



HAL
open science

Key-Schedule Security for the TLS 1.3 Standard

Chris Brzuska, Antoine Delignat-Lavaud, Christoph Egger, Cédric Fournet,
Konrad Kohbrok, Markulf Kohlweiss

► **To cite this version:**

Chris Brzuska, Antoine Delignat-Lavaud, Christoph Egger, Cédric Fournet, Konrad Kohbrok, et al..
Key-Schedule Security for the TLS 1.3 Standard. Asiacrypt 2022, IACR, Dec 2022, Taipei, Taiwan.
pp.621-650, 10.1007/978-3-031-22963-3_21 . hal-03969623

HAL Id: hal-03969623

<https://cnrs.hal.science/hal-03969623>

Submitted on 2 Feb 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Key-schedule Security for the TLS 1.3 Standard

Chris Brzuska¹, Antoine Delignat-Lavaud², Christoph Egger³,
Cédric Fournet², Konrad Kohbrok¹, and Markulf Kohlweiss⁴

¹ Aalto University, Finland {chris.brzuska,konrad.kohbrok}@aalto.fi

² Microsoft Research Cambridge, UK {fournet,antdl}@microsoft.com

³ IRIF, Université Paris Cité, France christoph.egger@alumni.fau.de

⁴ University of Edinburgh, UK mkohlwei@ed.ac.uk

Abstract. Transport Layer Security (TLS) is the cryptographic backbone of secure communication on the Internet. In its latest version 1.3, the standardization process has taken formal analysis into account both due to the importance of the protocol and the experience with conceptual attacks against previous versions. To manage the complexity of TLS (the specification exceeds 100 pages), prior reduction-based analyses have focused on some protocol features and omitted others, e.g., included session resumption and omitted agile algorithms or vice versa.

This article is a major step towards analysing the TLS 1.3 key establishment protocol as specified at the end of its rigorous standardization process. Namely, we provide a full proof of the TLS *key schedule*, a core protocol component which produces output keys and internal keys of the key exchange protocol. In particular, our model supports all key derivations featured in the standard, including its negotiated modes and algorithms that combine an optional Diffie-Hellman exchange for forward secrecy with optional pre-shared keys supplied by the application or recursively established in prior sessions.

Technically, we rely on *state-separating proofs* (Asiacrypt '18) and introduce techniques to model large and complex derivation graphs. Our key schedule analysis techniques have been used subsequently to analyse the key schedule of Draft 11 of the MLS protocol (S&P '22) and to propose improvements.

1 Introduction

Transport Layer Security (TLS) is the most widely used authenticated secure channel protocol on the Internet, protecting the communications of billions of users. Previous versions of TLS have suffered from impactful attacks against weaknesses in their design, including legacy algorithms (e.g. FREAK for export RSA [9], LogJam [2] for export Diffie-Hellman, WeakDH for ill-chosen groups, and exploits against Mantin biases of RC4 [21]); the RSA key encapsulation (e.g. the ROBOT [19] variant of Bleichenbacher's PKCS1 padding oracle); the fragile MAC-encode-encrypt construction leading to many variants of Vaudenay's padding oracles against CBC cipher suites (e.g. BEAST [38],

Lucky13 [3]); the weak signature over nonces allowing protocol version downgrades (e.g. DROWN [5] and POODLE); attacks on other negotiated parameters [11], the key exchange logic (e.g. the cross-protocol attack of [49] and 3SHAKE [12]); exploitations of collisions on the hash transcript (e.g. SLOTH [15]). TLS 1.3 intends both to fix the weaknesses of previous versions and to improve the protocol performance, notably by lowering the latency of connection establishment from two roundtrips down to one, or even zero when resuming a connection.

Historically, the IETF process to adopt a standard involves an open consortium of contributors mostly coming from industry, with a bias towards early implementers. The TLS working group at the IETF acknowledged that this process puts too much emphasis on deployment and implementation concerns, and tends to address security issues reactively [50]. For TLS 1.3, it decided to address security upfront by welcoming feedback from various cryptographic efforts, including symbolic [30,29] and computational protocol models [34,35,48], both on paper and implemented in tools such as Tamarin or CryptoVerif. Early drafts of TLS 1.3 also drew much inspiration from Krawczyk’s OPTLS protocol [47], which comes with a detailed security proof, although later versions diverged from it (in particular in the design of resumption). This proactive approach has certainly improved the overall design of TLS 1.3, and uncovered flaws along its 28 intermediate drafts. However, many of these efforts are incomplete (focusing, e.g., on fixed protocol configurations) or do not account for the final version published in RFC 8446, see Section 6 for a more detailed discussion of related work. Since final adoption, further questions have been raised about pre-shared keys, potential reflection attacks [37], and difficulties in separating resumption PSKs (produced internally by the key exchange) from external ones installed by the application. In short: we still miss provable security for the final Internet standard.

TLS can be decomposed into sub-protocols: the *record layer* manages the multiplexing, fragmentation, padding and encryption of data into packets (also called *records*) from three separate streams of handshake, alert, and application data. Incoming handshake messages are passed to the *handshake* sub-protocol, which in turn produces fresh record keys and outgoing handshake messages. Taking advantage of this well-understood modularity, other protocols re-use the TLS 1.3 handshake with different record layers: for instance, DTLS 1.3 is a variant based on UDP datagrams instead of TCP streams, while the IETF version of QUIC replaces the record layer with a much extended transport [42], adding features such as dynamic application streams and fine-grained flow control. Detailed security proofs for the TLS 1.3 record layer have been proposed by Patton et al. [51] (extending the work of Fischlin et al. [40] on stream-based channels), Badertscher et al. [6], and Bhargavan et al. [32], who also provide a verified reference implementation. Therefore, we defer to these works for the record layer, and focus on the handshake protocol.




Fig. 1: Overview over the TLS 1.3 Handshake (left) and its key schedule (right). $[m]_k$ denotes encryption of message m under key k . k_{ae1} and τ_c are derived from k_{cht} , k_{ae2} and τ_s are derived from k_{shs} , and k_{ae3} is derived from k_{sat} . We color digests and keys in alternating pink and blue to clarify digest-key dependency. E.g., label c e t r a f f i c and digest d_{as} is used to derive k_{cet} .

1.1 TLS 1.3 Handshake and Key Schedule

The top of Fig. 1 gives an abstract view of the TLS 1.3 protocol message flow. In the client hello message, the client sends a nonce n_c , its Diffie-Hellman (DH) share g^x , a PSK $label$ and a $binder$ value for domain separation and session resumption. As a means of negotiation, the client may offer shares for different groups and different PSK options (thus the indices i, j in $g_i^x, label_j, binder_j$). The server communicates its choice of the DH group and the PSK when sending the server hello message which contains the server nonce n_s , its share $g_{i_0}^y$ (including the group description) and the label $label_{j_0}$ of the chosen PSK. The remaining messages consist of server certificate, signature (C(pk), CV(σ)), key confirmation messages in the forms of messages authentication codes (MACs) τ_s and τ_c computed over the transcript, and a $ticket$ which is used on the client side to store a resumption key (later referred to as $resumption$ PSK) derived from the key material of the current key exchange session.

The key $schedule$ is the core part of the handshake that performs all key computations. It takes as main input PSK and DH key materials and, at each phase of the handshake, it derives keys, e.g., to encrypt c e a r l y t r a f f i c (k_{cet}), to compute the binder value (k_{binder}), to encrypt s e r v e r h a n d s h a k e t r a f f i c (k_{shs}) and to encrypt c l i e n t h a n d s h a k e t r a f f i c (k_{cht}).

The key schedule relies on the hashed key derivation function (HKDF) standard [45], which uses HMAC [7] to implement $extract$ (xtr) and $expand$ (xpd) operations. In addition, the key schedule makes calls to xpd to expand keys into further subkeys. The key schedule thus consists of a collection of xtr and xpd operations, organized in a graph. Each of the operations takes as input a $chaining$ key and/or new key material, (k_{psk} in the xtr in the early phase and k_{dh} in the xtr in the handshake phase), together with the latest digest and auxiliary inputs such as a resumption status r and a ticket nonce tn .

In this article, we consider eight output keys of the TLS key schedule: k_{cet} , k_{eem} , k_{binder} , k_{cht} , k_{sht} , k_{cat} , k_{sat} , k_{eam} . They constitute a natural boundary, inasmuch as all other TLS keys and IVs are further derived from them in a transcript-independent manner.

1.2 Key Schedule Model and Key Exchange Model

We model the security of the key schedule as an indistinguishability game between a real and an ideal game. The real game allows the adversary to use their own dishonest application PSKs and Diffie-Hellman shares. In addition, it allows the adversary to instruct the game to sample honest PSKs and Diffie-Hellman shares. From these base keys, the adversary can then instruct the model to derive further keys. The adversary cannot see internal keys, but it can obtain the 8 output keys from the model. In turn, in the ideal game, the output keys are replaced by unique, random keys which are sampled independently from the input key material.

The interface of this model captures how the key exchange protocol uses the key schedule. The key exchange protocol should, indeed, not use the internal keys, but instead only use the output keys. Moreover, the final session keys are to be used only by the Record Layer to implement a secure channel. In a companion paper [25], we show that key exchange security of the TLS 1.3 handshake protocol reduces to the key schedule security established in this paper. Note that authentication is proved based on *keys* and does not capture binding between keys and identities, as needed, e.g., for reflection attacks [29].

Outline We introduce our overall technical approach in Section 2. We define our assumptions for collision-resistance, pseudorandomness and pre-image resistance in Section 3. Section 4 defines syntax and security of the TLS key schedule. Section 5 states the main key schedule theorem and provides its proof. This article gives proof sketches of all lemmata, highlighting their conceptual insights. The complete proofs are provided in the full version [23]. Finally, Section 7 includes proposals for (late) changes to the TLS 1.3 standard.

2 Technical Approach

2.1 Handles

Complex derivation steps make it crucial to maintain administrative *handles* in the model state, both for internal bookkeeping and security modeling as well as for communication with the adversary. Namely, to instruct the model to perform further computations on keys, the adversary can point to the keys to be used via handles. Such handles are particularly important for honest keys, i.e., honest psks, honest Diffie-Hellman shares and honest internal keys derived via xtr and xpd from honest base keys, because the model cannot provide the adversary with the actual values of these secrets.

Our model constructs handles as nested data records where each nesting step keeps track of the inputs which were used to compute the associated key. We have base handles for PSKs and DH secrets, including handles for dummy zero values to be used in noDH and noPSK mode as well as base handles for a fixed 0_{salt} and fixed 0_{ikm} .

$\text{dh}\langle\text{sort}(X, Y)\rangle$	Diffie-Hellman secret
$h = \text{psk}\langle\text{ctr}, \text{alg}\rangle$	application PSK
$\text{noDH}\langle\text{alg}\rangle$	fixed $0^{\text{len}(\text{alg})}$ Diffie-Hellman secret
$\text{noPSK}\langle\text{alg}\rangle$	fixed $0^{\text{len}(\text{alg})}$ PSK
0_{salt}	fixed 0 salt
$0_{\text{ikm}}\langle\text{alg}\rangle$	fixed $0^{\text{len}(\text{alg})}$ initial key material (IKM)

The model then inductively applies the following constructors to build all other handles from the base handles:

$\text{xtr}\langle\text{name}, \text{left parent handle}, \text{right parent handle}\rangle.$
 $\text{xpd}\langle\text{name}, \text{label}, \text{parent handle}, \text{other arguments}\rangle.$

For example, given a handle to the early master secret h_{es} , the handle h_{cet} to the client early transport secret is defined as

$$h_{cet} = \text{xpd}\langle\text{cet}, \text{c e traffic}, h_{es}, t_{es}\rangle$$

where t_{es} is the transcript of the protocol messages exchanged so far, and ‘c e traffic’ is the constant byte string label prescribed in the RFC [52] for this derivation step.

Agility Our model is *agile*, i.e., it supports multiple algorithms. Thus, we tag the handles $h = \text{psk}\langle\text{ctr}, \text{alg}\rangle$, $\text{noPSK}\langle\text{alg}\rangle$ and $0_{\text{ikm}}\langle\text{alg}\rangle$ with the algorithm alg for which the keys are intended. Jumping ahead, we note that we also tag *keys* with their intended algorithm so that in the key derivation

$$k_{cet} = \text{xpd}(k_{es}, \text{c e traffic}, d_{es}),$$

the agile xpd function can retrieve the correct hash algorithm alg to use within hmac from the key’s tag. We write $\text{alg}(h_{cet})$ for the algorithm descriptor of h_{cet} and $\text{tag}_h(k)$ for key k tagged with this algorithm.

Length The handle determines the algorithm, and the algorithm determines the length of keys and outputs of a hash-algorithm alg . For convenience, we write $\text{len}(h_{cet})$ as an alias for $\text{len}(\text{alg}(h_{cet}))$.

Note that we introduced handles $0_{\text{ikm}}\langle\text{alg}\rangle$ for the dummy key value $0^{\text{len}(\text{alg})}$ as well as 0_{salt} for the 1-bit-long 0-key. This is because hmac pads keys with zeroes up to their block length and thus, storing multiple zero values would introduce redundancy in the model without a correspondence in real-life.

Name and level In addition to the algorithm and its key length, the handle determines the key name (cet) and a *level*. The level is the number of resumptions the handle records, counting from 0 and adding one for each node with a **resumption** label. We write $\text{level}(h_{cet})$ for this level. We will often need to refer to the *parent* names of a particular key (name) n , and write the pair of parent names as $\text{prntn}(n)$. In the case of xpd , the key is only derived from one key and thus, in this case, $\text{prntn}(n) = (n_1, \perp)$. Conversely, we refer by $\text{chldrnn}(n_1)$ to the set of all key names which are derived from n_1 . In particular, if $\text{prntn}(n) = (n_1, \perp)$, then $n \in \text{chldrnn}(n_1)$. We refer to all names which share a parent with n as $\text{sblngn}(n)$.

Handshake mode Jumping ahead, we note that we use handle data also to communicate the handshake mode to the key schedule model. A $\text{noDH}\langle alg \rangle$ Diffie-Hellman handle signals a psk_ke mode, while a $\text{noPSK}\langle alg \rangle$ PSK handle signals a dh_ke mode.

2.2 Application Key Registration & Honesty

Honesty of a handle is a crucial concept to model that the key associated with the handle, when returned to the adversary, looks pseudorandom. Honesty is inductively computed, starting from the base keys: All zero keys have dishonest handles. Handles of application PSKs are honest if their key was sampled by the security model and dishonest if their key was sampled by the security model. Diffie-Hellman handles are honest if both shares are honest. Derived handles are honest if and only if at least one of their input handles are honest. Considering the derivation graph (cf. right side of Fig. 1), we obtain that the h_{esalt} handles and the handles which appear *before* have the same honesty as the last PSK handle, while the handles after h_{esalt} are honest if the last PSK handle was honest *or* the last Diffie-Hellman handle was honest.

