

# A Unilateral-to-Mutual Authentication Compiler for Key Exchange (with Applications to Client Authentication in TLS 1.3) \*

Hugo Krawczyk<sup>†</sup>  
IBM Research

## ABSTRACT

We study the question of how to build “compilers” that transform a unilaterally authenticated (UA) key-exchange protocol into a mutually-authenticated (MA) one. We present a simple and efficient compiler and characterize the UA protocols that the compiler upgrades to the MA model, showing this to include a large and important class of UA protocols. The question, while natural, has not been studied widely. Our work is motivated in part by the ongoing work on the design of TLS 1.3, specifically the design of the client authentication mechanisms including the challenging case of post-handshake authentication. Our approach supports the analysis of these mechanisms in a general and modular way, in particular aided by the notion of “functional security” that we introduce as a generalization of key exchange models and which may be of independent interest.

## 1. INTRODUCTION

A natural question in the area of key-exchange (KE) protocols is how to transform a given KE protocol that is secure in the unilateral-authentication (UA) setting (where only one of the parties authenticates and the other remains anonymous) into a secure mutual-authentication (MA) protocol. Somewhat surprisingly this question has not received much attention in the complexity-based KE literature, in part since much of the work in this area has focused on the mutual authentication case. In practice, however, the most widely used KE protocol, TLS, is centered around unilateral (server-only) authentication, with optional upgrade to mutual authentication. Namely, one starts from a core protocol where only the server authenticates and then extends it with an optional mechanism for client authentication.

In this work we investigate the above question in the context of complexity-theoretic models of KE, striving for results that are general as well as practically relevant to real-

world protocols, in particular to the ongoing work on the design of TLS 1.3 [33] which provides particular motivation for this work. We think of this problem as developing *compilers* that extend secure unilateral key-exchange protocols into secure mutual key-exchange protocols. We focus on signature-based compilers, namely the case in which the second party to authenticate - which we refer to as the client - uses a digital signature as the means for public-key for authentication (even though our approach can potentially be extended to other forms of authentication).

**The SIGMAC Compiler.** Our compiler is simple: To upgrade a unilateral protocol  $\Pi_1$  into a mutually authenticated  $\Pi_2$ , upon completion of  $\Pi_1$  the client sends a single message comprised of: (i) the client’s signature on a portion of the  $\Pi_1$ ’s transcript; and (ii) a MAC value computed on the client and server’s identities with a MAC key computed by  $\Pi_1$ . The intuition of the design is simple too: First note that anyone can sign the transcript, hence just a signature by the client on the transcript is not sufficient. Just MACing the identities is clearly not sufficient either since any party that participated in the unilateral protocol can compute the MAC on any identities. Yet the combination of signature and MAC ensures a binding between an identity and the knowledge of a key (computed in the unilateral protocol). We call the compiler SIGMAC for SIGnature and MAC Compiler, and also since it is reminiscent of the SIGMA protocol [24].

While intuitively appealing, proving the compiler, namely, showing that it can upgrade a secure unilateral protocol into a secure mutual one turns out to be non-trivial, in particular regarding what information needs to be covered by the signature. One can show examples, even in natural and practical settings, where even signing the whole transcript is not enough to ensure client authentication. Fortunately, we show that the compiler works for important classes of protocols, including any protocol that derives its session key from a Diffie-Hellman exchange. A core issue is to characterize what needs to be included under the signature. For this, we find a general sufficient condition (through a notion we call *transport replication security*, abbreviated as “treplication security”) that if satisfied by a unilateral KE protocol  $\Pi$ , then applying the SIGMAC compiler to  $\Pi$  results in a provably secure mutually authenticated KE protocol. We then show extensive classes of unilateral protocols to possess this property, including KE mechanisms used in TLS 1.3.

By characterizing the *minimal* amount of information to be included under the client’s signature, we achieve more

\*Full version (includes proofs): [eprint.iacr.org/2016/711](http://eprint.iacr.org/2016/711)

<sup>†</sup>Email: [hugo@ee.technion.ac.il](mailto:hugo@ee.technion.ac.il).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

CCS '16, October 24–28, 2016, Vienna, Austria.

© 2016 ACM. ISBN 978-1-4503-4139-4/16/10...\$15.00

DOI: <http://dx.doi.org/10.1145/2976749.2978325>

general results that not only apply to a wider class of protocols but also allow to remove from the signature any information that the client may not want to sign with a non-repudiable signature. In particular, signing the server's identity leaves a transferable proof that the client communicated with the specific server. While full deniability of communication is hard to achieve against a malicious server [32], an explicit signature on the server's identity fails deniability even with a passive (or "honest-but-curious") server. We show that signing the server's identity can be safely omitted. On the other hand, we stress that if the protocol contains specific information that needs explicit authentication, e.g., the negotiation of security properties, these can and must be included under the signature even if these elements are not essential for the correct functioning of SIGMAC.

Finally, we comment on the use of a MAC function in the SIGMAC compiler. The essential observation (see Section 3.1) is that even if the client's signature covers both server's and client's identities, omitting the MAC - in particular, the MACed client identity - voids the ability of the compiler to produce a mutually authenticated protocol. As we note below, the use of a MAC function fits particularly well with TLS and its Finished message.

**TLS 1.3: Application and Motivation.** While the main results in this paper are of general applicability and intended to generically lift unilateral authentication protocols into mutual authentication ones, one timely motivation and application of this work is the analysis of TLS 1.3 [33]. The TLS key exchange (or *handshake*) protocol in the current and prior versions is built as a server-only authentication protocol with optional extension to client authentication. This approach calls exactly for a unilateral-to-mutual compiler as studied here. In particular, the specific SIGMAC compiler captures the type of mechanism used in TLS which is based on the combination of a signature and a MAC (the latter referred to as the client's Finished message in TLS).

A challenging aspect of the new TLS 1.3 design is that the protocol allows for *post-handshake* client authentication, namely, a setting where a session key, authenticated by the server only, is used to protect TLS data (record layer) and only *later* is the client authentication extension applied. This raises the question of what's the value of this late authentication; can the exchanged data be considered as mutually authenticated upon this late client authentication? Fortunately, our compiler approach turns out to be particularly useful for analyzing this setting. It allows us to show that upon client authentication and successful verification by the server, the data exchanged and protected up to this point (as well as subsequent data) enjoys mutual authentication security in the sense of "secure channels". That is, it is guaranteed that only the parties that passed authentication (the named server and client) are those generating the information and being able to decrypt it. Of course, the server obtains this assurance for past data only *a-posteriori*, i.e., only after verifying client authentication. Other questions our methodology helps addressing regard the effect of late authentication on keys derived prior to the authentication (e.g., a "resumption key") or the validity of using a session key to encrypt the very messages that carry client authentication (see Section 6).<sup>1</sup>

<sup>1</sup>An immediate conclusion from these results is that it is not necessary to introduce a dedicated key for encrypting post-

In all, our results are relevant to the analysis of client authentication for the PFS modes of TLS 1.3, namely, 1-RTT and PSK, and for PSK without PFS (hence, also for resumption mode). This includes the cases of post-handshake authentication and the use of the session key (or application key) for encrypting the client authentication message.

An important tool for analyzing the above complex questions is our "functional" generalization of the key exchange model from [9] which may be of independent interest. It provides an abstraction that frees the analysis from specific model and functionality details while at the same time adding generality. For example, one can use this abstraction to reason about whether a session key used to protect key exchange messages may still be secure for another purpose (say, to derive other keys) or whether a late client authentication implies authentication of keys derived prior to the authentication (e.g., a resumption key in TLS 1.3).

*Final note:* Familiarity with TLS 1.3 is not essential for understanding the results in this paper that, while applicable and motivated by the TLS 1.3 design, are more general and emphasize the conceptual aspects rather than the details of a specific protocol. Even in the context of the ongoing specification of TLS 1.3, generality is a benefit for informing basic aspects of the design (such as the complex issues related to late client authentication). At the same time, we stress that a full analysis of TLS 1.3 is well beyond the scope of this work.

**Related Work.** The literature on key-exchange protocols is extensive with many different models and many analyzed protocols. Yet, the number of papers dealing with unilateral authentication is relatively small. In particular, the basic earlier models, such as those of Bellare-Rogaway [1] and Canetti-Krawczyk (CK) [9], do not treat it. One exception is the work by Shoup [34] that deals with "anonymous protocols" explicitly. Halevi and Krawczyk [17] treat unilateral security in the context of password protocols. Our treatment is based on the CK model which we specialize to the unilateral case with minimal changes following the approach of [26, 25]. More recently, and motivated by the increased interest in analyzing TLS, several papers study unilateral authentication. Particularly, the works of Maurer et. al [29] and Goldberg et al. [16] center on this model but do not deal with the question of building generic transformations from unilateral to mutual authentication.

Several works have presented authentication compilers for *unauthenticated* key-exchange protocols including [2, 9] and [19]. Authenticators from [2, 9] can be applied to upgrade unilateral authentication (UA) to mutual authentication (MA) but since authenticators are applied separately to each UA protocol message they result in added (and possibly interleaved) interaction, hence necessitating of changes in the UA part of the protocol (rather than acting as an extension). The compilers from [19] are intended to lift unauthenticated protocols to mutually authenticated ones, but a closer look reveals that they do not achieve this goal except if identities are somehow included under their signatures and MACs. Even if fixed, applying such compilers to a UA protocol is "overkill" (in terms of communication and complexity, e.g., the compilers from [19] add 3-6 flows to the protocol) and

handshake messages, a design choice being contemplated by the TLS working group. Avoiding such dedicated key has the major advantage of dispensing with cumbersome mechanisms involving trial decryption and the like.

applying “half” compiler (in one direction only) does not necessarily work. For example, the SIGMA protocol from [24] shows that if *each party* signs the transcript (but not the identities) and also applies a MAC on its own identity then the protocol is secure. However, if we used the same technique for one party only to upgrade a UA protocol to MA we would end with an insecure protocol (see Section 3.1). The recent work of Kohlweiss et al. [23] analyzes a UA-to-MA transformation based on client passwords. We leave for future work the analysis of other compilers (e.g., one where both parties sign part of the transcript and only MAC their own identity as in [24] or a variant of SIGMAC where the client’s signature is replaced with KEM-based authentication).

