



A Tale of Two Models: Formal Verification of KEMTLS via Tamarin

Sofia Celi¹, Jonathan Hoyland¹, Douglas Stebila²,
and Thom Wiggers³

¹ Cloudflare, Inc., San Francisco, USA

cherenkov@riseup.net, jhoyland@cloudflare.com

² University of Waterloo, Waterloo, Canada

dstebila@uwaterloo.ca

³ Radboud University, Nijmegen, The Netherlands

thom@thomwiggers.nl

Abstract. KEMTLS is a proposal for changing the TLS handshake to authenticate the handshake using long-term key encapsulation mechanism keys instead of signatures, motivated by trade-offs in the characteristics of post-quantum algorithms. Prior proofs of security of KEMTLS and its variant KEMTLS-PDK have been hand-written proofs in the reductionist model under computational assumptions. In this paper, we present computer-verified symbolic analyses of KEMTLS and KEMTLS-PDK using two distinct Tamarin models. In the first analysis, we adapt the detailed Tamarin model of TLS 1.3 by Cremers et al. (ACM CCS 2017), which closely follows the wire-format of the protocol specification, to KEMTLS(-PDK). We show that KEMTLS(-PDK) has equivalent security properties to the main handshake of TLS 1.3 proven in this model. We were able to fully automate this Tamarin proof, compared with the previous TLS 1.3 Tamarin model, which required a big manual proving effort; we also uncovered some inconsistencies in the previous model. In the second analysis, we present a novel Tamarin model of KEMTLS(-PDK), which closely follows the multi-stage key exchange security model from prior pen-and-paper proofs of KEMTLS(-PDK). The second approach is further away from the wire-format of the protocol specification but captures more subtleties in security definitions, like deniability and different levels of forward secrecy; it also identifies some flaws in the security claims from the pen-and-paper proofs. Our positive security results increase the confidence in the design of KEMTLS(-PDK). Moreover, viewing these models side-by-side allows us to comment on the trade-off in symbolic analysis between detail in protocol specification and granularity of security properties.

Keywords: Post-quantum cryptography · TLS · Key encapsulation mechanisms · Formal analysis · Tamarin

S. Celi—Now works for Brave Software, Inc.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2022
V. Atluri et al. (Eds.): ESORICS 2022, LNCS 13556, pp. 63–83, 2022.
https://doi.org/10.1007/978-3-031-17143-7_4

1 Introduction

The Transport Layer Security (TLS) protocol is one of the most used cryptographic protocols. In its most recent version, the TLS 1.3 [31] handshake employs an ephemeral (elliptic-curve) Diffie–Hellman (DH) key exchange to establish session keys for confidentiality. In the regular handshake, TLS 1.3 authenticates the server and optionally the client using RSA or elliptic-curve signatures. It transmits the public keys to verify those signatures during the handshake, in certificates signed by a certificate authority (CA).

KEMTLS [32] is an alternative proposal for a post-quantum TLS 1.3 [31] handshake. It avoids using handshake signatures, which typically authenticate the TLS 1.3 handshake, replacing them with end-entity authentication based on key encapsulation mechanisms (KEMs) following well-established techniques for implicitly authenticated key exchange. As post-quantum KEMs are typically more efficient than the post-quantum signature schemes, either in bytes on the wire or computational efficiency, this saves resources. KEMTLS-PDK (“pre-distributed public key”) [34] is a variant of KEMTLS that offers a more efficient handshake if the client already has the server’s long-term public key. The authentication mechanisms from KEMTLS and KEMTLS-PDK have been proposed for standardisation to the Internet Engineering Task Force (IETF) TLS working group [13].

Figure 1 shows the cryptographic core of the unilaterally authenticated TLS 1.3 and KEMTLS handshakes. KEMTLS replaces the TLS 1.3 Diffie–Hellman-based ephemeral key exchange by KEM operations. Most importantly, whereas in TLS 1.3 the server authenticates by signing the transcript using the key from the server’s certificate, in KEMTLS the client encapsulates against the KEM public key in the server’s certificate. KEMTLS then combines both KEM shared secrets—one from the ephemeral key exchange and one from the server’s long-term key—to derive a key that is *implicitly authenticated*, meaning only the intended server will be able to derive the secret. The client can then use the derived key to transmit application data.

At many levels, the KEMTLS handshake is similar to the TLS 1.3 handshake. However, due to the usage of KEMs, the order of messages in TLS 1.3 has been significantly changed. Additionally, the server can no longer send data in its first response to the client. However, KEMTLS preserves the client’s ability to send its message after receiving the first flight from the server.

As KEMTLS is a novel way to achieve authentication in the TLS 1.3 handshake, the security of its design should be carefully checked not only with pen-and-paper proofs but with a computer-assisted formal analysis of it to provide stronger evidence of its soundness to adopters and standardization bodies like the IETF.

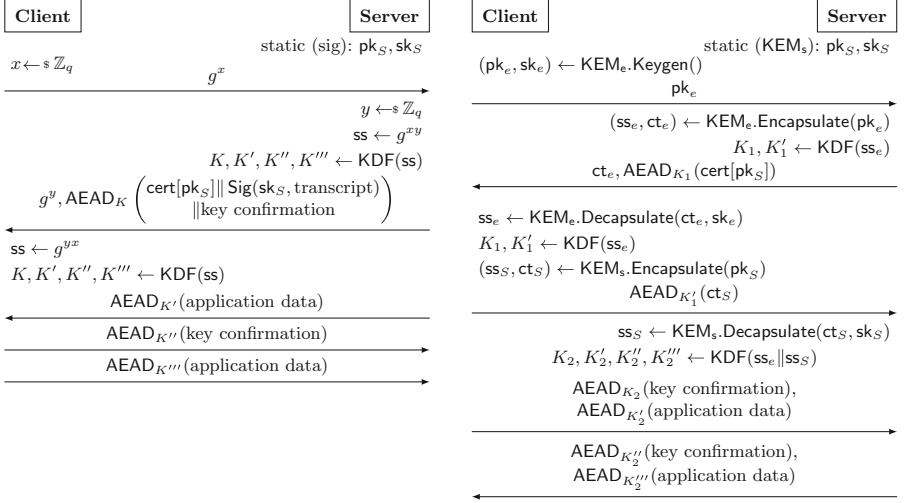


Fig. 1. Simplified protocol diagrams of server-only authenticated versions of: (left) the TLS 1.3 handshake, using signatures for authentication; and (right) the KEMTLS handshake, using KEMs for authentication.

1.1 Related Work

Analysis of TLS 1.3. During the development process of TLS 1.3, there was a strong collaboration between the standardisation community with the academic research community. Initial TLS 1.3 protocol designs were based on academic designs [28], and it was explicit goal of the TLS 1.3 process to incorporate academic security analysis of new designs before continuing with standardisation. Paterson and van der Merwe described this as a “design-break-fix-release” process rather than the “design-release-break-patch” cycle that was found on prior versions of the standardisation and usage of TLS [30]. Many of the security analyses of TLS 1.3 used the reductionist security paradigm [19–21, 27, 28]. Complementing this manual proof work, computer-aided cryptography [1] was also instrumental in checking TLS 1.3. Analyses were done using the Proverif [7] and Tamarin [14, 15] symbolic analysis tools, as well as a verified implementation in F* [16].