2.3 State-Separating Proofs (SSPs)

In the following we use the the pseudorandomness game $\text{Gxpd}_{n,\ell}^0$ for the xpd function (depicted in Fig. 2) as a running example to introduce core concepts. As is common in cryptography, security is modeled as an interaction between an adversary \mathcal{A}




Fig. 2: Game $\text{Gxpd}_{n,\ell}^b$ for $b \in \{0, 1\}$

(which can be thought of as sitting left of the picture) and a program which we call the *game*. This interaction happens via so-called *oracles*—which we describe in pseudo-code—corresponding to the arrows from the left side of the picture. The task of the adversary consists in *distinguishing* two variants of the

game G^0 and G^1 with identical interfaces and we measure the success probability of any such adversary \mathcal{A} and call it *advantage*.

Definition 1 (Advantage). For adversary \mathcal{A} , we define the advantage

$$\text{Adv}(\mathcal{A}; \mathsf{G}^0, \mathsf{G}^1) := |\Pr[1 = \mathcal{A} \rightarrow \mathsf{G}^0] - \Pr[1 = \mathcal{A} \rightarrow \mathsf{G}^1]|.$$

In particular, for the pseudorandomness game $\text{Gxpd}_{n,\ell}^b$ for xpd , the analogous definition is as follows.

Definition 2 (XPD). For adversary \mathcal{A} , we define the xpd pseudorandomness advantage $\text{Adv}(\mathcal{A}, \text{Gxpd}_{n,\ell}^0, \text{Gxpd}_{n,\ell}^1)$ as

$$|\Pr[1 = \mathcal{A} \rightarrow \text{Gxpd}_{n,\ell}^0] - \Pr[1 = \mathcal{A} \rightarrow \text{Gxpd}_{n,\ell}^1]|,$$

where Fig. 2 defines $\text{Gxpd}_{n,\ell}^0$.

The graphs specifying such a security game suggest a natural flow downwards. While we discuss the details of the game later in this section, one can extract a conceptual picture already from the graph alone. Concretely the intended usage (by the adversary) of $\text{Gxpd}_{n,\ell}^b$ consists on first registering input values using the $\text{SET}_{n_1,\ell}$ oracle, executing key derivation using the $\text{XPD}_{CN,\ell}$ oracle and finally retrieving and testing the output using the $\text{GET}_{n,\ell}$ oracle. In addition, the adversary gets access to auxiliary oracles, namely the HASH oracle modeling a cryptographic hash function as well as the Q and UNQ oracles.⁵ Finally, $\text{Gxpd}_{n,\ell}^b$ is structured in individual components which we call *packages*.

Definition 3 (Package). A package M consists of a set of oracles $[\rightarrow \mathsf{M}] = \{\text{O}1, \dots, \text{O}t\}$, specified by pseudo-code and operating on a set of state variables Σ , specified on the top of each package description. All other variables used by oracles are temporary and their values are forgotten after each call. The oracles of M may depend on oracles $[\mathsf{M} \rightarrow] = \{\text{O}'1, \dots, \text{O}'t'\}$, i.e., make calls to oracles in $[\mathsf{M} \rightarrow]$. We say that a package M is stateless if $\Sigma = \emptyset$. We say that a package M is a game if $[\mathsf{M} \rightarrow] = \emptyset$.

While some oracles of a package are exposed to the adversary, others are used only internally within the game. A monolithic version of a game such as $\text{Gxpd}_{n,\ell}^b$ can be obtained by *inlining* all internal oracle calls. With the concept of packages we can now discuss the individual parts of $\text{Gxpd}_{n,\ell}^b$. $\text{Xpd}_{CN,\ell}$ is a parallel composition of $\text{Xpd}_{n,\ell}$ for all children of n_1 exposing the oracles $\text{XPD}_{n,\ell}$ for $n \in CN$, we write $\text{XPD}_{CN,\ell}$ as shorthand for these oracles. The $\text{Xpd}_{CN,\ell}$ packages are the only stateless packages in the game, indicated by the white color as opposed to the gray of stateful packages.

⁵ These two oracles in particular are necessary for composition: Note that the main oracles the adversary interacts with are subscripted by a name n and a level ℓ while the Q and UNQ oracles only take the name n as subscript. We will share the same Q and UNQ oracles between many instances of $\text{Gxpd}_{n,\ell}^b$ and therefore need to allow reductions access to these oracles.

The $\text{XPD}_{CN,\ell}$ oracle of package $\text{Xpd}_{n,\ell}$ computes a new handle $h \leftarrow \text{xpd}(n, \text{label}, h_1, \text{args})$ alongside a new key $k \leftarrow \text{xpd}(k_1, (\text{label}, d))$ based on the parent handle h_1 , the arguments (e.g. transcript) and the bit r indicating whether this is a resumption session. The evaluation also includes a *label* which depends on the name of the package as well as the resumption bit. Note that the oracle only receives the *handle* of the input key from the adversary and only returns the newly constructed *handle* of the newly derived key. Concrete secrets are passed to $\text{Key}_{n,\ell}^b$ packages using the GET and SET oracles. Here we can distinguish the upper $\text{Key}_{n_1,\ell}^1$ package and the lower $\text{Key}_{CN,\ell}^b$ packages (for all n in CN). We defer discussion about the Q and UNQ oracle calls to the description of the **Log** package.

The upper $\text{Key}_{n_1,\ell}^1$ package offers oracle $\text{SET}_{n_1,\ell}(h, \text{hon}, k)$ to the adversary which allows it to register a key. The oracle first verifies that the handle h matches the name n and level ℓ of this key package and—modeling algorithmic agility—verifies that the algorithm tag matches the value of the key, and else, **assert** throws an *abort*. As this is an ideal key package (indicated by superscript $b=1$) for honest keys, instead of using the value provided by the adversary a fresh value is sampled—as indicated by using \leftarrow_s in contrast to \leftarrow used for assignments. Finally the key is stored in this package’s state and the handle returned to the caller. The GET oracle simply restores algorithm tagging on the key value and returns it to the caller (in this case the **Xpd** package). The lower $\text{Key}_{CN,\ell}^b$ packages work the other way round in that they expose the GET oracle to the adversary while the SET oracle is used by **Xpd**. We encode the distinguishing task for the adversary in the $\text{Key}_{CN,\ell}^b$ package: In $\text{Gxpd}_{n,\ell}^0$ ($b = 0$), the keys returned from the GET oracle of the $\text{Key}_{CN,\ell}^0$ is honestly computed based on the input keys while in the ideal game $\text{Gxpd}_{n,\ell}^1$ the values of honest keys are sampled in the **Key** package ignoring the value computed by **Xpd**.

Finally, queries Q_n and UNQ_n to the **Log** _{n} package (Fig. 4) model collisions. The Q query simply returns if a *handle* is re-used while UNQ concerns itself with collisions between keys via an abort pattern and a mapping method. In slightly nonstandard notation, we use existential quantors here to express searching for *indices* into tables. The pattern models conditions on states where the game aborts (i.e. terminates and outputs a special symbol), cf. Section 5.3 for their

```

Xpd $n,\ell$ 
-----
Parameters
-----
 $n$  : name
 $\ell$  : level

prntn :  $N \rightarrow (N_\perp \times N_\perp)$ 
label :  $N \times \{0, 1\} \rightarrow \{0, 1\}$ 96

State
-----
no state

XPD $n,\ell$ ( $h_1, r, \text{args}$ )
-----
 $n_1, \_ \leftarrow \text{prntn}(n)$ 
label  $\leftarrow \text{label}(n, r)$ 
 $h \leftarrow \text{xpd}(n, \text{label}, h_1, \text{args})$ 
( $k_1, \text{hon}$ )  $\leftarrow \text{GET}_{n_1,\ell}(h_1)$ 
if  $n = \text{psk}$  :
   $\ell \leftarrow \ell + 1$ 
   $k \leftarrow \text{xpd}(k_1, (\text{label}, \text{args}))$ 
else
   $\text{alg} \leftarrow \text{alg}(h_1)$ 
   $d \leftarrow \text{HASH}(\text{tag}_{\text{alg}}(\text{args}))$ 
   $k \leftarrow \text{xpd}(k_1, (\text{label}, d))$ 
 $h \leftarrow \text{SET}_{n,\ell}(h, \text{hon}, k)$ 
return  $h$ 

```

Fig. 3: Xpd package

<u>Key_{n,ℓ}^b</u>		<u>Log_n^{P,map}</u>
Parameters	State	Parameters
n : name	$K_{n,\ell}$: Keytable	n : name
ℓ : level		
		<u>State</u>
<u>SET_{n,ℓ}(h, hon, k^*)</u>		L_n : Log
assert name(h) = n		<u>Q_n(h)</u>
assert level(h) = ℓ		if $L_n[h] = \perp$:
assert alg(k^*) = alg(h)		return \perp
$k \leftarrow \text{untag}(k^*)$		else
assert len(h) = $ k $		($h', _, _$) $\leftarrow L_n[h]$
if Q _n (h) $\neq \perp$:		return h'
return Q _n (h)		<u>UNQ_n(h, hon, k)</u>
if b :		if ($\exists h' : L_n[h'] = (h', hon', k)$
if hon :		\wedge level(h) = $r \wedge$ level(h^*) = r') :
$k \leftarrow_{\$} \{0, 1\}^{\text{len}(h)}$		if map($r, hon, r', hon' J_n[k]$) :
$h' \leftarrow \text{UNQ}_n(h, hon, k)$		$L_n[h] \leftarrow (h', hon, k)$
if $h' \neq h$:		$J_n[k] \leftarrow 1$
return h'		return h'
$K_{n,\ell}[h] \leftarrow (k, hon)$		if ($\exists h^* : L_n[h^*] = (h', hon', k)$
return h		\wedge level(h) = $r \wedge$ level(h^*) = r') :
<u>GET_{n,ℓ}(h)</u>		$P(r, hon, r', hon')$
assert $K_{n,\ell}[h] \neq \perp$		$L_n[h] \leftarrow (h, hon, k)$
(k^*, hon) $\leftarrow K_{n,\ell}[h]$		return h
$k \leftarrow \text{tag}_h(k^*)$		
return (k, hon)		
<hr/>		
P	the command $P(r, hon, r', hon')$ is	
Z	\emptyset	
A	if $hon = hon' = 0 \wedge r = r' = 0$: throw abort	
D	if $hon = hon' = 0$: throw abort	
R	if $hon = hon' = 0$: throw abort else throw win	
F	throw abort	
<hr/>		
map	the command $\text{map}(r, hon, r', hon', J_n[k])$ is	
0	0	
1	$hon = hon' = 0 \wedge r \neq r' \wedge 0 \in \{r, r'\} \wedge J_n[k] \neq 1$	
∞	$hon = hon' = 0$	

Fig. 4: Code for the Key and Log. In addition we use Nkey for a single key package that answers queries for all levels from the same table and Okey for a NKey package which consistently answers with the constant all-zeros key.

use. We use the **throw** notation here to allow special symbols in addition to *abort* which is also used by **assert**. In the game $\text{Gxpd}_{n,\ell}^b$, the D pattern aborts on collisions between dishonest keys. The F and R pattern abort if there is a collision between key values, regardless of their honesty, and they return different abort messages. Z does not abort at all, and A aborts upon a collision of two dishonest level 0 keys (which we use to constrain the adversary’s psk registrations in the key schedule model).

Mapping methods filter certain collisions (preventing an *abort* event. ∞ allows collisions between Diffie-Hellman secrets (the adversary can construct colliding values via $X^zY = XY^z$) and the 1 method allows the adversary to register a dishonest application PSK colliding with an dishonest resumption PSK. The mapping methods are only used in the proof and not in the security model.

3 Assumptions

3.1 Collision-Resistance

Fig. 5 defines the collision-resistance game $\text{Gcr}^{\text{f-}alg,b}$ for a given function $\text{f-}alg$, where $\text{f} \in \{\text{hash}, \text{xtr}, \text{xpd}\}$ and $alg \in \mathcal{H}$ which TLS 1.3 currently defines as

$$\mathcal{H} = \{\text{sha256}, \text{sha384}, \text{sha512}\}$$

(see FIPS 180-2). The HASH oracle takes as input a text t from the domain of $\text{f-}alg$ and returns its digest d . If that text t has not been queried before, the digest is stored in table H at index t . In the ideal game ($b = 1$), the oracle first checks whether d already occurs in H , and if so, throws an abort. Hence, the adversary can distinguish between the real and the ideal game if and only if it can submit two different texts with the same digest. Our definition generalizes to n -ary functions by letting the text t be the tuple of their arguments.

Definition 4 (Collision-Resistance). For an adversary \mathcal{A} , a function $\text{f} \in \{\text{hash}, \text{xtr}, \text{xpd}\}$ and algorithm $alg \in \mathcal{H}$, define collision-resistance advantage $\text{Adv}(\mathcal{A}, \text{Gcr}^{\text{f-}alg,0}, \text{Gcr}^{\text{f-}alg,1})$ is

$$|\Pr[1 = \mathcal{A} \rightarrow \text{Gcr}^{\text{f-}alg,0}] - \Pr[1 = \mathcal{A} \rightarrow \text{Gcr}^{\text{f-}alg,1}]|.$$

Agile Collision-resistance It is convenient to define the *agile* collision-resistance game $\text{Gacr}^{\text{f},b}$ as well, where $\text{f} \in \{\text{hash}, \text{xtr}, \text{xpd}\}$ takes *tagged* inputs, i.e., **hash** takes a single input, tagged with the algorithm to use, **xpd** takes three inputs ($k, \text{label}, \text{args}$), where k is tagged, and **xtr** takes inputs (k_1, k_2) where one is tagged, and if both are tagged, they are tagged consistently. The adversary can then make queries to **HASH** with values in the domain of the *agile* functions. We write $\text{Hash}^b := \text{Gacr}^{\text{hash},b}$. See Section 2.1 for further discussion of tagging.

```

Gcrf-alg,b
HASH(t)
-----
assert t ∈ dom(f-alg)
d ← f-alg(t)
if H[t] = ⊥ :
  if b ∧ d ∈ range(H) :
    throw abort
  H[t] ← d
return d

```

Fig. 5: $\text{Gcr}^{\text{f-}alg,b}$ code.

3.2 Pseudorandomness of xpd

For most key names n , Definition 2 already captures pseudorandomness of xpd . We now cover two special cases.

XPD to derive PSK For $n = \text{psk}$ (cf. Fig. 6a), the *layer* index increases from ℓ to $\ell + 1$. Thus, the $\text{XPD}_{\text{psk},\ell}$ oracle reads keys via $\text{GET}_{\text{rm},\ell}$ queries, but writes keys using the level $\ell + 1$ query $\text{SET}_{\text{psk},\ell+1}$. Another difference in $\text{Gxpd}_{\text{psk},\ell}^b$ compared to the general $\text{Gxpd}_{n,\ell}^b$ is that the lower $\text{Log}_{\text{psk}}^{D1}$ package uses a $D1$ pattern for logging which ignores level 0 $\text{UNQ}_{\text{psk}}(h, \text{hon}, k)$ queries with $\text{hon} = 0$ whenever there already exists a dishonest handle h' for key value k at level 0. Since $\text{XPD}_{\text{psk},\ell}$ writes only on level $\ell + 1 > 0$, this difference in logging does not affect the strength of the assumption, but it makes the assumption code align with the key schedule game, cf. Section 4.1. Finally, for deriving the psk, no hash-operation is performed.