Analyses of TLS that address unilateral authentication include [30, 26, 21, 5, 22, 6, 3] and work directly relevant to TLS 1.3 [28, 14, 12, 23, 25, 13, 11]. Of particular relevance to the present paper is the work of Cremers et al. [11] who present a detailed formal analysis of TLS 1.3 (draft 10) using the Tamarin prover. It includes a formulation of delayed client authentication whose analysis uncovered a surprising attack on an early TLS 1.3 version of this mechanism which we use as a motivating example in Section 3.1. Additional work related to post-authentication in different protocol settings is the design of “channel binding” protocols whose track record of repeated breaks illustrates the misleading intuition behind many of these techniques - see Bhargavan et al. [7, 4]. It will be interesting to apply the treplication-security notion as a formal basis for studying channel binding and related notions. Finally, we point out that our “functional” approach from Section 6.1 bears resemblance to Shoup’s model [34] and to the “suitable for” notion from [8].

## 2. THE UA AND MA KE MODELS

This section presents an abridged description of the Canetti-Krawczyk security model for key-exchange protocols [9, 10] which serves as the basis for the formal security treatment in this paper. Please consult [9] for complete details.

### 2.1 Mutual Authentication (MA) KE Model

We start by presenting the original CK model for mutual authenticated KE, and then specialize this model to the case of unilateral authentication. For the purpose of this paper we will refer to this model as MA (for mutual authentication) and differentiate it from the unilateral authentication model denoted UA.

A key-exchange (KE) protocol is run in a network of interconnected parties (all of which, including adversaries, are modeled as probabilistic polynomial-time machines) where each party can be activated to run an instance of the protocol called a *session*. Within a session a party can be activated to initiate the session or to respond to an incoming message. As a result of these activations, and according to the specification of the protocol, the party creates and maintains a *session state*, generates outgoing messages, and eventually *completes* the session by outputting a *session key*  $sk$  and erasing the session state. A session may not complete (e.g., if authentication fails) in which case it is *aborted* without generating a session key (output  $sk = \perp$ ). A KE session is associated with its *owner* (the party at which the session exists), a *peer* (the party with which the session key is intended to be established), and a *session identifier*, de-

noted *sid*. A KE protocol defines the contents of *sid*, typically including nonces and other transcript information. We name sessions by their owner and *sid*, and if the session peer is identified then also by the peer identity. For example,  $(S, sid)$  for a session at party  $S$  or  $(S, C, sid)$  if the peer to the session is identified as  $C$ . It is assumed (and enforced by the protocol) that any two sessions at the same (honest) party have different identifiers.

**Credentials.** We consider the public key setting where parties possess public keys of their choice and parties can verify the authentic binding between identities and public keys, be it either a certification authority (CA) or other means such as out-of-band distribution. We also consider the “pre-shared key setting” where pairs of parties share a secret symmetric key that they use for authentication.

**Attacker model.** The attacker, denoted  $\mathcal{A}$ , is an active “man-in-the-middle” adversary with full control of the communication links between parties.  $\mathcal{A}$  can intercept and modify messages sent over these links, it can delay or prevent their delivery, inject its own messages, interleave messages from different sessions, etc. (Formally, it is  $\mathcal{A}$  to whom parties hand their outgoing messages for delivery.)  $\mathcal{A}$  also schedules all session activations and session-message delivery. In addition, the attacker is allowed access to secret information via the following attacks:

- A **Reveal** query can be performed against an individual session after completion and the result is that the attacker learns the corresponding session key  $sk$ .
- A **Corrupt** against a party means that the attacker learns the long-term secret of that party; in addition, from the moment a party is corrupted all its actions may be controlled by the attacker. Non-corrupted parties are referred to as uncorrupted or *honest*.

For simplicity, in the current presentation we do not consider the CK *StateReveal* query.

**Basic security (mutual authentication).** To define security, we establish a notion of *matching* defined via session identifiers. If session  $(C, S, sid)$  completes at party  $C$ , then session  $(S, C, sid)$ , if it exists at  $S$  and completes, is said to be *matching* to  $(C, S, sid)$  (note the correspondence of peers and equality of *sid* and that only complete sessions have matching sessions).

Sessions against which any one of the attacks **Reveal** or **Corrupt** is performed (including sessions compromised via party corruption) are called *exposed*. A session is called *fresh* if it is complete and not exposed, and its matching session (if it exists) is also not exposed.

The security of session keys generated in fresh sessions is captured via the inability of the attacker  $\mathcal{A}$  to distinguish the session key of a *test session*, chosen by  $\mathcal{A}$  among all *fresh* sessions in the protocol, from a random value. This is captured via a **Test** query. This query may only be asked once during the security game. It sets  $K_0 := sk$  (or random if  $sk = \perp$ ) as the real session key, and sets  $K_1 \leftarrow_{\mathcal{R}} \{0, 1\}^\lambda$ . Then, it picks  $b \leftarrow_{\mathcal{R}} \{0, 1\}$  and returns  $K_b$ . The attacker can continue with the regular actions against the protocol also after the **Test** query; at the end of its run  $\mathcal{A}$  outputs a bit  $b'$ , which is meant as a guess for the value of  $b$ .

A key-exchange protocol  $\pi$  is *secure* if for polynomial-time attackers  $\mathcal{A}$  running against  $\pi$  it holds:

SIGMAC applied to UA protocol  $\Pi_1$  uses the following components:

- A key derivation function KDF applied to session keys output by  $\Pi_1$  for producing two keys  $K_a$  and  $K_s$  (this can be as simple as computing  $K_a = f_K(0)$  and  $K_s = f_K(1)$  where  $f$  is a PRF, or any other derivation of mutually pseudorandom values  $K_a, K_s$ ).
- A signature algorithm and a MAC algorithm (with the usual chosen message security requirements for both).
- A value  $\text{sid}^*$  defined for each session in  $\Pi_1$  and consisting of the concatenation of the session id  $\text{sid}$  and a subset of transcript information, called a transcript core, specifically defined for each protocol  $\Pi_1$  (see Section 4).
- Identities and keys for clients: Clients are anonymous in  $\Pi_1$ , but in  $\Pi_2$  they have identities and signature keys whose public keys other parties can obtain and validate. Servers have same identities and keys as in  $\Pi_1$ .

MA protocol  $\Pi_2$  is defined identically to  $\Pi_1$  except for the following extension:

- When a client  $C$  completes a session  $(C, S, \text{sid})$  in  $\Pi_1$  with peer  $S$  and outputs a session key  $K$ , it performs the following operations in  $\Pi_2$ . It derives keys  $K_a$  and  $K_s$  from  $K$  using KDF; sends a message, denoted CSM, consisting of a signature on the value  $\text{sid}^*$  and a MAC value,  $\text{MAC}_{K_a}(C, S)$ , on the identities  $C$  and  $S$ ; it erases  $K_a$  and completes session  $(C, S, \text{sid})$  in  $\Pi_2$  with session key  $K_s$ .
- When a server  $S$  would establish a session  $(S, \text{sid})$  with session key  $K$  according to  $\Pi_1$ , in  $\Pi_2$  it does not complete the session but computes  $K$  and keeps it in the session's state. When a CSM message (allegedly) from client  $C$  is delivered to a session  $(S, \text{sid})$ ,  $S$  checks that a key  $K$  was previously computed for this session and if so it derives keys  $K_a$  and  $K_s$  from  $K$  using KDF. It then checks that the incoming signature is valid (under the public key of  $C$ ) and that it covers the appropriate value  $\text{sid}^*$  (including the session id  $\text{sid}$ ). Finally, it verifies that the MAC computed with key  $K_a$  covers the client identity  $C$  and its own identity  $S$ . If all checks succeed,  $S$  completes session  $(S, C, \text{sid})$  with session key  $K_s$ .

Figure 1: The SIGMAC Compiler.

1. If two uncorrupted parties complete matching sessions in a run of protocol  $\pi$  under attacker  $\mathcal{A}$  then they output the same key (except for a negligible probability).
2. The probability that  $\mathcal{A}$  wins the **Test** experiment, namely, it outputs a correct guess  $b = b'$  is at most  $1/2$  plus a negligible fraction.<sup>2</sup>

We note that this model is general enough to capture protocols with implicit or explicit authentication and with and without key confirmation (or liveness). This is important since it makes our results more general, showing that our compiler applies to a UA protocol with any of the above characteristics.

A key-exchange protocol is said to achieve **perfect forward secrecy** (PFS) if it satisfies the above definition when relaxing the notion of fresh sessions to allow a **Corrupt** query against the owner of the session but only upon completion of the session.

## 2.2 Unilateral Authentication (UA) KE Model

The model as described above is intended to capture protocols where both peers authenticate to each other. Here we specialize this model to the case of unilateral authentication (UA), namely, when only one party authenticates to its peer (but the second party remains “anonymous”). The treatment is similar to [25].

The unilateral setting is best described by client-server terminology. We denote by  $C, S$  the client-server parties to the protocol as well as their identities, and refer to the authenticating party as the server. In the public key setting servers have public keys and in the case of a server and client sharing a symmetric key (the pre-shared key setting) we assume that the client associates the key to a server identity but not (necessarily) the other way. Sessions at a client are denoted as triples  $(C, S, \text{sid})$  as in the regular MA model but

<sup>2</sup>We use asymptotic, polynomial-time, terminology throughout the paper as it simplifies presentation. Moving to a concrete security setting is mostly straightforward.

sessions at servers do not have a named peer hence they are denoted  $(S, \text{sid})$ .

The essential characteristic of the UA model is that *only (fresh) client sessions are allowed as test sessions*.<sup>3</sup> Thus, matching is defined only for client sessions, namely, if session  $(C, S, \text{sid})$  exists and completes, and session  $(S, \text{sid})$  exists and completes then  $(S, \text{sid})$  is called a matching session to  $(C, S, \text{sid})$ .

