

Verified Models and Reference Implementations for the TLS 1.3 Standard Candidate

Karthikeyan Bhargavan, Bruno Blanchet, Nadim Kobeissi

INRIA

{*karthik.bhargavan,bruno.blanchet,nadim.kobeissi*}@inria.fr

Abstract—TLS 1.3 is the next version of the Transport Layer Security (TLS) protocol. Its clean-slate design is a reaction both to the increasing demand for low-latency HTTPS connections and to a series of recent high-profile attacks on TLS. The hope is that a fresh protocol with modern cryptography will prevent legacy problems; the danger is that it will expose new kinds of attacks, or reintroduce old flaws that were fixed in previous versions of TLS. After 18 drafts, the protocol is nearing completion, and the working group has appealed to researchers to analyze the protocol before publication. This paper responds by presenting a comprehensive analysis of the TLS 1.3 Draft-18 protocol.

We seek to answer three questions that have not been fully addressed in previous work on TLS 1.3: (1) Does TLS 1.3 prevent well-known attacks on TLS 1.2, such as Logjam or the Triple Handshake, even if it is run in parallel with TLS 1.2? (2) Can we mechanically verify the computational security of TLS 1.3 under standard (strong) assumptions on its cryptographic primitives? (3) How can we extend the guarantees of the TLS 1.3 protocol to the details of its implementations?

To answer these questions, we propose a methodology for developing verified symbolic and computational models of TLS 1.3 hand-in-hand with a high-assurance reference implementation of the protocol. We present symbolic ProVerif models for various intermediate versions of TLS 1.3 and evaluate them against a rich class of attacks to reconstruct both known and previously unpublished vulnerabilities that influenced the current design of the protocol. We present a computational CryptoVerif model for TLS 1.3 Draft-18 and prove its security. We present RefTLS, an interoperable implementation of TLS 1.0-1.3 and automatically analyze its protocol core by extracting a ProVerif model from its typed JavaScript code.

I. INTRODUCTION

The Transport Layer Security (TLS) protocol is widely used to establish secure channels on the Internet. It was first proposed under the name SSL [45] in 1994, and has undergone a series of revisions since, leading up to the standardization of TLS 1.2 [37] in 2008. Each version adds new features, deprecates obsolete constructions, and introduces countermeasures for weaknesses found in previous versions. The behavior of the protocol can be further customized via *extensions*, some of which are mandatory to prevent known attacks on the protocol.

One may expect that TLS clients and servers would use only the latest version of the protocol with all security-critical extensions enabled. In practice, however, many legacy variants of the protocol continue to be supported for backwards compatibility, and the everyday use of TLS

depends crucially on clients and servers negotiating the most secure variant that they have in common. Securely composing and implementing the many different versions and features of TLS has proved to be surprisingly hard, leading to the continued discovery of high-profile vulnerabilities in the protocol.

A history of vulnerabilities. We identify four kinds of attacks that TLS has traditionally suffered from. *Downgrade* attacks enable a network adversary to fool a TLS client and server into using a weaker variant of the protocol than they would normally use with each other. In particular, version downgrade attacks were first demonstrated from SSL 3 to SSL 2 [72] and continue to be exploited in recent attacks like POODLE [60] and DROWN [7]. *Cryptographic* vulnerabilities rely on weaknesses in the protocol constructions used by TLS. Recent attacks have exploited key biases in RC4 [3], [71], padding oracles in MAC-then-Encrypt [4], [60], padding oracles in RSA PKCS#1 v1.5 [7], weak Diffie-Hellman groups [1], and weak hash functions [23]. *Protocol composition* flaws appear when multiple modes of the protocol interact in unexpected ways if enabled in parallel. For example, the renegotiation attack [65] exploits the sequential composition of two TLS handshakes, the Triple Handshake attack [15] composes three handshakes, and cross-protocol attacks [58], [72] use one kind of TLS handshake to attack another. *Implementation bugs* contribute to the fourth category of attacks on TLS, and are perhaps the hardest to avoid. They range from memory safety bugs like HeartBleed and coding errors like GotoFail to complex state machine flaws like SKIP and FREAK [12]. Such bugs can be exploited to bypass all the security guarantees of TLS, and their prevalence, even in widely-vetted code, indicates the challenges of implementing TLS securely.