Definition 5 (XPD for psk). For an adversary \mathcal{A} , we define the xpd pseudorandomness advantage for psk derivation $\text{Adv}(\mathcal{A}, \text{Gxpd}_{\text{psk},\ell}^0, \text{Gxpd}_{\text{psk},\ell}^1)$ as

$$|\Pr[1 = \mathcal{A} \rightarrow \text{Gxpd}_{\text{psk},\ell}^0] - \Pr[1 = \mathcal{A} \rightarrow \text{Gxpd}_{\text{psk},\ell}^1]|$$

XPD to derive esalt For $n = \text{esalt}$, the lower $\text{Log}_{\text{esalt}}^R$ package uses an R pattern instead of a D pattern, sending abort messages whenever the same key value k is registered as an *esalt* under two distinct handles h and h' (across all levels and regardless of honesty). Note that the adversary could simulate the R pattern itself (by retrieving all keys and checking for equality) and thus, the R pattern only *weakens* the adversary since it can no longer query the game after triggering an R abort and since the adversary does not learn the value of the collision which caused the abort.

Definition 6 (XPD for esalt). For an adversary \mathcal{A} , we define the xpd pseudorandomness advantage for esalt derivation $\text{Adv}(\mathcal{A}, \text{Gxpd}_{\text{esalt},\ell}^0, \text{Gxpd}_{\text{esalt},\ell}^1)$ as

$$|\Pr[1 = \mathcal{A} \rightarrow \text{Gxpd}_{\text{esalt},\ell}^0] - \Pr[1 = \mathcal{A} \rightarrow \text{Gxpd}_{\text{esalt},\ell}^1]|.$$




Fig. 6: xpd assumptions

$\underline{\text{Xtr}}_{n,\ell}^b$
Parameters

n : name
 ℓ : level
 b : bit
 $\text{prntn} : N \rightarrow (N_\perp \times N_\perp)$
 $\text{label} : N \times \{0, 1\} \rightarrow \{0, 1\}^{96}$


State

no state


$\text{XTR}_{n,\ell}(h_1, h_2)$

$n_1, n_2 \leftarrow \text{prntn}(n)$
if $\text{alg}(h_1) \neq \perp \wedge \text{alg}(h_2) \neq \perp$:
 assert $\text{alg}(h_1) = \text{alg}(h_2)$
 $h \leftarrow \text{xtr}(n, h_1, h_2)$
 $(k_1, \text{hon}_1) \leftarrow \text{GET}_{n_1,\ell}(h_1)$
 $(k_2, \text{hon}_2) \leftarrow \text{GET}_{n_2,\ell}(h_2)$
 $k \leftarrow \text{xtr}(k_1, k_2)$
 $\text{hon} \leftarrow \text{hon}_1 \vee \text{hon}_2$
if $b \wedge \text{hon}_2$:
 $k^* \leftarrow_{\$} \{0, 1\}^{\text{len}(k)}$
 $k \leftarrow \text{tag}_{\text{alg}(k)}(k^*)$
 $h \leftarrow \text{SET}_{n,\ell}(h, \text{hon}, k)$
return h


(a) Code of Xtr



(b) Game $\text{Gxtr1}^b_{es,\ell}$ for $b \in \{0, 1\}$



(c) Game $\text{Gxtr2}^b_{hs,\ell}$ for $b \in \{0, 1\}$



(d) Game $\text{Gxtr3}^b_{as,\ell}$ for $b \in \{0, 1\}$

Fig. 7: xtr Pseudorandomness Assumption

3.3 Pseudorandomness of xtr

The TLS 1.3 key schedule performs three xtr operations (cf. Fig. 1), and the modeling is analogous to the XPD assumptions, except that for the early secret es , xtr security relies on the psk which is the *right* input to xtr , and for the application secret as , xtr security relies on $esalt$ which is the *left* input to xtr . The derivation of the handshake secret hs is a special case, because its security is an *OR* of the honesty of its left and right input. We here state the xtr security assumption required for hs security based on its *left* input $esalt$ and turn to the security based in its right input (the Diffie-Hellman (DH) secret) shortly. Note that the security of $esalt$ will be applied *after* the security of the DH secret and thus, the bit b in the $\text{Xtr}^b_{hs,\ell}$ is already set to 1 and samples output keys uniformly at random whenever the Diffie-Hellman secret is honest. The security of $esalt$ thus only increases security for those keys where the Diffie-Hellman secret is dishonest.

Definition 7 (XTR advantages). For adversary \mathcal{A} , level $\ell \in \mathbb{N}_0$, we define the xtr pseudorandomness advantage for *es* as $\text{Adv}(\mathcal{A}, \text{Gxtr1}_{es,\ell}^0, \text{Gxtr1}_{es,\ell}^1)$, the pseudorandomness advantage for *hs* as $\text{Adv}(\mathcal{A}, \text{Gxtr2}_{hs,\ell}^0, \text{Gxtr2}_{hs,\ell}^1)$ and the pseudorandomness advantage for *as* as $\text{Adv}(\mathcal{A}, \text{Gxtr3}_{as,\ell}^0, \text{Gxtr3}_{as,\ell}^1)$, where Fig. 7b-7d define the games Gxtr1_{es}^b , Gxtr2_{hs}^b and Gxtr_{as}^b and Definition 1 defines advantage.

3.4 Salted ODH

Our salted oracle Diffie-Hellman assumption (SODH) is a stronger variant of the oracle Diffie-Hellman assumption introduced by Abdalla et al. [1] and the PRF oracle Diffie-Hellman assumption studied by Brendel et al. [20]. Most importantly, SODH is an *agile*, i.e., it requires pseudorandomness of the derived keys even when the adversary can see hash-values of the same Diffie-Hellman secret under *different* hash-functions and different, possibly adversarially chosen salts. In practice, different salts can emerge from disagreement between server and client about the PSK to use since the early salt *esalt* (and possibly also the *alg*) changes when the PSK changes (see Fig. 1). The Gsodh^b game (cf. Fig. 8) allows the adversary to generate honest Diffie-Hellman shares via DHGEN , to combine them (or an honest and a dishonest share) into a Diffie-Hellman secret via DHEXP and to derive keys from them via $\text{XTR}_{n,\ell}$ for an arbitrary level $\ell \in \{0, \dots, d\}$. Oracle $\text{GET}_{n,\ell}$ then allows to retrieve the derived keys. Note that pseudorandomness is modeled, this time, by a bit in the $\text{Xtr}_{n,\ell}^b$ package (Fig. 7a).

Definition 8 (SODH). For an adversary \mathcal{A} , we define the Salted Oracle Diffie Hellman (SODH) advantage $\text{Adv}(\mathcal{A}, \text{Gsodh}^0, \text{Gsodh}^1) :=$

$$|\Pr[1 = \mathcal{A} \rightarrow \text{Gsodh}^0] - \Pr[1 = \mathcal{A} \rightarrow \text{Gsodh}^1]|,$$

3.5 Pre-image resistance for xpd

Pseudorandomness and collision resistance of xpd also imply that it is hard to find pre-images for *honest* output keys. We prove this implication in the full




Fig. 8: Game Gsodh^b (top), package Dh (bottom)

version of this article [23, Lemma E.7] and in this conference version rely on pre-image resistance as a separate assumption for convenience.




Fig. 9: Pre-image resistance assumptions

Definition 9 (Pre-image resistance advantages). For an adversary \mathcal{A} and level $\ell \in \mathbb{N}_0$ we define the pre-image resistance advantage for deriving keys in O^* (a set to be specified later) $\text{Adv}(\mathcal{A}, \mathbf{Gpi}_{O^*}^D, \mathbf{Gpi}_{O^*}^F) :=$

$$|\Pr[1 = \mathcal{A} \rightarrow \mathbf{Gpi}_{O^*}^D] - \Pr[1 = \mathcal{A} \rightarrow \mathbf{Gpi}_{O^*}^F]|,$$

the pre-image resistance advantage for deriving keys with the same parent as esalt by $\text{Adv}(\mathcal{A}, \mathbf{Gpi}_{esalt}^D, \mathbf{Gpi}_{esalt}^F) :=$

$$|\Pr[1 = \mathcal{A} \rightarrow \mathbf{Gpi}_{esalt}^D] - \Pr[1 = \mathcal{A} \rightarrow \mathbf{Gpi}_{esalt}^F]|.$$

Fig. 9a and Fig. 9b define $\mathbf{Gpi}_{O^*}^P$ and \mathbf{Gpi}_{esalt}^P .

Our modular assumptions for xpd and xtr are agile, multi-instance security assumptions with registration of dishonest keys. They reduce to their non-agile, single-instance, monolithically written counterparts with a security loss equal to the number of honest keys. Since TLS 1.3 currently only supports hash-algorithms of different length, indeed, our agile assumptions for xtr and xpd reduce to *non-agile* assumptions. In turn, we can only reduce our modular agile SODH assumption to an *agile* monolithic SODH assumption, because TLS 1.3 indeed requires such a strong, agile SODH assumption (cf. Section 3.4 and Section 7) for further discussion. See full version [23, Appendix E] for the reduction proofs.




Fig. 10: Parent names `prntn` in TLS 1.3

4 Key Schedule

We reason about the TLS 1.3 key schedule in terms of its three elementary operations extract (`xtr`), expand (`xpd`) and computation of Diffie-Hellman secrets. This section first introduces an abstract key schedule syntax and refines it to capture TLS 1.3 as part of a bigger class of *TLS-like* key schedules. We then define key schedule security and state our theorem for all TLS-like key schedules.

4.1 Key Schedule Syntax

Our formalization interprets the key schedule as a directed graph where nodes describe *key names* (cf. Fig. 10 for the case of TLS 1.3). In addition to the set of names N and the graph description (encoded as `prntn` function, cf. Section 2.1), a key schedule has a function `label` which maps the name and a resumption bit to a derivation label. We conveniently model `hmac` operations by using `xpd` with *empty label* as an alias for `hmac`. By sound cryptographic practice, a key should be either used for `xpd` or for `hmac` but not both, so if a node has an empty label, it is not allowed to have siblings. Similarly, `xtr` operations only yield a single child, and the multiple children of `xpd` operations are derived using distinct labels.

Definition 10 (Key Schedule Syntax). *A key schedule $ks = (N, \text{label}, \text{prntn})$ consists of a set of names N and two functions*

$$\begin{aligned} \text{label} : & \quad N \times \{0, 1\} \rightarrow \{0, 1\}^{96} \cup \{\perp\} \\ \text{prntn} : & \quad N \rightarrow (N \cup \perp) \times (N \cup \perp) \end{aligned}$$

with the previously described restrictions.

Fig. 10 describes the `prntn` function of the TLS 1.3 key schedule as a graph. Stating and proving our theorem in terms of the concrete TLS key schedule would require listing and treating each `xpd` operation individually. Instead, we

prove our theorem for all *TLS-like* key schedules (of which the TLS key schedule is an instance). We consider a key schedule as *TLS-like* if it aligns with TLS in terms of base keys and xtr operations and treats the *psk* name as the main root from which all keys except for the base keys can be reached. Moreover, a *TLS-like* key schedule only has a single loop. This loop contains the edge from *rm* to *psk* and models resumptions. This edge has the special property of increasing the associated level as the *psk* is computed in an earlier session to be used in a later key schedule session. As such the cycle does not contradict an ordering on key computations.

Definition 11 (TLS-like Key Schedule Syntax). *A key schedule $ks = (N, \text{label}, \text{prntn})$ is TLS-like if its prntn graph satisfies the above restrictions, its set of names N contains at least the names $0\text{salt}, \text{psk}, \text{es}, \text{esalt}, \text{dh}, \text{hs}, \text{hsalt}, 0\text{ikm}, \text{as}, \text{rm}$ and the prntn function maps $0\text{salt}, \text{dh}$ and 0ikm to (\perp, \perp) , maps es, hs and as according to Fig. 10, maps psk to (rm, \perp) and each of the remaining names n to some pair (n_1, \perp) with $n_1 \neq \perp$.*

We use several subsets of N which we summarize in Table 1.

4.2 Key Schedule Security Model

Our key schedule security model captures that the key schedule produces keys which are pseudorandom and unique. We formulate security as indistinguishability between a real and an ideal game where the real game implements the actual key schedule derivations, while in the ideal game, output keys are unique, and honest keys are sampled uniformly at random. Concretely, we follow a simulation approach (somewhat similar to the Canetti and Krawczyk [26] approach to key exchange), where the ideal game is defined as a composition of a simulator \mathcal{S} and an ideal functionality. The simulator instructs the ideal functionality to produce output keys of certain length, however the *value* of the output keys is sampled independently from the simulator. As we require that no adversary can distinguish these two settings this captures security: The protocol determines when an output key becomes available and which type of key but no information about the concrete value is disclosed in the protocol (as the simulator does not have such information).

Concretely, in our ideal game $\text{Gks}^1(\mathcal{S})$ (Fig. 11b), the simulator \mathcal{S} is a parameter and the $\text{Key}_{O^*, 0..d}^1$ and Log_{O^*} packages (cf. Section 2.3) constitute the ideal functionality. Namely, the $\text{Key}_{O^*, 0..d}^1$ package samples a uniformly random key for handles which correspond to honest keys with a name $n \in O^*$ and some level $0 \leq \ell \leq d$. The Log_{O^*} package, in turn, ensures that each handle corresponds to a *different* key, modeling key uniqueness for both honest and dishonest keys.

Similarly, we describe the real execution of the key schedule as a game Gks^0 , written in pseudocode. Following the SSP methodology outlined in Section 2.3, we split the pseudocode of the game Gks^0 into several packages most of which (Xpd , Xtr , DH , Key , and Log) have been introduced before and Check is described in Section 4.3. Fig. 11a depicts the composed game Gks^0 —recall that this graph is not merely an illustration, it is part of the formal definition of Gks^0 .




Fig. 11: Key schedule security games with internal keys I^* , output keys O^* and XPN , the set of key names produced by xpd . We write OK_n as an abbreviation for $Nk_n \rightarrow L_n^Z$. We initialize K and Nk_n with suitable 0 values (cf. Section 2.1).

N	The set of all (key) names
N^*	$N \setminus \{psk, dh\}$
I^*	The set of internal keys $\{n \in N^* \mid \text{chldrnn}(n) = \emptyset\}$
O^*	The set of output keys $\{n \in N^* \mid \text{chldrnn}(n) = \emptyset\}$
O	$O^* \cup \{psk\}$
S	The set of separation points (Definition 13)
XPN	The set of expand names $\{n \in N : \text{prntn}(n) = (_, \perp)\}$
XPR	The set of representatives (Section 4.3)

Table 1: Notation

The game Gks^0 exposes $\text{SET}_{psk,0}$ and DHGEN oracles which sample honest Diffie-Hellman shares, honest application PSKs and enable the adversary to register dishonest application PSKs with a chosen value. The XTR and XPD oracles trigger key derivations. Finally, the adversary can access output keys via the GET oracle on the (real) key package $\text{Key}_{\mathcal{O}^*,0..d}$.

