than in reality. When final messages contain critical information, warnings, etc., this scenario should undoubtedly be considered as an attack. Classically, this is referred to as a *truncation* attack. In order to capture this attack in our secure channel environment, we define secure termination.

Truncation attacks have been shown to be a very real problem [35], including effects on voting. Usually these attacks follow from a failure of the implementation to check for the closure alert (if TLS), or from a misinterpretation of what constitutes a “terminate” message. While the importance of the former is highlighted by our model, the latter is at the heart of secure termination. Most analyses of truncation attacks are ad-hoc, essentially cryptanalyses based on weaknesses discovered in a particular protocol. Some of these have been against the TLS protocol [27, 35], leveraging and exploiting implementation faults. Other recent work on protocol termination, albeit in an unrelated aspect, highlights a growing interest in the final stages of a cryptographic protocol run [10]. Consequently, we define secure termination in the interest of providing a formal framework for modeling truncation attacks, by modeling finalization and completion guarantees on received communication flows.

Channels To discuss secure channels and channel termination it is vital to clearly define what a channel actually is. Past work using channels has provided mixed descriptions of this concept – Hoepman [18] describes *unidirectional* and *bidirectional* channels, hence conceptually equating a channel to a transport link between entities. In this sense, a bidirectional channel for a real-world protocol such as TLS would have separate keys for each direction, but all keys would be considered to be within the same channel. Meanwhile, another line of research [26, 25] defines a channel in a unidirectional sense, with messages input from a sender and output to a receiver. Ultimately, this formulation allows for a modular analysis of channel security, with channel keys being used for sending messages at one end of the channel and for receiving messages at the other end. This practice has met with wide-spread acceptance, with channels generally being modeled by three interfaces (sender, receiver, and adversary) based on various adversarial capabilities [2].

In 2001, Canetti and Krawczyk [9] defined a *secure channel* as “a link between a pair of parties” which provides message authentication and confidentiality via a key obtained via a key-exchange protocol. As in many other works on key exchange and secure channel analysis [20, 23, 4], the authors consider *sessions* at principals as participants in the key exchange protocol, and call the resulting key the *session key*. This session key is then used to secure the channel (e.g. encryption and authentication keys are derived from the session key).

One salient issue arises when the standard conceptualization of a session key is compared with the modeling of a channel, as discussed above. Namely, a key exchange protocol should minimally result in a set of two channel keys. This assumption follows from standard real-world protocols which maintain separate session keys for sending channels and receiving channels [11, 19, 31]. In consideration of this, we undertake to formalize channels in the context of sessions for bi-directional communication. With sent messages at a session being not necessarily independent from received messages (particularly, as we will see, in the context of secure termination), this new, “big picture” view of session modeling raises interesting questions for channel analysis. One recent work [24] similarly aims to address bi-directional channels, but does not consider sub-connections (parallel channels after key derivation) or channel termination.

While TLS and similar protocols have been analyzed under the assumption of discrete messages, important work has been done considering AEAD in the streaming setting [12]. We view our work as easily adaptable to the streaming security context, since the security of our channel protocol model is reliant upon underlying, “plug-in” primitives and their security games.

Cryptographic Agility While channels generally use separate keys for sending and receiving, these keys are often generated from the same master session key. This is indeed the case in many real-world protocols such as TLS. Moreover, it is possible that the same master session key could be re-used with different concrete cryptographic algorithms, which may be controllable by the adversary. We account for this possibility by requiring *agile* [1, 15] channel key derivation in our channel security analysis. Bhargavan et al. [6] analyzed cryptographic agility in TLS, focusing on the public key algorithms used in the handshake protocol. Agility definitions appear App. B (adapted from [1]), and we apply agile key derivation functions in our results in order to provide strong channel security.

Paper Outline Introducing keyed two-party controllable channel protocols, §2 defines sessions, connections, schemes, channels, and connection/channel closure, outlining the fundamental structure of the channel environment even before consideration of security. Combinations of security schemes and key derivation functions (for the channel key) comprise *suites*.

Subsequently, §3 describes security. We focus on channel stLHAE security due to it being a frequent real-world objective. Our security experiments here are broad, and encompass the situation where parallel and consecutive connections may employ various suites. We expect that this work can be extended to situations where parallel and consecutive connections may employ combinations of various authentication suites, stLHAE suites, AEAD suites, etc.

Additionally, we provide constructions of channel protocol environments where channels are constructed from various stLHAE schemes. Proving security for our constructions, we provide a reduction from the channel security to the security of the underlying schemes used in the construction.

Linking our definitions and constructions to the ACCE model, we furthermore demonstrate that channel stLHAE security can be reduced to ACCE security in the post-accept phase (channel phase of ACCE), when the session key is used directly as the channel key and only one connection for the session is initialized (§5.1). Thus our channel protocol can be used as a refined post-accept phase for the ACCE model, if such a model is necessary due to lack of key indistinguishability. Formally, the one-connection and session key restrictions on the reduction arise from the formalization of the ACCE model, as keyed two-party controllable channel protocols encompass a wider vision of protocol interaction, based on real-world protocol implementations.

In §4 secure termination and the secure termination experiment are defined. Additionally, we associate secure termination and channel protocol security by reductions between the former and the latter, under the generic case of parallel and consecutive connections constructed from various stLHAE suites. We demonstrate that TLS 1.2 achieves secure termination on receipt of Alert or `close_notify` messages by combining our work on secure channels with previous analyses of the TLS 1.2 protocol in the ACCE model.

While presented in detail in the following sections, the notation for channel protocols, schemes, and respective security games is summarized in Table 1. This provides a reference point, highlighting the differences between similar terms and notation, and is an extension of previously established notation.

ChannelSnd ChannelRcv	Send and receive algorithms for a keyed two-party controllable channel protocol $\text{CHNL}[\{(II, \text{KDF})\}]$.
ChnlEnc ChnlDec	Security experiment oracles for a keyed two-party controllable channel protocol $\text{CHNL}[\{(II, \text{KDF})\}]$, where $\{II\}$ are stLHAE schemes.
Send $_{\top}$ Receive $_{\top}$	Security experiment oracles for the secure termination of a channel protocol experiment.
Enc & Dec	Encrypt and decrypt algorithms for a stLHAE scheme.

Table 1. Notation reference for protocols, schemes, and experiments.

2 Channels

In this section we formalize the natural real-world protocol environment of sessions, connections, and channels. Fig. 1 provides a depiction of this environment.

2.1 Definitions

Every channel protocol will be associated with one or more cryptographic *schemes*. Typical schemes include authentication schemes, under a message authentication code, and stLHAE, but many other types are possible. We define abstractly the elements of any such scheme.




Fig. 1. Communication diagram with sessions π_s^P and $\pi_s^{P'}$, write and read connections cxn_W and cxn_R for each session, and channels between connections protected by the channel keys $\pi_s^P . cxn_W^i . CK$, etc. Each principal may be modeled by multiple sessions.

Definition 1 (Scheme). Let \mathcal{M} be a message space, \mathcal{K} a key space, and \mathcal{C} an output space, and let the elements of \mathcal{C} be called ciphertexts. A scheme Π is a tuple of algorithms:

- $\text{Kgn}() \xrightarrow{\S} k$: A probabilistic key gen. algorithm that outputs a key $k \in \mathcal{K}$.
- $\text{Snd}(k, m, st_S) \xrightarrow{\S} (c, st'_S)$: A probabilistic send algorithm that takes as input a key $k \in \mathcal{K}$, a message $m \in \mathcal{M}$, and a write state st_S , and outputs an outgoing ciphertext $c \in \mathcal{C}$ or an error symbol \perp , and an updated state st'_S .
- $\text{Rcv}(k, c, st_R) \rightarrow (m, st'_R)$: A deterministic receive algorithm that takes as input a key $k \in \mathcal{K}$, a ciphertext $c \in \mathcal{C}$, and a read state st_R , and outputs either a message $m \in \mathcal{M}$ or a error symbol \perp , and an updated state st'_R .

On first use, st_S and st_R are initialized to st_S^0 and st_R^0 , resp. If, for $\text{Snd}(k, m, st_S) \xrightarrow{\S} (c, st'_S)$ and $\text{Rcv}(k, c, st_R) \rightarrow (m, \alpha, st'_R)$, $st'_S = \perp$ and $st'_R = \perp$, then Π is said to be stateless. Otherwise Π is said to be stateful.

Correctness Consider the following: $i \geq 0$, all $m_i \in \mathcal{M}$, all $k \xleftarrow{\S} \text{Kgn}()$, initial states st_S^0 and st_R^0 , and a sequence $(c_i, st_S^{i+1}) \xleftarrow{\S} \text{Snd}(k, m_i, st_S^i)$, where $c_i \neq \perp$ for all i . Then, for a matching sequence of message receipts, we have $(m_i, st_R^{i+1}) \leftarrow \text{Rcv}(k, c_i, st_R^i)$. Further correctness requirements may hold, dependent on the scheme. In §3, stLHAE schemes will be considered; however, many scheme types are possible. Other types of schemes include signatures schemes and authentication schemes.

Definition 2 (Principals and Sessions). For a collection of principals $\{P_1, \dots, P_p\}$, where $p \in \mathbb{N}$, each $P_l \in \{P_1, \dots, P_p\}$ is (potentially) in possession of a long-term private/public key pair (sk_P, pk_P) , and is modeled by a collection of session oracles $\{\pi_{s_1}^{P_1}, \dots, \pi_{s_n}^{P_n}\}$.

Correspondingly and wlog, each session π_s^P is an oracle with access to the (potential) long-term key pair of P and the (potential) long-term public keys pk_1, \dots, pk_p of all other principals. Furthermore, a session π_s^P at a principal P maintains a collection of variables $(\{\pi_s^P . cxn_{role}^i\}, \pi_s^P . (P', \bar{s}), \pi_s^P . \alpha, \pi_s^P . SK)$:

- $\{cxn_{role}^i\}$: A collection of connections, where $role \in \{W, R\}$, representing ‘write’ and ‘read’ connections, respectively.
If $role = W$, then \overline{role} denotes R , and vice versa.
- $(P', [\bar{s}])$: An identifier for the partner and (optionally) the partner’s session.
- α : A status variable in $\{0, 1\}$, where the session is active if $\alpha = 1$ and inactive if $\alpha = 0$. A session must be active to send or receive messages.
- SK : A session key shared with the partner session.

Any P may maintain several sessions, both in parallel and consecutively.

The following definition uses a key derivation function (KDF). We follow standard assumptions by requiring KDF to be a PRF [8]; KDFs may also be defined more explicitly (see [22]).

Definition 3 (Connection). A connection cxn_{role} with $role \in \{W, R\}$ is defined by a set of variables $(cxn_{role}.CK, cxn_{role}.status, cxn_{role}.suite, cxn_{role}.substate)$. Let the notation e.g. $cxn_{role}.CK$ denote variables at a specific connection:

- CK : A channel key corresponding to the connection and role.
- $status$: A status variable in $\{\text{active}, \text{terminated}\}$.
- $suite$: A variable identifying the scheme/KDF pair (Π, KDF) implemented on the connection. If no pair is specified for the connection, then $cxn_{role}.suite \leftarrow \perp$.
- $substate$: Any additional scheme-specific connection state variables.