Security proofs. Historically, when an attack is found on TLS, practitioners propose a temporary fix that is implemented in all mainstream TLS libraries, then a longer-term countermeasure is incorporated into a protocol extension or in the next version of the protocol. This has led to a attack-patch-attack cycle that does not provide much assurance in any single version of the protocol, let alone its implementations.

An attractive alternative would have been to develop security proofs that systematically demonstrated the absence of large classes of attacks in TLS. However, developing proofs for an existing standard that was not designed with security models in mind is exceedingly hard [63]. After years of effort, the cryptographic community only recently

published proofs for the two main components of TLS: the *record* layer that implements authenticated encryption [57], [62], and the *handshake* layer that composes negotiation and key-exchange [46], [51]. These proofs required new security definitions and custom cryptographic assumptions, and even so, they apply only to abstract models of certain modes of the protocol. For example, the proofs do not account for low-level details of message formats, downgrade attacks, or composition flaws. Since such cryptographic proofs are typically carried out by hand, extending the proofs to cover all these details would require a prohibitive amount of work, and the resulting large proofs themselves would need to be carefully checked.

A different approach taken by the protocol verification community is to *symbolically* analyze cryptographic protocols using simpler, stronger assumptions on the underlying cryptography, commonly referred to as the Dolev-Yao model [39]. Such methods are easy to automate and can tackle large protocols like TLS in all their gory detail, and even aspects of TLS implementations [31], [18]. Symbolic protocol analyzers are better at finding attacks, but since they treat cryptographic constructions as perfect black boxes, they provide weaker security guarantees than classic cryptographic proofs that account for probabilistic and computational attacks.

The most advanced example of mechanized verification for TLS is the ongoing miTLS project [21], which uses dependent types to prove both the symbolic and cryptographic security of a TLS implementation that supports TLS 1.0-1.2, multiple key exchanges and encryption modes, session resumption, and renegotiation. This effort has uncovered weaknesses in both the TLS 1.2 standard [15] and its other implementations [12], and the proof is currently being extended towards TLS 1.3.

Towards Verified Security for TLS 1.3. In 2014, the TLS working group at the IETF commenced work on TLS 1.3, with the goal of designing a faster protocol inspired by the success of Google’s QUIC protocol [44]. Learning from the pitfalls of TLS 1.2, the working group invited the research community to contribute to the design of the protocol and help analyze its security even before the standard is published. A number of researchers, including the authors of this paper, responded by developing new security models and cryptographic proofs for various draft versions, and using their analyses to propose protocol changes. Cryptographic proofs were developed for Draft-5 [40], Draft-9 [52], and Draft-10 [55], which justified the core design of the protocol. A detailed symbolic model in Tamarin was developed for Draft-10 [35]. Other works studied specific aspects of TLS 1.3, such as key confirmation [41], client authentication [50], and downgrade resilience [14].

Some of these analyses also found attacks. The Tamarin analysis [35] uncovered a potential attack on the composition of pre-shared keys and certificate-based authentication, and this attack was prevented in Draft-11. A version downgrade attack was found in Draft-12 and its countermeasure in Draft-13 was proved secure [14]. A cross-protocol attack on RSA signatures was described in [47]. Even in this paper,

we describe two vulnerabilities in 0-RTT client authentication that we discovered and reported, which influenced the subsequent designs of Draft-7 and -13.

After 18 drafts, TLS 1.3 is entering the final phase of standardization. Although many of its design decisions have now been vetted by multiple security analyses, several unanswered questions remain. First, the protocol has continued to evolve rapidly with every draft version, so many of the cryptographic proofs cited above are already obsolete and do not apply to Draft-18. Since many of these are manual proofs, it is not easy to update them and check all the proof steps. Second, none of these symbolic or cryptographic analyses, with the exception of [14], consider the composition of TLS 1.3 with legacy versions like TLS 1.2. Hence, they do not account for attacks like [47] that exploit weak legacy crypto in TLS 1.2 to break the modern cryptographic constructions of TLS 1.3. Third, none of these works addresses TLS 1.3 implementations. In this paper, we seek to cover these gaps with a new comprehensive analysis of TLS 1.3 Draft-18.

Our Contributions. We propose a methodology for developing mechanically verified models of TLS 1.3 alongside a high-assurance reference implementation of the protocol.