Definition 12 (Key Schedule Advantage). For a key schedule $ks = (N, \text{label}, \text{prntn})$, a natural number d , a simulator \mathcal{S} and an adversary \mathcal{A} which makes queries for at most d levels we define the advantage $\text{Adv}(\mathcal{A}, \text{Gks}^0, \text{Gks}^1(\mathcal{S})) :=$

$$|\Pr[1 = \mathcal{A} \rightarrow \text{Gks}^0] - \Pr[1 = \mathcal{A} \rightarrow \text{Gks}^1(\mathcal{S})]|,$$

where Fig. 11b defines $\text{Gks}^1(\mathcal{S})$ and Fig. 11a defines Gks^0 .

4.3 Front-End Checks

The **Check** package acts as a restriction on the adversary since the **assert** conditions in the **Check** code force the adversary to use the correct Diffie-Hellman shares and binder value in its transcript when the transcript is included in a derivation step. In terms of composability, the **assert** conditions in **Check** force the key exchange to call the key schedule with consistent values, i.e., derive the Diffie-Hellman secret from a pair of shares that is included in the transcript and not from an unrelated pair of shares. The TLS 1.3 specification ensures these innocent conditions, and requiring them formally means that the proof breaks down when session memory in TLS 1.3 is unsafely implemented.

In addition to enforcing the use of consistent shares in the transcript, the XPD oracle of the **Check** package (Fig. 12) ensures that the resumption flag is consistent with the level of the PSK; and that the binder tag included in the transcript of later stages (at the end of the last `ClientHello` message) is the same that was computed and checked in the early stage. The transcript is not included into all xpd derivations, but only once on the path from psk to output key, and **Check** only filters queries on these particular derivation steps. Since including the transcript ensures domain separation between different protocol runs and derivation pathes, we refer to the derivation steps which include the transcript as a *separation point*.

```

Check
 $\text{XPD}_{n,\ell}(h_1, r, args)$ 


---


if  $n = \text{bind}$  :
  if  $r = 0$ , assert  $\text{level}(h_1) = 0$ 
  if  $r = 1$ , assert  $\text{level}(h_1) > 0$ 
elseif  $n \in S \cap \text{early}$  :
   $\text{binder} \leftarrow \text{BinderArgs}(args)$ 
   $h_{\text{bndr}} \leftarrow \text{BinderHand}(h_1, args)$ 
   $(k, \_) \leftarrow \text{GET}_{\text{binder},\ell}(h_{\text{bndr}})$ 
  assert  $\text{binder} = k$ 
elseif  $n \in S$  :
   $X, Y \leftarrow \text{DhArgs}(args)$ 
   $h_{\text{dh}} \leftarrow \text{DhHand}(h_1)$ 
  assert  $h_{\text{dh}} = \text{dh}(\text{sort}(X, Y))$ 
   $\text{binder} \leftarrow \text{BinderArgs}(args)$ 
   $h_{\text{bndr}} \leftarrow \text{BinderHand}(h_1, args)$ 
   $(k, \_) \leftarrow \text{GET}_{\text{binder},\ell}(h_{\text{bndr}})$ 
  assert  $\text{binder} = k$ 
 $h \leftarrow \text{XPD}_{n,\ell}(h_1, r, args)$ 
return  $h$ 

```

Fig. 12: Code of **Check**

Definition 13 (Separation Points). For a key schedule $ks = (N, \text{label}, \text{prntn})$, we call $S \subseteq N$ a set of separation points, if it satisfies the following two requirements:

- $\forall n \in O$: the path from psk to n contains an $n' \in S$.
- If there exists a path from dh to an $n \in O$, then it contains an $n' \in S$.

In addition, for each xpd operation, we choose one representative child. I.e., $XPR \subseteq N$ is a *representative set* for ks if $psk, esalt \in XPR$ and for each name $n \in N$ with only a single parent (these are the xpd nodes), either n or exactly one sibling of n is contained in XPR .

5 Key Schedule Theorem

Theorem 1. Let ks be a TLS-like key schedule with representative set XPR and separation points S . Let $d \in \mathbb{N}$. There is an efficient simulator \mathcal{S} such that for all adversaries \mathcal{A} which make queries for at most d resumption levels,

$$\begin{aligned}
\text{Adv}(\mathcal{A}, \text{Gks}^0, \text{Gks}^1(\mathcal{S})) &\leq \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{cr}^{\text{main}}, \text{Gacr}^{\text{hash},b}) \\
&+ \sum_{j \in \{Z, D\}, f \in \{\text{xtr}, \text{xpd}\}} \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{j,f}^{\text{main}}, \text{Gacr}^{f,b}) \\
&+ \max_{i \in \{0,1\}} \left[\text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{sodh}^{\text{main}}, \text{Gsodh}^b) \right. \\
&\quad \sum_{\ell=0}^{d-1} \left(\text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{es,\ell}^{\text{main}}, \text{Gxtr}_{es,\ell}^b) \right. \\
&\quad \quad + \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{hs,\ell}^{\text{main}}, \text{Gxtr}_{hs,\ell}^b) \\
&\quad \quad + \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{as,\ell}^{\text{main}}, \text{Gxtr}_{as,\ell}^b) \\
&\quad \left. + \sum_{n \in XPR} \left(\text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{n,\ell}^{\text{main}}, \text{Gxpd}_{n,\ell}^b) \right) \right. \\
&\quad \quad + \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{esalt,pi}^{\text{main}}, \text{Gpi}_{esalt}^b) \\
&\quad \left. + \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{O^*,pi}^{\text{main}}, \text{Gpi}_{O^*}^b) \right],
\end{aligned}$$

where \mathcal{A}_i behaves as \mathcal{A} except that it returns bit i on a so-called win abort (cf. Lemma 8); $\mathcal{R}_*^{\text{main}} := \mathcal{R}^{\text{ch-map}} \rightarrow \mathcal{R}_*$ when replacing $*$ by $cr, (Z, f), (D, f), sodh, es, hs, as, n, O^*, pi$ or $esalt, pi$, the simulator \mathcal{S} is marked in grey in Fig. 16b, Fig. 18a defines \mathcal{R}_{sodh} , Fig. 18e defines $\mathcal{R}_{es,\ell}$, $\mathcal{R}_{hs,\ell}$ and $\mathcal{R}_{as,\ell}$ are defined analogously, and Fig. 18f defines $\mathcal{R}_{n,\ell}$ for $n \in XPR$, $0 \leq \ell \leq d$, Fig. 18b defines $\mathcal{R}_{esalt,pi}$ and Fig. 18c defines $\mathcal{R}_{O^*,pi}$.

5.1 Proof Technique

A recurrent proof technique which we use are *reductions*, written in SSP style. As usually, we want to show that if there is an adversary \mathcal{A} which successfully




Fig. 13: Proof Structure

distinguishes between two games \mathbf{G}_{big}^0 and \mathbf{G}_{big}^1 , then based on \mathcal{A} , we can construct an adversary \mathcal{B} of similar complexity as \mathcal{A} which successfully distinguishes between two games \mathbf{G}_{sml}^0 and \mathbf{G}_{sml}^1 . Our reductions will have the following form.

Lemma 1 (Reduction Technique). *If we can define a reduction \mathcal{R} such that*

$$\mathbf{G}_{big}^0 \stackrel{code}{\equiv} \mathcal{R} \rightarrow \mathbf{G}_{sml}^0 \quad (1) \quad \text{and} \quad \mathbf{G}_{big}^1 \stackrel{code}{\equiv} \mathcal{R} \rightarrow \mathbf{G}_{sml}^1 \quad (2)$$

then

$$\text{Adv}(\mathcal{A}; \mathbf{G}_{big}^0, \mathbf{G}_{big}^1) = \text{Adv}(\mathcal{B}; \mathbf{G}_{sml}^0, \mathbf{G}_{sml}^1), \quad (3)$$

where

$$\mathcal{B} := \mathcal{A} \rightarrow \mathcal{R}. \quad (4)$$

Proof. Assuming Equation 1, 2 and 4, we deduce Equation 3 as follows:

$$\begin{aligned} & \text{Adv}(\mathcal{A}, \mathbf{G}_{big}^0, \mathbf{G}_{big}^1) \\ & \stackrel{\text{def.}}{=} |\Pr[1 = \mathcal{A} \rightarrow \mathbf{G}_{big}^0] - \Pr[1 = \mathcal{A} \rightarrow \mathcal{R} \rightarrow \mathbf{G}_{big}^1]| \\ & \stackrel{\text{Eq. 1 \& 2}}{=} |\Pr[1 = \mathcal{A} \rightarrow (\mathcal{R} \rightarrow \mathbf{G}_{sml}^0)] - \Pr[1 = \mathcal{A} \rightarrow (\mathcal{R} \rightarrow \mathbf{G}_{sml}^1)]| \\ & = |\Pr[1 = (\mathcal{A} \rightarrow \mathcal{R}) \rightarrow \mathbf{G}_{sml}^0] - \Pr[1 = (\mathcal{A} \rightarrow \mathcal{R}) \rightarrow \mathbf{G}_{sml}^1]| \\ & \stackrel{\text{def.}}{=} \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}, \mathbf{G}_{sml}^0, \mathbf{G}_{sml}^1) \stackrel{\text{Eq. 4}}{=} \text{Adv}(\mathcal{B}, \mathbf{G}_{sml}^0, \mathbf{G}_{sml}^1) \end{aligned}$$

Importantly, throughout this article, we define reductions graphically as composition of previously defined packages so that the reduction *re-uses* code, as opposed to the usual technique which introduces new code for a reduction. As a

<u>Map</u>			
$\text{SET}_{\text{psk},0}(h, \text{hon}, k)$	$\text{XPD}_{n \in \text{XPN}, \ell}(h_1, r, \text{args})$	$\text{DHGEN}()$	$\text{XTR}_{n \in \{\text{es}, \text{hs}, \text{as}\}, \ell}(h_1, h_2)$
$h' \leftarrow \text{SET}_{\text{psk},0}(h, \text{hon}, k)$	$i_1, _ \leftarrow \text{prntidx}(n, \ell)$	$\text{return DHGEN}()$	$i_1, i_2 \leftarrow \text{prntidx}(n, \ell)$
$M_{\text{psk}}[h] \leftarrow h'$	$\text{assert } M_{i_1}[h_1] \neq \perp$	$\text{DHEXP}(X, Y)$	$\text{assert } M_{i_1}[h_1] \neq \perp$
$\text{return } h$	$\text{label} \leftarrow \text{label}(n, r)$	$h \leftarrow \text{dh}(\text{sort}(X, Y))$	$\text{assert } M_{i_2}[h_2] \neq \perp$
$\text{GET}_{n \in O^*, \ell}(h)$	$\ell_1 \leftarrow \text{level}(M_{i_1}[h_1])$	$h' \leftarrow \text{DHEXP}(X, Y)$	$\ell' \xleftarrow{\text{choose}} \text{level}(M_{i_1}[h_1]),$ $\text{level}(M_{i_2}[h_2])$
$\text{assert } M_{n, \ell}[h] \neq \perp$	$h \leftarrow \text{xpd}(n, \text{label}, h_1, \text{args})$	$\text{if } M_{dh}[h] = \perp :$	$h \leftarrow \text{xtr}(n, h_1, h_2)$
$h' \leftarrow M_{n, \ell}[h]$	$h' \leftarrow \text{XPD}_{n, \ell_1}(M_{i_1}[h_1], r, \text{args})$	$M_{dh}[h] \leftarrow h'$	$h' \leftarrow \text{XTR}_{n, \ell'}(M_{i_1}[h_1], M_{i_2}[h_2])$
$\text{return } h'$	$\text{if } n = \text{psk} : \ell \leftarrow \ell + 1$	$\text{return } h$	$M_{n, \ell}[h] \leftarrow h'$
$\text{GET}_{n, \text{level}(h')}(h')$	$M_{n, \ell}[h] \leftarrow h'$		$\text{return } h$
	$\text{return } h$		

Fig. 14: Oracles of **Map**. Here, $\ell \in \{0 \dots d\}$. $\ell' \xleftarrow{\text{choose}} \text{level}(M_{n_1}[h_1]), \text{level}(M_{n_2}[h_2])$ assigns to ℓ' the value $\text{level}(M_{n_1}[h_1])$ if it is not \perp and $\text{level}(M_{n_2}[h_2])$, else.

result, we can argue Equations 1 and 2 graphically. E.g., in Fig. 17b we highlight the reduction in gray and observe that the only change from Fig. 17a is the collision resistance assumption—the $\mathbf{G}_{\text{sml}}^b$ in this case. Observing the graph of \mathbf{Gks}^0 (cf. Fig. 11a) closely and comparing it with the graphs of the assumptions introduced in Section 3, one can identify that the assumptions are almost sub-graphs of \mathbf{Gks}^0 , and by an appropriately chosen sequence of reduction arguments, the graphs of the assumptions will appear as actual subgraphs.

5.2 Proof of Theorem 1

We need to show the indistinguishability of the real game \mathbf{Gks}^0 and the ideal game $\mathbf{Gks}^1(\mathcal{S})$. Fig. 15a depicts the real game \mathbf{Gks}^0 (cf. Fig. 11a), with slightly different graph layouting. Fig. 16b depicts the ideal game $\mathbf{Gks}^1(\mathcal{S})$ (cf. Fig. 11b) where the simulator \mathcal{S} is described in concrete code. To show the indistinguishability between \mathbf{Gks}^0 (Fig. 15a) and $\mathbf{Gks}^1(\mathcal{S})$ (Fig. 16b), we make 4 *game hops*, depicted as the sequence of the five games depicted in Fig. 15a, 15b, 15c, 16a and 16b. We now describe each of the game hops and state the corresponding lemma.

First, recall that the key schedule security model stores keys in a redundant fashion (a) due to possible equal values of a dishonest resumption psk ($\text{level}(h) > 0$) and an adversarially registered application psk ($\text{level}(h) = 0$) and (b) due to the equal values of the (dishonest) DH keys corresponding to (X^a, Y) and (X, Y^a) .

Lemma 2 introduces a **Map** package (see Fig. 15b for the game and the left column of Fig. 14 for the code of **Map**) to remove the redundantly stored keys—note that the $\text{Log}_{\text{psk}}^{A1}$ and the $\text{Log}_{dh}^{Z\infty}$ package now use the $\text{map} = 1$ and the $\text{map} = \infty$ code of **Log** (see Fig. 4 for its code). As a result, any adversary playing against \mathbf{Gcore}^0 (defined in Fig. 15b) cannot create (this particular) redundancy anymore since the $\text{Key}_{\text{psk}, \ell}$ and DHKey_{dh} packages do not store the key again when the mapping code is triggered. We defer the proof of code equality proof of

Lemma 2 to the full version [23]. It relies on proving the invariant that whenever \mathbf{Gks}^0 stores key k with honesty hon under handle h , then game \mathbf{Gks}^{0Map} stores key k with honesty hon under the mapped handle $h' = M[h]$. The proof proceeds by induction over the oracle calls.