The connection $substate$ variable models other state information, which may be defined by $suite$. For example, this variable may handle protocol state if Π is a stLHAE scheme. Exact use of this variable can be seen in the concrete constructions provided in §3. We abuse notation and use cxn_{role} as an identifier by which an entity may refer to the collection of variables, without having access to them.

Definition 4 (Keyed Two-Party Controllable Channel Protocol). Let $\mathcal{M} = \mathcal{AD} \times \mathcal{AED}$ be a message space and \mathcal{L} an optional length space. A keyed two-party controllable channel protocol $\text{CHNL}[\{(II, \text{KDF})\}]$ over a set of scheme, key derivation function pairs $\{(II, \text{KDF})\}$ is a tuple of algorithms:

- $\text{SessionKeyGen}(1^\lambda, P, P', s, \bar{s}) \xrightarrow{\mathbb{S}} SK$: A probabilistic session key generation algorithm that takes as input a security parameter λ , identities P and P' , and session indices s and \bar{s} . It sets the respective partner identifiers and session status $\pi_s^P.\alpha \leftarrow 1$, and outputs a shared session key SK .
- $\text{ConnectionInit}(SK, \pi_s^P, [II], [\text{KDF}], i) \xrightarrow{\mathbb{S}} (\pi_s^P.cxn_{\overline{W}}^i, \pi_s^P.cxn_{\overline{R}}^i)$: A probabilistic connection initiation algorithm that takes as input a tuple of a shared session key SK , a session π_s^P , an optional scheme Π , an optional key derivation function KDF used to derive CK , and connection index i , and outputs the i -th ‘read’ and ‘write’ connections at the session, $\pi_s^P.cxn_{\overline{W}}^i$ and $\pi_s^P.cxn_{\overline{R}}^i$, or a distinguishing failure symbol \perp .

- $\text{ChannelSnd}(m, [\ell], \text{cxn}_{\mathbb{W}}^i) \xrightarrow{\mathbb{S}} (c, \text{cxn}_{\mathbb{W}}^{i'})$: A probabilistic channel sending algorithm that takes as input a message $m \in \mathcal{M} = \mathcal{AD} \times \mathcal{AED}$, an optional output length $\ell \in \mathcal{L}$, and a connection $\text{cxn}_{\mathbb{W}}^i$, and outputs a ciphertext $c \in \mathcal{C}$ or a distinguishing failure symbol \perp , and an updated connection $\text{cxn}_{\mathbb{W}}^{i'}$.
- $\text{ChannelRcv}(p, \text{cxn}_{\mathbb{R}}^i) \rightarrow (m, \text{cxn}_{\mathbb{R}}^{i'})$: A deterministic channel receiving algorithm that takes as input a packet $p \in \mathcal{P} = \mathcal{AD} \times \mathcal{C}$ and a connection $\text{cxn}_{\mathbb{R}}^i$, and outputs a message $m \in \mathcal{M} = \mathcal{AD} \times \mathcal{AED}$ or a distinguishing failure symbol \perp , and an updated connection $\text{cxn}_{\mathbb{R}}^{i'}$.

The packet space \mathcal{P} is a set induced by $\text{CHNL}[\Pi, \text{KDF}]$ for each $\Pi \in \{\Pi\}$, the length space \mathcal{L} , and the message space \mathcal{M} , where \mathcal{M} is a tuple of the data space of authenticated transmissions \mathcal{AD} , and a data space of authenticated and encrypted transmissions \mathcal{AED} . Consequently, \mathcal{P} is a tuple of \mathcal{AD} and a ciphertext space \mathcal{C} , where \mathcal{C} is defined by the scheme used on \mathcal{M} . The i -th read-write connection pair at a session π_s^P is denoted $(\pi_s^P.\text{cxn}_{\mathbb{W}}^i, \pi_s^P.\text{cxn}_{\mathbb{R}}^i)$.

We define correctness in the logical way where ConnectionInit outputs the failure symbol if the session is not active, and ChannelSnd and ChannelRcv , respectively, output the failure symbol if the connection status variable is not active.

Correctness is defined as follows:

- For all $SK \leftarrow \text{SessionKeyGen}(1^\lambda, P, P', s, \bar{s})$:
 - If $\text{ConnectionInit}(\cdot, \pi_s^P, \cdot, \cdot, i)$ has previously been called, then
 - * $\perp \leftarrow \text{ConnectionInit}(SK, \pi_s^P, [\Pi], [\text{KDF}], i)$.
 - If $\mathcal{CK} \not\subseteq \mathcal{K}$, i.e. the output key space of KDF is not a subset of the input key space of Π , then
 - * $\perp \leftarrow \text{ConnectionInit}(SK, \pi_s^P, [\Pi], [\text{KDF}], i)$.
 - Else, $(\pi_s^P.\text{cxn}_{\mathbb{W}}^i, \pi_s^P.\text{cxn}_{\mathbb{R}}^i) \leftarrow \text{ConnectionInit}(SK, \pi_s^P, [\Pi], [\text{KDF}], i)$ if
 - * $\pi_s^P.\alpha = 1$ and $\pi_s^P.\text{cxn}_{\mathbb{W}}^i.CK \leftarrow \text{KDF}(SK, (P, s), (P', \bar{s}), i)$ and $\pi_s^P.\text{cxn}_{\mathbb{R}}^i.CK \leftarrow \text{KDF}(SK, (P', \bar{s}), (P, s), i)$, where $\pi_s^P.\text{cxn}_{\mathbb{W}}^i.CK, \pi_s^P.\text{cxn}_{\mathbb{R}}^i.CK \in \mathcal{K}$ the key space of Π , and
 - * $\perp \leftarrow \text{ConnectionInit}(SK, \pi_s^P, [\Pi], [\text{KDF}], i)$ otherwise.
- If SK is not the output of any previous run of $\text{SessionKeyGen}(1^\lambda, P, P', s, \bar{s})$, then $\perp \leftarrow \text{ConnectionInit}(SK, \pi_s^P, [\Pi], [\text{KDF}], i)$.
- For all
 - $SK \leftarrow \text{SessionKeyGen}(1^\lambda, P, P', s, \bar{s})$,
 - $(\pi_s^P.\text{cxn}_{\mathbb{W}}^i, \pi_s^P.\text{cxn}_{\mathbb{R}}^i) \leftarrow \text{ConnectionInit}(SK, \pi_s^P, [\Pi], [\text{KDF}], i)$, and
 - $(\pi_{\bar{s}}^{P'}.\text{cxn}_{\mathbb{W}}^i, \pi_{\bar{s}}^{P'}.\text{cxn}_{\mathbb{R}}^i) \leftarrow \text{ConnectionInit}(SK, \pi_{\bar{s}}^{P'}, [\Pi], [\text{KDF}], i)$,
we have that $\pi_s^P.\text{cxn}_{\text{role}}^i.CK = \pi_{\bar{s}}^{P'}.\text{cxn}_{\text{role}}^i.CK$.
- For sending:
 - For all $m \in \mathcal{M}$,
 - * $(c, \pi_s^P.\text{cxn}_{\mathbb{W}}^{i'}) \leftarrow \text{ChannelSnd}(m, [\ell], \pi_s^P.\text{cxn}_{\mathbb{W}}^i)$, provided that $\pi_s.\text{cxn}_{\mathbb{W}}^i.\text{status} = \text{active}$, and
 - * $(\perp, \pi_s^P.\text{cxn}_{\mathbb{W}}^{i'}) \leftarrow \text{ChannelSnd}(m, [\ell], \pi_s^P.\text{cxn}_{\mathbb{W}}^i)$ otherwise.
 - If $\pi_s.\text{cxn}_{\mathbb{W}}^i.\text{status} = \text{active}$, then for all $m \in \mathcal{M}$:
 - * $(c, \pi_s^P.\text{cxn}_{\mathbb{W}}^{i'}) \leftarrow \text{ChannelSnd}(m, \ell, \pi_s^P.\text{cxn}_{\mathbb{W}}^i)$, where $|c| = \ell$, or

- * $(\perp, \pi_s^P . \text{cxn}_W^i) \leftarrow \text{ChannelSnd}(m, \ell, \pi_s^P . \text{cxn}_W^i)$.
- For receiving, for all $m = (\mathbf{ad}, m) \in \mathcal{M}$ and $p = (\mathbf{ad}, c)$, where
 - $(c, \pi_s^P . \text{cxn}_W^i) \leftarrow \text{ChannelSnd}(m, [\ell], \pi_s^P . \text{cxn}_W^i)$ and
 - $\pi_s^P . \text{cxn}_W^i . CK = \pi_s^{P'} . \text{cxn}_R^i . CK$,

we have

- $(m, \pi_s^{P'} . \text{cxn}_R^i) \leftarrow \text{ChannelRcv}(p, \pi_s^{P'} . \text{cxn}_R^i)$, provided that
 - * $\pi_s^{P'} . \text{cxn}_R^i . \text{status} = \mathbf{active}$, and
 - * $\pi_s^{P'} . \text{cxn}_R^i . \text{substate}$ and $\pi_s^P . \text{cxn}_W^i . \text{substate}$ satisfy any substate conditions (e.g. Π scheme-specific statefulness conditions), and
- $(\perp, \pi_s^{P'} . \text{cxn}_R^i) \leftarrow \text{ChannelRcv}(p, \pi_s^{P'} . \text{cxn}_R^i)$ otherwise.

Remark 1. Note that there is no restriction on the number of calls made to $\text{ConnectionInit}(SK, \pi_s^P, \pi_s^{P'}, [\Pi], [\text{KDF}], i)$, for a given session key SK and sessions π_s^P and $\pi_s^{P'}$. This models real-world protocols where session resumption is possible. Particularly, connections $\pi_s^P . \text{cxn}_{role}^i$ and $\pi_s^P . \text{cxn}_{role}^i$ may be *terminated*, resulting in the connections being destroyed completely; later resumption based on the session keying material is possible by calling ConnectionInit again.

Remark 2. We explicitly allow ConnectionInit to be used to create connections at only one session. Matching connections can be created by using the algorithm again: $\text{ConnectionInit}(SK, \pi_s^{P'}, \pi_s^P, [\Pi], [\text{KDF}], i)$. This matches real-world protocols where each session derives its connection keys independently, regardless of whether or not the partner session is still active.

CHNL may take as input several schemes which may each be used to initialize different connections. If no KDF is used to derive channel keys, then either the session key is used directly in the channel, or no security scheme Π is implemented (i.e. the $([], [])$ suite). In either case $[]$ is included in the list of functions $\{\text{KDF}\}$, and represents the identity function.