We present symbolic protocol models for TLS 1.3 written in ProVerif [27]. They incorporate a novel security model (described in §II) that accounts for all recent attacks on TLS, including those relying on weak cryptographic algorithms. In §III-V, we use ProVerif to evaluate various modes and drafts of TLS 1.3 culminating in the first symbolic analysis of Draft-18 and the first composite analysis of TLS 1.3+1.2. Our analyses uncover known and new vulnerabilities that influenced the final design of Draft-18. Some of the features we study no longer appear in the protocol, but our analysis is still useful for posterity, to warn protocol designers and developers who may be tempted to reintroduce these problematic features in the future.

In §VI, we develop the first machine-checked cryptographic proof for TLS 1.3 using the verification tool CryptoVerif [24]. Our proof reduces the security of TLS 1.3 Draft-18 to standard cryptographic assumptions over its primitives. In contrast to manual proofs, our CryptoVerif script can be more easily updated from draft-to-draft, and as the protocol evolves.

Our ProVerif and CryptoVerif models capture the protocol core of TLS 1.3, but they elide many implementation details such as the protocol API and state machine. To demonstrate that our security results apply to carefully-written implementations of TLS 1.3, we present RefTLS (§VII), the first reference implementation of TLS 1.0-1.3 whose core protocol code has been formally analyzed for security. RefTLS is written in Flow, a statically typed variant of JavaScript, and is structured so that all its protocol code is isolated in a single module that can be automatically translated to ProVerif and symbolically analyzed against our rich threat model.

The full version of this paper is published as a technical report [13], and our models and code are available at:

<https://github.com/inria-prosecco/reftls>

II. A SECURITY MODEL FOR TLS

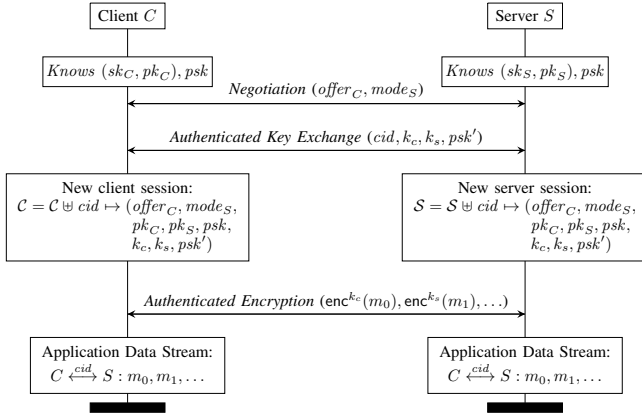


Figure 1: TLS Protocol Structure: Negotiation, then Authenticated Key Exchange (AKE), then Authenticated Encryption (AE) for application data streams.

Figure 1 depicts the progression of a typical TLS connection. Since a client and server may support different sets of features, they first *negotiate* a protocol mode that they have in common. In TLS, the client C makes an $offer_C$ and the server chooses its preferred $mode_S$, which includes the protocol version, the key exchange protocol, the authenticated encryption scheme, the Diffie-Hellman group (if applicable), and the signature and hash algorithms.

Then, C and S execute the negotiated *authenticated key exchange* protocol (e.g. Ephemeral Elliptic-Curve Diffie Hellman), which may use some combination of the long-term keys (e.g. public/private key pairs, symmetric pre-shared keys) known to the client and server. The key exchange ends by computing fresh symmetric keys (k_c, k_s) for a new session (with identifier cid) between C and S , and potentially a new pre-shared key (psk') that can be used to authenticate future connections between them.

In TLS, the negotiation and key exchange phases are together called the *handshake* protocol. Once the handshake is complete, C and S can start exchanging application data, protected by an authenticated encryption scheme (e.g. AES-GCM) with the session keys (k_c, k_s). The TLS protocol layer that handles authenticated encryption for application data is called the *record* protocol.

Security Goals for TLS. Each phase of a TLS connection has its own correctness and security goals. For example, during negotiation, the server must choose a $mode_S$ that is consistent with the client's $offer_C$; the key exchange must produce a secret session key, and so on. Although these intermediate security goals are important building blocks towards the security of the full TLS protocol, they are less meaningful to applications that typically use TLS via a TCP-socket-like API and are unaware of the protocol's internal structure. Consequently, we state the security goals of TLS from the viewpoint of the application, in terms of messages it sends and receives over a protocol session.

All goals are for messages between honest and authenticated clients and servers, that is, for those whose long-term

keys (sk_C, sk_S, psk) are unknown to the attacker. If only the server is authenticated, then the goals are stated solely from the viewpoint of the client, since the server does not know whether it is talking to an honest client or the attacker.