Lemma 2 (Map-Intro). *For all adversaries \mathcal{A} which make queries for at most d resumption levels,*

$$\Pr[1 = \mathcal{A} \rightarrow \mathbf{Gks}^0] = \Pr[1 = \mathcal{A} \rightarrow \mathbf{Gks}^{0Map}].$$

In particular $\mathbf{Gks}^0 \stackrel{func}{\equiv} \mathbf{Gks}^{0Map}$.

Lemma 3 then reduces the indistinguishability of \mathbf{Gks}^{0Map} (Fig. 15b) and \mathbf{Gks}^{1Map} (Fig. 15c) to the indistinguishability of \mathbf{Gcore}^0 and $\mathbf{Gcore}^1(\mathcal{S}^{core})$ using reduction \mathcal{R}^{core} . The indistinguishability of \mathbf{Gcore}^0 and $\mathbf{Gcore}^1(\mathcal{S}^{core})$ will be established in Theorem 2 in Appendix 5.3 and contains the main technical argument of this article.

Lemma 3 (Main). *For all PPT adversaries \mathcal{A} which make queries for at most d resumption levels,*

$$\begin{aligned} & \text{Adv}(\mathcal{A}, \mathbf{Gks}^{0Map}, \mathbf{Gks}^{1Map}) \\ &= \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}^{ch-map}, \mathbf{Gcore}^0, \mathbf{Gcore}^1(\mathcal{S}^{core})), \end{aligned}$$

where Fig. 15b defines \mathbf{Gks}^{0Map} , Fig. 15c defines \mathbf{Gks}^{1Map} , \mathcal{R}^{ch-map} and \mathcal{S}^{core} are marked in grey in Fig. 15c, and Fig. 17a and Fig. 17i define \mathbf{Gcore}^0 and $\mathbf{Gcore}^1(\mathcal{S}^{core})$, respectively.

Proof. The proof of Lemma 3 is an instance of Lemma 1 with $\mathbf{G}_{big}^0 = \mathbf{Gks}^{0Map}$, $\mathbf{G}_{big}^1 = \mathbf{Gks}^{1Map}$, $\mathbf{G}_{sml}^0 = \mathbf{Gcore}^0$, $\mathbf{G}_{sml}^1 = \mathbf{Gcore}^1(\mathcal{S}^{core})$ and $\mathcal{R} = \mathcal{R}^{ch-map}$.

By Lemma 1, it suffices to show that

$$\mathbf{Gks}^{0Map} \stackrel{code}{\equiv} \mathcal{R}^{ch-map} \rightarrow \mathbf{Gcore}^0 \tag{5}$$

$$\mathbf{Gks}^{1Map} \stackrel{code}{\equiv} \mathcal{R}^{ch-map} \rightarrow \mathbf{Gcore}^1(\mathcal{S}^{core}) \tag{6}$$

Equation 5 follows by definition, since Fig. 15b defines \mathbf{Gks}^{0Map} as the composition of \mathcal{R}^{ch-map} and \mathbf{Gcore}^0 . Similarly, for Equation 6, Fig. 15c defines \mathbf{Gks}^{1Map} as the composition of \mathcal{R}^{ch-map} and $\mathbf{Gcore}^1(\mathcal{S}^{core})$.

In Lemma 4, we inline the $\mathbf{Xpd}_{n,0..d}$ code into \mathbf{Map} for $n \in O^*$ and call the result $\mathbf{Map-Xpd}$ (see Fig. 15c and Fig. 16a for the two games). The proof is a simple inlining argument and included into the full version [23] for completeness.

Lemma 4 (Xpd-Inlining). *For all PPT adversaries \mathcal{A} which make queries for at most d resumption levels,*

$$\Pr[1 = \mathcal{A} \rightarrow \mathbf{Gks}^{1Map}] = \Pr[1 = \mathcal{A} \rightarrow \mathbf{Gks}^{Mapxpd}].$$

In particular $\mathbf{Gks}^{1Map} \stackrel{code}{\equiv} \mathbf{Gks}^{Mapxpd}$.

Finally, Lemma 5 establishes the (perfect) indistinguishability of $\mathbf{Gks}^{\text{Map-Xpd}}$ and $\mathbf{Gks}^1(\mathcal{S})$. The proof of Lemma 5, essentially, removes or rather *inverts* the mapping on the output keys in order to recover the ideal functionality. Inverting the handle mapping, however, requires that it is *injective*. Conceptually, it is also clear that injectivity of the handle mapping needs to play a role in the proof: We prove uniqueness of output keys which means that equal keys imply equal handles. The injectivity proof ensures that the mapping did not introduce additional collisions and that the proof of Theorem 2 indeed suffices to establish the uniqueness of output keys in $\mathbf{Gks}^1(\mathcal{S})$.

Lemma 5 (Map-Outro). *For all PPT adversaries \mathcal{A} which make queries for at most d resumption levels,*

$$\Pr[1 = \mathcal{A} \rightarrow \mathbf{Gks}^{\text{Mapxpd}}] = \Pr[1 = \mathcal{A} \rightarrow \mathbf{Gks}^1(\mathcal{S})].$$

In particular, $\mathbf{Gks}^{\text{Mapxpd}} \stackrel{\text{func}}{\equiv} \mathbf{Gks}^1(\mathcal{S})$.

In summary, Lemma 3 is the core argument, Lemma 2 is proven via a mechanical invariant proof, Lemma 5 is proven via a conceptually interesting invariant proof and Lemma 4 is a straightforward inlining argument.

Theorem 1 directly follows from Lemma 2-Lemma 5 and Theorem 2 (stated in Section 5.3).

$$\begin{aligned} \text{Adv}(\mathcal{A}, \mathbf{Gks}^0, \mathbf{Gks}^1(\mathcal{S})) &\stackrel{\text{Lm. 2}}{=} \text{Adv}(\mathcal{A}, \mathbf{Gks}^{0\text{Map}}, \mathbf{Gks}^1(\mathcal{S})) \\ &\stackrel{\text{Lm. 5}}{=} \text{Adv}(\mathcal{A}, \mathbf{Gks}^{0\text{Map}}, \mathbf{Gks}^{\text{Mapxpd}}) \\ &\stackrel{\text{Lm. 4}}{=} \text{Adv}(\mathcal{A}, \mathbf{Gks}^{0\text{Map}}, \mathbf{Gks}^{1\text{Map}}) \\ &\stackrel{\text{Lm. 3}}{=} \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}^{\text{ch-map}}, \mathbf{Gks}^{0\text{core}}, \mathbf{Gks}^{1\text{core}}(\mathcal{S}^{\text{core}})) \\ &\stackrel{\text{Th. 2}}{\leq} \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{\text{cr}}^{\text{main}}, \mathbf{Gacr}^{\text{hash},b}) \\ &\quad + \sum_{j \in \{Z, D\}, f \in \{\text{xtr}, \text{xpd}\}} \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{j,f}^{\text{main}}, \mathbf{Gacr}^{\text{hash},b}) \\ &\quad + \max_{i \in \{0,1\}} \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{sodh}}^{\text{main}}, \mathbf{Gsodh}^b) \\ &\quad + \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{esalt}, pi}^{\text{main}}, \mathbf{Gpi}_{\text{esalt}}^b) \\ &\quad + \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{O^*, pi}^{\text{main}}, \mathbf{Gpi}_{O^*}^b) \\ &\quad + \sum_{\ell=0}^{d-1} (\text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{es}, \ell}^{\text{main}}, \mathbf{Gxtr}_{\text{es}, \ell}^b) \\ &\quad \quad + \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{hs}, \ell}^{\text{main}}, \mathbf{Gxtr}_{\text{hs}, \ell}^b) \\ &\quad \quad + \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{as}}^{\text{main}}, \mathbf{Gxtr}_{\text{as}, \ell}^b) \\ &\quad \quad + \sum_{n \in \text{XPR}} (\text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{n, \ell}^{\text{main}}, \mathbf{Gxpd}_{n, \ell}^b))), \end{aligned}$$

where XPR is the representative set required by the theorem, $\mathcal{R}_*^{\text{main}} := \mathcal{R}^{\text{ch-map}} \rightarrow \mathcal{R}_*$ when replacing $*$ by cr , (Z, f) , (D, f) sodh , es , hs , as , n , O^* , pi or esalt , pi .

5.3 Core Key Schedule Theorem

It remains to show that the *core* key schedule game Gcore^0 without the `Map` and `Check` package in front (Fig. 17a) is indistinguishable from an ideal game $\text{Gcore}^1(\mathcal{S}^{\text{core}})$ which consists of an ideal functionality with a simulator $\mathcal{S}^{\text{core}}$ (Fig. 17i). We prove Theorem 2 in Appendix A.

Theorem 2 (Core). *Let ks be a TLS-like key schedule with XPR . Let d be an integer. Let $\mathcal{S}^{\text{core}}$ be the efficient simulator defined in Fig. 16b. Then, for all adversaries \mathcal{A} which make queries for at most d resumption levels, we have that*

$$\begin{aligned}
& \text{Adv}(\mathcal{A}, \text{Gcore}^0, \text{Gcore}^1(\mathcal{S}^{\text{core}})) \\
\leq & \sum_{\mathcal{R} \in \{\mathcal{R}_{\text{cr}}, \mathcal{R}_Z, \mathcal{R}_D\}} \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}, \text{Gacr}^b) \\
& + \max_{i \in \{0,1\}} \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{sodh}}, \text{Gsodh}^b) \\
& + \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{esalt}, \text{pi}}, \text{Gpi}_{\text{esalt}}^b) \\
& + \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{O^*, \text{pi}}, \text{Gpi}_{O^*}^b) \\
& + \sum_{\ell=0}^{d-1} (\text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{es}, \ell}, \text{Gxtr}_{\text{es}, \ell}^b) \\
& \quad + \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{hs}, \ell}, \text{Gxtr}_{\text{hs}, \ell}^b) \\
& \quad + \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{as}, \ell}, \text{Gxtr}_{\text{as}, \ell}^b) \\
& \quad + \sum_{n \in XPR} (\text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{n, \ell}, \text{Gxpd}_{n, \ell}^b))),
\end{aligned}$$

where XPR is the required representation set (cf. Table 1), Fig. 17a defines Gcore^0 and Fig. 17i defines $\text{Gcore}^1(\mathcal{S}^{\text{core}})$, Fig. 17b defines \mathcal{R}_{cr} , Fig. 18a defines $\mathcal{R}_{\text{sodh}}$, Fig. 18e defines $\mathcal{R}_{\text{es}, \ell}$, $\mathcal{R}_{\text{hs}, \ell}$ and $\mathcal{R}_{\text{as}, \ell}$ are defined analogously, and $\mathcal{R}_{n, \ell}$ for $n \in XPR$ and $0 \leq \ell \leq d$ is defined in Fig. 18f, $\mathcal{R}_{\text{esalt}, \text{pi}}$ is defined in Fig. 18b and $\mathcal{R}_{O^*, \text{pi}}$ is defined in Fig. 18c.

6 Related Work

The following discussion focuses on attacker capabilities and security guarantees, and glosses over the exact encoding into security games and the use of multiple keys and stages.

Dowling et al. [34,35,36] present a multi-stage security model of `draft-05`, `draft-10`, and the final version of the standard. Their multi-stage model considers `psk_ke`, `dh_ke`, and `psk_dhe_ke` modes in isolation. Li et al. [48] adapt the multi-stage security model to also capture the recursive nature of the TLS

1.3 key schedule, by accounting for the re-use of resumption secrets between different modes (`psk_ke`, `psk_dhe_ke`, and the now removed semi-static share 0-RTT).

Cremers et al. [30,29] investigate the security of `draft-10` and `draft-21`, using the automated Tamarin prover (in the symbolic model). Their work investigates the proposed post-handshake client authentication and finds an attack that exploited a missing binding between PSKs and transcripts that led to the addition of binders to the standard. To our knowledge ours is the first reduction proof that models the additional security afforded by binder values.

Bhargavan et al. [10] also model TLS 1.3, decomposed into 3 separate pieces: `dh_ke` 1-RTT handshake, the 0-RTT handshake, and the record protocol. They verify these models using both ProVerif [18] and CryptoVerif [16]. A limitation of their model is the informal way in which the separate guarantees for the three components are combined to justify the overall security of the protocol.

Blanchet [17] introduces a new proof modularization framework in CryptoVerif, which bears significant similarities with the state-separating proof framework [24] that our work is based on. The work also updates some of the model from `draft-18` to `draft-28`; however, the model still assumes that all pre-shared keys are derived from resumption secrets and does not capture adaptively-created dishonest application PSKs, or the security of PSK binders.

Many other works focus on analysing certain properties of the TLS 1.3 handshake protocol. For instance, Arfaou et al. [4] specifically analyse the privacy of the TLS 1.3 `psk_ke`, `dh_ke`, and `psk_dhe_ke` handshakes. Fischlin et. al. [41] analyse the `draft-06` TLS 1.3 handshake, and show that its modes achieve key confirmation in isolation. Fischlin et. al. [39] considers replay attacks against various drafts of TLS 1.3 0-RTT handshakes such as `draft-14`'s `psk_ke` mode, similarly considering versions and modes in isolation. Other relevant papers on TLS handshake analysis are [46,37,27].

The idea of analyzing a key schedule (rather than a key exchange protocol) is conceptually similar to the SIGMA-I pattern of Krawczyk [44] and Krawczyk and Wee [47]. These works prove a reduction from key exchange security to key schedule security analogously to our companion paper [25].

Recent work also looked at the tightness of TLS 1.3 security proofs [33,31]. Besides natural birthday bounds for collision resistance, our reductions avoid the common quadratic loss in the number of sessions. We remark however, that tightness was not the principal focus of our analysis.

Subsequent work to the present article [22] uses our methodology, e.g., our recursive handle structure and the style of encoding security guarantees in `Log` packages to analyse the key schedule security of the Messaging Layer Security (MLS) protocol whose conclusions were integrated into the IETF standard, e.g., [28]. In the present paper, in addition to key techniques which were picked up by [22], we introduce a plethora of techniques to tackle *indirect* domain separation by late hashing of Diffie-Hellman shares and binders such as the notion of separation points and the `Check` component introduced in Section 4.3. In a similar way, the additional mapping step (Lemma 2, 4 and 5) handle redundancy

not present in MLS. See Section 7 for simplifications of the TLS protocol which would allow for a much simpler analysis than the one presented in this article.

7 Lessons Learned & Afterthoughts on the Key Schedule

We now discuss changes to the key schedule that would improve its security and simplify its analysis and may be of independent interest for other protocols.