Definition 5 (Channel). *If CHNL is a keyed two-party controllable channel protocol, and $\pi_s^P . \text{cxn}_W^i$ and $\pi_s^{P'} . \text{cxn}_R^j$*

- *share a channel key $\pi_s^P . \text{cxn}_W^i . CK = \pi_s^{P'} . \text{cxn}_R^j . CK$ and*
- *implement the same scheme/KDF suite such that $\pi_s^P . \text{cxn}_W^i . \text{suite} = \pi_s^{P'} . \text{cxn}_R^j . \text{suite}$,*

then we say that $\pi_s^P . \text{cxn}_W^i$ has a channel to $\pi_s^{P'} . \text{cxn}_R^j$. Moreover, if $\pi_s^P . \text{cxn}_W^i . \text{suite} = (\Pi, \text{KDF}) = \pi_s^{P'} . \text{cxn}_R^j . \text{suite}$, then we say that $\pi_s^P . \text{cxn}_W^i$ has a Π -channel to $\pi_s^{P'} . \text{cxn}_R^j$.

By not demanding that $i = j$ in Def. 5, we enable modeling of unknown key share. Namely, *any* two connections using the same channel key and suite share a channel, regardless if they were correctly initiated via ConnectionInit .

3 Keyed Two-Party stLHAE Channel Protocol Security

As discussed in the introduction and §2, channel protocols may take as input Π -schemes of various types as well as various KDF-functions. In this section, we introduce stateful length-hiding authenticated encryption (stLHAE) channel security and provide concrete constructions of channels where the set $\{\Pi\}$ is comprised of stLHAE schemes (see App. A.2).

While the following definition, Def. 6, is not necessary for defining the channel protocol environment, it is essential for consideration of its security. Termination messages affect connection status, and therefore the ability to send or receive messages. Consequently, an adversary could use the encryption of such messages to distinguish between ciphertexts (see Fig. 2).

Definition 6 (Terminate Message). *Let CHNL be a keyed two-party controllable channel protocol, and let $\mathcal{T} \subset \mathcal{M}$, the message space of CHNL . Then, for $m \in \mathcal{T}$, m is called a **terminate** message and, for all $m \in \mathcal{T}$, we add the following to the correctness requirements of CHNL :*

- if $(c, \pi_s^P.cxn_W^{i'}) \leftarrow \text{ChannelSnd}(m, [\ell], \pi_s^P.cxn_W^i)$ and $\pi_s^P.cxn_W^i.status = \text{active}$, then $\pi_s^P.cxn_W^{i'}.status \leftarrow \text{terminated}$, and
- if $(m, \pi_s^P.cxn_R^{i'}) \leftarrow \text{ChannelRcv}(p, \pi_s^P.cxn_R^i)$ and $\pi_s^P.cxn_R^i.status = \text{active}$, then $\pi_s^P.cxn_R^{i'}.status \leftarrow \text{terminated}$.

We define *secure stLHAE channels* in Def. 7. Note that we do not present a single, generic *secure channel* definition since security for channel protocols must be considered with respect to the Π -scheme goals (stLHAE, authentication, signatures, etc.). We envisage protocols possibly implementing several Π -schemes. For example, an encryption scheme Π_1 and authentication scheme Π_2 . Connections would be initiated using (Π_1, KDF_1) , (Π_2, KDF_2) , or no security scheme $([], [])$, denoted $\text{CHNL}[(\Pi_1, \text{KDF}_1), (\Pi_2, \text{KDF}_2), ([], [])]$. In terms of real-world communication, a session could have, and close, connections running HTTPS, while maintaining other connections that send and receive information unprotected via HTTP (i.e. the $([], [])$ suite). It may then, via a TLS key-renegotiation under the existing session key, initiate new connections for channels that will be protected under an stLHAE Π -scheme. Such possibilities extend the current security considerations and are left for future work.

Definition 7 (stLHAE Channel Security). *Let $\text{CHNL}[\{(II, \text{KDF})\}]$ be a keyed two-party controllable channel protocol such that $\{II\}$ are stLHAE schemes, and let \mathcal{A} be a PPT adversarial algorithm. The stLHAE experiment for $\text{CHNL}[\{(II, \text{KDF})\}]$ is given by $\text{Exp}_{\text{CHNL}[\{(II, \text{KDF})\}]}^{\text{stlhae-chnl}}$ in Fig. 2. We define*

$$\text{Adv}_{\text{CHNL}[\{(II, \text{KDF})\}]}^{\text{stlhae-chnl}}(\mathcal{A}) = 2 \Pr \left[\text{Exp}_{\text{CHNL}[\{(II, \text{KDF})\}], \mathcal{A}}^{\text{stlhae-chnl}}(\lambda) \right] - 1 .$$

An channel protocol $\text{CHNL}[\{(II, \text{KDF})\}]$ is a secure channel stLHAE protocol if $\text{Adv}_{\text{CHNL}[\{(II, \text{KDF})\}]}^{\text{stlhae-chnl}}(\mathcal{A})$ is a negligible function in λ for all PPT adversaries \mathcal{A} .

$\text{Exp}_{\text{CHNL}[\{(II, \text{KDF})\}], \mathcal{A}}^{\text{stLHAE-chnl}}(\lambda):$	$\text{SnPairInit}((P, s), (P', \bar{s})):$
1: list $\leftarrow \perp$ 2: $b'_{(P^*, s^*, i^*)}$ $\leftarrow \mathcal{A}^{\text{SnPairInit}(\cdot), \text{CxnInit}(\cdot), \text{ChnlEnc}(\cdot), \text{ChnlDec}(\cdot)}()$ 3: return $(b'_{(P^*, s^*, i^*)} = b_{(P^*, s^*, i^*)})$	1: $SK \leftarrow \text{SessionKeyGen}(1^\lambda, P, P', s, \bar{s})$ 2: return \perp
$\text{CxnInit}(\pi_s^P, II^*, \text{KDF}^*, i):$	$\text{ChnlDec}((\text{ad}, c), \pi_s^P . \text{cxn}_R^i):$
1: if $(II^*, \text{KDF}^*) \notin \{(II, \text{KDF})\}$ then 2: return \perp 3: if $(\pi_s^P . \text{cxn}_W^i \in \text{list}) \vee (\pi_s^P . \text{cxn}_R^i \in \text{list})$ then 4: return \perp 5: $(\pi_s^P . \text{cxn}_W^i, \pi_s^P . \text{cxn}_R^i) \leftarrow \text{ConnectionInit}(SK, \pi_s^P, II^*, \text{KDF}^*, i)$ 6: list $\leftarrow \text{list} \cup \{(\pi_s^P . \text{cxn}_W^i, \pi_s^P . \text{cxn}_R^i)\}$ 7: $u_{(P, s, i)} \leftarrow 0, v_{(P, s, i)} \leftarrow 0, \text{phase}_{(P, s, i)} \leftarrow 0$ 8: $b_{(P, s, i)} \xleftarrow{\$} \{0, 1\}$ 9: return $(\pi_s^P . \text{cxn}_W^i, \pi_s^P . \text{cxn}_R^i)$	1: if $b = 0$ then 2: return $(\perp, \pi_s^P . \text{cxn}_R^i . \text{status})$ 3: $v \leftarrow v + 1$ 4: $(\text{ad}, m), \pi_s^P . \text{cxn}_R^i$ $\leftarrow \text{ChannelRcv}((\text{ad}, c), \pi_s^P . \text{cxn}_R^i)$ 5: if $(\exists \pi_{\bar{s}}^{P'} . \text{cxn}_W^i: \pi_{\bar{s}}^{P'} . \text{cxn}_W^i \text{ has a channel to } \pi_s^P . \text{cxn}_R^i) \wedge (\nexists j, j \neq i: \pi_{\bar{s}}^{P'} . \text{cxn}_W^j \text{ has a channel to } \pi_s^P . \text{cxn}_R^i)$ then 6: $u \leftarrow u_{(P', \bar{s}, i)}$ 7: else 8: $u \leftarrow 0$ 9: if $(u < v) \vee (c \neq \text{sent}.c_v) \vee (\text{ad} \neq \text{sent}.ad_v)$ then 10: phase $\leftarrow 1$ 11: if phase = 1 then 12: return $(\text{ad}, m), \pi_s^P . \text{cxn}_R^i . \text{status}$ 13: return $(\perp, \pi_s^P . \text{cxn}_R^i . \text{status})$
$\text{ChnlEnc}((\text{ad}, (m_0, m_1)), \ell, \pi_s^P . \text{cxn}_W^i):$	
1: $u \leftarrow u + 1$ 2: if $((\text{ad}, m_0) \in \mathcal{T}) \wedge ((\text{ad}, m_1) \notin \mathcal{T})$ or $((\text{ad}, m_1) \in \mathcal{T}) \wedge ((\text{ad}, m_0) \notin \mathcal{T})$ then 3: return $(\perp, \pi_s^P . \text{cxn}_W^i . \text{status})$ 4: $(c^{(0)}, \pi_s^P . \text{cxn}_W^i . (0))$ $\leftarrow \text{ChannelSnd}((\text{ad}, m_0), [\ell], \pi_s^P . \text{cxn}_W^i)$ 5: $(c^{(1)}, \pi_s^P . \text{cxn}_W^i . (1))$ $\leftarrow \text{ChannelSnd}((\text{ad}, m_1), [\ell], \pi_s^P . \text{cxn}_W^i)$ 6: if $c^{(0)} = \perp$ or $c^{(1)} = \perp$ then 7: return $(\perp, \pi_s^P . \text{cxn}_W^i . \text{status})$ 8: $(\text{sent}.ad_u, \text{sent}.c_u) := (\text{ad}, c^{(b)})$ 9: $\pi_s^P . \text{cxn}_W^i \leftarrow \pi_s^P . \text{cxn}_W^i . (b)$ 10: return $(\text{sent}.c_u, \pi_s^P . \text{cxn}_W^i . \text{status})$	

Fig. 2. Oracles of the stLHAE $\text{CHNL}[\{(II, \text{KDF})\}]$ security experiment, where II is specified in $\pi_s^P . \text{cxn}_W^i . \text{suite}$ (resp. $\pi_s^P . \text{cxn}_R^i . \text{suite}$). For conciseness, the synchronization variable $u_{(P, s, i)}$ is referenced as u in ChnlEnc (resp. for $v_{(P, s, i)}$). Similarly $b_{(P, s, i)}$ is referenced as b .

In the security game for $\text{CHNL}[\{(II, \text{KDF})\}]$ in Fig. 2, the adversary may select (II, KDF) pairs used to initiate the channel connections. However, if the pair is not valid – in the set $\{(II, \text{KDF})\}$ – no connection will be initiated. ChnlSnd may be called on any connection as ChnlSnd adaptively uses ChannelSnd constructed from the scheme II from the connection’s internal variable $\pi_s^P . \text{cxn}_W^i . \text{suite}$ (analogously ChnlRcv). For agility between suites, every connection pair is initialized with the same session key SK .