Analysis of KEMTLS. The initial KEMTLS and KEMTLS-PDK papers included reductionist security proofs [32, 34], adapting the multi-stage key exchange approach used by Dowling et al. [19, 20] for TLS 1.3. Subsequently, Towa et al. proposed and proved an alternative abbreviated handshake, with additional short-lived static keys [25], and found a few minor mistakes in the original security proofs, which were subsequently fixed in online versions of the original papers [33, 35]. All these proofs treat protocol modes independently—one-at-a-time—and do not consider the presence of the other protocol modes.

1.2 Contributions

In this work, we present two security analyses of all four variants of KEMTLS (the base KEMTLS protocol, with server-only or mutual authentication, and the pre-distributed public keys variant KEMTLS-PDK, also with server-only or mutual authentication) using Tamarin [3, 29]. The source code of our models is available at <https://github.com/kemtls/>.

Our first model, presented in Sect. 3, is based on the Tamarin analysis of TLS 1.3 by Cremers et al. [14]. This is a highly detailed model in terms of the protocol specification, closely following the TLS 1.3 wire format. In this model, we show that all four KEMTLS variants have equivalent security properties to the main handshake of TLS 1.3 without extensions. In implementing this model for KEMTLS, we were able to fully automate the proof, unlike the original model which required significant manual effort.

Our second model, presented in Sect. 4, is a novel Tamarin model developed from scratch that closely follows the multi-stage key exchange security model used in the pen-and-paper proofs [32, 34]. This model focuses on the “cryptographic core”, meaning that it is further away from the wire specification and does not model details like message encryption or the record layer. However, it captures more details in the security definitions, using the more granular definitions of forward secrecy from [32, 34] as well as including an analysis of deniability. This model allows us to symbolically verify the reductionist security claims from the pen-and-paper proofs, but goes further by considering all four KEMTLS variants simultaneously. This Tamarin model allowed us to identify some minor flaws in the properties stated based on pen-and-paper proofs.

In Sect. 5, we compare the features of our two Tamarin models. Having these two models side-by-side illustrates the trade-off between detail of protocol specification and granularity of security properties. Ideally, of course, one would achieve both levels of detail simultaneously, but such complexity is challenging both for the humans reading and writing pen-and-paper proofs or authoring Tamarin models, and for computers checking such Tamarin models (where runtime typically scales exponentially with the complexity of the model). Our side-by-side approach with two very different perspectives still yields significant confidence in the soundness of the KEMTLS protocol design and each provides insight into flaws in the earlier models that it was based on.

2 Background on Symbolic Analysis

One approach to proving the security properties of protocols is *symbolic analysis*, which uses formal logic to reason about the properties of an algebraic model of a protocol. Computational tools, such as Tamarin [3, 29] or ProVerif [9], can then be used to check whether certain properties hold in the symbolic model.

In symbolic analysis, generic symbols replace specific values. Operations like encryption are also modelled symbolically: for example, $\mathbf{senc}(a, b)$ represents the value a being symmetrically encrypted with the key b . In symbolic analysis,

cryptographic operations are *perfect*, meaning the adversary can learn nothing about an encrypted message without the correct key. The operations that describe a protocol in a symbolic model take messages and state information, and transform them into the next state or emit another protocol message. A tool can then use all operations and symbols to generate every possible protocol run.

Many symbolic analyses of protocols use the Dolev–Yao [18] attacker model, in which an attacker can manipulate all messages at will, e.g. by redirecting them, replaying them, dropping them, or manipulating their contents. It can also construct new messages from information previously learnt. However, as the cryptography is assumed to be perfect, the attacker can not read or modify encrypted or authenticated messages if it does not have the right keys.

Symbolic models can also be extended to give the attacker special extra abilities. For example, one can allow the attacker to reveal private keys or state information of parties by performing queries to a reveal oracle. We record when the attacker uses this oracle, so reveal queries become part of the trace of execution.

Security properties are modelled as predicates over execution traces. In Tamarin, during the execution of the rules of the protocol, we can emit *action facts*. We use these action facts to record, for example, the session’s impression of the authentication status or the current keys. We then write lemmas representing security properties as predicates over action facts: for example, that any key recorded in a certain type of action fact must not be known to the adversary, unless the adversary cheated by revealing keys. A model checker like Tamarin can then be used to check if the protocol maintains the required security property. Assuming soundness of the tool, either the tool will give a proof that the protocol has the required property, find a counter-example, or fail to terminate.

3 Model #1: High-Resolution Protocol Specification

In this section, we discuss the natural approach of taking one of the TLS 1.3 models and adapting it to KEMTLS(-PDK). Our work demonstrates that KEMTLS provides security guarantees at least equivalent to those proven by Cremers et al. for the main handshake of TLS 1.3.

3.1 Cremers et al.’s Tamarin TLS 1.3 Model

The Tamarin model of TLS 1.3 [14] is very high-resolution in terms of its modelling of protocol details and adherence to the protocol specification. It covers the cryptographic computations such as the key exchange and the key schedule; for example, calls to HKDF are decomposed into hash function calls. This model also includes the extensions to the basic TLS 1.3 handshake, such as the `HelloRetryRequest` mechanism, pre-shared keys, and resumption via session tickets. Additionally, it models the encryption of handshake messages, the syntax of the protocol messages, and mechanics such as TLS 1.3 extensions.

In terms of security properties, the Cremers et al. model extends Tamarin’s basic Dolev–Yao attacker with the ability to recover secrets from Diffie–Hellman key shares and reveal the long-term keys of participants. TLS 1.3 is not secure against an attacker who can use these attacks freely, but aims to provide confidentiality and integrity against an attacker who is restricted from revealing secrets of the target session. Cremers et al. were able to encode lemmas capturing most of the security properties claimed by the TLS 1.3 specification [31, Appendix E.1]. They report that proving all lemmas in their model took about a week. Much of this time was spent on manual interaction with Tamarin’s prover to guide it to prove some of the more complex lemmas. Verifying the generated proof requires “about a day” and “a vast amount of RAM” [14].

3.2 Representing KEMTLS in the Model

We now describe how we modified the existing TLS 1.3 model to represent both KEMTLS and its variant with pre-distributed keys, KEMTLS-PDK. The original model is highly modular, which made it relatively easy to modify.

The model of Cremers et al. represents TLS 1.3 through rules that manipulate a specific state object, which keeps track of many protocol variables, such as keys, authentication status, and the currently active handshake mode. Tamarin rules create transitions between these states. Where the protocol branches, such as when the server requests client authentication by sending `CertificateRequest`, there are two rules that end up in the same next state; for example, they would set the `cert_req` variable differently. The server later uses this variable to decide which of the rules `recv_client_auth` or `recv_client_auth_cert` to use; the latter expects the `Certificate`, `CertificateVerify`, and `Finished` messages, while the former only expects `Finished`. We handle the public key infrastructure (PKI) for KEM public keys in the same way as [14]: we do not model CA certificates, and assume an out-of-band binding between public keys and identities.