## 3. THE SIGMAC COMPILER

We present our signature-based compiler that we call SIGMAC and which augments a UA-secure protocol  $\Pi_1$  into a MA-secure protocol  $\Pi_2$ . Roughly, protocol  $\Pi_2$  is obtained via the SIGMAC compiler by adding to  $\Pi_1$  a single message sent from client  $C$  to server  $S$  upon completion of a session in  $\Pi_1$ , where the message, denoted CSM (for “client sign-and-mac”), comprises a signature of  $C$  and a MAC (with a key derived from  $\Pi_1$ ’s session key). The signature is applied to a value  $\text{sid}^*$  defined for  $\Pi_1$  while the MAC is applied on the identities of  $C$  and  $S$ . The compiler is presented in Figure 1.

We stress that the value  $\text{sid}^*$  plays an essential role in the security of SIGMAC and its requirements are presented in detail in Section 4 via the notions of a transcript core and treplication security. We simplify some of our treatment by assuming that the session identifiers  $\text{sid}$  in protocol  $\Pi_1$  include unique random nonces contributed by  $C$  and  $S$ .

Another simplification in our treatment is that we assume that the signature keys used by clients for the SIGMAC signature are not used for other purposes or as part of protocol  $\Pi_1$ . This limitation is not essential as long as different uses are properly insulated via domain separation or similar methods.

<sup>3</sup>Restricting test sessions to client sessions only reflects the fact that there is no requirement on the security of session keys at a server as these may have the attacker as the unidentified peer. One could allow to also test server sessions that have a matching fresh client session, but this would be equivalent to our simpler definition.

More generally, we require that no two sessions at a party will have the same `sid` (except for negligible probability); also, we assume that `sid*` includes a fresh session-specific value from  $S$  (needed, in particular, against replay attacks). *Note on preserving the session key from  $\Pi_1$ .* The above specification of SIGMAC defines the keys  $K_a, K_s$  used by  $\Pi_2$  as derived from the session key  $K$  output by  $\Pi_1$ . Alternatively, one can specify that  $\Pi_1$  outputs  $K_a, K_s$ , where  $K_s$  serves as the session key in both  $\Pi_1$  and  $\Pi_2$  while  $K_a$  is only used in  $\Pi_2$  as a MAC key. We chose the former approach as it allows us to keep  $\Pi_1$  unmodified and is somewhat more convenient in the formal treatment. On the other hand, if one requires that  $\Pi_1$  and  $\Pi_2$  use the exact same session key, one needs to accommodate the generation of key  $K_a$  already as part of  $\Pi_1$  (this would be the case in TLS 1.3).

*Note on (not) signing the server's identity.* If one wants to provide some deniability against a passive adversary as discussed in the introduction, one should avoid entering the identity of  $S$  into the signature (i.e., into `sid*`).

**On (hashing) the identities included under SIGMAC's MAC.** The identities  $C$  and  $S$  included under the MAC computation in SIGMAC can have different forms, e.g., a subject name in a public key certificate, a whole certificate, an email address, etc. However, we note that an identity of a party  $X$  can be replaced, for the sake of including it under the SIGMAC MAC, by  $H(X)$  if  $H$  is a collision resistant function (this is similar to the use of hash functions for generating inputs to digital signatures). Moreover, hashing  $X$  with additional information is also acceptable as long as the parsing of the input to the hash function is defined as to uniquely determine the identity included in the hash. Such hashing of identities is used in many protocols including TLS. As a subtle example, resumption mode of TLS 1.3, replaces the server's identity included in the client's Finished message (that implements the SIGMAC's MAC functionality) with a unique value, `resumption_context`, computed as a hash of the server's identity and carried from a prior exchange between this client and server. For this to be secure one has to assume that the function used to compute `resumption_context` is collision resistant. In the case of TLS 1.3 this function is based on HKDF-Expand which can be shown to be collision resistant if the underlying hash is collision resistant. We stress that it is imperative in all uses of hashed identities that the output from the hash function not be truncated in a way that weakens the collision resistance property.

### 3.1 Rationale and examples

As explained in the introduction, the intuition behind SIGMAC is that the matching of the same client identity as the signer as well as under the MAC creates the needed binding between the client and the session key. Adding the server's identity to the MAC is needed to prevent UKS attacks (see below), namely, ensuring the server that the client's view of the peer is correct. Alternatively, one can include the server's identity under the signature which raises some privacy issues as discussed earlier.

The actual considerations in analyzing the compiler are more subtle particularly with respect to what needs to be covered by the signature. In some cases, even covering the whole transcript is insufficient to ensure mutual authentication via SIGMAC while in other cases covering part of the transcript is enough. We strive to identify the essential

parts of a transcript that need to be signed (this adds to the generality of the compiler and to our understanding of this mechanism, and can help achieving some level of deniability).

We illustrate some of the subtleties through some important examples serving as motivation for the notion of *transcript replication attacks* introduced in the next section and central to our analysis (and with direct implications to the security of client authentication in TLS 1.3).

**MACing the server identity  $S$ .** We show a UA protocol that when upgraded to MA via SIGMAC but without including the server's identity under the client's MAC, results in a MA protocol that is open to a UKS attack (the necessity of including  $C$  under the MAC is more obvious since anyone can sign the transcript hence such a signature has no binding to the UA run). The UA protocol is a simple server-authenticated DH exchange where the client sends a DH value  $g^x$  to which the server responds with  $g^y$  and a signature on  $g^x$  and  $g^y$ . The session key  $K$  is derived from  $g^{xy}$ . If we extend this to a MA protocol via SIGMAC but removing  $S$  from the MAC, namely,  $C$  sends  $MAC_{K_a}(C)$  and its signature on both  $g^x$  and  $g^y$  (whose concatenation defines `sid`), the resultant protocol is open to a UKS attack as follows. A malicious server  $S'$  relays  $g^y$  to  $C$  but signs it with its own signature. When  $C$  responds with  $MAC_{K_a}(C)$ ,  $S'$  sends this message to  $S$ . In this case  $K_s$  is computed by both  $C$  and  $S$ . But while  $S$  has  $C$  as the peer to the exchange,  $C$  has  $S'$  as the peer, hence resulting in a UKS attack.

**Pre-shared key protocols and why sometimes even signing the whole transcript is not enough.** We show an example of a UA protocol where applying the SIGMAC compiler with the *entire transcript* under the client's signature does not result in a MA-secure protocol. For this we consider a UA-secure protocol  $\Pi_1$  where the client uses as the server's credentials the server's identity and a symmetric pre-shared key (PSK). We assume that this key, that might have been previously shared by  $C$  with  $S$  through another UA key exchange protocol, has been authenticated only by  $S$  (as in the resumption mode of TLS 1.3). We define  $\Pi_1$  as an implicitly-authenticated protocol where the parties just exchange nonces (which form the whole transcript) and the session key is derived by applying a PRF keyed with PSK to the nonces. It is not hard to show that  $\Pi_1$  with `sid` set to the concatenation of the nonces is UA-secure. Now, consider the protocol  $\Pi_2$  resulting from applying SIGMAC to  $\Pi_1$  where we set the transcript core included in `sid*` to the entire transcript. Let  $S'$  be an attacker that has a shared key with  $C$  and one with  $S$  (the two keys may be random and independent) and interacts as a MitM between  $C$  and  $S$ , acting as the server with  $C$  and as a client with  $S$ .  $S'$  simply relays without modification the nonces chosen by  $C$  and  $S$ . In addition,  $S'$  forwards  $C$ 's signature on `sid*` to  $S$  but replaces the MAC generated by  $C$  with a MAC computed by  $S'$  on identities  $(C, S)$  using a MAC key derived from the PSK that  $S'$  shares with  $S$ . As a result  $S$  ends its session with peer  $C$  but with a session key that is known to  $S'$ ! Note that if we add explicit authentication to  $\Pi_1$  by sending a MAC from the server to client (using a key derived from the PSK) but this MAC is not included in `sid*` then the above attack works as well. Moreover, as we will see in Section 4.2, in some cases even including the server's MAC under the signature is still not enough to provide MA secu-

rity (in particular, this indicates that the source of problem was not in the implicit authentication nature of  $\Pi_1$ ).

Interestingly, this attack scenario was discovered by the Tamarin team [11] in their formal analysis of an early draft of TLS 1.3, showing that this is not just an hypothetical concern but something that can arise in practice.

**SIGMAC security for DH-based protocols (and some more subtleties).** An important and extensive class of protocols where SIGMAC can be applied to lift a protocol from UA to MA are protocols that provide forward secrecy by deriving the session key from a DH value  $g^{xy}$ , where  $g^x, g^y$  are exchanged between  $C$  and  $S$  and  $g^y$  is authenticated by  $S$ . We show in Section 4.1 that defining  $\text{sid}^*$  to include both values  $g^x, g^y$  is sufficient for SIGMAC to ensure MA security.

As yet another interesting subtlety regarding SIGMAC, we can show that without including  $g^y$  under the client's signature, MA-security is not guaranteed (this is somewhat counter-intuitive as one could expect that signing the client's message, with some freshness information from the server, would suffice). For this we consider a simple UA protocol, namely, a DH exchange between  $C$  and  $S$  where the server signs the whole transcript. If  $g^y$  is not included under the client's signature, an attacker  $S'$  can change  $g^y$  to  $g^{2y}$  leading  $C$  to compute its MAC using a key derived from  $g^{2xy}$  while  $S$  derives the verification key from  $g^{xy}$ . Using ideas similar to the example used in [26] to show the necessity of the PRF-ODH assumption in the analysis of TLS 1.2, one can build a KDF and MAC so that the MAC computed on (chosen) values  $C$  and  $S'$  with a key derived from  $g^{2xy}$  results in the same MAC tag as when computed on values  $C$  and  $S$  with key derived from  $g^{xy}$ . Also, the session key computed through the KDF is the same when derived from  $g^{xy}$  and  $g^{2xy}$ . This results in a UKS attack where both  $C$  and  $S$  compute the same session key but while  $S$  binds it to peer  $C$ ,  $C$  binds it to peer  $S'$  (which signed  $g^{2y}$ ). While clearly requiring contrived MAC and KDF functions, this argument shows that if one does not include  $g^y$  in  $\text{sid}^*$ , the SIGMAC compiler *cannot be proved* to work without assuming additional properties of the KDF and MAC functions.