Remark 3. Since multiple Π schemes may be implemented by $\text{CHNL}[\{(II, \text{KDF})\}]$, we envisage that it is possible to run various but simultaneous (or consecutive) experiments on the different connections – all linked to the master session key. Note the applications to real-world protocols: in 802.11 [19], TKIP uses a pairwise master key to derive 2 sets of MAC and 2 sets of encryption keys for EAPOL and application data protection. Different, simultaneous experiments can also be considered for the MAC and encryption goals. This is left for future work.

Definition 8 (Channel Construction from stLHAE). *Let $\{\Pi\}$, where $\Pi = (\text{Kgn}, \text{Enc}, \text{Dec})$, be stLHAE scheme(s) and let $\{\text{KDF}\}$ be key derivation function(s). A keyed two-party controllable channel protocol $\text{CHNL}[\{(II, \text{KDF})\}]$ is constructed to achieve a pair of linked stLHAE channels, with message space $\mathcal{M} = \mathcal{AD} \times \mathcal{AE}\mathcal{D}$ and packet space $\mathcal{P} = \mathcal{AD} \times \mathcal{C}$, as follows:*

- $\text{SessionKeyGen}(1^\lambda, P, P', s, \bar{s})$:
 - Selects a shared session key SK according to the distribution of key space,
 - sets $\pi_s^P.\alpha \leftarrow 1$, $\pi_{\bar{s}}^{P'}.\alpha \leftarrow 1$,
 - sets respective partner identifiers $\pi_s^P.(P', \bar{s})$ and $\pi_{\bar{s}}^{P'}.(P, s)$,
 - sets $\pi_s^P.SK \leftarrow SK$ and $\pi_{\bar{s}}^{P'}.SK \leftarrow SK$.

Return SK .

- $\text{ConnectionInit}(SK, \pi_s^P, \Pi, \text{KDF}, i)$:

If $\text{ConnectionInit}(\cdot, \pi_s^P, \cdot, \cdot, i)$ has previously been called, return \perp .

 - Compute $\pi_s^P.\text{c}x\text{n}_{\bar{w}}^i.CK \leftarrow \text{KDF}(SK, (P, s), (P', \bar{s}), i)$ and
 - $\pi_{\bar{s}}^{P'}.\text{c}x\text{n}_{\bar{r}}^i.CK \leftarrow \text{KDF}(SK, (P', \bar{s}), (P, s), i)$.

If $\pi_s^P.\text{c}x\text{n}_{\bar{w}}^i.CK$ or $\pi_{\bar{s}}^{P'}.\text{c}x\text{n}_{\bar{r}}^i.CK$ are not in key space \mathcal{K} of Π , return \perp .

 - Set status $\pi_s^P.\text{c}x\text{n}_{\bar{w}}^i.\text{status} = \text{active}$ and $\pi_{\bar{s}}^{P'}.\text{c}x\text{n}_{\bar{r}}^i.\text{status} = \text{active}$,
 - set suite $\pi_s^P.\text{c}x\text{n}_{\bar{w}}^i.\text{suite} \leftarrow (\Pi, \text{KDF})$, and $\pi_{\bar{s}}^{P'}.\text{c}x\text{n}_{\bar{r}}^i.\text{suite} \leftarrow (\Pi, \text{KDF})$,
 - $\pi_s^P.\text{c}x\text{n}_{\bar{w}}^i.\text{substate} \leftarrow 0$, and $\pi_{\bar{s}}^{P'}.\text{c}x\text{n}_{\bar{r}}^i.\text{substate} \leftarrow 0$.

Return $(\pi_s^P.\text{c}x\text{n}_{\bar{w}}^i, \pi_{\bar{s}}^{P'}.\text{c}x\text{n}_{\bar{r}}^i)$.

- $\text{ChannelSnd}(\text{ad}, m, \ell, \pi_s^P.\text{c}x\text{n}_{\bar{w}}^i)$:

If $\pi_s^P.\text{c}x\text{n}_{\bar{w}}^i.\text{status} \neq \text{active}$, return $(\perp, \pi_s^P.\text{c}x\text{n}_{\bar{w}}^i)$.

Otherwise,

 - compute:

$(c, \pi_s^P.\text{c}x\text{n}_{\bar{w}}^i.\text{substate}') \stackrel{\$}{\leftarrow} \text{Enc}(\pi_s^P.\text{c}x\text{n}_{\bar{w}}^i.CK, \ell, \text{ad}, m, \pi_s^P.\text{c}x\text{n}_{\bar{w}}^i.\text{substate})$, where Enc is specified by the scheme Π defined in $\pi_s^P.\text{c}x\text{n}_{\bar{w}}^i.\text{suite}$.
 - If $(\text{ad}, m) \in \mathcal{T}$, set $\pi_s^P.\text{c}x\text{n}_{\bar{w}}^i.\text{status} \leftarrow \text{terminated}$.

Return $(c, \pi_s^P.\text{c}x\text{n}_{\bar{w}}^i')$.

- $\text{ChannelRcv}(\text{ad}, c, \pi_{\bar{s}}^{P'}.\text{c}x\text{n}_{\bar{r}}^i)$:

If $\pi_{\bar{s}}^{P'}.\text{c}x\text{n}_{\bar{r}}^i.\text{status} \neq \text{active}$, return $(\perp, \pi_{\bar{s}}^{P'}.\text{c}x\text{n}_{\bar{r}}^i)$.

Otherwise,

 - compute:

$(m, \pi_{\bar{s}}^{P'}.\text{c}x\text{n}_{\bar{r}}^i.\text{substate}') \leftarrow \text{Dec}(\pi_{\bar{s}}^{P'}.\text{c}x\text{n}_{\bar{r}}^i.CK, \text{ad}, c, \pi_{\bar{s}}^{P'}.\text{c}x\text{n}_{\bar{r}}^i.\text{substate})$, where Dec is specified by the scheme Π in $\pi_{\bar{s}}^{P'}.\text{c}x\text{n}_{\bar{r}}^i.\text{suite}$.

If $m = \perp$, return $(\perp, \pi_{\bar{s}}^{P'}.\text{c}x\text{n}_{\bar{r}}^i)$.

If $(\text{ad}, m) \in \mathcal{T}$, set $\pi_{\bar{s}}^{P'}.\text{c}x\text{n}_{\bar{r}}^i.\text{status} \leftarrow \text{terminated}$. Return $((\text{ad}, m), \pi_{\bar{s}}^{P'}.\text{c}x\text{n}_{\bar{r}}^i')$.

Naturally, stLHAE security of a channel protocol construction ought to be reducible to the security of the stLHAE scheme(s) $\{II\}$ underlying the construction. Yet there is an additional consideration. While each II uses connection keys derived via the KDF, all KDFs use the same shared master session key SK . Consequently, we require *agility* for $\{KDF\}$. Agility for the entire set $\{KDF\}$ implies that the individual primitives can share SK securely [1].

Theorem 1. *Let $\text{CHNL}[\{(II_j, \text{KDF})\}]$ be a keyed two-party controllable channel protocol constructed from stLHAE scheme(s) $\{II_j\}$ and key derivation function(s) $\{\text{KDF}\}$, such that $\{\text{KDF}\}$ is a compatible, finite set which is agile with respect to pseudo-randomness.*

Let \mathcal{A} be an adversarial algorithm against the $\text{Exp}_{\text{CHNL}[\{(II, \text{KDF})\}]}^{\text{stlhae-chnl}, \mathcal{A}}(\lambda)$ experiment in Fig. 2. Let p be the number of identities and n be the maximum number of sessions at an identity. Then we can construct a $\{\text{KDF}\}$ -restricted adversarial algorithm \mathcal{B}_0 against the pseudo-randomness agility of $\{\text{KDF}\}$ in $\text{Exp}_{\{\text{KDF.F}\}}^{\text{agile-prf}}$, and adversarial algorithms \mathcal{B}_j against $\text{Exp}_{II_j}^{\text{stlhae}}$ from Fig. 6, such that

$$\text{Adv}_{\text{CHNL}[\{(II, \text{KDF})\}]}^{\text{stlhae-chnl}}(\mathcal{A}) \leq p^2 \cdot n^2 \cdot (\text{Adv}_{\{\text{KDF.F}\}}^{\text{agile-prf}}(\mathcal{B}_0) + \max_j \{\text{Adv}_{II_j}^{\text{stlhae}}(\mathcal{B}_j)\}) .$$

Proof. Consider a channel protocol $\text{CHNL}[\{(II_j, \text{KDF})\}]$, where $\{II_j\}$ is a set of stLHAE scheme(s) and $\{\text{KDF}\}$ is a finite, compatible set of KDF functions.

We will follow a series of game hops, where the adversarial advantage in Game i is denoted Adv_i .

Game 0 This is the original experiment,

$$\text{Adv}_0 = \text{Adv}_{\text{CHNL}[\{(II, \text{KDF})\}]}^{\text{stlhae-chnl}}(\mathcal{A}) .$$

Game 1 Initializing the simulation, the challenger selects its guess of identities P and P' from the set of all identities $\{P_1, \dots, P_p\}$, as well as send/receive target sessions π_s^P and $\pi_{\bar{s}}^{P'}$, respectively, via random indices s and \bar{s} from $\{1, \dots, n\}$ where n is the maximum number of sessions at an identity. If \mathcal{A} does not call Cxnlnit on the guessed target session pair, or attempt to win with a bit guess $b'_{(P^*, s^*, i^*)}$ for $(P^*, s^*) \in \{(P, s), (P', \bar{s})\}$, then \mathcal{B} aborts. Thus,

$$\text{Adv}_0 = p^2 \cdot n^2 \cdot \text{Adv}_1 .$$

Game 2 In this game we replace the channel keys $\{CK\}$ with randomly sampled keys for all connections at the target sessions π_s^P and $\pi_{\bar{s}}^{P'}$. Distinguishing between Games 1 and 2 implies that we can build an adversary \mathcal{B}_0 which succeeds against the PRF agility of the set $\{\text{KDF}\}$. It follows that

$$\text{Adv}_1 = \text{Adv}_{\{\text{KDF.F}\}}^{\text{agile-prf}}(\mathcal{B}_0) + \text{Adv}_2 .$$

Following from Game 2, we have that all $\text{KDF} \in \{\text{KDF}\}$ can securely share the same session key SK . Thus, all channel keys $\{CK\}$ used in CHNL connections output by CxnlInit are indistinguishable from random.

Now we build adversaries against the stLHAE schemes $\{\Pi_j\}$, in each connection initiated by CxnlInit. We assign \mathcal{B}_j as an adversary against the Π_j -th scheme. If more than one connection is initiated with a given scheme Π_j , we build additional adversaries $\mathcal{B}_{j,1}, \mathcal{B}_{j,2}$, etc. Thus, we construct an adversary against the stateful authentication scheme for each connection at the target sessions.