Ephemeral key exchange in the TLS 1.3 model uses Tamarin’s Diffie–Hellman functionality. It also allows the negotiation of two different DH groups. During the handshake, the client and server generate ephemeral DH secrets for the chosen group. If the server rejects the client’s choice of DH group, it falls back to another group through the `HelloRetryRequest` mechanism. To model the post-quantum ephemeral key exchange in KEMTLS, we replaced the Diffie–Hellman operations by `kemencaps` (KEM encapsulation) in place of the server’s DH key generation. The client then computes the shared secret via `kemdecaps` (KEM decapsulation).

The authentication rules and states required more careful consideration. In the TLS 1.3 model, the `Certificate`, `CertificateVerify`, and `Finished` messages were sent and received simultaneously. In KEMTLS, we split the handling of these messages, as the peer that is authenticating needs to first receive a ciphertext to decapsulate. Doing this requires more states. Additionally, in KEMTLS the client sends `Finished` before the server, which deviates from TLS 1.3.

To finish our integration of KEMTLS, we made changes to the key schedule to include the computation of KEMTLS’ Authenticated Handshake Secret (AHS)

```

lemma secret_session_keys:
  "All tid actor peer kw kr pas #i.
   SessionKey(tid, actor, peer, <pas, 'auth'>, <kw, kr>@#i &
   not (Ex #r. RevLtk(peer)@#r & #r < #i) &
   not (Ex tid3 esk #r. RevEKemSk(tid3, peer, esk)@#r & #r < #i) &
   not (Ex tid4 esk #r. RevEKemSk(tid4, actor, esk)@#r & #r < #i)
   ==> not Ex #j. K(kr)@#j"

```

Listing 1. The `secret_session_keys` lemma proves application traffic keys are secret.

and use the correct handshake traffic encryption keys. We also modified the action facts emitted in the various rules to match our KEM operations; lemmas that made use of these action facts were also updated. We disabled the PSK and session ticket features of the original model.

Modeling KEMTLS-PDK. In KEMTLS-PDK, the client has the server’s long-term public key beforehand. Access to the public key allows the client to send a ciphertext in the initial `ClientHello` message. Additionally, the client may attempt client authentication proactively and thus transmit its `Certificate` before receiving `ServerHello` from the server. We model this through an additional initial state for the KEMTLS-PDK client. From this state, there are two rules which set the state variable that will decide if the client will send its certificate. KEMTLS-PDK is otherwise implemented as a mostly separate sequence of states and rules, as the key schedule and order of messages are quite different. The client and server still transition through a state shared with KEMTLS, so they can fall back to the “full” handshake.

3.3 Security Properties

We adapt the lemmas from the Cremers et al. model for TLS 1.3. Many core lemmas are constructed around the `SessionKey` fact: the client and the server record this fact when the handshake concludes. `SessionKey` contains the actor’s final understanding of its and its peer’s identities, authentication statuses, and the application traffic keys. We prove all security properties discussed in [14], and briefly explain the most important of these below.

Adversary Compromise of Secrets. First, we note the extent to which the adversary can compromise ephemeral or long-term secrets. KEMTLS uses ephemeral KEM keys for ephemeral secrecy and long-term KEM keys for authentication. The adversary can reveal actors’ long-term secret keys; this records the `RevLtk($actor)` fact. We also allow revealing the ephemeral secret key in individual sessions, recording the `RevEKemSk(tid, $actor, esk)` fact. Variables `tid` (“thread identifier”) and `esk` track the specific session and secret key.

KEMs are not “symmetric” in the same way that Diffie–Hellman key exchange is. Only one party in each KEM key exchange has a secret key that can be

targeted by a reveal query. We do not model revealing the shared secret from the ciphertext.

Intermediate session keys, like the Main Secret (MS), can not be revealed directly. This follows from the design of the original model: in TLS 1.3, these secrets only depend on the ephemeral key exchange, so revealing the ephemeral key exchange in sessions not targeted by a lemma still allows the adversary to obtain those sessions’ intermediate session keys. In KEMTLS, this is no longer the case: we mix the shared secrets encapsulated against long-term keys into the key schedule; as a result, our attacker is slightly weaker. The model discussed in Sect. 4 does directly allow session key reveal.

(Forward) Secrecy of Session Keys. The outputs of the handshake, as recorded in the `SessionKey` fact, are the application traffic read and write keys `kr` and `kw`. We require these keys to remain secret against various forms of attacks. Forward secrecy requires that if the long-term keys (but not the ephemeral keys) were compromised after the session completes, the session keys remain secure.

We model this in the `secret_session_keys` lemma as shown in Listing 1. This lemma considers a client or server that believes it has authenticated its peer, where the attacker has not revealed the ephemeral KEM secret keys. We allow the attacker to reveal the peer’s long-term secret key, but only after the `SessionKey` fact was emitted; this is the “forward” secrecy aspect. The attacker should not be able to learn (`not Ex #j. K(kr)@#j`) the target’s read key `kr` under these constraints. We similarly prove forward secrecy for each of the intermediate keys in the key schedule: the Handshake Secret (HS), AHS and MS.

Note that in KEMTLS, the session keys are derived from not just the ephemeral key exchange as in TLS 1.3, but also include the secret encapsulated during the authentication phase of the handshake. This implies that both the ephemeral key and the server’s long-term key need to be compromised in client sessions, and the ephemeral key and the server’s long-term key in server sessions with mutual authentication. We prove this in our model through a variant of the `secret_session_keys` lemma that allows ephemeral key compromise, as long as the peer’s long-term key is never revealed.

Authentication. We model the authentication properties of KEMTLS in the same way as they were modelled for TLS 1.3. The client and server are partnered via the nonces exchanged in the initial messages. The `entity_authentication` lemma captures that if the client, at the end of the handshake protocol, has authenticated their peer, and the peer’s long-term keys have not been revealed, then there must be a peer session that started with the same nonces. The `mutual_` variant of this lemma states the same, but with the roles of client and server reversed. As these lemmas allow revealing the targeted actor’s long-term keys, these properties also cover key-compromise impersonation attacks. Similarly, in the lemma `transcript_agreement`, we prove that when the client, after receiving the server’s `Finished` message, commits to a transcript, there

exists a server that is running with the same transcript (or their long-term keys have been revealed). The `mutual_` variant lemma states the same but with the roles reversed.

3.4 Results

After adding relevant helper lemmas, Tamarin was able to auto-prove all the correctness and security lemmas for Model #1, with all four KEMTLS variants supported simultaneously. Run-times are shown in Appendix B.