Simplify SODH The salted Diffie-Hellman computation extracts entropy from the DH secret and mixes it with the PSK-derived salt (which is under adversarial influence). A separate DH extraction, preferably hashing the (sorted) public shares together with the secret, followed by a dual PRF, would enable a proof based on the simpler and better understood Oracle Diffie-Hellman assumption. The hashing of shares would also remove the need to map DH secrets (currently computable from multiple pairs of shares), and would enable the use of a more abstract functionality such as a CCA-secure KEM (as in TLS 1.2 [14]). These changes would thus also ease the integration of post-quantum secure primitives.

Eliminate PSK mapping Similarly, *directly* applying domain-separation for computations based on application and resumption PSKs via distinct labels would remove the need to map PSKs and argue via inclusion of binders at separation points indirectly. Both proposals follow the same design pattern: first sanitize input key materials to prevent malleability (DH secrets) and collisions (dishonest resumption PSKs and adversarially-chosen application PSKs).

Avoid Agile Assumptions Our development supports multiple hash algorithms without requiring any hash-agile assumptions, by observing that the hash functions currently used by TLS 1.3 have pairwise-distinct digest lengths. This is brittle, e.g. adding support for SHA3 with the same lengths as SHA2 would prevent formally accounting for cross-algorithm collisions. This may be prevented by tagging the outputs of all extractors and KDFs with hash algorithms. Similarly, we may avoid the current need for agile (S)ODH assumptions by tagging group elements with both a group descriptor and a single extraction algorithm.

Prevent PSK Reflections Drucker and Gueron note that TLS 1.3 is subject to reflection attacks due to its symmetric use of PSKs [37]. Hence, in our model, the same PSK handle may either be used by two parties, as intended, or by the same party acting both as a client and as a server. This is a security risk, inasmuch as applications may embed identity information in PSK identifiers to benefit from their early authentication. It may also enable key synchronization attacks and other variants of key compromise impersonation [13] when identities are also symmetrical. When using PSKs, the standard unfortunately forbids certificate-based authentication, which would otherwise provide more detailed, role-specific identity information. At the key schedule level, it may be possible to enforce better separation by tagging PSK identifiers with roles.

Enforce Stronger Modularity Applied cryptographers often complain that, in TLS 1.2, the subtle interleaving of the handshake with the record layer hinders its analysis based on the well-established Bellare-Rogaway [8] security model [43]. While TLS 1.3 tries to enforce cleaner separation between handshake and record keys, it still fails in some important places. Notably, the handshake traffic secrets, meant to be released to the record layer (be it TLS, DTLS, or QUIC) are also used by the handshake to derive finished keys. Similarly, some handshake messages are encrypted under keys derived from application traffic secrets (e.g. New Session Ticket, carrying resumption PSKs, late client authentication, and key updates). This complicates the modeling of data stream security, as application data may be interleaved with handshake messages (e.g. the same QUIC packet may contain both data and session tickets). To prevent such issues, and many others, we suggest the RFC documents more explicitly its application interface and, in particular, recommends not to derive keys from keys released to the record layer.

References

1. Abdalla, M., Bellare, M., Rogaway, P.: The oracle Diffie-Hellman assumptions and an analysis of DHIES. In: CT-RSA 2001. LNCS, vol. 2020, pp. 143–158. Springer, Heidelberg (Apr 2001). https://doi.org/10.1007/3-540-45353-9_12
2. Adrian, D., Bhargavan, K., Durumeric, Z., Gaudry, P., Green, M., Halderman, J.A., Heninger, N., Springall, D., Thomé, E., Valenta, L., VanderSloot, B., Wustrow, E., Zanella-Béguelin, S., Zimmermann, P.: Imperfect forward secrecy: How Diffie-Hellman fails in practice. In: ACM CCS 2015. pp. 5–17. ACM Press (Oct 2015). <https://doi.org/10.1145/2810103.2813707>
3. AlFardan, N.J., Paterson, K.G.: Lucky thirteen: Breaking the TLS and DTLS record protocols. In: 2013 S&P. pp. 526–540. IEEE (May 2013). <https://doi.org/10.1109/SP.2013.42>
4. Arfaoui, G., Bultel, X., Fouque, P.A., Nedelcu, A., Onete, C.: The privacy of the TLS 1.3 protocol. PoPETS **2019**(4), 190–210 (Oct 2019). <https://doi.org/10.2478/popets-2019-0065>
5. Aviram, N., Schinzel, S., Somorovsky, J., Heninger, N., Dankel, M., Steube, J., Valenta, L., Adrian, D., Halderman, J.A., Dukhovni, V., Käsper, E., Cohny, S., Engels, S., Paar, C., Shavitt, Y.: DROWN: Breaking TLS using SSLv2. In: USENIX Security 2016. pp. 689–706. USENIX (Aug 2016)
6. Badertscher, C., Matt, C., Maurer, U., Rogaway, P., Tackmann, B.: Augmented secure channels and the goal of the TLS 1.3 record layer. In: ProvSec 2015. LNCS, vol. 9451, pp. 85–104. Springer, Heidelberg (Nov 2015). https://doi.org/10.1007/978-3-319-26059-4_5
7. Bellare, M.: New proofs for NMAC and HMAC: Security without collision resistance. *Journal of Cryptology* **28**(4), 844–878 (Oct 2015). <https://doi.org/10.1007/s00145-014-9185-x>
8. Bellare, M., Rogaway, P.: Entity authentication and key distribution. In: CRYPTO'93. LNCS, vol. 773, pp. 232–249. Springer, Heidelberg (Aug 1994). https://doi.org/10.1007/3-540-48329-2_21
9. Beurdouche, B., Bhargavan, K., Delignat-Lavaud, A., Fournet, C., Kohlweiss, M., Pironti, A., Strub, P.Y., Zinzindohoue, J.K.: A messy state of the union: Taming

- the composite state machines of TLS. In: 2015 S&P. pp. 535–552. IEEE (May 2015). <https://doi.org/10.1109/SP.2015.39>
10. Bhargavan, K., Blanchet, B., Kobeissi, N.: Verified models and reference implementations for the TLS 1.3 standard candidate. In: 2017 S&P. pp. 483–502. IEEE (May 2017). <https://doi.org/10.1109/SP.2017.26>
 11. Bhargavan, K., Brzuska, C., Fournet, C., Green, M., Kohlweiss, M., Zanella-Béguelin, S.: Downgrade resilience in key-exchange protocols. In: 2016 S&P. pp. 506–525. IEEE (May 2016). <https://doi.org/10.1109/SP.2016.37>
 12. Bhargavan, K., Delignat-Lavaud, A., Fournet, C., Pironti, A., Strub, P.Y.: Triple handshakes and cookie cutters: Breaking and fixing authentication over tls. In: IEEE Symposium on Security & Privacy (Oakland) (2014), <pubs/triple-handshakes-and-cookie-cutters-sp14.pdf>
 13. Bhargavan, K., Delignat-Lavaud, A., Pironti, A.: Verified contributive channel bindings for compound authentication. In: NDSS 2015. ISOC (Feb 2015)
 14. Bhargavan, K., Fournet, C., Kohlweiss, M., Pironti, A., Strub, P.Y., Zanella Béguelin, S.: Proving the TLS handshake secure (as it is). In: CRYPTO 2014, Part II. LNCS, vol. 8617, pp. 235–255. Springer, Heidelberg (Aug 2014). https://doi.org/10.1007/978-3-662-44381-1_14
 15. Bhargavan, K., Leurent, G.: Transcript collision attacks: Breaking authentication in TLS, IKE and SSH. In: NDSS 2016. ISOC (Feb 2016)
 16. Blanchet, B.: CryptoVerif: Computationally sound mechanized prover for cryptographic protocols. In: Formal Protocol Verification. vol. 117, p. 156 (2007)
 17. Blanchet, B.: Composition theorems for CryptoVerif and application to TLS 1.3. In: CSF. pp. 16–30 (July 2018). <https://doi.org/10.1109/CSF.2018.00009>
 18. Blanchet, B., Smyth, B., Cheval, V., Sylvestre, M.: ProVerif 2.00: automatic cryptographic protocol verifier. User Manual (2018)
 19. Böck, H., Somorovsky, J., Young, C.: Return of bleichenbacher’s oracle threat (ROBOT). In: USENIX Security 2018. pp. 817–849. USENIX (Aug 2018)
 20. Brendel, J., Fischlin, M., Günther, F., Janson, C.: PRF-ODH: Relations, instantiations, and impossibility results. In: CRYPTO 2017, Part III. LNCS, vol. 10403, pp. 651–681. Springer, Heidelberg (Aug 2017). https://doi.org/10.1007/978-3-319-63697-9_22
 21. Bricout, R., Murphy, S., Paterson, K.G., van der Merwe, T.: Analysing and exploiting the mantin biases in RC4. Cryptology ePrint Archive, Report 2016/063 (2016), <http://eprint.iacr.org/2016/063>
 22. Brzuska, C., Cornelissen, E., Kohbrok, K.: Security analysis of the mls key derivation. In: 2022 IEEE Symposium on Security and Privacy. pp. 595–613. IEEE Computer Society, Los Alamitos, CA, USA (may 2022). <https://doi.org/10.1109/SP46214.2022.00035>
 23. Brzuska, C., Delignat-Lavaud, A., Egger, C., Fournet, C., Kohbrok, K., Kohlweiss, M.: Key-schedule security for the tls 1.3 standard. Cryptology ePrint Archive, Report 2021/467 (2021), <https://eprint.iacr.org/2021/467>
 24. Brzuska, C., Delignat-Lavaud, A., Fournet, C., Kohbrok, K., Kohlweiss, M.: State separation for code-based game-playing proofs. In: ASIACRYPT 2018, Part III. LNCS, vol. 11274, pp. 222–249. Springer, Heidelberg (Dec 2018). https://doi.org/10.1007/978-3-030-03332-3_9
 25. Brzuska, C., Egger, C.: Key exchange to key schedule reduction for TLS 1.3 (2022), preprint
 26. Canetti, R., Krawczyk, H.: Universally composable notions of key exchange and secure channels. In: EUROCRYPT 2002. LNCS, vol. 2332, pp. 337–351. Springer, Heidelberg (Apr / May 2002). https://doi.org/10.1007/3-540-46035-7_22

27. Chen, S., Jero, S., Jagielski, M., Boldyreva, A., Nita-Rotaru, C.: Secure communication channel establishment: TLS 1.3 (over TCP fast open) vs. QUIC. In: ESORICS 2019, Part I. LNCS, vol. 11735, pp. 404–426. Springer, Heidelberg (Sep 2019). https://doi.org/10.1007/978-3-030-29959-0_20
28. Cornelissen, E.: Pull request 453: Use the GroupContext to derive the joiner_secret, <https://github.com/mlswg/mls-protocol/pull/453>
29. Cremers, C., Horvat, M., Hoyland, J., Scott, S., van der Merwe, T.: A comprehensive symbolic analysis of TLS 1.3. In: ACM CCS 2017. pp. 1773–1788. ACM Press (2017)
30. Cremers, C., Horvat, M., Scott, S., van der Merwe, T.: Automated analysis and verification of TLS 1.3: 0-RTT, resumption and delayed authentication. In: 2016 S&P. pp. 470–485. IEEE (May 2016). <https://doi.org/10.1109/SP.2016.35>
31. Davis, H., Günther, F.: Tighter proofs for the SIGMA and TLS 1.3 key exchange protocols. In: ACNS 2021, Japan, June 21–24, 2021, Proceedings, Part II. Lecture Notes in Computer Science, vol. 12727, pp. 448–479. Springer (2021). https://doi.org/10.1007/978-3-030-78375-4_18, https://doi.org/10.1007/978-3-030-78375-4_18
32. Delignat-Lavaud, A., Fournet, C., Kohlweiss, M., Protzenko, J., Rastogi, A., Swamy, N., Zanella-Beguelin, S., Bhargavan, K., Pan, J., Zinzindohoue, J.K.: Implementing and proving the TLS 1.3 record layer. In: IEEE Security & Privacy. IEEE (2017)
33. Diemert, D., Jager, T.: On the tight security of TLS 1.3: Theoretically sound cryptographic parameters for real-world deployments. *Journal of Cryptology* **34**(3), 30 (Jul 2021). <https://doi.org/10.1007/s00145-021-09388-x>
34. Dowling, B., Fischlin, M., Günther, F., Stebila, D.: A cryptographic analysis of the TLS 1.3 handshake protocol candidates. In: ACM CCS 2015. pp. 1197–1210. ACM Press (Oct 2015). <https://doi.org/10.1145/2810103.2813653>
35. Dowling, B., Fischlin, M., Günther, F., Stebila, D.: A cryptographic analysis of the TLS 1.3 draft-10 full and pre-shared key handshake protocol. *Cryptology ePrint Archive*, Report 2016/081 (2016), <http://eprint.iacr.org/2016/081>
36. Dowling, B., Fischlin, M., Günther, F., Stebila, D.: A cryptographic analysis of the TLS 1.3 handshake protocol. *Journal of Cryptology* **34**(4), 37 (Oct 2021). <https://doi.org/10.1007/s00145-021-09384-1>
37. Drucker, N., Gueron, S.: Selfie: reflections on TLS 1.3 with PSK. *Journal of Cryptology* **34**(3), 27 (Jul 2021). <https://doi.org/10.1007/s00145-021-09387-y>
38. Duong, T., Rizzo, J.: Here come the \oplus ninjas (2011), http://nerdoholic.org/uploads/dergln/beast_part2/ssl_jun21.pdf
39. Fischlin, M., Günther, F.: Replay attacks on zero round-trip time: The case of the TLS 1.3 handshake candidates. In: 2017 IEEE European Symposium on Security and Privacy (EuroS&P). pp. 60–75. IEEE (2017)
40. Fischlin, M., Günther, F., Marson, G.A., Paterson, K.G.: Data is a stream: Security of stream-based channels. In: CRYPTO 2015, Part II. LNCS, vol. 9216, pp. 545–564. Springer, Heidelberg (Aug 2015). https://doi.org/10.1007/978-3-662-48000-7_27
41. Fischlin, M., Günther, F., Schmidt, B., Warinschi, B.: Key confirmation in key exchange: A formal treatment and implications for TLS 1.3. In: 2016 S&P. pp. 452–469. IEEE (May 2016). <https://doi.org/10.1109/SP.2016.34>
42. Iyengar, J., Thomson, M.: QUIC. IETF draft (2019)
43. Jager, T., Kohlar, F., Schäge, S., Schwenk, J.: Authenticated confidential channel establishment and the security of TLS-DHE. *Journal of Cryptology* **30**(4), 1276–1324 (Oct 2017). <https://doi.org/10.1007/s00145-016-9248-2>