All ChnlEnc and ChnlDec-queries for connections at the target sessions are answered via the Enc and Dec oracles maintained by each \mathcal{B} (see Figure ??). All ChnlEnc and ChnlDec-queries at other connections (i.e. at non-target sessions) can be correctly answered by the challenger via ChannelSnd and ChannelRcv algorithms. Thus we can correctly simulate queries for all connections.

By the success of \mathcal{A} , it must output a bit $b'_{(P^*, s^*, i^*)}$ such that $(P^*, s^*) \in \{(P, s), (P', \bar{s})\}$ and $\text{Exp}_{\text{CHNL}[(\Pi_j, \text{KDF}), \mathcal{A}]}^{\text{stlhae-chnl}}(\lambda)$ returns *true*. Since (P^*, s^*, i^*) corresponds to a specific connection at one of the target sessions, and all queries at such connections are answered via oracles for various \mathcal{B} adversaries, an adversary \mathcal{B}_{j^*} also wins with the bit guess $b'_{(P, s, i^*)}$ in the experiment $\text{Exp}_{\Pi_{j^*}, \mathcal{B}_{j^*}}^{\text{stlhae}}(\lambda)$. \mathcal{B}_{j^*} may be one of several adversaries against Π_{j^*} if more than one connection is initiated with Π_{j^*} . We bound the winning advantage by the maximum over all possible advantages, with $1 \leq j \leq |\{\Pi\}|$. \square

4 Secure Termination

Ultimately, the goal of secure termination is a guarantee to the receiver connection that no further messages are being sent. We define *closure* in the contexts of both connections and channels before presenting the secure termination adversarial advantage and experiment.

4.1 Closure Alerts and Channel Closure

Inherently, secure termination is the closure of connections, controlled by the sending and receipt of signifying messages from a subset of the message space.

Definition 9 (Connection Closure). *Let CHNL be a keyed two-party controllable channel protocol and π_s^P be a session of CHNL.*

- *The i -th write channel connection at π_s^P is said to be closed if $\pi_s^P.\text{cxn}_W^i.\text{status} = \text{terminated}$.*
- *The i -th read channel connection at π_s^P is said to be closed if $\pi_s^P.\text{cxn}_R^i.\text{status} = \text{terminated}$.*

If all channels at a given session are closed, the session is said to be closed.

Definition 10 (Channel Closure). *We say that a channel from $\pi_s^P.\text{cxn}_W^i$ to $\pi_{\bar{s}}^{P'}.\text{cxn}_R^i$ closes using a message $m^* \in \mathcal{M}$ if*

- $(c, \pi_s^P.cxn_W^i) \stackrel{\S}{\leftarrow} \text{ChannelSnd}(m^*, [\ell], \pi_s^P.cxn_W^i)$, for some $c \in \mathcal{C}$,
where $\pi_s^P.cxn_W^i.status = \mathbf{active}$, and $\pi_s^P.cxn_W^{i'}.status = \mathbf{terminated}$, and
- $(m^*, \pi_s^{P'}.cxn_R^i) \leftarrow \text{ChannelRcv}(p, \pi_s^{P'}.cxn_R^i)$, for some $p \in \mathcal{P}$,
where $\pi_s^{P'}.cxn_R^i.status = \mathbf{active}$ and $\pi_s^{P'}.cxn_R^{i'}.status = \mathbf{terminated}$.

The sending, resp. receiving, of any **terminate** in the set of termination messages $\mathcal{T} \subset \mathcal{M}$ results in channel closure. Moreover, if the channel from π_s^P to $\pi_s^{P'}$ closes using a **terminate** message, we say that π_s^P initiated the channel closure.

Note that the sending, resp. receiving, of a **terminate** message $m \in \mathcal{T}$ results in *connection closure*; however, *channel closure* demands a causal relationship between both end connections based on m .

Remark 4. While a session may consist of multiple channel connections at any given time, it may also consist of no connections (this may occur when no channel parameters have been negotiated, or all channels have been *closed*). The number of channel connections may change during a session’s lifespan, namely by means of session *closure* and *resumption*. Resumption is realized via `ConnectionInit`.

4.2 Secure Termination Experiment

Definition 11 (Secure Termination Experiment). Let `CHNL` be a keyed two-party controllable channel protocol with message space \mathcal{M} , and let \mathcal{A} be an adversary algorithm. Let **terminate** be an element of $\mathcal{T} \subset \mathcal{M}$, where \mathcal{T} is the set of all termination messages. With the secure termination experiment for `CHNL` given by $\text{Exp}_{\text{CHNL}}^{\text{sc.t}}$ in Figure 3, define

$$\text{Adv}_{\text{CHNL}}^{\text{sc.t}}(\mathcal{A}) = \Pr [\text{Exp}_{\text{CHNL}, \mathcal{A}}^{\text{sc.t}}(\lambda) = 1] \ .$$

The existence of a signifying **terminate** message (i.e. $\mathcal{T} \neq \emptyset$) is insufficient for claiming that a protocol *always* securely terminates connections. Only connection closure with such a **terminate** message yields secure termination of the connection. Intrinsically, secure termination is a property of the read side of a channel, which is achieved upon receipt of **terminate**.

Remark 5. Secure termination is a per-connection goal, achieved on message receipt, and handles adversarial intervention in channel closure – what is commonly referred to as a *truncation attack*. Truncation, malicious or otherwise, can happen via various means; for example, closure of the underlying TCP connection as a means of closing a TLS channel. The adversary wins the secure termination security game only if it is able to make an honest party accept that the connection has been correctly terminated when it has not been. Thus malicious closure of the TCP connection is not a valid attack.

Remark 6. Some protocols have “terminate”-looking messages that may not be authenticated at all or authenticated properly, and our model aids in understanding and comparing the security of such protocols against truncation attacks. For

$\text{Exp}_{\text{CHNL}[\{(II, \text{KDF})\}], \mathcal{A}}^{\text{sc.t}}(\lambda):$ <ol style="list-style-type: none"> 1: phase $\leftarrow 0$, list $\leftarrow \perp$ 2: $\mathcal{A}^{\text{SnPairInit}(\cdot), \text{ConnectionInit}(\cdot), \text{Send}_{\top}(\cdot), \text{Receive}_{\top}(\cdot)}()$ 3: return phase <hr/> $\text{ConnectionInit}(\pi_s^P, \Pi^*, \text{KDF}^*, i):$ <ol style="list-style-type: none"> 1: if $(\Pi^*, \text{KDF}^*) \notin \{(II, \text{KDF})\}$ then 2: return \perp 3: if $(\pi_s^P.\text{cxn}_w^i \in \text{list}) \vee (\pi_s^P.\text{cxn}_r^i \in \text{list})$ then 4: return \perp 5: $(\pi_s^P.\text{cxn}_w^i, \pi_s^P.\text{cxn}_r^i) \leftarrow \text{ConnectionInit}(SK, \pi_s^P, \Pi^*, \text{KDF}^*, i)$ 6: list $\leftarrow \text{list} \cup \{\pi_s^P.\text{cxn}_w^i, \pi_s^P.\text{cxn}_r^i\}$ 7: return $(\pi_s^P.\text{cxn}_w^i, \pi_s^P.\text{cxn}_r^i)$ <hr/> $\text{Send}_{\top}(m, [\ell], \pi_s^P.\text{cxn}_w^i):$ <ol style="list-style-type: none"> 1: $(c, \pi_s^P.\text{cxn}_w^i) \leftarrow \text{ChannelSnd}(m, [\ell], \pi_s^P.\text{cxn}_w^i)$ 2: return $(c, \pi_s^P.\text{cxn}_w^i.\text{status})$ 	$\text{SnPairInit}((P, s), (P', \bar{s})):$ <ol style="list-style-type: none"> 1: $SK \leftarrow \text{SessionKeyGen}(1^\lambda, P, P', s, \bar{s})$ 2: return \perp <hr/> $\text{Receive}_{\top}(p, \pi_s^P.\text{cxn}_r^i):$ <ol style="list-style-type: none"> 1: $(m, \pi_s^P.\text{cxn}_r^i) \leftarrow \text{ChannelRcv}(p, \pi_s^P.\text{cxn}_r^i)$ 2: if $(m \in \mathcal{T}) \wedge (\pi_s^P.\text{cxn}_r^i.\text{status} = \text{terminated}) \wedge ((\exists \pi_{\bar{s}}^{P'}.\text{cxn}_w^i : \pi_{\bar{s}}^{P'}.\text{cxn}_w^i \text{ has a channel to } \pi_s^P.\text{cxn}_r^i) \implies (\pi_{\bar{s}}^{P'}.\text{cxn}_w^i.\text{status} \neq \text{terminated}))$ then 3: phase $\leftarrow 1$ 4: return phase from experiment 5: return $\pi_s^P.\text{cxn}_r^i.\text{status}$
---	---

Fig. 3. Secure termination experiment for a protocol $\text{CHNL}[\{(II, \text{KDF})\}] = (\text{SessionGen}, \text{ConnectionInit}, \text{ChannelSnd}, \text{ChannelRcv})$ with message space \mathcal{M} , $\mathcal{T} \subset \mathcal{M}$, and adversary \mathcal{A} , where Send_{\top} and Receive_{\top} are constructed from (II, KDF) as defined by $\pi_s^P.\text{cxn}_w^i.\text{suite}$ and $\pi_s^P.\text{cxn}_r^i.\text{suite}$.

example, the `DeauthenticationRequest` of 802.11 *may* meet termination message requirements, but that is dependent on statefulness and whether or not the requests are protected as Robust Management Frames (RMF). A session that does not negotiate to send such requests as RMFs is susceptible to truncation attacks. Similarly, the security implications of 802.11's `DisassociationRequest` should be questioned. Many protocols have specified messages which appear to indicate termination; however, exactly what cryptographic guarantees are provided on receipt has not been well understood, in the absence of a secure termination model. Hence the success of truncation attacks.

Definition 12 (Secure Termination of a Protocol). *Let CHNL be a two-party controllable channel protocol, let π_s^P and $\pi_{\bar{s}}^{P'}$ be any two sessions of CHNL , and let π_s have a channel to $\pi_{\bar{s}}$ according to Def. 5.*

We say that CHNL achieves secure termination on the channel from π_s to $\pi_{\bar{s}}$ with `terminate` if

- *the channel from π_s^P to $\pi_{\bar{s}}^{P'}$ closes using a termination message `terminate` $\in \mathcal{T}$, according to Def. 10, and*
- $\text{Adv}_{\text{CHNL}}^{\text{sc.t}}(\mathcal{A})$ *is a negligible function in λ for all PPT adversaries \mathcal{A} .*

Naturally, the fact that the read connection of the closure-initiator's session may not necessarily be required to close when the write connection sends a `terminate` message gives rise to the concepts of *fatal* and *graceful secure termination*. When a session π_s^P initiates closure with `terminate` $\in \mathcal{T}$ such that only the write connection closes, graceful closure *can* be achieved by waiting for a corresponding `terminate` message to be sent to the read connection of π_s^P before the read connection closes. Thus both sessions sharing the channels *may* achieve secure termination. Comparatively, if π_s^P initiates closure with a fatal `terminate` message, it *cannot* achieve secure termination – only the receiving session may achieve it.