## 4. TRANSCRIPT-REPLICATION (TREPLICATION) SECURITY

As shown in Section 3.1, the information that needs to be signed by the client in the SIGMAC compiler, namely  $\text{sid}^*$ , plays an essential role in the security of the compiler. The value  $\text{sid}^*$  is comprised of the session id  $\text{sid}$  and a subset of transcript elements as defined next via the notion of treplication security.

**DEFINITION 1 (TREPLICATION SECURITY).** *We say that a UA protocol  $\Pi$  is secure against transcript replication (treplication) attacks with respect to a transcript subset  $\tau$  (called a transcript core)<sup>4</sup> if any efficient UA-attacker  $S'$  against  $\Pi$  has only negligible probability to win the Test experiment on a fresh session  $(S, \text{sid})$  for which there is a session  $(C, S', \text{sid})$  with same session identifier  $\text{sid}$  and same transcript core  $\tau$  as  $(S, \text{sid})$ , and where  $C$  is a honest client.*

<sup>4</sup>That is, we view the protocol transcript as an ordered set of (possibly optional) elements and define the transcript core as a specified subset or projection.

Note that while  $S'$  is given the same capabilities of a UA attacker and the same indistinguishability Test, the conditions for the choice of the test session are very different than in the regular UA model (in particular, this is a server's session, not a client's one). Essentially, the requirement is that if  $S'$  acts as a man-in-the-middle between honest parties  $C$  and  $S$  but cannot change or choose the session id  $\text{sid}$  and transcript core  $\tau$  sent between  $C$  and  $S$ , then  $S'$  cannot learn information on the session key computed by  $S$ . This requirement is incomparable to UA security, namely, it is not implied by nor it implies UA security.

The importance of the treplication security notion is demonstrated by Theorem 3 where we show that joint UA and treplication security suffices for the successful use of the SIGMAC compiler. Here we prove some important classes of UA protocols to be treplication-secure (with respect to specific defined transcript cores), hence upgradeable to mutual authentication via SIGMAC. These protocols include those discussed in Section 3.1 and those with direct application to TLS 1.3 client authentication (particularly to the PFS modes of the protocol as well as the non-PFS PSK mode which also functions as session resumption).

### 4.1 Treplication security of DH Protocols

In Section 3.1 we showed some of the subtleties arising in proving treplication security of DH protocols. Here we show how such security can be achieved. We consider three cases: (i) plain DH where  $g^{xy}$  is only used to derive the session key; (ii) in addition to deriving the session key from  $g^{xy}$ , additional keys used during the key exchange protocol itself are derived from  $g^{xy}$  - however, we assume that the session key is secure (indistinguishable from uniform) even if these additional keys are provided to the attacker; (iii) the exponent  $x$  (and/or  $y$ ) is used for additional computation (not just for computing  $g^{xy}$ ). We note that case (i), while being the simplest and more natural, is generally insufficient for addressing real-world protocols. Case (ii) is directly relevant to TLS 1.3 in its 1-RTT mode and in the pre-shared key mode with PFS (the additional  $g^{xy}$ -derived keys are used for handshake traffic protection). Case (iii) would show up in a protocol like QUIC [27], or a similar extension to TLS 1.3, where two DH keys involving  $g^x$  are derived: an ephemeral  $g^{xy}$  as well as a key  $g^{xs}$  that combines  $g^x$  with a public key  $g^s$  of the server.

**LEMMA 2.** *Let  $\Pi$  be a UA-secure Diffie-Hellman protocol of the above types over a group  $G$ . Then, under the DDH assumption on  $G$ ,  $\Pi$  is treplication-secure if one defines the transcript core  $\tau$  to include the session identifier  $\text{sid}$  of  $\Pi$  as well as the two DH values  $g^x$  and  $g^y$ .*

### 4.2 Treplication security of PSK Protocols

Here we consider *pre-shared key (PSK) key-exchange protocols*, namely, protocols where the parties use as their credentials a symmetric key they previously shared by some method (e.g., out-of-band or a previous key exchange execution). We are interested in protocols that are UA-secure, hence we assume the shared key to have been authenticated only by the server (TLS 1.3 resumption presents such a case).

PSK protocols that offer forward secrecy (i.e., derive their session key through a DH exchange) are treplication-secure under the conditions demonstrated earlier. Thus, here we focus on UA PSK protocols without forward secrecy and investigate under which conditions they are treplication-secure.

We’ve seen in Section 3.1 that the basic implicitly-authenticated PSK protocol where the parties exchange nonces and derive a session key from their PSK and nonces (without further authentication) *cannot* be made treplication-secure even if the full transcript (the nonces in this case) is included in  $\text{sid}^*$ . Thus we consider an extension of this protocol that is UA-secure with explicit authentication where  $S$  sends a MAC computed on the nonces with a MAC key derived from the PSK (this MAC would correspond to the server’s Finished message in TLS). For the moment we restrict the transcript to only include the random nonces of  $C$  and  $S$ , denoted  $n_C, n_S$ , and the server’s MAC. We refer to this protocol as Basic-UA-PSK and define the whole transcript as the transcript core  $\tau$ .

Interestingly, this protocol *is not treplication-secure for generic MAC functions*. However, we can prove its security under a mild assumption on MAC security. Please see full version for the full treatment of treplication security of PSK protocols.

### 4.3 One-pass (0-RTT) Protocols

TLS 1.3 supports a so-called 0-RTT handshake consisting of a one-pass authenticated key-exchange protocol where authentication is carried through the use of a pre-shared key that client and server exchange in a previous regular handshake, typically with server-only authentication. In particular, this mode is used to instantiate session resumption in TLS 1.3. The 0-RTT handshake essentially consists of the client sending a key identifier (that allows the server to identify or compute the pre-shared key  $K$ ) and a nonce. The session key (used to immediately encrypt client’s data) is derived from a combination of the key  $K$  and the client’s nonce. Note that there is no server’s nonce assumed in the protocol which opens this mode to replay attacks. Here we discuss how client authentication could be added to such a mode, although TLS 1.3 does not currently support client authentication for 0-RTT.

We first note that one-pass protocols can be framed in the CK model through an adaptation as presented in [18], and the above nonce-based protocol can be proven to be UA-secure in this model. Thus, we can apply SIGMAC to this protocol MAC-ing the identities of  $C$  and  $S$  (in the case of TLS 1.3, the MAC would be implemented via the client’s Finished message that is part of the 0-RTT exchange) and signing the client’s nonce (which acts as  $\text{sid}$ ). Note, however, that SIGMAC assumes a server’s nonce which is not present here but this requirement can be relaxed if one accepts replay attacks in the model (as in [18]). More fundamentally, the above one-pass protocol is *not* treplication secure for reasons similar to (and more straightforward than) the interactive PSK case discussed in Section 3.1. Specifically, the attacker  $S'$  would establish separate pre-shared keys with both  $C$  and  $S$  (with  $S'$  acting as server and client, respectively) and later run a 0-RTT handshake with both parties, relaying the client-chosen nonce to  $S$ . Clearly,  $S'$  learns the key computed by  $S$  as it knows the pre-shared key used by  $S$ , hence breaking treplication security.<sup>5</sup> Unfortunately, and in contrast to the interactive PSK case treated earlier, this treplication insecurity cannot be fixed without including the server’s identity (or a value derived from it, such

<sup>5</sup>The attack is possible even if a nonce by  $S$  is involved and may even work if the key identifier is part of the transcript core as long as  $S'$  can choose the same key identifier as  $S$ .

as `resumption_context` discussed before) in the transcript core. Thus, client authentication for 0-RTT would need to include the server’s identity (and nonce) under the client’s signature. Fortunately, in the next section we show that when the server’s identity is signed by  $C$ , SIGMAC ensures MA security even if the base UA protocol is not treplication secure.

## 5. PROOF OF THE SIGMAC COMPILER

We prove the security of the SIGMAC compiler defined in Section 3 for treplication-secure protocols. We assume the MAC and signature functions used in SIGMAC to be secure in the standard chosen-message unforgeability sense.

**THEOREM 3.** *Let  $\Pi_1$  be a secure unilaterally authenticated KE protocol that is also treplication-secure with transcript core  $\tau$  and set  $\text{sid}^*$  to the concatenation of the session id  $\text{sid}$  and  $\tau$ . Then protocol  $\Pi_2$  resulting from the application of the SIGMAC compiler with this  $\text{sid}^*$  is a secure mutually authenticated KE protocol.*

Before proving the theorem we recall some of the adversarial actions in each model and how they differ in the case of  $\Pi_1$  and  $\Pi_2$ .

- Party activation for session initiation and message receiving/sending: This is same in  $\Pi_1$  and  $\Pi_2$  except for the CSM message sent by  $C$  upon receiving the last message from  $S$  (in case  $C$  sends the last message in  $\Pi_1$  then CSM is sent immediately after this last message).
- Party corruption: In both cases this includes revealing to the attacker all secret material of the corrupted party (and subsequent full control of that party by the attacker). Note that in the case of  $\Pi_2$ , client corruption includes revealing the signing key of  $C$ .
- Reveal queries: This can be applied to any completed session (other than the test session) at a honest party and returns the session key (note that the timing of completion for  $S$  may be different in  $\Pi_1$  and  $\Pi_2$ ).
- Session tests: In  $\Pi_1$  session tests are only allowed against  $(C, S, \text{sid})$  while in  $\Pi_2$  they are also allowed against  $(S, C, \text{sid})$ .