Auto-proving and Helper Lemmas. Many of the lemmas in Cremers et al.’s model of TLS 1.3 were not able to be auto-proved by Tamarin; instead, the authors had to manually guide Tamarin through parts of the proof. Our goal was to improve the model so that it could be proved without manual intervention.

To help the automated prover, Cremers et al. introduced many intermediate lemmas, many of which state properties of earlier keys or more limited message exchanges. Inheriting these lemmas proved to be both helpful as well as distracting. Incrementally proving and adjusting the intermediate lemmas to apply to KEMTLS(-PDK) helped us spot bugs and make progress. But starting from their helper lemmas often left us unclear as to why particular intermediate lemmas were necessary to prove the final security properties.

In our experience, Tamarin does not find counterexamples very easily in big models. As a result, we wrote increasingly “smaller” lemmas whenever we ran into a lemma that was hard to prove. This greatly expanded the number of helper lemmas available. While we believe that this helped auto-prove the model, it also resulted in cases where the helper lemmas interacted in bad ways and had to be ignored. (Replacing Diffie–Hellman by KEM, thus avoiding Tamarin’s algebraic analysis of DH group operations, may also have eased analysis.) Additionally, the model of [14] is carefully split over different files to avoid certain helper lemmas from interacting. With much less experience, we joined together most of those files, which in many cases lead to Tamarin getting distracted by helper lemmas.

A Bug in Cremers et al.’s TLS 1.3 Lemmas. While working on the proof, we found that one of the core lemmas in [14]’s TLS 1.3 model seems to have changed after creating the proof. The lemma `session_key_agreement` tried to prove that the client’s and servers values of `keys` in the `SessionKey` fact matched. However, variable `keys` is a tuple $\langle kr, kw \rangle$ of the reading and writing keys. As the server’s writing key should match the client’s reading key and not the client’s writing key, this lemma did not hold. The rendered proofs included in the repository alongside the model and lemmas revealed that in the executed proof, `keys` was split into its elements and equated correctly. We disclosed the bug to the authors.

3.5 Limitations

Although the model is very granular in its description of KEMTLS(-PDK), we do have some limitations. As discussed in Sect. 3.3, we do not model intermediate

session key reveal. We also have not modelled session resumption or pre-shared key modes with KEMTLS. Finally, we have not attempted to model deniability, which we will model in Sect. 4.

4 Model #2: Multi-stage Key Exchange Model

The security properties shown in the original KEMTLS paper [32] and the KEMTLS-PDK paper [34] are stated using the reductionist security paradigm, via the *multi-stage key exchange model* [24], which was adapted for proofs of the TLS 1.3 handshake [19, 20]. Our goal in this section is to translate the reductionist security properties in this model—match security, session key indistinguishability, and authentication—from a pen-and-paper model to being encoded in Tamarin, then have the Tamarin prover confirm these properties hold. Notably, this model discriminates between the several keys established within a single KEMTLS handshake, associating distinct security properties with individual stage keys.

4.1 Reductionist Security Model for TLS 1.3 and KEMTLS

The multi-stage key exchange security model, first introduced by Fischlin and Günther [24], is an extension of the Bellare–Rogaway (BR) model [6] for proving security of authenticated key exchange in the reductionist security paradigm. In the BR model, the adversary is in control of all communications between honest parties, so the adversary can activate honest parties to send their next protocol message, and can also modify, delay, drop, replay, or create messages. Each honest party can run multiple simultaneous or sequential executions of the protocol (each execution at a party is called a session) sharing a single long-term key pair across their sessions. Within each session, a party maintains several variables, including the execution status, a session identifier, an identifier for the peer (if the peer is to be authenticated), and a session key. The adversary interacts with the honest parties via oracles, including oracles for starting a new session at a party (the `NewSession` oracle) and message delivery and response (`Send`), as well as letting the adversary learn an honest party’s long-term key (`Corrupt`) or the session key of a particular session (`Reveal`). The adversary may choose one session as a challenge session and, via a call to the `Test` oracle, receive an indistinguishability game challenge.

There are many extensions to the BR model to capture different functionality and security properties; see [10, Ch. 2] for a summary. One important extension is the formalisation by Brzuska et al. [11, 12] which introduces a property called *match security*. This checks the technical condition that the session identifiers specified by the protocol effectively match the partnered sessions. Among other benefits, match security helps with composition theorems involving AKE protocols. Our Tamarin model does address match security fully, but due to space constraints, we omit discussion of match security in the proceedings version of the paper.

In real-world protocols like QUIC and TLS, as well as KEMTLS, multiple keys are established in each session for different purposes. Fischlin and Günther [24] created the multi-stage key exchange model, an extension of the BR model in which a single session can have multiple stages, each of which establishes a key with certain security characteristics; they used this approach to analyse QUIC. It was also used by Dowling et al. [19, 20] to analyse the TLS 1.3 handshake. As KEMTLS is an alternative realization of the TLS 1.3 handshake, it is natural to similarly use this model for analyzing KEMTLS, as done in [32].

We now present the technical components of the multi-stage key exchange security model as used in KEMTLS [32] and KEMTLS-PDK [34]. Our presentation here will be somewhat abbreviated; for full details, see the full versions of [32, 34].

Partnering. For the proof of KEMTLS(-PDK), we need to keep track of the pairs of sessions that are (supposedly) communicating. Each session keeps track of per-stage session identifiers, each of which is a distinct label for the stage followed by all plaintext messages transmitted up until that point in the protocol; for KEMTLS-PDK, this also includes the implicit `ServerCertificate` message. We call two sessions *partners* if their session identifiers match.

Adversary Interaction. The oracles and variables stated in Sect. 4.1 suffice for modelling the various properties of match-security. To model key indistinguishability, the multi-stage model includes an oracle $\text{Test}(\pi, i)$ which challenges the adversary to distinguish the i th stage key of session π from random.

Multi-stage Security and Malicious Acceptance. Multi-stage security models secrecy of each stage key under specific forward secrecy properties. These properties include *implicit* and *explicit authentication*. The model is parameterized by values indicating the expected security properties of particular stage keys. [32, 34] define four levels of forward secrecy:

- No forward secrecy (0);
- *Weak forward secrecy level 1* (wfs1): the key is confidential against passive adversaries. This level allows the adversary to access the peer’s long-term keys. Keys with this level of forward secrecy have no authentication.
- *Weak forwards secrecy level 2* (wfs2): the key is confidential against passive adversaries (wfs1) and against active adversaries who never corrupted the peer’s long-term key. In the latter case, the key is implicitly authenticated.
- *Forward secrecy* (fs): the key is confidential against passive adversaries (wfs1) and against active adversaries who did not corrupt the peer’s long-term key before the stage accepted. Keys with level fs are implicitly authenticated.

As the protocol is executed, the security level of a particular stage key may be upgraded once a later stage accepts. The specific security levels for each client and the server are indicated in [32, 34]; the server’s security levels are different if mutual authentication is used.

Explicit authentication, which is e.g. achieved by the `Finished` messages, is modelled through *malicious acceptance*: an adversary should not be able to cause a supposedly explicitly authenticated stage to accept without a partner stage.