Definition 13 (Fatal and Graceful Secure Termination). *Let π_s^P and $\pi_{\bar{s}}^{P'}$ be sessions of a two-party controllable channel protocol CHNL such that π_s^P has an channel to $\pi_{\bar{s}}^{P'}$ according to Def. 5. Let `terminate` $\in \mathcal{T}$.*

- *If, upon running $(c, \pi_s^P.cxn_{\bar{w}}^i) \stackrel{\S}{\leftarrow} \text{ChannelSnd}(\text{terminate}, \pi_s^P.cxn_{\bar{w}}^i)$, both the write and read connections at π_s^P are closed according to Def. 9, then `terminate` is said to be a fatal termination message.*
- *If, upon running $(c, \pi_s^P.cxn_{\bar{w}}^i) \stackrel{\S}{\leftarrow} \text{ChannelSnd}(\text{terminate}, \pi_s^P.cxn_{\bar{w}}^i)$, only the write connection at π_s^P is closed according to Def. 9, then `terminate` is said to be a graceful termination message.*

4.3 Reduction to stLHAE Security

Secure termination depends upon the relay of messages with unaltered content and therefore its security is reducible to that of the authentication guarantees of the channel.

Theorem 2. *Let $\text{CHNL}[\{(II, \text{KDF})\}]$ be a keyed two-party controllable channel protocol constructed from authentication scheme(s) $\{II\}$, with a message space \mathcal{M} and $\mathcal{T} \subset \mathcal{M}$, and let \mathcal{A} be an adversarial algorithm in the $\text{Exp}_{\text{CHNL}[\{(II, \text{KDF})\}]}^{\text{sc.t}}$ experiment. Then we can construct an adversarial algorithm \mathcal{B} against $\text{Exp}_{\text{CHNL}[\{(II, \text{KDF})\}]}^{\text{stlhae-chnl}}$ such that*

$$\mathbf{Adv}_{\text{CHNL}[\{(II, \text{KDF})\}]}^{\text{sc.t}}(\mathcal{A}) \leq \mathbf{Adv}_{\text{CHNL}[\{(II, \text{KDF})\}]}^{\text{stlhae-chnl}}(\mathcal{B}) .$$

Proof. Let $\text{CHNL}[\{(II, \text{KDF})\}]$ be a keyed two-party controllable channel protocol according to Definition 5 and Definition 8, with message space \mathcal{M} where $\mathcal{T} \subset \mathcal{M}$, and let \mathcal{A} be an adversary algorithm against the secure termination security of $\text{CHNL}[\{(II, \text{KDF})\}]$.

The challenger answers `SnPairInit` queries with its `SnPairInit` oracle, and `ConnectionInit` queries using its `CxnInit` oracle. Similarly, the challenger answers `SendT` with his `ChnlEnc` oracle, letting $m = m_0 = m_1$. The challenger answers `ReceiveT` queries with his `ChnlDec` oracle.

\mathcal{A} must eventually call `ReceiveT`($p, \pi_{\bar{s}}^{P'}.cxn_{\bar{r}}^i$) on a packet $p = (\text{ad}, c)$, for some $\pi_{\bar{s}}^{P'}.cxn_{\bar{r}}^i$ such that $(m, \pi_{\bar{s}}^{P'}.cxn_{\bar{r}}^i) \leftarrow \text{ChannelRcv}(p, \pi_{\bar{s}}^{P'}.cxn_{\bar{r}}^i)$ where $m \in \mathcal{T}$ and

phase $\leftarrow 1$ in $\text{Exp}_{\text{CHNL}\{\{(H, \text{KDF})\}\}, \mathcal{A}}^{\text{sc.t}}(\lambda)$. However, by the success of \mathcal{A} and the correctness of the ChannelSnd algorithm, we have that $(c \neq \text{sent}.c_v) \vee (\text{ad} \neq \text{sent}.ad_v)$. Thus \mathcal{B} successfully wins in the $\text{Exp}_{\text{CHNL}\{\{(H, \text{KDF})\}\}, \mathcal{B}}^{\text{stlhae-chnl}}(\lambda)$ experiment. \square

5 Secure Channels and Termination in TLS 1.2

TLS is one of the most important security protocols in the world today and serves as the backbone of internet security. Due to the lack of key indistinguishability in TLS 1.2, many works analyzing it rely on the ACCE model [20]. We show that the *post-accept* phase channel ACCE security model can be viewed as a highly restricted case of channel protocol stateful length-hiding AE (stLHAE) security, and provide positive secure termination results for TLS 1.2 under this restricted case. Note that the use of the *post-accept* ACCE model, together with the coinciding channel restrictions, is not inherent; we demonstrate the correlation to bypass a reanalysis of TLS and directly consider secure termination. Analyzing TLS in our channel protocol model would allow for consideration of parallel sessions and connections, as well as resumption, key derivation functions, etc.

5.1 Comparing Channel Protocols and ACCE

ACCE employs a *pre-accept* phase and a *post-accept* phase. While the former handles all protocol interactions before a session key has been accepted, the latter is correlated to channel security under the agreed-upon session key. Breaking ACCE security is described in terms of either getting a session oracle to accept maliciously, or correctly answering a stLHAE encryption challenge.

Since CHNL session keys need not be sampled uniformly at random, it is possible to adapt the CHNL stLHAE construction (Def. 8) and allow for SK to be derived by other means (e.g. as the output of a handshake in the ACCE *pre-accept* phase). Thus we can link the CHNL stLHAE security experiment to the ACCE *post-accept* phase experiment (see App. ??). The ACCE model uses the session key directly for the stLHAE primitive, with key derivation consisting of splitting the session key into two in order to obtain write/read keys. Since only one session key exists and is used directly as the connection keys (e.g. on cxn_w, cxn_r), only one connection pair is allowed per session. Another session key would be required for further connections if key separation is to be achieved. This assumption is inherent to the ACCE model TLS analyses to date, where SK is assumed to be a concatenation of the write and read keys (e.g. [23]).

However, in the $\text{CHNL}\{\{(H, \text{KDF})\}\}$ security game, connections are input into ChnlSnd and ChnlRcv , allowing for the game to be run on various connections of a given session, and an adversary is allowed to initialize connections of its choice under various AEAD schemes, with connection keys derived from the master session key. In comparison, the ACCE model uses the session key directly for the AEAD primitive. Thus, the correlation between channel protocol security and ACCE is highly constrained to cases where a single AEAD scheme is employed and in the absence of a specified KDF (e.g. the singleton suite $(H, [])$).

Despite being a highly restricted case, we compare the security guarantees of the channel phase of the ACCE experiment and $\text{CHNL}[(\Pi, \square)]$. This is to illustrate the applicability of the CHNL environment, although those protocols currently analyzed in the ACCE model need to be reanalyzed in the CHNL model to truly consider security under parallel connections, etc.

Theorem 3. *Let $\text{CHNL}[(\Pi, \square)]$ be a keyed two-party controllable channel protocol constructed from a single stLHAE scheme Π , using the identity KDF \square . Let p be the number of principals, n be the maximum number of sessions at a principal, and let \mathcal{A} be an adversarial algorithm against the $\text{Exp}_{\text{CHNL}[(\Pi, \square)]}^{\text{stlhae-chnl}}()$ experiment in Fig. 2. Then we can construct an adversarial algorithm \mathcal{B} against $\text{Exp}_{\Pi}^{\text{ACCE}}()$ such that*

$$\text{Adv}_{\text{CHNL}[(\Pi, \square)]}^{\text{stlhae-chnl}}(\mathcal{A}) \leq p^2 \cdot n^2 \cdot \text{Adv}_{\Pi}^{\text{ACCE}}(\mathcal{B}) .$$

Proof. Again we follow a series of game hops where the adversarial advantage in Game i is denoted Adv_i . For all identities and sessions, session keys are sampled according to the distribution of the key space \mathcal{K} .

Game 0 This is the original experiment,

$$\text{Adv}_0 = \text{Adv}_{\text{CHNL}[(\Pi, \square)]}^{\text{stlhae-chnl}}(\mathcal{A}) .$$

Game 1 Initializing the simulation, the challenger selects its guess of identities P and P' from the set of all identities $\{P_1, \dots, P_p\}$, as well as send/receive target sessions π_s^P and $\pi_{\bar{s}}^{P'}$, respectively, via random indices s and \bar{s} from $\{1, \dots, n\}$ where n is the maximum number of sessions at an identity. If \mathcal{A} does not call Cxnlnit on the guessed target session pair, calls Cxnlnit more than once on either session, or does not attempt to win with a bit guess $b'_{(P^*, s^*, i^*)}$ for $(P^*, s^*) \in \{(P, s), (P', \bar{s})\}$, then \mathcal{B} aborts. Thus,

$$\text{Adv}_0 = p^2 \cdot n^2 \cdot \text{Adv}_1 .$$

The challenger maintains *status* variables, and answers ChnlEnc and ChnlDec queries with \mathcal{B} 's Encrypt and Decrypt oracles, respectively, run on connection keys from the single initiated connection pair for the target session pair. All ChnlEnc and ChnlDec -queries at other connections (i.e. at non-target sessions) can be correctly answered by the challenger via ChannelSnd and ChannelRcv algorithms.

\mathcal{A} will eventually output a bit $b'_{(P^*, s^*, i^*)}$ for $(P^*, s^*) \in \{(P, s), (P', \bar{s})\}$, and \mathcal{B} also outputs the same bit $b'_{(P^*, s^*, i^*)}$. Thus \mathcal{B} wins in the post-accept experiment $\text{Exp}_{\Pi, \mathcal{B}}^{\text{ACCE}}(\lambda)$ of the ACCE game whenever \mathcal{A} wins the $\text{Exp}_{\text{CHNL}[(\Pi, \square)], \mathcal{A}}^{\text{stlhae-chnl}}(\lambda)$ experiment. □




Fig. 4. TLS fatal alert behavior, where *initiator* is the session initiating a channel closure. Both the read and write *cxn* of the initiator close immediately upon sending.




Fig. 5. TLS `close_notify` behavior. The write *cxn* of the initiator closes upon sending of `close_notify`. Upon receipt, the read *cxn* at the responder closes. Subsequently, the responder sends a corresponding `close_notify` alert and closes its write *cxn*.

5.2 Secure Termination in TLS

According to specification, the Alert Protocol in TLS falls into three categories: fatal alerts, warning alerts, and `close_notify` alerts. Unlike fatal alerts, which upon sending/receipt close both write and read connections at a session, `close_notify` alerts do not necessarily close the initiator’s receive connection immediately, but *may* wait until receipt of the reciprocal `close_notify` alert. Fig. 4 and 5 illustrate fatal and `close_notify` alert behavior. According to the TLS 1.2 standard, the determination of whether or not the initiator’s read connection should be closed when a `close_notify` is sent is left to the usage profile [11, p. 28]. However, upon receipt of a `close_notify` alert, the responder *must* close its read connection and *must* send a corresponding `close_notify` alert on its send connection, before closing it also. Receipt of the reciprocal alert results in the closure of the original initiator’s read connection.