44. Krawczyk, H.: SIGMA: The “SIGn-and-MAC” approach to authenticated Diffie-Hellman and its use in the IKE protocols. In: CRYPTO 2003. LNCS, vol. 2729, pp. 400–425. Springer, Heidelberg (Aug 2003). https://doi.org/10.1007/978-3-540-45146-4_24
45. Krawczyk, H.: Cryptographic extraction and key derivation: The HKDF scheme. In: CRYPTO 2010. LNCS, vol. 6223, pp. 631–648. Springer, Heidelberg (Aug 2010). https://doi.org/10.1007/978-3-642-14623-7_34
46. Krawczyk, H.: A unilateral-to-mutual authentication compiler for key exchange (with applications to client authentication in TLS 1.3). In: ACM CCS 2016. pp. 1438–1450. ACM Press (Oct 2016). <https://doi.org/10.1145/2976749.2978325>
47. Krawczyk, H., Wee, H.: The OPTLS protocol and TLS 1.3. Cryptology ePrint Archive, Report 2015/978 (2015), <http://eprint.iacr.org/2015/978>
48. Li, X., Xu, J., Zhang, Z., Feng, D., Hu, H.: Multiple handshakes security of TLS 1.3 candidates. In: 2016 S&P. pp. 486–505. IEEE (May 2016). <https://doi.org/10.1109/SP.2016.36>
49. Mavrogianopoulos, N., Vercauteren, F., Velichkov, V., Preneel, B.: A cross-protocol attack on the TLS protocol. In: ACM CCS 2012. pp. 62–72. ACM Press (Oct 2012). <https://doi.org/10.1145/2382196.2382206>
50. Paterson, K.G., van der Merwe, T.: Reactive and proactive standardisation of TLS. In: Security Standardisation Research. pp. 160–186 (2016)
51. Patton, C., Shrimpton, T.: Partially specified channels: The TLS 1.3 record layer without elision. Cryptology ePrint Archive, Report 2018/634 (2018)
52. Rescorla, E.: The Transport Layer Security (TLS) Protocol Version 1.3. <https://tools.ietf.org/html/rfc8446> (Aug 2018)

A Proof of Theorem 2

Fig. 17 provides an overview over the proof of Theorem 2, which consists of 5 game-hops: From Gcore^0 (Fig. 17a) to $\text{Gcore}^{\text{Hash}}$ (Fig. 17b), to Gcore^D (Fig. 17c), to $\text{Gcore}^{\text{Resalt}}$ (Fig. 17d) and finally to $\text{Gcore}^1(\mathcal{S})$ (Fig. 17i), as required by Theorem 2. The last step from $\text{Gcore}^{\text{Resalt}}$ and to $\text{Gcore}^1(\mathcal{S})$ is implemented by a hybrid argument which we prove in Appendix A, the hybrid games are illustrated in Fig. 17f, Fig. 17g and Fig. 17h. The first game hop (Lemma 6) idealizes collision resistance of the $\text{Hash}_{\mathcal{H}}^b$ package (which is used to hash transcripts).

Lemma 6 (Collision-Resistance).

$$\text{Adv}(\mathcal{A}, \text{Gcore}^0, \text{Gcore}^{\text{Hash}}) \leq \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{cr}, \text{Hash}_{\mathcal{H}}^0, \text{Hash}_{\mathcal{H}}^1),$$

where Fig. 17a and Fig. 17b define Gcore^0 and $\text{Gcore}^{\text{Hash}}$, respectively, and \mathcal{R}_{cr} is marked in grey in Fig. 17b.

Proof. The proof of Lemma 6 is an instance of Lemma 1 with $\mathbf{G}_{\text{big}}^0 = \text{Gcore}^0$, $\mathbf{G}_{\text{big}}^1 = \text{Gcore}^{\text{Hash}}$, $\mathbf{G}_{\text{sml}}^0 = \text{Gacr}^{\text{hash},0}$, $\mathbf{G}_{\text{sml}}^1 = \text{Gacr}^{\text{hash},1}$ and $\mathcal{R} = \mathcal{R}_{cr}$.

By Lemma 1, it suffices to show that

$$\text{Gcore}^0 \stackrel{\text{code}}{\equiv} \mathcal{R}_{cr} \rightarrow \text{Gacr}^{\text{hash},0} \tag{7}$$

$$\text{Gcore}^{\text{Hash}} \stackrel{\text{code}}{\equiv} \mathcal{R}_{cr} \rightarrow \text{Gacr}^{\text{hash},1} \tag{8}$$

Equation 7 follows by definition, since Fig. 15b defines $\mathbf{Gks}^{\text{0Map}}$ as the composition of $\mathcal{R}^{\text{ch-map}}$ and \mathbf{Gcore}^0 . Similarly, for Equation 6, Fig. 15c defines $\mathbf{Gks}^{\text{1Map}}$ as the composition of $\mathcal{R}^{\text{ch-map}}$ and $\mathbf{Gcore}^1(\mathcal{S}^{\text{core}})$.

The second game hop (Lemma 7) changes the pattern in the Log packages from Z to D and reduces to collision-resistance.

Lemma 7 (D-Pattern).

$$\begin{aligned} & \text{Adv}(\mathcal{A}, \mathbf{Gcore}^{\text{Hash}}, \mathbf{Gcore}^D) \leq \\ & \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{Z,\text{xtr}}, \mathbf{Gacr}^{\text{xtr},b}) + \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{Z,\text{xpd}}, \mathbf{Gacr}^{\text{xpd},b}) + \\ & \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{D,\text{xtr}}, \mathbf{Gacr}^{\text{xtr},b}) + \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{D,\text{xpd}}, \mathbf{Gacr}^{\text{xpd},b}), \end{aligned}$$

where $\mathbf{Gcore}^{\text{Hash}}$ is defined in Fig. 17b, \mathbf{Gcore}^D is defined in Fig. 17c, and the reductions $\mathcal{R}_{Z,\text{xtr}}$, $\mathcal{R}_{Z,\text{xpd}}$, $\mathcal{R}_{D,\text{xpd}}$ and $\mathcal{R}_{D,\text{xtr}}$ are provided in the full version [23].

Proof (Proof Sketch). We (a) first idealize collision-resistance for xtr and xpd using the \mathbf{Rxtr} and \mathbf{Rxp} implementation (cf. Fig. 18d) instead of \mathbf{Xtr} and \mathbf{Xpd} to construct reductions $\mathcal{R}_{Z,\text{xtr}}$ and $\mathcal{R}_{Z,\text{xpd}}$ to agile collision-resistance (cf. Section 3.1) of xtr and xpd , respectively. (b) Then, we introduce D abort patterns in a perfect equivalence step and then (c) de-idealize collision-resistance of xtr and xpd analogously to (a). Step (b) relies on an *inductive* argument which we refer to as the *co-dependance* lemma. Namely, we show that once collision-resistance of the hash-function is idealized, a D abort on one key requires that there has been a collision before and therefore, since every collision requires a previous collision, no collision can occur. We defer the details of this argument to the full version.

The third game-hop (Lemma 8) now introduces a special abort symbol *win* whenever two different *esalt* handles point to the same key value, *regardless* of the honesty of the handles.

We show that the introduction of this special *win* abort can only increase the success probability of an adversary \mathcal{A} by providing two adversaries \mathcal{A}_0 and \mathcal{A}_1 which behave almost identically to \mathcal{A} , but in addition to the advantage which \mathcal{A} has anyway, at least one out of \mathcal{A}_0 and \mathcal{A}_1 succeeds in benefitting from the *win* aborts. We defer the details of this proof to the full version. In a later game hop, we show that *win* aborts can be removed. Note that in this lemma, we do not bound the difference of an adversary between \mathbf{Gcore}^D and $\mathbf{Gcore}^{\text{Resalt}}$, but instead bound the difference between *advantages*, when comparing either of the games to $\mathbf{Gcore}^1(\mathcal{S}^{\text{core}})$.

Lemma 8 (Resalt). *For every adversary \mathcal{A} and $i \in \{0,1\}$, let \mathcal{A}_i be the adversary which behaves as \mathcal{A} , but returns i when receiving a win abort. Then,*

$$\begin{aligned} & \text{Adv}(\mathcal{A}, \mathbf{Gcore}^D, \mathbf{Gcore}^1(\mathcal{S}^{\text{core}})) \\ & \leq \max_{i \in \{0,1\}} \text{Adv}(\mathcal{A}, \mathbf{Gcore}^{\text{Resalt}}, \mathbf{Gcore}^1(\mathcal{S}^{\text{core}})) \end{aligned} \tag{9}$$

where $\mathbf{Gcore}^{\text{Resalt}}$ is defined in Fig. 17c.

Proof. First, observe that

$$\Pr[1 = \mathcal{A} \rightarrow \text{Gcore}^D] \leq \Pr[1 = \mathcal{A}_1 \rightarrow \text{Gcore}^{\text{Resalt}}] \quad (10)$$

because \mathcal{A}_1 returns 1 whenever \mathcal{A} returns 1 *plus* when receiving a *win* abort. Moreover,

$$\Pr[1 = \mathcal{A} \rightarrow \text{Gcore}^1(\mathcal{S})] = \Pr[1 = \mathcal{A}_1 \rightarrow \text{Gcore}^1(\mathcal{S})], \quad (11)$$

because $\text{Gcore}^1(\mathcal{S})$ never returns *win*. From here, the proof proceeds by case distinction. **Case I.**

$$\Pr[1 = \mathcal{A} \rightarrow \text{Gcore}^D] \geq \Pr[1 = \mathcal{A} \rightarrow \text{Gcore}^1(\mathcal{S}^{\text{core}})] \quad (12)$$

Inequality 9 holds because

$$\begin{aligned} & \text{Adv}(\mathcal{A}, \text{Gcore}^D, \text{Gcore}^1(\mathcal{S}^{\text{core}})) \\ & \stackrel{\text{Eq. 12}}{=} \Pr[1 = \mathcal{A} \rightarrow \text{Gcore}^D] - \Pr[1 = \mathcal{A} \rightarrow \text{Gcore}^1(\mathcal{S}^{\text{core}})] \\ & \stackrel{\text{Eq. 10\&11}}{\leq} \Pr[1 = \mathcal{A}_1 \rightarrow \text{Gcore}^D] - \Pr[1 = \mathcal{A}_1 \rightarrow \text{Gcore}^1(\mathcal{S}^{\text{core}})] \end{aligned}$$

Case II consists of Inequality 12 being the other way round. The proof is analogous.

With the collision-freeness of *esalt* at hand, we can (globally) apply the Gsodh assumption to idealize Diffie-Hellman secrets (Lemma 9). The proof is brief and straightforward and included after the lemma statement.

Lemma 9 (SODH Lemma).

$$\text{Adv}(\mathcal{A}, \text{Gcore}^{\text{Resalt}}, \text{Gcore}^{\text{SODH}}) = \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{\text{sodh}}, \text{Gsodh}^b),$$

where Fig. 17c and Fig. 17d define $\text{Gcore}^{\text{Resalt}}$ and $\text{Gcore}^{\text{SODH}}$, respectively, and $\mathcal{R}_{\text{sodh}}$ is marked in grey in Fig. 18a.

Proof. The proof of Lemma 9 is an instance of Lemma 1 with $\text{G}_{\text{big}}^0 = \text{Gcore}^{\text{Resalt}}$, $\text{G}_{\text{big}}^1 = \text{Gcore}^{\text{SODH}}$, $\text{G}_{\text{sml}}^0 = \text{Gsodh}^0$, $\text{G}_{\text{sml}}^1 = \text{Gsodh}^1$ and $\mathcal{R} = \mathcal{R}_{\text{sodh}}$.

By Lemma 1, it suffices to show that

$$\text{Gcore}^{\text{Resalt}} \stackrel{\text{code}}{\equiv} \mathcal{R}_{\text{sodh}} \rightarrow \text{Gsodh}^0 \quad (13)$$

$$\text{Gcore}^{\text{SODH}} \stackrel{\text{code}}{\equiv} \mathcal{R}_{\text{sodh}} \rightarrow \text{Gsodh}^1 \quad (14)$$

Equation 13 follows by observing that the graph of $\mathcal{R}_{\text{sodh}} \rightarrow \text{Gsodh}^0$ (Fig. 18a with $b = 0$) is identical to the graph of $\text{Gcore}^{\text{Resalt}}$ (Fig. 17c) by inspecting the different layouts of the graphs chosen in the two figures representing them. Similarly, for Equation 14, one can observe that the graph of $\mathcal{R}_{\text{sodh}} \rightarrow \text{Gsodh}^1$ (Fig. 18a with $b = 1$) is equal to the graph of $\text{Gcore}^{\text{SODH}}$ (Fig. 17d).

At this point we use a hybrid argument (Lemma 10) to idealize pseudorandomness of all xtr and xpd blocks. We provide the proof of Lemma 10 in Appendix B.

Lemma 10 (Hybrid Lemma).

$$\begin{aligned}
\text{Adv}(\mathcal{A}, \text{Gcore}^{\text{SODH}}, \text{Gcore}^{ki}) &\leq \sum_{\ell=0}^d \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{es,\ell}, \text{Gxtr}_{es,\ell}^b) \\
&+ \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{hs,\ell}, \text{Gxtr}_{hs,\ell}^b) + \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{as,\ell}, \text{Gxtr}_{as,\ell}^b) \\
&+ \sum_{n \in \text{XPR}} (\text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{n,\ell}, \text{Gxpd}_{n,\ell}^b)),
\end{aligned} \tag{15}$$

where Fig. 18f defines $\mathcal{R}_{n,\ell}$ for $n \in \text{XPN} \setminus \{psk, esalt\}$ and Fig. 18e defines $\mathcal{R}_{es,\ell}$.

Finally, we use preimage-resistance (Lemma 11) to remove the rewarding aborts introduced in Lemma 8.

Lemma 11 (Pre-image-resistance).

$$\begin{aligned}
&\text{Adv}(\mathcal{A}, \text{Gcore}^{ki}, \text{Gcore}^1(\mathcal{S}^{\text{core}})) \\
&\leq \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{esalt,pi}, \text{Gpi}_{esalt}^R, \text{Gpi}_{esalt}^D) \\
&\quad + \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{O^*,pi}, \text{Gpi}_{O^*}^D, \text{Gpi}_{O^*}^F),
\end{aligned} \tag{16}$$

where $\mathcal{R}_{esalt,pi}$ is defined in Fig. 18b and $\mathcal{R}_{O^*,pi}$ is defined in Fig. 18c.