Secrecy: If an application data message m is sent over a session cid between an honest client C and honest server S , then this message is kept confidential from an attacker who cannot break the cryptographic constructions used in the session cid .

Forward Secrecy: Secrecy (above) holds even if the long-term keys of the client and server (sk_C, pk_C, psk) are given to the adversary after the session cid has been completed and the session keys k_c, k_s are deleted by C and S .

Authentication: If an application data message m is received over a session cid from an honest and authenticated peer, then the peer must have sent the same application data m in a matching session (with the same parameters $cid, offer_C, mode_S, pk_C, pk_S, psk, k_c, k_s, psk'$).

Replay Prevention: Any application data m sent over a session cid may be accepted at most once by the peer.

Unique Channel Identifier: If a client session and a server session have the same identifier cid , then all other parameters in these sessions must match (same $cid, offer_C, mode_S, pk_C, pk_S, psk, k_c, k_s, psk'$).

These security goals encompass most of the standard security goals for secure channel protocols such as TLS. For example, secrecy for application data implicitly requires that the authenticated key exchange must generate secret keys. Authentication incorporates the requirement that the client and server must have matching sessions, and in particular, that they agree on each others' identities as well as the inputs and outputs of negotiation. Hence, it prohibits client and server impersonation, and man-in-the-middle downgrade attacks.

The requirement for a unique channel identifier is a bit more unusual, but it allows multiple TLS sessions to be securely composed, for example via session resumption or renegotiation, without exposing them to credential forwarding attacks like Triple Handshake [15]. The channel identifier could itself be a session key or a value generated from it, but is more usually a public value that is derived from session data contributed by both the client and server [17].

Symbolic vs. Computational Models. Before we can model and verify TLS 1.3 against the security goals given above, we need to specify our protocol execution model. There are two different styles in which protocols have classically been modeled, and in this paper, we employ both of them. *Symbolic* models were developed by the security protocol verification community for ease of automated analysis. Cryptographers, on the other hand, prefer to use *computational* models and do their proofs by hand. A full comparison between these styles is beyond the scope of this paper (see e.g. [26]); here we briefly outline their differences in terms of the two tools we will use.

ProVerif [25], [27] analyzes symbolic protocol models, whereas CryptoVerif [24] verifies computational models.

The input languages of both tools are similar. For each protocol role (e.g. client or server) we write a *process* that can send and receive messages over public channels, trigger security events, and store messages in persistent databases.

In ProVerif, messages are modeled as abstract terms. Processes can generate new nonces and keys, which are treated as atomic opaque terms that are fresh and unguessable. Functions map terms to terms. For example, encryption constructs a complex term from its arguments (key and plaintext) that can only be deconstructed by decryption (with the same key). The attacker is an arbitrary ProVerif process running in parallel with the protocol, which can read and write messages on public channels, and can manipulate them symbolically.

In CryptoVerif, messages are concrete bitstrings. Freshly generated nonces and keys are randomly sampled bitstrings that the attacker can guess with some probability (depending on their length). Encryption and decryption are functions on bitstrings to which we may associate standard cryptographic assumptions such as IND-CCA. The attacker is a probabilistic polynomial-time CryptoVerif process running in parallel.

Authentication goals in both ProVerif and CryptoVerif are written as correspondences between events: for example, if the client triggers a certain event, then the server must have triggered a matching event in the past. Secrecy is treated differently in the two tools; in ProVerif, we typically ask whether the attacker can compute a secret, whereas in CryptoVerif, we ask whether it can distinguish a secret from a random bitstring.

The analysis techniques employed by the two tools are quite different. ProVerif searches for a protocol trace that violates the security goal, whereas CryptoVerif tries to construct a cryptographic proof that the protocol is equivalent (with high probability) to a trivially secure protocol. ProVerif is a push-button tool that may return that the security goal is true in the symbolic model, or that the goal is false with a counterexample, or that it is unable to conclude, or may fail to terminate. CryptoVerif is semi-automated, it can search for proofs but requires human guidance for non-trivial protocols.

We use both ProVerif and CryptoVerif for their complementary strengths. CryptoVerif can prove stronger security properties of the protocol under precise cryptographic assumptions, but the proofs require more work. ProVerif can quickly analyze large protocols to automatically find attacks, but a positive result does not immediately provide a cryptographic proof of security. Deriving sound cryptographic proofs using symbolic analysis is still an open problem for real-world protocols [34].