### 5.1 The simulator

For the proof of Theorem 3 we specify a simulator that given an MA-attacker  $\mathcal{A}_2$  against protocol  $\Pi_2$  it defines a UA-attacker  $\mathcal{A}_1$  against  $\Pi_1$  such that if  $\Pi_2$  wins a test session in  $\Pi_2$  so does  $\mathcal{A}_1$  in  $\Pi_1$ . The simulator SIM acts as the challenger (i.e., the orchestrator of the protocol run) for a given MA-attacker  $\mathcal{A}_2$  against protocol  $\Pi_2$  and uses  $\mathcal{A}_2$  actions to implement the UA-attacker  $\mathcal{A}_1$  against  $\Pi_1$ . We think of the parties running  $\Pi_1$  as real parties, with their own secret keys, and those running  $\Pi_2$  as simulated by SIM with the same keys as in  $\Pi_1$  except for clients’ signature keys (that exist in  $\Pi_2$  but not in  $\Pi_1$ ). The latter are chosen by SIM for honest clients in  $\Pi_2$  while public keys for corrupted parties are chosen by the attacker.

SIM implements the CSM message in  $\Pi_2$  as follows. Upon completion of a session  $(C, S, \text{sid})$  by a client  $C$  in  $\Pi_1$ , SIM invokes  $\mathcal{A}_1$  to run a reveal query on that session to obtain

the corresponding session key  $K$ . From it, SIM derives keys  $K_a$  and  $K_s$ , and uses  $K_a$  and the signature key of  $C$  to compute the CSM message that it hands to  $\mathcal{A}_2$ . Completion of sessions by honest parties in  $\Pi_2$  is decided by SIM. In the case of a client, a session is completed if and only if the corresponding session completes in  $\Pi_1$ . Completion of a session  $(S, C, \text{sid})$  at a honest server in  $\Pi_2$  depends on receiving a valid CSM message, namely with a valid signature on  $\text{sid}^*$  and the MAC under  $K_a$ . If this message is not originated with a honest client (e.g., it is sent by a corrupted client) then SIM verifies it by invoking  $\mathcal{A}_1$  to reveal the corresponding completed server's session  $(S, \text{sid})$  in  $\Pi_1$  and deriving  $K_a$  from it. The validity of this action follows from Lemma 4 (c).

The actions of  $\mathcal{A}_1$  as decided by SIM mostly mimic those of  $\mathcal{A}_2$  with some important differences in the handling of the corrupt, reveal and test queries that we specify next. After each action we argue informally the validity of the action (showing that the attackers' actions and views are consistent with the protocol and the attack model). The formal proof of validity follows from Lemma 4 below. Throughout, we say that keys  $K_a, K_s$  in a session in  $\Pi_2$  are **consistent** with a session key  $K$  in  $\Pi_1$  if  $K_a, K_s$  are derived from  $K$  using the KDF function.

1. When  $\mathcal{A}_2$  corrupts a party, SIM invokes  $\mathcal{A}_1$  to do the same.  $\mathcal{A}_2$  receives the secret information from this party as received by  $\mathcal{A}_1$  and, in the case of clients, it receives from SIM the party's signature key.
2. When  $\mathcal{A}_2$  reveals a session  $(C, S, \text{sid})$  at honest client  $C$  in  $\Pi_2$ , SIM invokes  $\mathcal{A}_1$  to reveal  $(C, S, \text{sid})$  in  $\Pi_1$ ;  $\mathcal{A}_1$  is handed the session key  $K$  from  $\Pi_1$  while  $\mathcal{A}_2$  is handed the derived session key  $K_s$  from  $\Pi_2$ .

- Validity (Lemma 4 (b)): We use the fact that these two sessions exist and complete together and have consistent keys.

3. When  $\mathcal{A}_2$  reveals a session  $(S, C, \text{sid})$  at honest server  $S$  in  $\Pi_2$ , SIM invokes  $\mathcal{A}_1$  to reveal session  $(S, \text{sid})$  in  $\Pi_1$ ;  $\mathcal{A}_1$  is handed session key  $K$  while  $\mathcal{A}_2$  is handed the derived key  $K_s$ .

- Validity (Lemma 4 (c)): We use the fact that if  $S$  is honest and session  $(S, C, \text{sid})$  completes in  $\Pi_2$  then session  $(S, \text{sid})$  exists and completed in  $\Pi_1$ , and the two sessions have consistent keys.

4. The processing of the **Test** query and test sessions is handled by SIM as follows. At the onset of the protocol run, SIM selects a random session, called the *guess session*, from all sessions to be created during the run of protocol  $\Pi_2$  under  $\mathcal{A}_2$  (this assumes an a-priori bound on the number of sessions to be created by  $\mathcal{A}_2$ ). If at any point,  $\mathcal{A}_2$  issues an action that disqualifies the guess session as a test session, SIM instructs  $\mathcal{A}_1$  to output a random bit  $b$  and aborts. The guess session may be at a client or server:

- (a) If the session is at a client, call it  $(C, S, \text{sid})$ , then as soon as  $\mathcal{A}_2$  delivers the last incoming message to this session (from server  $S$ ), SIM invokes  $\mathcal{A}_1$  to deliver this last message (which, if the guess session is correct, results in  $C$  completing) and to issue a **Test** query on  $(C, S, \text{sid})$  in  $\Pi_1$ .

- (b) If the session is at a server, call it  $(S, C, \text{sid})$ , then as soon as  $\mathcal{A}_2$  delivers the last message from  $S$  to session  $(C, S, \text{sid})$ , SIM invokes  $\mathcal{A}_1$  to deliver this last message and to issue a **Test** query on  $(C, S, \text{sid})$  in  $\Pi_1$ .

In both cases SIM learns from  $\mathcal{A}_1$  the value of the real-or-random key  $K$  returned by the **Test** query in  $\Pi_1$ . SIM derives from  $K$  keys  $K_a$  and  $K_s$  and uses  $K_a$  to compute the outgoing CSM message from  $C$  in  $\Pi_2$  and for checking the validity of a CSM message delivered by  $\mathcal{A}_2$  to session  $(S, C, \text{sid})$ . When  $\mathcal{A}_2$  issues a **Test** query against the guess session (at  $C$  or  $S$ ), SIM provides  $K_s$  as the answer to that query.

- Validity: We will show that the test sessions invoked by  $\mathcal{A}_1$  are valid test sessions under  $\mathcal{A}_1$ 's run and that  $K_s$  is real if  $K$  is real and (pseudo) random otherwise. See Lemma 4 (f).
5. When  $\mathcal{A}_2$  stops its run with an output bit  $b$ , SIM instructs  $\mathcal{A}_1$  to stop with the same bit.
- We need to show that  $\mathcal{A}_1$  wins whenever  $\mathcal{A}_2$  wins, hence proving that the advantage of  $\mathcal{A}_2$  against  $\Pi_2$  is bounded (up to a polynomial factor induced by the probability of a correct test session guess) by the advantage of  $\mathcal{A}_1$  against  $\Pi_1$ .

## 5.2 Validity of SIM's actions

The proof of Theorem 3 uses the following properties of the simulator's actions and those induced on attacker  $\mathcal{A}_1$ , in particular showing the validity of these actions in the corresponding models.

LEMMA 4. *The following properties hold based on the simulation actions.*

- (a)  $\Pi_1$  and  $\Pi_2$  as run under  $\mathcal{A}_1$  and  $\mathcal{A}_2$ , respectively, have the same set of parties, including the same honest and same corrupted parties.
- (b) If  $C$  is honest, then when a session  $(C, S, \text{sid})$  completes in  $\Pi_2$ , a session  $(C, S, \text{sid})$  exists and has completed in  $\Pi_1$ ; in particular, the two sessions have consistent session keys.
- (c) If  $S$  is honest, if a session  $(S, C, \text{sid})$  completes in  $\Pi_2$  then session  $(S, \text{sid})$  exists and has completed in  $\Pi_1$ ; in particular, the two sessions have consistent session keys.
- (d) If  $(C, S, \text{sid})$  is a valid test session chosen by  $\mathcal{A}_2$  in  $\Pi_2$  then, except for negligible probability,  $(C, S, \text{sid})$  is fresh in  $\Pi_1$  and (by part (b)) the two sessions have consistent session keys.
- (e) Assuming  $\Pi_1$  is treplication-secure, if session  $(S, C, \text{sid})$  is a valid test session chosen by  $\mathcal{A}_2$  in  $\Pi_2$ , then, except for negligible probability, session  $(C, S, \text{sid})$  exists, completes and is fresh in  $\Pi_1$  and both sessions have consistent session keys.
- (f) If the test session chosen by  $\mathcal{A}_2$  in  $\Pi_2$  is a valid test session and SIM guessed this session correctly then, except for negligible probability, the test session chosen



by SIM is a valid test session under  $\mathcal{A}_1$ 's run, and if the key  $K$  returned to  $\mathcal{A}_1$  is real (resp. random) then  $K_s$  is real (resp. pseudo-random).

The next two lemmas, needed to complete the proof of Lemma 4 (parts (d) and (e)) and Theorem 3, show, respectively, the need to include the identities  $C$  and  $S$  under the CSM MAC.