Ostensibly, TLS fatal alerts are *fatal termination messages* per Def. 13 and the TLS 1.2 specification [11]. Comparatively, `close_notify` alerts are either *fatal* or *graceful termination messages*, depending on the implementation. Per specification, “It is not required for the initiator of the close to wait for the responding `close_notify` alert before closing the read side of the connection.”

Remark 7. It is crucial to note that Alert Protocol messages are in fact, within the message space of the TLS protocol according to Def. 5 and Def. 8, despite the Alert Protocol running on top of the Record Layer Protocol. This is due to the fact that TLS 1.2 encrypts and sequences these messages in the same manner as application messages [11, section 7.2].

We conclude with the following theorem on secure termination in TLS, under the restricted (ACCE-induced) channel environment of prior TLS analyses.

Theorem 4. *For TLS 1.2, TLS-RSA, TLS-CCA, TLS-DH, and TLS-DHE achieve secure termination under any fatal alert or close_notify alert, where each TLS session consists of a single connection pair with connection keys derived from the session key via the identity KDF.*

Theorem 4 follows by combining previous ACCE analyses of TLS 1.2 [29, 20, 23, 7], Theorem 3, and Theorem 2, and applying the observation above that *fatal* and *close_notify* alerts satisfy the definition of `terminate` messages within $\mathcal{T} \subset \mathcal{M}$, where \mathcal{M} is the protocol message space.

From previous analyses of TLS 1.2 in the ACCE model, Thm. 4 inherits the restriction that sessions consist of a single connection pair. However this is not intrinsic to the design of TLS and it may be analyzed in the keyed two-party controllable channel model for better consideration of the full protocol.

References

1. Acar, T., Belenkiy, M., Bellare, M., Cash, D.: Cryptographic agility and its relation to circular encryption. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 403–422. Springer, Heidelberg (May 2010)
2. Badertscher, C., Matt, C., Maurer, U., Rogaway, P., Tackmann, B.: Augmented secure channels and the goal of the TLS 1.3 record layer. In: Au, M.H., Miyaji, A. (eds.) ProvSec 2015. LNCS, vol. 9451, pp. 85–104. Springer, Heidelberg (Nov 2015)
3. Bellare, M., Namprempre, C.: Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. *Journal of Cryptology* 21(4), 469–491 (Oct 2008)
4. Bergsma, F., Dowling, B., Kohlar, F., Schwenk, J., Stebila, D.: Multi-Ciphersuite Security of the Secure Shell (SSH) Protocol. In: Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security. pp. 369–381. CCS '14, ACM (2014), <http://doi.acm.org/10.1145/2660267.2660286>
5. Bhargavan, K., Delignat-Lavaud, A., Fournet, C., Pironti, A., Strub, P.Y.: Triple handshakes and cookie cutters: Breaking and fixing authentication over TLS. In: 2014 IEEE Symposium on Security and Privacy. pp. 98–113. IEEE Computer Society Press (May 2014)
6. Bhargavan, K., Fournet, C., Kohlweiss, M., Pironti, A., Strub, P.Y., Zanella Béguelin, S.: Proving the TLS handshake secure (as it is). In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part II. LNCS, vol. 8617, pp. 235–255. Springer, Heidelberg (Aug 2014)
7. Boyd, C., Hale, B., Mjølsnes, S.F., Stebila, D.: From stateless to stateful: Generic authentication and authenticated encryption constructions with application to TLS. In: Sako, K. (ed.) CT-RSA 2016. LNCS, vol. 9610, pp. 55–71. Springer, Heidelberg (Feb / Mar 2016)
8. Brier, E., Peyrin, T.: A forward-secure symmetric-key derivation protocol - how to improve classical DUKPT. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 250–267. Springer, Heidelberg (Dec 2010)
9. Canetti, R., Krawczyk, H.: Analysis of key-exchange protocols and their use for building secure channels. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 453–474. Springer, Heidelberg (May 2001)
10. Cohen, R., Coretti, S., Garay, J., Zikas, V.: Probabilistic Termination and Composability of Cryptographic Protocols (2016), <http://eprint.iacr.org/>
11. Dierks, T., Rescorla, E.: The Transport Layer Security (TLS) Protocol Version 1.2 (2008), <https://tools.ietf.org/html/rfc5426>, RFC 5426
12. Fischlin, M., Günther, F., Marson, G.A., Paterson, K.G.: Data is a stream: Security of stream-based channels. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015, Part II. LNCS, vol. 9216, pp. 545–564. Springer, Heidelberg (Aug 2015)
13. Giesen, F., Kohlar, F., Stebila, D.: On the security of TLS renegotiation. In: Sadeghi, A.R., Gligor, V.D., Yung, M. (eds.) ACM CCS 13. pp. 387–398. ACM Press (Nov 2013)
14. Günther, F., Mazaheri, S.: A Formal Treatment of Multi-key Channels. In: Proceedings of the 37th International Cryptology Conference. CRYPTO'17, Springer (2017)
15. Haber, S., Pinkas, B.: Securely combining public-key cryptosystems. In: ACM CCS 01. pp. 215–224. ACM Press (Nov 2001)

16. Hoang, V.T., Krovetz, T., Rogaway, P.: Robust authenticated-encryption AEZ and the problem that it solves. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part I. LNCS, vol. 9056, pp. 15–44. Springer, Heidelberg (Apr 2015)
17. Hoang, V.T., Reyhanitabar, R., Rogaway, P., Vizár, D.: Online authenticated-encryption and its nonce-reuse misuse-resistance. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015, Part I. LNCS, vol. 9215, pp. 493–517. Springer, Heidelberg (Aug 2015)
18. Hoepman, J.H.: The ephemeral pairing problem. In: Juels, A. (ed.) FC 2004. LNCS, vol. 3110, pp. 212–226. Springer, Heidelberg (Feb 2004)
19. IEEE 802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications (2012), <http://dx.doi.org/10.1109/IEEESTD.2012.6178212>
20. Jager, T., Kohlar, F., Schäge, S., Schwenk, J.: On the security of TLS-DHE in the standard model. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 273–293. Springer, Heidelberg (Aug 2012)
21. Katz, J., Yung, M.: Unforgeable encryption and chosen ciphertext secure modes of operation. In: Schneier, B. (ed.) FSE 2000. LNCS, vol. 1978, pp. 284–299. Springer, Heidelberg (Apr 2001)
22. Krawczyk, H.: Cryptographic extraction and key derivation: The HKDF scheme. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 631–648. Springer, Heidelberg (Aug 2010)
23. Krawczyk, H., Paterson, K.G., Wee, H.: On the security of the TLS protocol: A systematic analysis. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 429–448. Springer, Heidelberg (Aug 2013)
24. Marson, G., Poettering, B.: Security Notions for Bidirectional Channels. In: IACR Transactions on Symmetric Cryptology. vol. 2017, pp. 405–426. <http://ojs.ub.rub.de/index.php/ToSC/article/view/602>
25. Maurer, U., Tackmann, B.: On the soundness of authenticate-then-encrypt: formalizing the malleability of symmetric encryption. In: Al-Shaer, E., Keromytis, A.D., Shmatikov, V. (eds.) ACM CCS 10. pp. 505–515. ACM Press (Oct 2010)
26. Maurer, U.M.: Perfect cryptographic security from partially independent channels. In: 23rd ACM STOC. pp. 561–571. ACM Press (May 1991)
27. Microsoft-Inria Joint Centre: miTLS: A verified reference TLS implementation (2012), <https://www.mitls.org/pages/attacks>
28. Nampreppe, C.: Secure channels based on authenticated encryption schemes: A simple characterization. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 515–532. Springer, Heidelberg (Dec 2002)
29. Paterson, K.G., Ristenpart, T., Shrimpton, T.: Tag size does matter: Attacks and proofs for the TLS record protocol. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 372–389. Springer, Heidelberg (Dec 2011)
30. Rescorla, E.: The Transport Layer Security (TLS) Protocol Version 1.3: draft-ietf-tls-tls13-20 (April 2017), <https://tools.ietf.org/pdf/draft-ietf-tls-tls13-20.pdf>, expires October 30, 2017
31. Rescorla, E., Modadugu, N.: Datagram Transport Layer Security (2006), <https://tools.ietf.org/html/rfc4347>, RFC 4347
32. Rogaway, P.: Authenticated-encryption with associated-data. In: Atluri, V. (ed.) ACM CCS 02. pp. 98–107. ACM Press (Nov 2002)
33. Rogaway, P., Shrimpton, T.: Deterministic authenticated-encryption: A provable-security treatment of the key-wrap problem. Cryptology ePrint Archive, Report 2006/221 (2006), <http://eprint.iacr.org/2006/221>

34. Shrimpton, T.: 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>
35. Smyth, B., Pironti, A.: Truncating TLS connections to violate beliefs in web applications. In: 7th USENIX Workshop on Offensive Technologies, WOOT '13 (2013)

A Definitions

A.1 KDF

Definition 14 (KDF).

A key derivation function, $\text{KDF}(sk, \text{id}) \rightarrow ck$, is a pseudo-random function $SK \times \mathcal{ID} \rightarrow \mathcal{CK}$, which takes as input original keying material sk and a derived key identifier id , and outputs a key ck .

A.2 Authenticated Encryption with Associated Data (AEAD)

Definition 15 (Stateful Length-Hiding AEAD). A stateful length-hiding AEAD scheme Π for a message space \mathcal{M} , an associated data space \mathcal{AD} , a key space \mathcal{K} , and a ciphertext space \mathcal{C} , is a tuple of algorithms:

- $\text{Kgn}() \xrightarrow{\mathbb{S}} k$: A probabilistic key generation algorithm that outputs a key k .
- $\text{Enc}(k, \ell, \mathbf{ad}, m, st_{\mathbb{S}}) \xrightarrow{\mathbb{S}} (c, st'_{\mathbb{S}})$: A probabilistic encryption algorithm that takes as input a key $k \in \mathcal{K}$, a length $\ell \in \mathbb{Z} \cup \{\perp\}$, associated data $\mathbf{ad} \in \mathcal{AD}$, a message $m \in \mathcal{M}$, and an write state $st_{\mathbb{S}}$, and outputs a ciphertext $c \in \mathcal{C}$ or an error symbol \perp , and updated state $st'_{\mathbb{S}}$.
- $\text{Dec}(k, \mathbf{ad}, c, st_{\mathbb{R}}) \rightarrow (m, st'_{\mathbb{R}})$: A deterministic decryption algorithm that takes as input a key $k \in \mathcal{K}$, associated data $\mathbf{ad} \in \mathcal{AD}$, a ciphertext c , and a read state $st_{\mathbb{R}}$, and outputs a message $m \in \mathcal{M}$ or an error symbol \perp , and an updated state $st'_{\mathbb{R}}$.