A Realistic Threat Model for TLS. We seek to analyze TLS 1.3 for the above security goals against a rich threat model that includes both classic protocol adversaries as well as new ones that apply specifically to multi-mode protocols like TLS. In particular, we model recent downgrade attacks on TLS by allowing the use of weak cryptographic algorithms in older versions of TLS. In our analyses, the attacker can use any of the following attack vectors to disrupt the protocol.

- **Network Adversary:** As usual, we assume that the

attacker can intercept, modify, and send all messages sent on public network channels.

- **Compromised Principals:** The attacker can compromise any client or server principal P by asking for its long-term secrets, such as its private key (sk_P) or pre-shared key (psk). We do not restrict which principals can be compromised, but whenever such a compromise occurs, we mark it with a security event: `Compromised(pk_p)` or `CompromisedPSK(psk)`. If the compromise event occurs after a session is complete, we issue a different security event: `PostSessionCompromise(cid, pk_p)`.

- **Weak Long-term Keys:** If the client or server has a weak key that the attacker may be able to break with sufficient computation, we treat such keys the same way as compromised keys and we issue a more general event: `WeakOrCompromised(pk_p)`. This conservative model of weak keys is enough to uncover attacks like FREAK [12] that rely on the use of 512-bit RSA keys by TLS servers.

- **RSA Decryption Oracles:** TLS versions up to 1.2 use RSA PKCS#1 v1.5 encryption, which is known to be vulnerable to a form of padding oracle attack on decryption originally discovered by Bleichenbacher [28]. Although countermeasures to this attack have been incorporated into TLS, they remains hard to implement securely [59] resulting in continued attacks such as DROWN [7]. Furthermore, such padding oracles can sometimes even be converted to signature oracles for the corresponding private key [47].

We assume that any TLS server (at any version) that enables RSA decryption may potentially be vulnerable to such attacks. We distinguish between two kinds of RSA key exchange: `RSA(StrongRSAEncryption)` and `RSA(WeakRSAEncryption)`. In any session, if the server chooses the latter, we provide the attacker with a decryption and signature oracle for that private key.

- **Weak Diffie-Hellman Groups:** To account for attacks like Logjam [1], we allow servers to choose between strong and weak Diffie-Hellman groups (or elliptic curves), and mark the corresponding key exchange mode as `DHE(StrongDH)` or `DHE(WeakDH)`. We conservatively assume that weak groups have size 1, so all Diffie-Hellman exponentiations in these groups return the same distinguished element `BadElement`.

Even strong Diffie-Hellman groups typically have small subgroups that should be avoided. We model these subgroups by allowing a weak subgroup (of size 1) even within a strong group. A malicious client or server may choose `BadElement` as its public value, and then all exponentiations with this element as the base will also return `BadElement`. To avoid generating keys in this subgroup, clients and servers must validate the received public value.

- **Weak Hash Functions:** TLS uses hash functions for key derivation, HMAC, and for signatures. Versions up to TLS 1.2 use various combinations of MD5 and SHA-1, both of which are considered weak today, leading to exploitable attacks on TLS such as SLOTH [23].

We model both strong and weak hash functions, and the client and server get to negotiate which function they will use in signatures. Strong hash functions are treated as one-way functions in our symbolic model, whereas weak hash functions are treated as point functions that map all inputs to a constant value: `Collision`. Hence, in our model, it is trivial for the attacker to find collisions as well as second preimages for weak hash functions.

- **Weak Authenticated Encryption:** To model recent attacks on RC4 [3], [71] and TripleDES [22], we allow both weak and strong authenticated encryption schemes. For data encrypted with a weak scheme, irrespective of the key, we provide the adversary with a decryption oracle. A number of attacks on the TLS Record protocol stem from its use of a MAC-Encode-Encrypt construction for CBC-mode ciphersuites. This construction is known to be vulnerable to padding oracle attacks such as POODLE [60] and Lucky13 [4], and countermeasures have proved hard to implement correctly [2]. We model such attacks using a leaky decryption function. Whenever a client or server decrypts a message with this function, the function returns the right result but also leaks the plaintext to the adversary.