LEMMA 5. *If  $C$  and  $S$  are honest, session  $(C, S, \text{sid})$  is complete and unrevealed in  $\Pi_1$ , and session  $(S, \text{sid})$  is complete in  $\Pi_1$ , then in  $\Pi_2$  either  $S$  does not complete session  $\text{sid}$  or, if it does, it completes it as  $(S, C, \text{sid})$  (except for a negligible probability).*

LEMMA 6. *If  $\Pi_1$  is treplication-secure and  $(S, C, \text{sid})$  is fresh in  $\Pi_2$  then, except for negligible probability,  $(C, S, \text{sid})$  exists and completes in  $\Pi_2$ .*

### 5.3 Proof of Theorem 3

Consider simulation executions where SIM does not abort (i.e., SIM guesses correctly the test session) and in which the negligible probability events from the proofs of Lemmas 5 and 6 do not happen. In this case, Lemma 4 shows the validity of the actions of both adversaries,  $\mathcal{A}_1$  and  $\mathcal{A}_2$ , as set by the simulator actions, and their consistency with the UA and MA models, respectively. In particular, the views of  $\mathcal{A}_1$  and  $\mathcal{A}_2$  under the simulated games are exactly as those of a real run of protocols  $\Pi_1$  and  $\Pi_2$ , respectively. Thus, the winning probability of  $\mathcal{A}_2$  is the same as in a real run of protocol  $\Pi_2$  and by Lemma 4 (part f)  $\mathcal{A}_1$  wins its test session whenever  $\mathcal{A}_2$  does. We conclude that the winning advantage of  $\mathcal{A}_2$  is upper bounded by that of  $\mathcal{A}_1$ . In all, we have that if  $m$  is an upper bound on the number of sessions initiated by  $\mathcal{A}_2$ , then the winning advantage of  $\mathcal{A}_2$  against  $\Pi_2$  is at most  $m$  times the winning advantage of  $\mathcal{A}_1$  against  $\Pi_1$  plus a negligible fraction. Hence, if any UA-attacker  $\mathcal{A}_1$  against  $\Pi_1$  has only negligible advantage, then so is the case for any MA-attacker  $\mathcal{A}_2$  against  $\Pi_2$ . This completes the proof of MA-security of the compiled protocol  $\Pi_2$ .

**Note on security quantification.** The proofs of Lemmas 5 and 6 show events where the selection of a test session for  $\mathcal{A}_1$  fails due to some adversarial action. The lemmas show these events to have negligible probability. An examination of the proof shows this probability to be at most  $2 \cdot (n \cdot \epsilon_{\text{sig}} + \epsilon_{\text{mac}} + \epsilon_{\text{kdf}}) + \epsilon_{\text{UA}} + \epsilon_{\text{tr}} + \epsilon_{\text{nc}}$  where the  $\epsilon$  values denote the security of various elements in the protocol: the unforgeability of the signature and MAC schemes used in SIGMAC, the pseudorandomness of the KDF, the UA-security of protocol  $\Pi_1$ , and the treplication security of  $\Pi_1$  (these bounds represent the attackers' advantage for a given time bound);  $\epsilon_{\text{nc}}$  bounds the probability of nonce collision among honest clients and  $n$  bounds the number of clients (the latter is needed in the client's signature-forging reduction for which one needs to guess the client in the test session in order to embed the attacked signature key). The above expression is further multiplied by factor  $m$  (denoting an upper bound on the number of sessions established by adversary  $\mathcal{A}_2$  running against  $\Pi_2$ ) to account for the probability  $1/m$  that SIM does not abort its run, namely, it guesses correctly the session that  $\mathcal{A}_2$  will choose as its test session.<sup>6</sup>

<sup>6</sup>This may result in a  $m \cdot n$  factor multiplying  $\epsilon_{\text{sig}}$  but in this case one can set  $m$  to be the maximal number of sessions

**Streamlining the reduction.** The above bound contains the additive term  $m \cdot \epsilon_{\text{UA}}$  where  $\epsilon_{\text{UA}}$  bounds the advantage of the UA-attacker against protocol  $\Pi_1$ . It is likely that the bound  $\epsilon_{\text{UA}}$  resulting from the analysis of a particular protocol  $\Pi_1$  already contains a factor of  $m$  equal to the number of sessions, in which case  $m \cdot \epsilon_{\text{UA}}$  contains a factor  $m^2$ . The following analysis strategy avoids this unnecessary squared degradation. We consider a “semi-adaptive” adversarial model, for both the UA and MA settings, where the attacker has to *declare* its choice of test session as soon as the session is initiated. In this setting, SIM does not need to guess  $\mathcal{A}_2$ 's test session but it rather learns the session when  $\mathcal{A}_2$  initiates (and declares) it. At this point, SIM instructs the semi-adaptive  $\mathcal{A}_1$  to declare the choice of the corresponding test session in  $\Pi_1$ . Everything else stays the same as in the current simulation and proof except that here the reduction is between semi-adaptive UA and MA attackers without incurring the loss of the  $m$  factor. This factor *only comes once* through a generic reduction from a regular (adaptive) attacker to a semi-adaptive one. In this generic reduction the regular attacker acts semi-adaptive by choosing the test session at random among all  $m$  sessions and declaring it upon initiation - hence incurring a (one-time)  $m$  factor advantage degradation.

**Note on signing  $S$  and  $C$ .** Inspection of the proof of Lemma 6 shows that if the server's identity  $S$  is included under  $\text{sid}^*$ , i.e., signed by the client, then there is no need to include that identity under the MAC or to resort to treplication security. Yet, our results show that if signing this identity is omitted (for privacy reasons or simply because the identity is not part of the signed transcript - as in the resumption mode of TLS 1.3) its inclusion under the MAC is sufficient. As for the client's identity, we note that including this identity under the client's signature but not under a MAC would fail to provide MA security.

## 6. POST-HANDSHAKE AND ENCRYPTED AUTHENTICATION

In this section we adapt the analysis of the SIGMAC compiler to some specific settings arising in the ongoing specification of TLS 1.3 [33]. We also introduce the functional CK model which may be of independent interest.

The SIGMAC compiler as analyzed in previous sections can be readily applied in the context of TLS 1.3 in the case that the client sends its authentication message (the signature and Finished message, or CSM in our terminology<sup>7</sup>) in the third handshake flight, i.e., immediately upon verification of server authentication. However, TLS 1.3 also includes some cases where the SIGMAC analysis from Section 5 does not apply directly. We identify three such cases.

The first considers the fact that TLS 1.3 always encrypts the CSM message for reasons of identity protection. In the regular handshake case, the encryption key, or “handshake transport key (HTK)”, is derived from the same intermediate key from which  $K_a$  and  $K_s$  are derived, and by the properties of the key derivation function HTK is (computationally) independent from  $K_a$  and  $K_s$ . Extending the

established by a single client rather than representing the total number of sessions.

<sup>7</sup>We take the liberty of referring to the signature and MAC combination used in TLS 1.3 as “the CSM message”.

SIGMAC proof for this case is immediate: the simulator derives the HTK key from the UA session key  $K$ , as it does with  $K_a$  and  $K_s$ , and uses it to encrypt/decrypt the CSM message. Due to the independence of HTK from  $K_a$  and  $K_s$ , the proof holds with minor adjustments.

In the second case the CSM message is also encrypted but in this case the encryption uses the *same key*  $K_s$  that the protocol outputs as the session key. This immediately violates key indistinguishability since by the time the client establishes the session key (which is when the key can be tested for indistinguishability), this key has been used already. We show (Section 6.2) how to adapt the analysis of SIGMAC to this case, at the cost of reducing the security guarantee. Rather than ensuring generic security of  $K_s$  (as induced by the regular indistinguishability property of key-exchange protocols), we show that  $K_s$  can be used as a key to an authenticated encryption scheme for implementing *secure channels*. Our analysis is general and can be applied to different modelings of secure channels (e.g. [20, 15]) as well as to other applications of the session key<sup>8</sup> as long as the security of such application is not voided by the use of  $K_s$  for encrypting CSM. Our formalization introduces the notion of *functional tests* (Section 6.1) that abstracts out details of applications and implementations.

The third case is the so called “*post-handshake client authentication*”. Here, the client authenticates with a CSM message but only after the server and client have already exchanged application data (i.e., record layer communication in the TLS context) protected with the session key  $K_s$ . The encrypted data exchange starts after the parties compute the session key but before the client authenticates. This raises questions about the notion of security that can be claimed and the level of protection provided to the exchanged data (to which we will refer as *pre-authentication data*). One can see that since the handshake protocol without client authentication provides unilateral authentication, then the data gets the assurance of unilateral authentication (server authentication in the case of TLS). However, the SIGMAC analysis from previous sections showing that the CSM message upgrades the protocol to MA security does not hold anymore once the session key  $K_s$  is used to protect data before CSM is delivered and verified. Yet, we are able to show (Section 6.3) that the CSM message does provide mutual authentication of the exchanged data but only in the sense of secure channels, and only after the CSM message is verified by the server. Namely, the parties get the guarantee (delayed for the server) that only the named party that passed authentication can authenticate and decrypt the exchanged information (including data sent before client authentication). Here too, our treatment via functional tests adds generality and allows us to abstract the details of the secure channels modeling and implementation.

Finally, combining the last two cases one obtains secure channels security for the case (also in TLS 1.3) where in addition to the exchange of pre-authentication data, the post-handshake CSM message itself is encrypted with the session key (Section 6.3.1).

<sup>8</sup>In the case of TLS 1.3, such applications may include the derivation/update of keys or the issuing of resumption keys via a New Session Ticket Message. Here the session key can be the application key or other keys computed in the protocol such as the exporter master secret.

Fortunately, we can address all these scenarios with simple adaptations of the analysis of the SIGMAC compiler. We do so in the following subsections.

## 6.1 Functional Queries and Functional Tests

Since our goal is to show that the above authentication variants are sufficient to ensure “secure channels” functionality, it is worth recalling how such functionality is modeled. Take as example the ACCE model of Jager et al. [20] that has been used to prove security of several variants of TLS 1.2. Such model extends the usual key exchange security formalism with the ability of the attacker to run Encrypt and Decrypt queries on messages of its choice where the session key (or keys derived from it) serves as the authenticated encryption key. In addition, to capture the inability of the attacker to subvert the desired secure channels functionality, the attacker is tested through a dedicated game that replaces the standard indistinguishability test of key exchange models. In the ACCE model this test is adapted from the treatment of stateful authenticated encryption by Paterson et al. [31]. Other secure channels models differ in their details and scope but they follow a similar approach.