If $\ell \neq \perp$, then we say that Π is length-hiding.

Correctness is defined in the obvious way, based on scheme correctness from Def. 1.

Definition 16 (Stateful Length-Hiding AEAD Security). Let Π be a stateful length-hiding AEAD scheme and let \mathcal{A} be an PPT adversarial algorithm. The stateful length-hiding AEAD experiment for Π with bit b is given by $\text{Exp}_{\Pi}^{\text{stlhae}}(\mathcal{A})$ in Fig. 6. We define

$$\text{Adv}_{\Pi}^{\text{stlhae}}(\mathcal{A}) = 2 \Pr \left[\text{Exp}_{\Pi}^{\text{stlhae}}(\mathcal{A}) \right] - 1 .$$

Note that the state variables $st_{\mathbb{S}}$, $st_{\mathbb{R}}$ in Fig. 6 are considered substate variables in terms of the channel environment (e.g. $\pi_s^P.\text{cxn}_V^i.\text{substate}$, $\pi_s^P.\text{cxn}_R^i.\text{substate}$). This is due to the increased state considerations of the environment.

$\text{Exp}_{\Pi}^{\text{stlhae}}(\mathcal{A})$:	$\text{Dec}(\text{ad}, c)$:
1: $k \xleftarrow{\$} \text{Kgn}()$ 2: $b \xleftarrow{\$} \{0, 1\}$ 3: $st_S \leftarrow \perp, st_R \leftarrow \perp$ 4: $u \leftarrow 0, v \leftarrow 0$ 5: phase $\leftarrow 0$ 6: $b' \xleftarrow{\$} \mathcal{A}^{\text{Encrypt}(\cdot), \text{Decrypt}(\cdot)}()$ 7: return $(b' = b)$	1: if $b = 0$ then 2: return \perp 3: $v \leftarrow v + 1$ 4: $(m, st_R) \leftarrow \text{Dec}(k, \text{ad}, c, st_R)$ 5: if $(u < v) \vee (c \neq \text{sent}.c_v) \vee (\text{ad} \neq \text{sent}.ad_v)$ then 6: phase $\leftarrow 1$ 7: if phase = 1 then 8: return m 9: return \perp
$\text{Enc}(\ell, \text{ad}, (m_0, m_1))$:	
1: $u \leftarrow u + 1$ 2: $(\text{sent}.c^{(0)}, st_S^{(0)}) \leftarrow \text{Enc}(k, \ell, \text{ad}, m_0, st_S)$ 3: $(\text{sent}.c^{(1)}, st_S^{(1)}) \leftarrow \text{Enc}(k, \ell, \text{ad}, m_1, st_S)$ 4: if $\text{sent}.c^{(0)} = \perp$ or $\text{sent}.c^{(1)} = \perp$ then 5: return \perp 6: $(\text{sent}.ad_u, \text{sent}.c_u, st_S) := (\text{ad}, \text{sent}.c^{(b)}, st_S^{(b)})$ 7: return $\text{sent}.c_u$	

Fig. 6. Stateful length-hiding AEAD experiment stlhae for stateful length-hiding AEAD scheme $\Pi = (\text{Kgn}, \text{Enc}, \text{Dec})$ and adversary \mathcal{A} .

A.3 Authenticated and Confidential Channel Establishment (ACCE)

The following game, $\text{Exp}_{\Pi}^{\text{ACCE}}()$, is for the *post-accept* phase of the authenticated and confidential channel establishment (ACCE) model. More detail about the full ACCE experiment, including the post-accept phase, can be found in [20].

$\text{Encrypt}(\pi_s^P, l, \text{ad}, m_0, m_1)$:	$\text{Decrypt}(\pi_s^P, \text{ad}, c)$:
1: $u \leftarrow u + 1$ 2: $(c^{(0)}, st_S^{(0)}) \xleftarrow{\$} \text{Enc}(SK, l, \text{ad}, m_0, st_S)$ 3: $(c^{(1)}, st_S^{(1)}) \xleftarrow{\$} \text{Enc}(SK, l, \text{ad}, m_1, st_S)$ 4: if $c^{(0)} = \perp$ or $c^{(1)} = \perp$ then 5: return \perp 6: $(\text{sent}.ad_u, \text{sent}.c_u, st_S) \leftarrow (\text{ad}, c^{(b_s^P)}, st_S^{(b_s^P)})$ 7: return $\text{sent}.c_u$	1: if $b_s^P = 0$ then 2: return \perp 3: $v \leftarrow v + 1$ 4: $(m, st_R) \leftarrow \text{Dec}(SK, \text{ad}, c, st_R)$ 5: if $(u < v) \vee (c \neq \text{sent}.c_v) \vee (\text{ad} \neq \text{sent}.ad_v)$ then 6: phase $\leftarrow 1$ 7: if phase = 1 then 8: return m 9: return \perp

Fig. 7. ACCE game oracles, for principal P and session π_s^P , with session key SK .

B Agility

Here we summarize pseudo-random function agility, based on the notions of Acar et al. [1] which use more generalized scheme algorithms than we employ.

Let $\text{PRF.F} = (\text{Pg}, \text{Kg}, \text{PRF}, \text{Dom}, \text{Rng})$ be a scheme consisting of the following PPT algorithms: a probabilistic parameter generator Pg , a probabilistic key generator Kg , a deterministic evaluator PRF , a deterministic domain Dom , and a deterministic range Rng .

$\text{Exp}_{\text{PRF.F}}^{\text{agile-prf}}(\mathcal{A})$:	$\text{PRF}(\text{PRF.F}, x)$:
1: $\text{pars} \stackrel{\$}{\leftarrow} \text{Pg}(1^k)$	1: if $b = 1$ then
2: $K \stackrel{\$}{\leftarrow} \text{Kg}()$	2: $y \leftarrow \text{PRF}(\text{pars}, K, x)$
3: $b \stackrel{\$}{\leftarrow} \{0, 1\}$	3: else
4: $b' \leftarrow \mathcal{A}^{\text{PRF}(\cdot)}()$	4: $y \stackrel{\$}{\leftarrow} \text{Rng}(\text{pars})$
5: return $(b' = b)$	5: return y

Fig. 8. PRF agility game.

Here *compatibility* is defined for the set of PRF functions, but easily generalizes to other scheme types.

Definition 17 (Compatible). *A set $\{\text{PRF.F}\}$ of PRF.F schemes is compatible if all $\text{PRF} \in \{\text{PRF.F}\}$ have the same parameter generator and the same key generator.*

Definition 18 (Set-Restricted). *An adversary \mathcal{A} is $\{\text{PRF.F}\}$ -restricted if:*

- *the scheme arguments to its PRF queries, per Fig. 8, all come from $\{\text{PRF.F}\}$,*
- *it does not repeat an oracle query on the same inputs, and*
- *for all queries $\text{PRF}(\text{PRF.F}, x)$, $x \in \text{Dom}(\text{pars})$.*

Definition 19 (Agility). *Let $\text{Adv}_{\{\text{PRF.F}\}}^{\text{agile-prf}}(\mathcal{A}) = 2 \Pr[\text{Exp}_{\text{PRF.F}}^{\text{agile-prf}}(\mathcal{A})] - 1$, where $\{\text{PRF.F}\}$ is a finite, compatible set, and let \mathcal{A} be any PPT, $\{\text{PRF.F}\}$ -restricted adversary. Then $\{\text{PRF.F}\}$ is agile if $\text{Adv}_{\{\text{PRF.F}\}}^{\text{agile-prf}}(\mathcal{A})$ is negligible.*

Remark 8. While it has been shown by Acar et al. [1] that the set of all PRF functions is not agile, it is still possible for smaller subsets of PRF functions to achieve agility. For example, the set of KDF schemes available to a particular channel protocol could achieve agility. Special cases mentioned by Acar et al. include the singleton set, and any set of collision-resistant hash functions when formalized as keyed families. In practice, we expect that the sets of available KDF functions for real-world protocols may be small enough to achieve agility (trivially, if only one KDF is negotiated), but leave it to future work in particular protocol analyses.

C Composition of Protocol Epochs

This work addresses secure termination for TLS 1.2 under Alert Protocol messages sent after a completed handshake – it is assumed that no Alert Protocol message fragments are buffered for processing from an earlier protocol epoch (i.e. the Handshake Protocol). The importance of this observation is as follows: mentioned briefly above, in 2012 an attack on TLS via the Alert Protocol was discovered by the miTLS reference implementation team [27] which exploited alert message fragmentation due to the 2-byte message length. Basically, the attack was performed by sending an unauthenticated alert message fragment during the handshake which, due to its incompleteness, would be buffered for later processing. Once the channel was established and application data sent authenticated, the receipt of a subsequent alert message would result in the combination of the previous fragment and the first byte of the current alert message being processed as a single alert message. One possible consequence of this attack would be a downgrade of the intended alert level from *fatal* to *warning*, undermining the specified use of *fatal* alerts. The attack could also be used to upgrade the alert level – the sending of a *fatal* alert fragment during the handshake could result in the processing of a *warning* alert, sent later, at a fatal level. This would trivially result in closure of the receiver’s channel with a `terminate` message while the sender’s channel remains open; hence achieving a successful adversarial win according to the secure termination experiment of Figure 3.

By message format, TLS alert messages sent *after* the handshake completion fall under the ciphertext protections defined for Record Layer messages [11], and thereby satisfy the constraints of `terminate` messages. (i.e. messages that signal closure to the channel connections). However, those sent *during* the handshake do not (Theorem 3 addresses only channels satisfying stLHAE channel security). Fundamentally, this raises an important issue: if m is a `terminate` message in one protocol epoch, it may not be in others. The logical linking and buffering of m across protocol epochs with different security guarantees is a dangerous game, and lays the protocol open to attacks such as the one above. Similar issues could potentially arise during session resumption if data from the first connection is buffered into the second. Note that *resumption* can have different meanings dependent on protocol use. For protocols which allow session resumption via a key derived from the original master session key, we consider initiation of additional/subsequent connections. Comparatively, for those protocols that reuse the same key, we consider that the original connection was never actually closed.¹

To counter the miTLS Alert attack vulnerability, the miTLS reference implementation resets buffers at the sending/receipt of a ChangeCipherSpec message, effectively allowing a fresh start to all messages in the epoch, including `terminate` messages. However, the viability of the above attack is dependent

¹ Compared to resumption, *renegotiation* refers to sessions where keys for subsequent connections are derived from new keying material (possibly in addition to the original session key or a previous connection key); we leave this as a possible extension for future work.

on the implementation, with OpenSSL, for example, being subject to the attack while some other implementations are not [27]. Thus, proofs for protocol secure termination must clearly define assumptions as we have done, to dissever protocol implementations and delineate those that achieve secure termination.

According to miTLS, the viability of the above attack is dependent on the implementation. OpenSSL, for example, is subject to the attack while miTLS reference implementation, as well as some other implementations, are not [27].