The series of threats described above comprise our conservative threat model for TLS 1.3, and incorporates entire classes of attacks that have been shown to be effective against older versions of the protocol, including Triple Handshake, POODLE, Lucky 13, RC4 NOMORE, FREAK, Logjam, SLOTH, DROWN. In most cases, we assume strictly stronger adversaries than have been demonstrated in practice, but since attacks only get better over time, our model seeks to be defensive against future attacks. It is worth noting that, even though TLS 1.3 does not itself support any weak ciphers, TLS 1.3 clients and servers will need to support legacy protocol versions for backwards compatibility. Our model enables a fine-grained analysis of vulnerabilities: we can ask whether TLS 1.3 connections between a client and a server are secure even if TLS 1.2 connections between them are broken.

Verifying TLS 1.2 in ProVerif. We encode our threat model as a generic ProVerif crypto library that can be used with any protocol. To evaluate this model, and in preparation for our analysis of TLS 1.3, we symbolically analyze a model of TLS 1.2 using ProVerif. Our model includes TLS 1.2 clients and servers that support both RSA and Diffie-Hellman key exchanges, and are willing to use both weak and strong cryptography. We assume that clients are unauthenticated.

We write ProVerif processes for TLS 1.2 clients and servers that exchange messages according to the protocol standard, and issue a sequence of events—`ClientOffers`, `ServerChooses`, `ClientFinished`, `ServerFinished`, `ClientSends`, `ServerReceives`—indicating their progress through the protocol. We then compose these processes with our threat model and add queries for message authenticity and secrecy. For example, a secrecy query may ask whether the attacker can learn some application data message m sent by the client over a TLS 1.2 session with identifier cid .

When we run ProVerif for this query, it finds a counter-

example: the attacker can learn m if it can compromise server’s private key (`WeakOrCompromised(pkS)`). To check whether this is the only case in which m is leaked, we refine the secrecy query and run ProVerif again. ProVerif again finds a counter-example: the attacker can learn m if the server chooses a weak Diffie-Hellman group (`ServerChoosesKex(DHE(WeakDH))`). In this way, we keep refining our queries until we obtain the strongest security properties that hold for TLS 1.2 in our model:

- **TLS 1.2 Secrecy:** A message m sent by an honest client in a session cid to a server S cannot be known to the adversary unless one of the following conditions holds:
 - (1) the server’s public key is weak or compromised, or
 - (2) the session uses a weak Diffie-Hellman group, or
 - (3) the session uses weak authenticated encryption, or
 - (4) the server uses weak RSA decryption with the same public key (in this or any other session), or
 - (5) the server uses a weak hash function for signing with the same public key (in any session).
- **TLS 1.2 Authenticity & Replay Protection:** Every message m accepted by an honest client in a session cid with some server S corresponds to a unique message sent by S on a matching session, unless one of the conditions (1)-(5) above holds.

Both these queries are verified by ProVerif in a few seconds. All the disjuncts (1)-(5) in these queries are necessary, removing any of them results in a counterexample discovered by ProVerif, corresponding to some well-known attack on badly configured TLS 1.2 connections.

Interestingly, the conditions (2) and (3) are session specific, that is, only the sessions where these weak constructions are used are affected. In contrast, (4) and (5) indicate that the use of weak RSA decryption or a weak hash function in any session affects all other sessions that use the same server public key. As we shall see, this has an impact on the security of TLS 1.3 when it is composed with TLS 1.2.

We can also verify our TLS 1.2 model for more advanced properties. Forward secrecy does not hold in general for TLS 1.2, but can be proved for DHE sessions that use strong groups. Channel identifiers like $cid = k_c$ are not unique, and ProVerif finds a variant of the Triple Handshake attack, unless we implement the recommended countermeasure [64].

Verification Effort. The work of verifying TLS 1.2 can be divided into three tasks. We first modeled the threat model as a 400 line ProVerif library, but this library can now be reused for other protocols, including TLS 1.3. We then modeled the TLS 1.2 protocol in about 200 lines of ProVerif. Finally, we wrote about 50 lines of queries, both to validate our model (e.g. checking that the protocol completes in the absence of an attacker) and to prove our desired security goals. Most of the effort is in formalizing, refining, and discovering the right security queries. Although ProVerif is fully automated, verification gets more expensive as the protocol grows more complex. So, as we extend our models to cover multiple modes of TLS 1.3 composed with TLS 1.2, we sometimes need to simplify or restructure our models to aid verification.