To increase generality and simplicity, our treatment abstracts out the details of implementation of secure channels, framing the above model ingredients in an abstract way. That is, we extend the capabilities of a key exchange attacker with the ability to query abstract “functional queries” on session keys and we replace the usual indistinguishability test with an abstract “functional test” that is run on a session chosen by the attacker (the “test session”) and where the input to the test is the test session key. The only condition on such a test is that if one replaces the input session key with a random independent key, then the advantage of the attacker in winning the test is negligible. The ACCE model is an instantiation of this abstract framework where the functional queries represent the Encrypt/Decrypt queries and the session key experiment is modeled as a stateful authenticated encryption test (note that this test has the required property from a functional test, namely, if the authenticated encryption scheme is keyed with a random value then the attacker has negligible advantage in winning the ACCE game). This abstraction frees our analysis from dealing with specific details and models while at the same time adding generality. For example, one can use this abstraction to reason about whether a session key used to protect key exchange messages may still be secure for another purpose (say, to derive other keys) or whether a late client authentication implies authentication of keys derived prior to the authentication (e.g., a resumption key in TLS 1.3).

**Functional family  $F$ .** Let  $F$  be a parametrized family of (possibly randomized) functions that accept a single input to which we refer as a *session key*. Given a parameter  $p$  and a session key  $k$ ,  $f_p$  denotes a member of  $F$  and  $f_p(k)$  denotes an output distribution (or, more commonly, it denotes a specific value sampled from that distribution).

**Functional queries and tests.** We present a modification of the CK model of key exchange security [9] (which we recalled in Section 2). The only changes to the model are as follows (we refer by  $\mathcal{A}_F$  to the attacker in this setting):

- We add a new adversarial query type, called a *functional query*  $FQuery$ , associated to a functional family

$F$ .  $\mathcal{A}_F$  can issue **FQuery** against any session eligible for a **Reveal** query and also against the test session. In it,  $\mathcal{A}_F$  provides a session identifier, a parameter  $p$  and gets back  $f_p(k)$  where  $f_p$  is an element of  $F$  and  $k$  is the session key of the named session.

- We change the test session experiment: A test session can be chosen under the same rules as in the CK model but the regular indistinguishability-based **Test** query is replaced with a *functional test* **FTest**. This can be an arbitrary test taking the form of an interactive game between a challenger and attacker  $\mathcal{A}_F$ , where the challenger's input is the test session key and a random bit  $b$ , and the output of  $\mathcal{A}_F$  is a bit  $b'$ . We say that  $\mathcal{A}_F$  wins if  $b = b'$ . The only condition on a functional test is that if one replaces the input to the challenger with a random independent key then the advantage of  $\mathcal{A}_F$  in winning the game is negligible.

We allow the functional queries to use a shared state (formally accommodated through the parameter  $p$  that can be seen as a generalized second input to the function) and also allow the functional test to share such state. This adds generality to our treatment and is needed, in particular, for the formalization of stateful authenticated encryption as used in the definition of ACCE security.

**Functional CK model.** The security definition in this functional model remains the same as in the original CK model (end of Section 2.1) except that the notion of winning for the attacker changes from the indistinguishability **Test** experiment to a given functional test **FTest**. We will refer to this notion as Functional CK security and we will denote the MA and UA variants by F-MA and F-UA (as in the regular CK model of Section 2, the main difference between UA and MA is that only clients have a defined peer and only client sessions can be chosen as test sessions).

*Remark:* For sessions other than the test session, the functional queries seem redundant as the attacker can always query these sessions via a **SK reveal** query and compute the functional query by itself. However, note that information learned via functional queries on a session can help the attacker choose that session as the test session (running a reveal query on a session would disqualify that session for testing). Moreover, quantifying security via the number of functional queries in an attack (e.g., number of encrypt and decrypt queries) can be an important measure. On the other hand, one may simplify the model by assuming that the attacker chooses its test session at random (at the cost of a number-of-sessions factor loss in the attacker's advantage) in which case functional queries are not needed, except for the test session and as part of defining a functional test.

*Remark:* We note that the above formalism via a bit-guessing game can be generalized. First, one can replace such a game with any experiment (not necessarily bit-guessing) where the winning probability of  $\mathcal{A}_F$  on a run over a random independent key is negligible. Second, it would be enough to require that the difference between the success of the attacker on a run over the real session key be negligibly higher than on a run over a random key (even if the latter case is not in itself negligible). Finally, note that the functional approach can be applied to other key exchange models such as the Bellare-Rogaway [1]; a notion similar to functional

queries was considered by Shoup [34]. One may also be able to use the "suitable for" notion from Brzuska et al. [8] as a replacement for our functional tests. The two notions differ in their approaches but have a motivation in common, namely, replacing pure key indistinguishability with more limited functionalities where the use of the session key during the KE process does not void the security of the task at hand.

**Secure channels via functional security.** As discussed above, we are interested in freeing our results from the specific details of a secure channels model or its implementation. The idea is that any such formalism can be framed via the above functional model. Specifically, when referring to secure channels we will define two abstract functional queries, **Encrypt** and **Decrypt**, and will assume a given functional test. **Encrypt** and **Decrypt** will have the regular functionality of encrypting and decrypting messages but their implementation can take different forms (e.g., they can be defined via general authenticated encryption or with a more specific scheme, they can be stateless or stateful, apply to full messages or fragments, etc.). The functional test can also take different forms, e.g. [20, 15]. For example, [20] defines its model, ACCE security, through a semantic security game that mimics the game used to define the security of the underlying (stateful) authenticated encryption scheme [31] and relies on the fact that no attacker can win the latter game if the key in use is fresh and random. Our formalism can represent such a model by instantiating the functional test via this particular ACCE game.

## 6.2 SIGMAC with encrypted CSM

We consider a variant of the SIGMAC compiler in which the CSM message is encrypted with the session key  $K_s$  from the same session to which the CSM message belongs. We show that any application whose security can be defined via a functional test as in Section 6.1, and where the encryption and decryption of CSM are included in the functional queries of the model, enjoys mutual authentication in the sense of F-MA. As an application, instantiating our abstract formalism with a secure channels model (e.g., [20, 15]), we get that the session key validated via SIGMAC with encrypted CSM implements secure channels with mutual authentication (it is assumed that the secure channels mechanism enforces domain separation between the encryption of CSM and encryption of application data, for example as the ACCE model does via stateful encryption).

Let  $\Pi_2'$  be a protocol resulting from applying SIGMAC with encrypted CSM to a UA-secure  $\Pi_1$  protocol. (It is possible to include further information under this encryption such as the client's identity or certificate to protect it from eavesdroppers.) The encryption of CSM is generated by the client and the corresponding decryption and verification is performed by the receiving server. We prove this modified compiler following the original proof of Theorem 3 but we replace the regular indistinguishability-based security test with a functional test as defined in Section 6.1; we also assume **Encrypt** and **Decrypt** operations as part of the allowed functional queries<sup>9</sup>. We prove:

<sup>9</sup>Note that the attacker does not provide the message CSM for encryption, but rather the functional query is defined as the application of **Encrypt** to the value CSM generated by the client and possibly unknown to the attacker.

LEMMA 7. *The SIGMAC compiler with message CSM encrypted under session key  $K_s$  and applied to a treplication-secure UA-secure protocol (under the conditions of Theorem 3) results in a F-MA-secure protocol where the session indistinguishability test is replaced with a functional test (defined by the protocol) and the Encrypt and Decrypt operations applied to CSM are considered as valid functional queries on session key  $K_s$ .*

**Application to ACCE security.** As previously stated our functional security model allows us to frame secure channel formalisms. One specific instance is the ACCE model from [20] that has been successfully applied to the analysis of TLS protocols (but see the refined definitions of [15]). We omit the ACCE details here and informally claim the following corollary. Please consult [20] for the model (including a definition of pre-accept and post-accept phases), the definition of the Encrypt and Decrypt operations (that we represent as functional queries), and their security game that instantiates our functional test.

COROLLARY 8 (INFORMAL). *Let  $\Pi$  be a protocol resulting from the application of the SIGMAC compiler with encrypted CSM to a treplication-secure UA-secure protocol, and let Encrypt and Decrypt denote a stateful encryption scheme as considered by the ACCE model [20]. Consider protocol  $\Pi'$  that runs  $\Pi$  in the pre-accept phase of ACCE and uses the resultant session key as the key for the Encrypt and Decrypt operations in the post-accept phase as well as for the encryption of CSM<sup>10</sup>. Then,  $\Pi'$  is secure according to the (mutual authentication) ACCE model.*

### 6.3 Security under post-handshake authentication

Here we address the third case discussed in this section's introduction, namely, where data protected under the session key is exchanged between server and client after the server computed the session key but before the client sends the CSM message. This case, known as “*post-handshake client authentication*” in the context of TLS 1.3, requires a weakening of the security guarantees as the early use of the session key voids key indistinguishability. Clearly, since the protocol is UA-secure even without the CSM message then the data exchanged prior to the sending of CSM is UA-secure. But what notion of security can one claim once the client authenticates? Fortunately, using a simple adaptation of the proof of the SIGMAC compiler to the present setting, we are able to show that the data exchanged before the sending of CSM (as well as later data) can be considered to be protected by mutually authenticated secure channels, with the assurance for the server delayed until it verifies the client's CSM message. We frame this result in our more general functional setting (Section 6.1).

Let  $\Pi_2$  be a protocol resulting from applying SIGMAC to a UA-secure protocol  $\Pi_1$ , namely,  $\Pi_2$  consists of running  $\Pi_1$  and adding to it the CSM message from client to server with a defined `sid*` value. Let  $\Pi_2'$  be a modification of  $\Pi_2$  where between the last message of  $\Pi_1$  and the CSM message from the client, data (referred to as *pre-authentication data*)

is exchanged encrypted under the session key  $K_s$  generated at the end of the  $\Pi_1$  run. We frame this protocol in our functional setting by modeling the encryption and decryption operations applied to the pre-authentication data with abstract Encrypt and Decrypt functional queries.

We say that a *functional test is consistent with a set of functional queries* if the advantage of any adversary to win the test with a random key remains negligible even when allowed to issue functional queries from this set.