Deniability. Roughly speaking, deniability is the property that a party cannot provide proof to a judge that a peer participated in a particular protocol execution, even if they did. First introduced in general by Dwork, Naor, and Sahai [23] and in the context of key exchange by Di Raimondo, Gennaro, and Krawczyk [17], there are many flavours and variations of deniability; see e.g. [26] for a classification. *Offline deniability* is the inability of a judge to distinguish between a transcript generated by honest parties and a transcript generated by a simulator. The form of deniability offered by KEMTLS and KEMTLS-PDK (following the terminology of [26]) is that it provides offline deniability in the *universal deniability* setting (meaning the simulator only has access to parties’ long-term public keys) against an unbounded judge with full corruption powers (meaning the judge gets the parties’ long-term secret keys as well as any per-session coins).

Pen-and-Paper Proofs. The KEMTLS and KEMTLS-PDK papers [32,34] provide theorems and give proofs that their respective protocols satisfy the match-security and multi-stage security properties; they do not include any proofs for offline deniability. The match-security properties are shown information-theoretically, with terms depending on the number of sessions, the correctness probability of the KEMs, and the size of the TLS nonce space. The multi-stage security properties are shown under the following computational assumptions: hash function collision resistance, IND-1CCA security of KEM_e , PRF and dual-PRF security of HKDF.Extract , PRF-security of HKDF.Expand , EUF-CMA security of HMAC , and IND-CCA security of KEM_c and KEM_s . There is a tightness loss proportional to the number of sessions squared.

4.2 Formalizing the Reductionist Security Model in Tamarin

We formalized all four KEMTLS variants (regular and PDK, server-only and mutually authenticated) in Tamarin, along with lemmas capturing correctness, match security, multi-stage security, and deniability, analogous to the definitions from Sect. 4.1. We now describe the formalization in more detail.

Protocol Description. This Tamarin formulation of the four KEMTLS variants focuses on the “cryptographic core” of the protocol. Roughly speaking, this is the protocol as formulated in figures in the original papers [32,34], which includes cryptographic operations involved in the key exchange, but does not include extra fields and operations arising from the integration of the cryptographic operations into a network protocol. We only address the handshake protocol and exclude TLS message formatting, algorithm negotiation, and data structures such as certificates. We exclude other modes such

as TLS 1.3 session resumption or pre-shared key handshakes. We assume that long-term public keys are reliably distributed out-of-band. We omit modelling handshake encryption: while the various handshake traffic secrets are established and recorded as accepted in each stage of the protocol, subsequent handshake messages are sent in plaintext. The various primitives based on hash functions (HMAC, HKDF.Extract, HKDF.Expand) are modelled as independent opaque functions, rather than relying on each other and ultimately on a common hash function. As in the pen-and-paper proofs, there are three KEMs: KEM_e , KEM_c , and KEM_s , for ephemeral key exchange, client authentication, and server authentication, respectively. The KEMs are modelled as distinct primitives, meaning that a party cannot use its long-term credential to act as both a client and a server.

Adversary Interaction. Among the oracles stated in Sect. 4.1, the `NewSession` and `Send` oracles are not needed, since the Tamarin model includes rules for each protocol step. The Tamarin model does include `Corrupt` and `Reveal` oracles. Because key security in Tamarin is modelled not using indistinguishability but key recovery (the $K(\dots)$ fact in Tamarin lemmas), there is no need for the `Test` query in the Tamarin model.

Correctness Lemmas. We include a collection of “reachability” lemmas which check that, for every stage in all 4 protocol variants, it is possible to arrive at that stage, with honest client and server sessions having correct owner and peer information, matching contributive and session identifiers, and correct expectations on authentication, forward secrecy, and replayability; the reachability lemmas include checking retroactive upgrading of properties. These lemmas are implemented using Tamarin’s `exist-trace` feature. There are 47 reachability lemmas in total, generated from a template using the M4 macro language.

Security and Authentication Lemmas. The match security lemmas from Definition B.1 of [32], plus the adjustments for replayability in [34], are directly translated into Tamarin. The lemmas are basically predicates over the session-specific variables defined in the model syntax, and can be stated analogously since the Tamarin model includes action facts for each session-specific variable.

Session key security in Tamarin is modelled based on infeasibility of session key recovery, rather than indistinguishability of a session key from random. We have lemmas for each type of forward secrecy a stage key can have, directly translating the freshness conditions of [32, Defn. B.3] and [34, Defn. B.5].

We have a lemma for explicit authentication analogous to [34, Defn. B.5 3], including not requiring uniqueness of the replayable KEMTLS-PDK stage 1.

Deniability Lemmas. Whereas the lemmas for the above properties all share the same Tamarin protocol description as explained above, the deniability lemmas use a re-statement of the protocol description. To formulate a deniability lemma, we need two versions of the protocol description: honest execution of

the protocol using long-term secrets, and simulation using only public keys. The judge in the offline deniability game is passive and receives only transcripts, so we can collapse the multiple rules for each client and server action into a single rule that generates a full transcript including both client and server operations. The deniability lemmas use Tamarin’s observational equivalence feature [4] to check that the real and simulated transcripts are indistinguishable. Using observational equivalence causes a substantial increase in state space, so for efficiency reasons, we have a mode that omits portions of the transcript that are deterministically generated from earlier parts of the transcript.

4.3 Comparison of Pen-and-Paper and Tamarin Models

In principle, if the same security properties have been encoded in both a pen-and-paper reductionist security model and in a Tamarin model, a full and correct proof in the reductionist security model yields everything that a Tamarin proof could, and potentially more. In particular, reductionist security proofs do not idealize cryptographic primitives as much as Tamarin does. Moreover, a reductionist security proof can be done in the “concrete setting” [5], yielding a precise (non-asymptotic) relationship between the runtime and success probability of an adversary against the protocol versus the runtime and success probability of breaking the underlying cryptographic assumption. While it would be possible to encode the pen-and-paper proofs of KEMTLS from the original papers into a computer verification tool such as EasyCrypt [2], that would also require the cryptographer to manually write all game hops and reductions, a massive undertaking. To date, there are no proofs of KEMTLS using a computer-aided verification tool for reductionist proofs.

Tamarin does not lend itself to writing security properties in exactly the same way as would be used in reductionist security models. Although there is no way to objectively justify how close the pen-and-paper and Tamarin models of this section are to each other, subjectively we think they are quite close:

- The protocol specification in Tamarin maps nearly line-for-line onto the protocol figures in the original papers, using the same function interfaces, same key schedule, and same session identifiers.
- The session-specific variables in the pen-and-paper model correspond nearly one-for-one to action facts in the Tamarin model.
- There are Tamarin lemmas for each security property in the pen-and-paper model, and there is a clear mapping between the clauses in the predicates in the pen-and-paper model and the Tamarin model.