Proof. The proof of Lemma 11 consists of two instances of Lemma 1. For the first part of the proof, we use $\text{G}_{big}^0 = \text{Gcore}^{ki}$, $\text{G}_{big}^1 = \mathcal{R}_{esalt,pi} \rightarrow \text{Gpi}_{esalt}^D$, $\text{G}_{sml}^0 = \text{Gpi}_{esalt}^R$, $\text{G}_{sml}^1 = \text{Gpi}_{esalt}^D$ and $\mathcal{R}_{esalt,pi}$. Now, $\text{Gcore}^{ki} \stackrel{\text{code}}{\equiv} \mathcal{R}_{esalt,pi} \rightarrow \text{Gpi}_{esalt}^R$ follows by carefully inspecting the graph layouts in Fig. 17h and Fig. 18b. $\text{G}_{big}^1 \stackrel{\text{code}}{\equiv} \mathcal{R} \rightarrow \text{G}_{sml}^1$ is true by definition. By Lemma 1, we thus obtain

$$\begin{aligned}
&\text{Adv}(\mathcal{A}, \text{Gcore}^{ki}, \mathcal{R}_{esalt,pi} \rightarrow \text{Gpi}_{esalt}^D) \\
&= \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{esalt,pi}, \text{Gpi}_{esalt}^R, \text{Gpi}_{esalt}^D)
\end{aligned} \tag{17}$$


For the second instance of Lemma 1, we use $\text{G}_{big}^0 = \mathcal{R}_{esalt,pi} \rightarrow \text{Gpi}_{esalt}^R$, $\text{G}_{big}^1 = \text{Gcore}^1(\mathcal{S}^{\text{core}})$, $\text{G}_{sml}^0 = \text{Gpi}_{O^*}^D$, $\text{G}_{sml}^1 = \text{Gpi}_{O^*}^F$ and $\mathcal{R}_{O^*,pi}$. $\mathcal{R}_{esalt,pi} \rightarrow \text{Gpi}_{esalt}^R$ is code equivalent to $\mathcal{R}_{O^*,pi} \rightarrow \text{Gpi}_{O^*}^F$ by carefully comparing Fig. 18b (with $P = D$) and Fig. 18c (with $\mathbf{P} = \mathbf{D}$). $\text{Gcore}^1(\mathcal{S}^{\text{core}})$ is code-equivalent to $\mathcal{R}_{O^*,pi} \rightarrow \text{Gpi}_{O^*}^D$ by comparing Fig. 18c (with $\mathbf{P} = \mathbf{F}$) and Fig. 17i. By Lemma 1, we obtain

$$\begin{aligned}
&\text{Adv}(\mathcal{A}, \mathcal{R}_{esalt,pi} \rightarrow \text{Gpi}_{esalt}^D, \text{Gcore}^1(\mathcal{S}^{\text{core}})) \\
&= \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{O^*,pi}, \text{Gpi}_{O^*}^D, \text{Gpi}_{O^*}^F)
\end{aligned} \tag{18}$$


From Eq. 17 and 18, we obtain Inequality 16 via triangle inequality.

From Lemma 6-11, we obtain Theorem 2 by triangle inequalities, included below for completeness. All sums over \mathcal{R} sum over $\mathcal{R} \in \{\mathcal{R}_{\text{cr}}, \mathcal{R}_{Z,\text{xtr}}, \mathcal{R}_{Z,\text{xpd}}, \mathcal{R}_{D,\text{xtr}}, \mathcal{R}_{D,\text{xpd}}\}$ and use matching $f \in \{\text{hash}, \text{xtr}, \text{xpd}\}$


$$\begin{aligned}
& \text{Adv}(\mathcal{A}, \text{Gcore}^0, \text{Gcore}^1(\mathcal{S}_{\text{core}})) \\
& \stackrel{\text{Lm. 6}}{\leq} \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{\text{cr}}, \text{Gacr}^{\text{hash},b}) + \text{Adv}(\mathcal{A}, \text{Gks}^{\text{Hash}}, \text{Gcore}^1(\mathcal{S}^{\text{core}})) \\
& \stackrel{\text{Lm. 7}}{\leq} \sum_{\mathcal{R}} \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}, \text{Gacr}^{f,b}) + \text{Adv}(\mathcal{A}, \text{Gcore}^{\text{D}}, \text{Gcore}^1(\mathcal{S}^{\text{core}})) \\
& \stackrel{\text{Lm. 8}}{\leq} \sum_{\mathcal{R}} \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}, \text{Gacr}^{f,b}) + \max_{i \in \{0,1\}} \text{Adv}(\mathcal{A}_i, \text{Gcore}^{\text{Resalt}}, \text{Gcore}^1(\mathcal{S}^{\text{core}})) \\
& \stackrel{\text{Lm. 9}}{=} \sum_{\mathcal{R}} \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}, \text{Gacr}^{f,b}) + \max_{i \in \{0,1\}} \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{sodh}}, \text{Gsodh}^b) \\
& \quad + \text{Adv}(\mathcal{A}_i, \text{Gcore}^{\text{SODH}}, \text{Gcore}^1(\mathcal{S}^{\text{core}})) \\
& \stackrel{\text{Lm. 10}}{\leq} \sum_{\mathcal{R}} \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}, \text{Gacr}^{f,b}) + \max_{i \in \{0,1\}} \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{sodh}}, \text{Gsodh}^b) \\
& \quad + \text{Adv}(\mathcal{A}_i, \text{Gcore}^{\text{ki}}, \text{Gcore}^1(\mathcal{S}^{\text{core}})) + \\
& \quad \sum_{\ell=0}^{d-1} (\text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{es},\ell}, \text{Gxtr}_{\text{es},\ell}^b) + \max_{i \in \{0,1\}} \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{hs},\ell}, \text{Gxtr}_{\text{hs},\ell}^b) + \\
& \quad \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{as}}, \text{Gxtr}_{\text{as},\ell}^b) + \sum_{n \in \text{XPR}} (\text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{n,\ell}, \text{Gxpd}_{n,\ell}^b))) \\
& \stackrel{\text{Lm. 11}}{\leq} \sum_{\mathcal{R}} \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}, \text{Gacr}^{f,b}) + \max_{i \in \{0,1\}} \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{sodh}}, \text{Gsodh}^b) \\
& \quad + \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{esalt},pi}, \text{Gpi}_{\text{esalt}}^b) + \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{O^*,pi}, \text{Gpi}_{O^*}^b) + \\
& \quad \sum_{\ell=0}^{d-1} (\text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{es},\ell}, \text{Gxtr}_{\text{es},\ell}^b) + \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{hs},\ell}, \text{Gxtr}_{\text{hs},\ell}^b) + \\
& \quad \text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{\text{as}}, \text{Gxtr}_{\text{as},\ell}^b) + \sum_{n \in \text{XPR}} (\text{Adv}(\mathcal{A}_i \rightarrow \mathcal{R}_{n,\ell}, \text{Gxpd}_{n,\ell}^b))),
\end{aligned}$$



(a) Game Gks^0 .




(b) Game $Gks^{0Map} \stackrel{code}{\equiv} \mathcal{R}^{ch-map} \rightarrow Gcore^0$.




(c) Game $Gks^{1Map} \stackrel{code}{\equiv} \mathcal{R}^{ch-map} \rightarrow Gcore^1(Score)$.

Fig. 15: Overview over proof of Theorem 1



(a) Game G_{ks}^{Mapxpd} .



(b) Game G_{ks}^1 .

Fig. 16: Overview over proof of Theorem 1







Fig. 17: Proof overview for Theorem 2




(a) Game $G_{\text{core}}^{\text{SODH}}$ if $b = 0$ and game $G_{\text{core}}^{\text{SODH}}$ if $b = 1$. Reduction $\mathcal{R}_{\text{SODH}}$ is marked in grey.




(b) Game $G_{\text{core}}^{\text{ki}}$ if $P = R$. Reduction $\mathcal{R}_{\text{Piesalt}}$ is marked in grey.




(c) Game $G_{\text{core}}^1(S^{\text{core}})$ if $P = F$. Reduction \mathcal{R}_{PIO} is marked in grey.



(d) Rxt_n for $n \in \{es, hs, as\}$ and Rxp_n for $n \in \text{XPN}$.



(e) Reduction $\mathcal{R}_{es,\ell}$ is marked in grey. Reductions $\mathcal{R}_{psk,\ell}$ and $\mathcal{R}_{esalt,\ell}$ are defined analogously.



(f) Reduction $\mathcal{R}_{n,\ell}$ is marked in grey for $n \in \text{XPN} \setminus \{psk, esalt\}$.

Fig. 18: Reductions for the lemmas related to Theorem 2

B Proof of Lemma 10

B.1 Idealization Order

The proof of Lemma 10 is a hybrid argument over all levels as well as *over the derivation graph*. Here, we need to specify the order in which keys are idealized, i.e., replaced by random keys. This order needs to satisfy the (a) all parents have been idealized before and (b) that when n is idealized in a step then all siblings are idealized, too. We refer to such an order as *idealization order* and show that each TLS-like key schedule indeed has an idealization order (as is required for the proof of Lemma 10).

Definition 14 (Idealization Order). Let $ks = (N, \text{label}, \text{prntn})$ be a TLS-like key schedule and $m \in \mathbb{N}$. An Idealization Order io for ks is a total order on m subsets of N together with a sequence of names n^1, \dots, n^{m-1} such that

- $\text{io}[1] = \{psk, 0salt, dh, 0ikm\}$
- $\text{io}[1] \subset \dots \subset \text{io}[m] = N$
- $\forall 1 \leq c < m : n^c \in \text{io}[c], \text{io}[c+1] = \text{io}[c] \cup \text{SblngN}(n^c)$

For $1 \leq c \leq m$ we denote by $N < c$ the set of names before c , $N \cap \text{io}[c]$ and by $c \leq N$ the set of names after c , $N \setminus \text{io}[c]$

Claim (Idealization Order). Every TLS-like key schedule has an idealization order.

Proof. Let $<$ be a total order on N and consider the algorithm in Fig. 19 which constructs the idealization computing min with respect to $<$. Firstly, the assignment in line 4 is well-defined: Let n be in $N \setminus \text{io}[c]$. Since n is reachable from psk (Definition 11), there is a path from psk to n . As there is a first element on this path not in $\text{io}[c]$, the set $\{n \in N \setminus \text{io}[c] : (n_1, n_2) \leftarrow \text{prntn}(n) \wedge \{n_1, n_2\} \subseteq \text{io}[c]_{\perp} \wedge \text{SblngN}(n) \not\subseteq \text{io}[c]\}$ is non-empty. As $\text{io}[c] \subsetneq \text{io}[c+1]$ and N is finite, the algorithm terminates and $\text{io}[m] = N$

B.2 Proof of Lemma 10

Using the existence of idealization order Claim B.1, the proof of Lemma 10 consists of the following lemmata whose proofs are a direct application of Lemma 1.

Lemma 12 (XTR1). For $n^c = psk$, we have that

$$\text{Adv}(\mathcal{A}, \text{Gcore}^{\text{hyb}:\ell, c}, \text{Gcore}^{\text{hyb}:\ell, c+1}) \leq \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{es, \ell}, \text{Gxtr}_{es, \ell}^b),$$

where $\mathcal{R}_{es, \ell}$ is defined in Fig. 18e.

```

IdealizationOrder( $ks, <$ )
hr-----
io[1] := {psk, 0salt, dh, 0ikm}
c ← 1
while io[c] ⊂neq N
  nc ← min{n ∈ N \ io[c] :
    (n1, n2) ← prntn(n) ∧
    {n1, n2} ⊆ io[c]⊥ ∧
    SblngN(n) ⊄ io[c]}
  io[c + 1] ← io[c] ∪ SblngN(nc)
  c ← c + 1
m ← c
return io, m, (n1, ..., nm-1)

```

Fig. 19: Idealization Order

Lemma 13 (XTR2). For $n^c = \text{esalt}$

$$\text{Adv}(\mathcal{A}, \text{Gcore}^{\text{hyb}:\ell,c}, \text{Gcore}^{\text{hyb}:\ell,c+1}) \leq \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{\text{hs},\ell}, \text{Gxtr}_{\text{hs},\ell}^b),$$

where $\mathcal{R}_{\text{hs},\ell}$ is defined analogously to reduction $\mathcal{R}_{\text{es},\ell}$ in Fig. 18e.

Lemma 14 (XTR3). For $n^c = \text{hsalt}$, we have that

$$\text{Adv}(\mathcal{A}, \text{Gcore}^{\text{hyb}:\ell,c}, \text{Gcore}^{\text{hyb}:\ell,c+1}) \leq \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{\text{as},\ell}, \text{Gxtr}_{\text{as},\ell}^b),$$

where $\mathcal{R}_{\text{as},\ell}$ is defined analogously to reduction $\mathcal{R}_{\text{es},\ell}$ in Fig. 18e.

Lemma 15 (XPD). For $\text{chldrnop}(n^c) = \text{xpd}$, we have that

$$\begin{aligned} & \text{Adv}(\mathcal{A}, \text{Gcore}^{\text{hyb}:\ell,c}, \text{Gcore}^{\text{hyb}:\ell,c+1}) \\ & \leq \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{n,\ell}, \text{Gxpd}_{n,\ell}^b), \end{aligned}$$

where Game $\text{Gxpd}_{n,\ell}^b$ for $b \in \{0,1\}$ is defined in Fig. 2 for $n \notin \{\text{psk}, \text{esalt}\}$, in Fig. 6b for $n = \text{esalt}$ and in Fig. 6a for $n = \text{psk}$. For $n \notin \{\text{esalt}, \text{esalt}\}$, where $\mathcal{R}_{n,\ell}$ is defined in Fig. 18f. The reductions for $n \in \{\text{psk}, \text{esalt}\}$ are defined analogously.

From Lemma 12, Lemma 13, Lemma 14, and Lemma 15 we obtain Lemma 10 by a 2-dimensional hybrid argument. First observe:

$$\text{Gcore}^{\text{SODH}} = \text{Gcore}^{\text{hyb}:0} \quad \text{Gcore}^{\text{hyb}:d} = \text{Gcore}^{\text{ki}} \quad (19)$$

$$\text{Gcore}^{\text{hyb}:\ell} = \text{Gcore}^{\text{hyb}:\ell,1} \quad \text{Gcore}^{\text{hyb}:\ell+1} = \text{Gcore}^{\text{hyb}:\ell,m} \quad (20)$$

From there, we obtain

$$\begin{aligned} & \text{Adv}(\mathcal{A}, \text{Gcore}^{\text{SODH}}, \text{Gcore}^{\text{ki}}) \\ & \stackrel{\text{Eq. 19}}{=} \text{Adv}(\mathcal{A}, \text{Gcore}^{\text{hyb}:0}, \text{Gcore}^{\text{hyb}:d}) \\ & \stackrel{\text{tele.sum.}}{\leq} \sum_{\ell=0}^d \text{Adv}(\mathcal{A}, \text{Gcore}^{\text{hyb}:\ell}, \text{Gcore}^{\text{hyb}:\ell+1}) \\ & \stackrel{\text{Eq. 20}}{=} \sum_{\ell=0}^d \text{Adv}(\mathcal{A}, \text{Gcore}^{\text{hyb}:\ell,1}, \text{Gcore}^{\text{hyb}:\ell,m}) \\ & \stackrel{\text{tele.sum.}}{\leq} \sum_{\ell=0}^d \sum_{c=1}^{m-1} \text{Adv}(\mathcal{A}, \text{Gcore}^{\text{hyb}:\ell,c}, \text{Gcore}^{\text{hyb}:\ell,c+1}) \\ & = \sum_{\ell=0}^d \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{\text{es},\ell}, \text{Gxtr}_{\text{es},\ell}^b) + \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{\text{hs},\ell}, \text{Gxtr}_{\text{hs},\ell}^b) \\ & \quad + \text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{\text{as},\ell}, \text{Gxtr}_{\text{as},\ell}^b) \\ & \quad + \sum_{n \in \text{XPR}} (\text{Adv}(\mathcal{A} \rightarrow \mathcal{R}_{n,\ell}, \text{Gxpd}_{n,\ell}^b)), \end{aligned}$$