Figure 2: TLS 1.3 Draft-18 1-RTT Protocol (left) and Key Schedule (right). The protocol uses an (EC)DHE key exchange with server certificate authentication: client authentication and the `RetryRequest` negotiation steps are optional. The hash function $H()$ used in the key schedule is typically SHA-256, which has length $len_{H()} = 32$ bytes. The PSK-based key derivations in the key schedule are not used in the 1-RTT protocol here; they will be used later in Figure 4.

III. TLS 1.3 1-RTT: SIMPLER, FASTER HANDSHAKES

In its simplest form, TLS 1.3 consists of a Diffie-Hellman handshake, typically using an elliptic curve, followed by application data encryption using an AEAD scheme like AES-GCM. The essential structure of 1-RTT has remained stable since early drafts of TLS 1.3. It departs from the TLS 1.2 handshake in two ways. First, the key exchange is executed alongside the negotiation protocol so the client can start sending application data along with its second flight of messages (after one round-trip, hence 1-RTT), unlike

TLS 1.2 where the client had to wait for two message flights from the server. Second, TLS 1.3 eliminates a number of problematic features in TLS 1.2; it removes RSA key transport, weak encryption schemes (RC4, TripleDES, AES-CBC), and renegotiation; it requires group negotiation with strong standardized Diffie-Hellman groups, and it systematically binds session keys to the handshake log to prevent attacks like the Triple Handshake. In this section, we detail the protocol flow, we model it in ProVerif, and we analyze it alongside TLS 1.2 in the security model of §II.

1-RTT Protocol Flow. A typical 1-RTT connection in Draft 18 proceeds as shown in Figure 2. The first four messages form the negotiation phase. The client sends a `ClientHello` message containing a nonce n_C and an $offer_C$ that lists the versions, groups, hash functions, and authenticated encryption algorithms that it supports. For each group G that the client supports, it may include a Diffie-Hellman key share g^x . On receiving this message, the server chooses a $mode_S$ that fixes the version, group, and all other session parameters. Typically, the server chooses a group G for which the client already provided a public value, and so it can send its `ServerHello` containing a nonce n_S , $mode_S$ and g^y to the client. If none of the client’s groups are acceptable, the server may ask the client (via `RetryRequest`) to resend the client hello with a key share $g^{x'}$ for the server’s preferred group G' . (In this case, the handshake requires two round trips.)

Once the client receives the `ServerHello`, the negotiation is complete and both participants derive handshake encryption keys from $g^{x'y}$, one in each direction (k_c^h, k_s^h), with which they encrypt all subsequent handshake messages. The client and server also generate two MAC keys (k_c^m, k_s^m) for use in the `Finished` messages described below. The server then sends a flight of up to 5 encrypted messages: `Extensions` contains any protocol extensions that were not sent in the `ServerHello`; `CertRequest` contains an optional request for a client certificate; `Certificate` contains the server’s X.509 public-key certificate; `CertVerify` contains a signature with server’s private key sk_S over the log of the transcript so far (log_2); `Finished` contains a MAC with k_s^m over the current log (log_3). Then the server computes the 1-RTT traffic keys k_c, k_s and may immediately start using k_s to encrypt application data to the client.

Upon receiving the server’s encrypted handshake flight, the client verifies the certificate, the signature, and the MAC, and if all verifications succeed, the client sends its own second flight consisting of an optional certificate `Certificate` and signature `CertVerify`, followed by a mandatory `Finished` with a MAC over the full handshake log. Then the client starts sending its own application data encrypted under k_c . Once the server receives the client’s second flight, we consider the handshake complete and put all the session parameters into the local session databases at both client and server (\mathcal{C}, \mathcal{S}).

In addition to the traffic keys for the current session, the 1-RTT handshake generates two extra keys: ems is an exporter master secret that may be used by the application to bind authentication credentials to the TLS channel; psk' is a resumption master secret that may be used as a pre-shared key in future TLS connections between C and S .

The derivation of keys in the protocol follows a linear key schedule, as depicted on the right of Figure 2. The first version of this key schedule was inspired by OPTLS [52] and introduced into TLS 1.3 in Draft-7. The key idea in this design is to accumulate key material and handshake context into the derived keys using a series of HKDF invocations as the protocol progresses. For example, in connections that

use pre-shared keys (see §V), the key schedule begins by deriving es from psk , but after the `ServerHello`, we add in $g^{x'y}$ to obtain the handshake secret hs . Whenever we extract encryption keys, we mix in the current handshake log, in order to avoid key synchronization attacks like the Triple Handshake.