The main gap in modelling, as mentioned earlier, is that session key security is modelled via indistinguishability in the pen-and-paper models but via infeasibility of key recovery in the Tamarin model. Though it is possible to verify indistinguishability through Tamarin’s observational equivalence features, the effect on the state space as discussed in Sect. 4.2 makes this impractical.

4.4 Results

Tamarin was able to auto-prove all the lemmas for correctness, reachability, match security, multi-stage session key security, authentication, and deniability in Model #2, with all four KEMTLS variants supported simultaneously. We did not need to create any helper lemmas for Tamarin. Run-times are shown in Appendix B.

Bugs in the Original Papers’ Security Properties. When translating the models into Tamarin, we identified minor mistakes in some of the security properties listed in the KEMTLS [32] and KEMTLS-PDK [34] papers, highlighting the value of formal verification. We summarize the corrected properties in Appendix A.

4.5 Limitations

As noted above, the design of the model in this section imposes some limitations. Unlike in Sect. 3, we generally did not model non-cryptographic details of the handshake, such as TLS handshake messages, extensions, or the record layer. We also did not model handshake encryption or algorithm negotiation.

We also had, unlike in Sect. 3, three distinct KEMs for ephemeral key exchange, server authentication and client authentication. This implicitly assumes the same certificate is not used for both purposes, which was the basis of the Selfie attack [22]. Without this limitation, we observe a state-space explosion with a major impact on performance. For example, if $\text{KEM}_c = \text{KEM}_s$, the first 10 out of 11 `reachable_*` lemmas take over 8 h, and the last `reachable_*` did not terminate after 45 h, compared to all 11 `reachable_*` lemmas taking just over 1 min with distinct KEM_c and KEM_s .

Our deniability lemmas are for abbreviated transcripts and omit ephemeral coins. Without this limitation there is a major impact on performance. Including full transcripts for KEMTLS-sauth increases run-time from 1 min to 16 h, and including ephemeral coins increases runtime from 1 min to 110 min.

Table 1. Comparison of features in our two Tamarin models of KEMTLS

| Feature | Model #1 | Model #2 |
|---|----------|----------|
| <i>Protocol modelling</i> | | |
| Encrypted handshake messages | ✓ | ✗ |
| HKDF and HMAC decomposed into hash calls | ✓ | ✗ |
| Key exch. and auth. KEMs are the same algorithm | ✓ | ✗ |
| TLS message structure | ✓ | ✗ |
| Algorithm negotiation | ✓ | ✗ |
| <i>Security properties</i> | | |
| Adversary can reveal long-term keys | ✓ | ✓ |
| Adversary can reveal ephemeral keys | ✓ | ✗ |
| Adversary can reveal intermediate session keys | ✗ | ✓ |
| Multiple flavours of forward secrecy | ✗ | ✓ |
| Deniability | ✗ | ✓ |

5 Comparison of Models

We discussed two very different models of KEMTLS(-PDK) in the previous sections. These models are examples of how we can view modelling as the art of replacing specifics with generalities. Model #1 stays very close to the wire format of TLS 1.3 and phrases the security properties in terms of attacks on the ephemeral and long-term keys. It contains more implementation details such as algorithm negotiation, message framing, encryption of handshake messages, and even application data. Model #2 is more abstract in its representation of protocol messages. However, it models the cryptographic properties in a more granular fashion. This more abstract description closely follows the multi-stage pen-and-paper proofs of KEMTLS and KEMTLS-PDK, and allowed verifying the properties claimed in the pen-and-paper proofs. Table 1 summarizes differences between the two models, a few aspects of which we discuss further below.

Modelling KEMs. The two models differ in the way that they model the KEMs in the protocol. Model #1 uses the same functions for all KEM modes in the protocol (key exchange, server authentication and client authentication). Model #2 has three separate sets of functions for the three different KEM modes; this means the attacker can not copy ciphertexts or public keys from one of the modes to another, which should make proving the protocol easier. Interestingly, we saw significantly different performance between these two approaches. The second model proves in very short time with the three separate KEMs, but runtime blows up if we define all three KEM modes with the same functions; we did not attempt to generate the full proof because it took so long as discussed in Sect. 4.5. This suggests that splitting the three KEM modes in the first model could result in a speed-up. However, splitting the KEMs in Model #1 did not improve the time to auto-prove lemmas; in fact, a few lemmas even stopped being auto-

provable. Ideally, this puzzle would be resolved with a justification that there is a way of safely separating uses of KEMs, allowing us to use whichever form happens to be easier for Tamarin to prove.

Threat Model. Both models use Dolev–Yao attackers, but give the attackers slightly different extra abilities as noted in the bottom half of Table 1. Consequently, the results hold in slightly different circumstances. The attacker in Model #1 can compromise ephemeral keys and long-term keys, but not session keys, whereas the attacker in Model #2 can compromise intermediate session keys and long-term keys, but not ephemeral keys. Revealing the HS intermediate session key allows the second attacker to simulate the abilities of the first, but the reverse does not hold; the attacker in Model #2 is thus slightly stronger.

Ease of Use. Work on each of our two models was done by separate authors of this paper, neither of whom had written a paper using Tamarin before and who had only had a basic introduction to Tamarin prior to this work. Surprisingly to us, creation of Model #2 from scratch was simpler and proceeded faster than the work in Model #1 adapting Cremers et al.’s TLS 1.3 model to KEMTLS. We attribute this to the higher fidelity of the protocol model in Cremers et al., requiring more code to model our changes, and the higher difficulty in proving.

6 Conclusion

We presented two Tamarin models checking security properties of KEMTLS and its variant protocol KEMTLS-PDK. Model #1 is highly detailed in implementation characteristics, close to the wire-format of the protocol. Model #2 presents the protocol at a higher level but provides a more precise characterization of security properties. We prove that KEMTLS(-PDK) is secure in both models; importantly these analyses include all four KEMTLS variants supported simultaneously. Additionally, we proved offline deniability of KEMTLS(-PDK) in Model #2.

Overall, comparing these two analyses is something of an apples-to-oranges comparison. The two very different approaches allow us to model and test different properties of the protocol. Model #1 is closer to what an implementation would be like, and verifies the security properties in such a scenario. Adopting the Cremers et al. TLS 1.3 model [14] also allowed us to quickly adapt the security claims of TLS 1.3 to our protocols. Model #2, on the other hand, is an adaptation of the multi-stage authenticated key exchange model from the pen-and-paper proofs in [32,34]. As such, Model #2 in a sense checks the claims in the pen-and-paper proofs, and in fact uncovered some minor mistakes in those proofs.

Our two models illustrate a common trade-off in formal analysis between the detail of the protocol specification and the granularity of the security properties we can prove. A similar observation was also made by Cremers et al. [14], who

commented computational analyses could only look at parts of TLS 1.3, rather than considering all the modes at once.

While we proved certain privacy properties, such as deniability, our models can be further expanded to include other privacy properties, such as the proposed Encrypted Client Hello extension (previously called ESNI). These properties have only been proven by using the symbolic protocol analyzer ProVerif [8].