LEMMA 9. *Given a treplication-secure UA protocol  $\Pi_1$  (i.e., one that satisfies the conditions of Theorem 3), Protocol  $\Pi_2'$  as defined above is F-MA-secure with respect to any functional test consistent with the Encrypt and Decrypt functional queries defined by the protocol.*

**Application to ACCE security.** A corollary of this lemma is that if we instantiate the functional test with a security experiment for a given “secure channels” model then we obtain that the encrypted pre-authentication data attains the same protection as guaranteed by the secure channels model. Illustrating this in the ACCE model [20] needs some adjustment since ACCE separates key exchange (pre-accept phase) and encrypted message exchange (post-accept phase) while in our case these two are interleaved. What we are really interested in is the security of the encrypted message exchange protocol. To capture the latter form of security we could model the pre-accept phase with an (idealized) trusted distribution of pairwise session keys between honest parties. If a post-accept message exchange protocol has the property that when coupled with the idealized key exchange results in an ACCE secure protocol then we say that the message exchange protocol is *post-accept ACCE secure*. In this setting, what the above lemma says is that the encrypted pre-authentication data in protocol  $\Pi_2'$  is post-accept ACCE secure. (Note that the ACCE model is defined with mutual authentication which is ensured here, although the assurance for the server is delayed until it verifies the client's CSM message.) We omit a more formal treatment here.

#### 6.3.1 Security under post-handshake authentication with encrypted CSM.

One can extend the result from Lemma 9 to the case where in addition to the encryption of pre-authentication data, the CSM message itself is encrypted as in Section 6.2 (this corresponds to the TLS 1.3 specification). The simulation combines the simulators from the proofs of Lemma 7 and Lemma 9 using the key  $K_s$  to encrypt both pre-authentication data and the CSM message. Note that the encrypted CSM is delivered only after the encryption of pre-authentication data, hence in a stateful encryption its ciphertext will depend on the presence of prior encryptions (TLS 1.3's encryption is stateful in that it uses a sequence number as the per-record nonce). The simulation will also deliver the encrypted CSM at this time, hence it will encrypt it correctly (a crucial point is that there is no need to send CSM out of order in the simulation).

<sup>10</sup>It is assumed that the Encrypt state is initialized and used first for encrypting the CSM message and then applied to messages exchanged in the post-accept phase.

## Acknowledgments.

This paper has benefited enormously from collaboration with Eric Rescorla and Hoeteck Wee on the ongoing design of TLS 1.3, and from enlightening inputs from Björn Tackmann, Douglas Stebila, Felix Günther, Karthik Bhargavan, Kenny Paterson, Markulf Kohlweiss and many of the other participants in the TRON workshop and the TLS working group. Special thanks to Björn, Douglas, Eric and the anonymous reviewers for their excellent and detailed comments.

## References

- [1] M. Bellare and P. Rogaway. Entity authentication and key distribution. In D. R. Stinson, editor, *CRYPTO*, volume 773 of *Lecture Notes in Computer Science*, pages 232–249. Springer, 1993. ISBN 3-540-57766-1.
- [2] M. Bellare, R. Canetti, and H. Krawczyk. A modular approach to the design and analysis of authentication and key exchange protocols (extended abstract). In *30th ACM STOC*, pages 419–428. ACM Press, May 1998.
- [3] B. Beurdouche, K. Bhargavan, A. Delignat-Lavaud, C. Fournet, M. Kohlweiss, A. Pironti, P. Strub, and J. K. Zinzindohoue. A messy state of the union: Taming the composite state machines of TLS. In *IEEE Symposium on Security and Privacy*, 2015.
- [4] K. Bhargavan and G. Leurent. Transcript collision attacks: Breaking authentication in tls, IKE and SSH. In *23rd Annual Network and Distributed System Security Symposium, NDSS 2016, San Diego, California, USA, February 21-24, 2016*. The Internet Society, 2016. URL <http://www.internetsociety.org/events/ndss-symposium-2016>.
- [5] K. Bhargavan, C. Fournet, M. Kohlweiss, A. Pironti, and P.-Y. Strub. Implementing TLS with verified cryptographic security. In *IEEE Symposium on Security and Privacy*, 2013. URL <http://mitls.rocq.inria.fr/>.
- [6] K. Bhargavan, A. Delignat-Lavaud, C. Fournet, A. Pironti, and P. Strub. Triple handshakes and cookie cutters: Breaking and fixing authentication over TLS. In *2014 IEEE Symposium on Security and Privacy, SP*, pages 98–113, 2014.
- [7] K. Bhargavan, A. Delignat-Lavaud, and A. Pironti. Verified contributive channel bindings for compound authentication. In *NDSS*, 2015.
- [8] C. Brzuska, M. Fischlin, N. P. Smart, B. Warinschi, and S. C. Williams. Less is more: relaxed yet composable security notions for key exchange. *Int. J. Inf. Sec.*, 12(4):267–297, 2013. Cryptology ePrint Archive, Report 2012/242.
- [9] R. Canetti and H. Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In *EUROCRYPT*, pages 453–474, 2001. See also Cryptology ePrint Archive, Report 2001/040.
- [10] R. Canetti and H. Krawczyk. Universally composable notions of key exchange and secure channels. In *EUROCRYPT*, pages 337–351, 2002. See also Cryptology ePrint Archive, Report 2002/059.
- [11] C. Cremers, M. Horvat, S. Scott, and T. van der Merwe. Automated verification of TLS 1.3: 0-RTT, resumption and delayed authentication. In *IEEE S&P 2016.*, 2016.
- [12] B. Dowling, M. Fischlin, F. Günther, and D. Stebila. A cryptographic analysis of the TLS 1.3 handshake protocol candidates. In *ACM CCS*, 2015. Also, Cryptology ePrint Archive, Report 2015/914.
- [13] B. Dowling, M. Fischlin, F. Günther, and D. Stebila. A cryptographic analysis of the TLS 1.3 draft-10 full and pre-shared key handshake protocol. Cryptology ePrint Archive, Report 2016/081, 2016.
- [14] M. Fischlin and F. Günther. Multi-stage key exchange and the case of Google’s QUIC protocol. In *ACM CCS*, 2014.
- [15] M. Fischlin, F. Günther, G. A. Marson, and K. G. Paterson. Data is a stream: Security of stream-based channels. In R. Gennaro and M. J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 545–564. Springer, Heidelberg, Aug. 2015. .
- [16] I. Goldberg, D. Stebila, and B. Ustaoglu. Anonymity and one-way authentication in key exchange protocols. *Des. Codes Cryptography*, 67(2):245–269, 2013. . URL <http://dx.doi.org/10.1007/s10623-011-9604-z>.
- [17] S. Halevi and H. Krawczyk. Public-key cryptography and password protocols. *ACM Transactions on Information and System Security*, 2(3):230–268, Aug. 1999.
- [18] S. Halevi and H. Krawczyk. One-pass HMQV and asymmetric key-wrapping. In *PKC 2011*, pages 317–334, 2011.
- [19] T. Jager, F. Kohlar, S. Schäge, and J. Schwenk. Generic compilers for authenticated key exchange. In M. Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 232–249. Springer, Heidelberg, Dec. 2010.
- [20] T. Jager, F. Kohlar, S. Schäge, and J. Schwenk. On the security of TLS-DHE in the standard model. In *CRYPTO*, pages 273–293, 2012. Also Cryptology ePrint Archive, Report 2011/219.
- [21] F. Kohlar, S. Schäge, and J. Schwenk. On the security of TLS-DH and TLS-RSA in the standard model. Cryptology ePrint Archive, Report 2013/367, 2013. <http://eprint.iacr.org/>.
- [22] M. Kohlweiss, U. Maurer, C. Onete, B. Tackmann, and D. Venturi. (De-)constructing TLS. Cryptology ePrint Archive, Report 2014/020, 2014. revised Apr 2015.
- [23] M. Kohlweiss, U. Maurer, C. Onete, B. Tackmann, and D. Venturi. (de-)constructing TLS 1.3. In *Progress in Cryptology - INDOCRYPT 2015 - 16th International Conference on Cryptology in India, Bangalore, India, December 6-9, 2015, Proceedings*, pages 85–102, 2015.
- [24] H. Krawczyk. SIGMA: The “SIGn-and-Mac” approach to authenticated Diffie-Hellman and its use in the IKE protocols. In *CRYPTO*, pages 400–425, 2003.
- [25] H. Krawczyk and H. Wee. The OPTLS protocol and TLS 1.3. In *EuroS&P*, 2016.
- [26] H. Krawczyk, K. G. Paterson, and H. Wee. On the security of the TLS protocol: A systematic analysis. In *CRYPTO (1)*, pages 429–448, 2013. Also, Cryptology ePrint Archive, Report 2013/339.
- [27] A. Langley and W.-T. Chang. QUIC crypto, 2013. URL <http://tinyurl.com/lrrjyjs>.
- [28] R. Lychev, S. Jero, A. Boldyreva, and C. Nita-Rotaru. How secure and quick is QUIC? Provable security and performance analyses. In *IEEE Symposium on Security and Privacy*, pages 214–231, 2015.
- [29] U. Maurer, B. Tackmann, and S. Coretti. Key exchange with unilateral authentication: Composable security definition and modular protocol design. *IACR Cryptology ePrint Archive*, 2013:555, 2013. URL <http://eprint.iacr.org/2013/555>.
- [30] P. Morrissey, N. P. Smart, and B. Warinschi. A modular security analysis of the TLS handshake protocol. In *ASIACRYPT*, pages 55–73, 2008.
- [31] K. G. Paterson, T. Ristenpart, and T. Shrimpton. Tag size does matter: Attacks and proofs for the TLS record protocol. In *ASIACRYPT*, pages 372–389, 2011.
- [32] M. D. Raimondo, R. Gennaro, and H. Krawczyk. Deniable authentication and key exchange. In *ACM CCS*, 2006.
- [33] E. Rescorla. The transport layer security (TLS) protocol version 1.3 (draft 13), Dec. 2015. URL <https://tools.ietf.org/html/draft-ietf-tls-tls13-13>.
- [34] V. Shoup. On formal models for secure key exchange. Cryptology ePrint Archive, Report 1999/012, 1999.