Acknowledgements. The authors gratefully acknowledge helpful suggestions from Peter Schwabe. D. Stebila was supported by Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery grants RGPIN-2016–05146 and RGPIN-2022–03187. D. Stebila and T. Wiggers were supported by an NLnet Assure grant for the project “Standardizing KEMTLS”. T. Wiggers was supported by the European Research Council through Starting Grant No. 805031 (EPOQUE).

A Errors Identified in the Stated Properties of KEMTLS(-PDK)

Using Model #2, we identified minor mistakes in some of the forward secrecy and authentication properties listed in the original KEMTLS [32] and KEMTLS-PDK [34] papers. See the original papers for the definition of the symbols.

- In KEMTLS-mutual: $\text{auth}_3^S = 3$ and $\text{auth}_4^S = 4$ both should have been set to 5; $\text{FS}_{3,3}^S = \text{FS}_{3,4}^S = \text{FS}_{4,4}^S = \text{wfs2}$ should all have been wfs1 ; and $\text{auth}_6^S = 6$ should have been $\text{auth}_6^S = \infty$.
- In KEMTLS-PDK-sauth: $\text{FS}_{1,j}^C$ and $\text{FS}_{1,j}^S$ should have been 0 for all j ; $\text{auth}_5^C = 5$ should have been $\text{auth}_5^C = \infty$; and $\text{FS}_{i,4}^S$ should have been wfs1 for $i = 2, 3, 4$.
- In KEMTLS-PDK-mutual: the message SKC should have been included in the SF MAC computation and SF should have been included in the CF MAC computation; $\text{FS}_{1,j}^C$ and $\text{FS}_{1,j}^S$ should have been 0 for all j ; $\text{auth}_5^C = 5$ should have been $\text{auth}_5^C = \infty$; and $\text{FS}_{4,4}^S = \text{wfs1}$ should have been wfs2 .

The source papers have been updated online [33,35] with our corrections.

B Performance

We ran our two models using `tamarin-prover` version 1.16.1 on a server that has two 20-core Intel Xeon Gold 6230 CPUs, which after hyperthreading gives us 80 threads; the server has 192 GB of RAM. We note that communication bottlenecks between cores prevent fully utilising all resources.

Model #1 Table 2 shows run-times of the most time-consuming lemmas from Model #1. All four KEMTLS variants were supported simultaneously.

Table 2. Wall-clock run-time (hh:mm:ss) and memory usage of selected lemmas for Model #1

| Lemma | Steps | Time | Memory |
|----------------------------|--------|----------|--------|
| session_key_auth_agreement | 29 116 | 6:42:01 | 16 GB |
| session_key_agreement | 57 680 | 13:56:04 | 32 GB |
| handshake_secret | 29 390 | 4:40:52 | 12 GB |
| master_secret_pfs | 29 535 | 2:53:11 | 76 GB |
| All lemmas | — | 28 h | 121 GB |

Model #2 Table 3 shows the run-time for the various lemmas, for each KEMTLS variant on its own, and when all four KEMTLS variants are supported simultaneously; Tamarin was restricted to using 16 cores.

Table 3. Wall-clock run-time (hh:mm:ss) of lemmas for Model #2

| Lemma | KEMTLS | | | KEMTLS-PDK | | | All 4 variants |
|---------------------|---------|---------|---------|------------|---------|---------|----------------|
| | sauth | mutual | both | sauth | mutual | both | |
| reachable_* | 0:01:17 | 0:01:20 | 0:04:32 | 0:01:46 | 0:01:36 | 0:04:40 | 0:13:25 |
| attacker_works_* | 0:00:17 | 0:00:46 | 0:01:16 | 0:00:17 | 0:00:23 | 0:00:53 | 0:12:04 |
| match_* | 0:01:02 | 0:01:22 | 0:02:55 | 0:00:55 | 0:01:14 | 0:02:46 | 0:09:53 |
| sk_sec_nofs_client | 0:00:05 | 0:00:07 | 0:00:16 | 0:00:05 | 0:00:05 | 0:00:14 | 0:00:41 |
| sk_sec_nofs_server | 0:00:05 | 0:00:06 | 0:00:12 | 0:00:05 | 0:00:06 | 0:00:14 | 0:00:40 |
| sk_sec_wfs1 | 0:00:21 | 0:00:10 | 0:01:05 | 0:00:17 | 0:00:18 | 0:00:41 | 0:03:00 |
| sk_sec_wfs2 | 0:00:36 | 0:00:28 | 0:01:30 | 0:00:28 | 0:00:22 | 0:01:23 | 0:24:28 |
| sk_sec_fs | 0:01:20 | 0:03:05 | 0:06:38 | 0:01:21 | 0:01:33 | 0:05:07 | 1:39:58 |
| malicious_accept. | 0:00:13 | 0:01:40 | 0:04:13 | 0:00:17 | 0:00:22 | 0:01:39 | 27:29:37 |
| deniability (abbr.) | 0:01:02 | 0:12:15 | — | 0:00:24 | 0:29:10 | — | — |
| Total (excl. den.) | 0:05:16 | 0:09:05 | 0:22:38 | 0:05:30 | 0:06:00 | 0:17:38 | 30:13:46 |

References

1. Barbosa, M., et al.: SoK: computer-aided cryptography. In: 2021 IEEE Symposium on Security and Privacy, pp. 777–795. IEEE Computer Society Press (2021). <https://doi.org/10.1109/SP40001.2021.00008>
2. Barthe, G., Grégoire, B., Héraud, S., Béguelin, S.Z.: Computer-aided security proofs for the working cryptographer. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 71–90. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22792-9_5
3. Basin, D., Cremers, C., Dreier, J., Meier, S., Sasse, R., Schmidt, B.: Tamarin prover (2022). <https://tamarin-prover.github.io>
4. Basin, D.A., Dreier, J., Sasse, R.: Automated symbolic proofs of observational equivalence. In: Ray, I., Li, N., Kruegel, C. (eds.) ACM CCS 2015, pp. 1144–1155. ACM Press (2015). <https://doi.org/10.1145/2810103.2813662>

5. Bellare, M.: Practice-oriented provable-security. In: Okamoto, E., Davida, G., Mambo, M. (eds.) ISW 1997. LNCS, vol. 1396, pp. 221–231. Springer, Heidelberg (1998). <https://doi.org/10.1007/BFb0030423>
6. Bellare, M., Rogaway, P.: Entity authentication and key distribution. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 232–249. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-48329-2_21
7. Bhargavan, K., Blanchet, B., Kobeissi, N.: Verified models and reference implementations for the TLS 1.3 standard candidate. In: 2017 IEEE Symposium on Security and Privacy, pp. 483–502. IEEE Computer Society Press (2017). <https://doi.org/10.1109/SP.2017.26>
8. Bhargavan, K., Cheval, V., Wood, C.: Handshake privacy for TLS 1.3 - technical report. Research report, Inria Paris, Cloudflare (2022). <https://hal.inria.fr/hal-03594482>
9. Blanchet, B.: An efficient cryptographic protocol verifier based on Prolog rules. In: 14th IEEE Computer Security Foundations Workshop (CSFW-14), pp. 82–96. IEEE Computer Society (2001)
10. Boyd, C., Mathuria, A., Stebila, D.: Protocols for Authentication and Key Establishment. Springer, Heidelberg (2019). <https://doi.org/10.1007/978-3-662-58146-9>
11. Brzuska, C.: On the Foundations of Key Exchange, Ph.D. thesis, Technische Universität Darmstadt (2013). <http://tuprints.ulb-tu-darmstadt.de/3414/>
12. Brzuska, C., Fischlin, M., Warinschi, B., Williams, S.C.: Composability of Bellare-Rogaway key exchange protocols. In: Chen, Y., Danezis, G., Shmatikov, V. (eds.) ACM CCS 2011, pp. 51–62. ACM Press (2011). <https://doi.org/10.1145/2046707.2046716>
13. Celi, S., Schwabe, P., Stebila, D., Sullivan, N., Wiggers, T.: KEM-based Authentication for TLS 1.3. Internet-Draft draft-celi-wiggers-tls-authkem-01, Internet Engineering Task Force (2022). <https://datatracker.ietf.org/doc/html/draft-celi-wiggers-tls-authkem-01>, Work in Progress
14. Cremers, C., Horvat, M., Hoyland, J., Scott, S., van der Merwe, T.: A comprehensive symbolic analysis of TLS 1.3. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017, pp. 1773–1788. ACM Press (2017). <https://doi.org/10.1145/3133956.3134063>
15. 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 IEEE Symposium on Security and Privacy, pp. 470–485. IEEE Computer Society Press (2016). <https://doi.org/10.1109/SP.2016.35>
16. Delignat-Lavaud, A., et al.: Implementing and proving the TLS 1.3 record layer. In: 2017 IEEE Symposium on Security and Privacy, pp. 463–482. IEEE Computer Society Press (2017). <https://doi.org/10.1109/SP.2017.58>
17. Di Raimondo, M., Gennaro, R., Krawczyk, H.: Deniable authentication and key exchange. In: Juels, A., Wright, R.N., De Capitani di Vimercati, S. (eds.) ACM CCS 2006, pp. 400–409. ACM Press (2006). <https://doi.org/10.1145/1180405.1180454>
18. Dolev, D., Yao, A.C.C.: On the security of public key protocols (extended abstract). In: 22nd FOCS, pp. 350–357. IEEE Computer Society Press (1981). <https://doi.org/10.1109/SFCS.1981.32>
19. Dowling, B., Fischlin, M., Günther, F., Stebila, D.: A cryptographic analysis of the TLS 1.3 handshake protocol candidates. In: Ray, I., Li, N., Kruegel, C. (eds.) ACM CCS 2015, pp. 1197–1210. ACM Press (2015). <https://doi.org/10.1145/2810103.2813653>

20. Dowling, B., Fischlin, M., Günther, F., Stebila, D.: A cryptographic analysis of the TLS 1.3 handshake protocol. *J. Cryptol.* **34**(4), 1–69 (2021). <https://doi.org/10.1007/s00145-021-09384-1>
21. Dowling, B., Stebila, D.: Modelling Ciphersuite and version negotiation in the TLS protocol. In: Foo, E., Stebila, D. (eds.) *ACISP 2015*. LNCS, vol. 9144, pp. 270–288. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-19962-7_16
22. Drucker, N., Gueron, S.: Selfie: reflections on TLS 1.3 with PSK. *J. Cryptol.* **34**(3), 1–18 (2021). <https://doi.org/10.1007/s00145-021-09387-y>
23. Dwork, C., Naor, M., Sahai, A.: Concurrent zero-knowledge. In: *30th ACM STOC*, pp. 409–418. ACM Press (1998). <https://doi.org/10.1145/276698.276853>
24. Fischlin, M., Günther, F.: Multi-stage key exchange and the case of Google’s QUIC protocol. In: Ahn, G.J., Yung, M., Li, N. (eds.) *ACM CCS 2014*, pp. 1193–1204. ACM Press (2014). <https://doi.org/10.1145/2660267.2660308>
25. Günther, F., Rastikian, S., Towa, P., Wiggers, T.: KEMTLS with delayed forward identity protection in (Almost) a single round trip. In: *ACNS 2022* (2022). <https://eprint.iacr.org/2021/725>
26. Hülsing, A., Weber, F.: Epochal signatures for deniable group chats. In: *2021 IEEE Symposium on Security and Privacy*, pp. 1677–1695. IEEE Computer Society Press (2021). <https://doi.org/10.1109/SP40001.2021.00058>
27. Kohlweiss, M., Maurer, U., Onete, C., Tackmann, B., Venturi, D.: (De-)Constructing TLS 1.3. In: Biryukov, A., Goyal, V. (eds.) *INDOCRYPT 2015*. LNCS, vol. 9462, pp. 85–102. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-26617-6_5
28. Krawczyk, H., Wee, H.: The OPTLS protocol and TLS 1.3. In: *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*, pp. 81–96 (2016). <https://doi.org/10.1109/EuroSP.2016.18>
29. Meier, S., Schmidt, B., Cremers, C., Basin, D.: The TAMARIN prover for the symbolic analysis of security protocols. In: Sharygina, N., Veith, H. (eds.) *CAV 2013*. LNCS, vol. 8044, pp. 696–701. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39799-8_48
30. Paterson, K.G., van der Merwe, T.: Reactive and proactive standardisation of TLS. In: Chen, L., McGrew, D., Mitchell, C. (eds.) *SSR 2016*. LNCS, vol. 10074, pp. 160–186. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-49100-4_7
31. Rescorla, E.: The Transport Layer Security TLS Protocol Version 1.3. RFC 8446, RFC Editor (2018). <https://doi.org/10.17487/RFC8446>
32. Schwabe, P., Stebila, D., Wiggers, T.: Post-quantum TLS without handshake signatures. In: Ligatti, J., Ou, X., Katz, J., Vigna, G. (eds.) *ACM CCS 2020*, pp. 1461–1480. ACM Press (2020). <https://doi.org/10.1145/3372297.3423350>
33. Schwabe, P., Stebila, D., Wiggers, T.: Post-quantum TLS without handshake signatures. *Cryptology ePrint Archive*, Report 2020/534 (2020). <https://eprint.iacr.org/2020/534>
34. Schwabe, P., Stebila, D., Wiggers, T.: More efficient post-quantum KEMTLS with pre-distributed public keys. In: Bertino, E., Shulman, H., Waidner, M. (eds.) *ESORICS 2021*. LNCS, vol. 12972, pp. 3–22. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-88418-5_1
35. Schwabe, P., Stebila, D., Wiggers, T.: More efficient post-quantum KEMTLS with pre-distributed public keys. *Cryptology ePrint Archive*, Report 2021/779 (2021). <https://eprint.iacr.org/2021/779>