Since its introduction in Draft-7, the key schedule has undergone many changes, with a significant round of simplifications in Draft-13. Since all previously published analyses of 1-RTT predate Draft-13, this leaves open the question whether the current Draft-18 1-RTT protocol is still secure.

Modeling 1-RTT in ProVerif. We write client and server processes in ProVerif that implement the message sequence and key schedule of Figure 2.

Our models are abstract with respect to the message formats, treating each message (e.g. `ClientHello(...)`) as a symbolic constructor, with message parsing modeled as a pattern-match with this constructor. This means that our analysis assumes that message serialization and parsing is correct; it won’t find any attacks that rely on parsing ambiguities or bugs. This abstract treatment of protocol messages is typical of symbolic models; the same approach is taken by Tamarin [35]. In contrast, miTLS [21] includes a fully verified parser for TLS messages.

The key schedule is written as a sequence of ProVerif functions built using an HMAC function, $hmac(H, m)$, which takes a hash function H as argument and is assumed to be a one-way function as long as $H = \text{StrongHash}$. All other cryptographic functions are modeled as described in §II, with both strong and weak variants.

Persistent state is encoded using tables. To model principals and their long-term keys, we use a global private table that maps principals (A) to their key pairs ((sk_A, pk_A)). To begin with, the adversary does not know any of the private keys in this table, but it can compromise any principal and obtain her private key. As described in §II, this compromise is recorded in ProVerif by an event `WeakOrCompromised(pk_A)`.

As the client and server proceed through the handshake they record security events indicating their progress. We treat the negotiation logic abstractly; the adversary gets to choose $offer_C$ and $mode_S$, and we record these choices as events (`ClientOffers`, `ServerChooses`) at the client and server. When the handshake is complete, the client and server issue events `ServerFinished`, `ClientFinished`, and store their newly established sessions in two private tables `clientSession` and `serverSession` (corresponding to \mathcal{C} and \mathcal{S}). These tables are used by the record layer to retrieve the traffic keys k_c, k_s for authenticated encryption. Whenever the client or server sends or receives an application data message, it issues further events (`ClientSends`, `ServerReceives`, etc.) We use all these events along with the client and server session tables to state our security goals.

1-RTT Security Goals. We encode our security goals as ProVerif *queries* as follows:

- **Secrecy** for a message, such as m_1 , is encoded using an auxiliary process that asks the adversary to guess the

value of m_1 ; if the adversary succeeds, the process issues an event `MessageLeaked(cid, m_1)`. We then write a query to ask ProVerif whether this event is reachable.

- **Forward Secrecy** is encoded using the same query, but we explicitly leak the client and server’s long-term keys (sk_C, sk_S) at the end of the session cid . ProVerif separately analyzes pre-compromise and post-compromise sessions as different *phases*; the forward secrecy query asks that messages sent in the first phase are kept secret even from attackers who learn the long-term keys in the second phase.
- **Authentication** for a message m_1 received by the server is written as a query that states that whenever the event `ServerReceives(cid, m_1)` occurs, it must be preceded by three matching events: `ServerFinished(cid, \dots)`, `ClientFinished(cid, \dots)`, and `ClientSends(cid, m_1)`, which means that some honest client must have sent m_1 on a matching session. The authentication query for messages received by clients is similar.
- **Replay protection** is written as a stronger variant of the authentication query that requires *injectivity*: each `ServerReceives` event must correspond to a unique, matching, preceding `ClientSends` event.
- **Unique Channel Identifiers** are verified using another auxiliary process that looks up sessions from the `clientSession` and `serverSession` tables and checks that if the cid in both is the same, then all other parameters match. Otherwise it raises an event, and we ask ProVerif to prove that this event is not reachable.

When we first ask ProVerif to verify these queries, it fails and provides counterexamples; for example, client message authentication does not hold if the client is compromised `Compromised(pk_C)` or unauthenticated in the session. We then refine the query by adding this failure condition as a disjunct, and run ProVerif again and repeat the process until the query is proved. Consequently, our final verification results are often stated as a long series of disjuncts listing the cases where the desired security goal does not hold.

Verifying 1-RTT in Isolation. For our model of Draft-18 1-RTT, ProVerif can prove the following secrecy query about all messages ($m_{0.5}, m_1, m_2$):

- **1-RTT (Forward) Secrecy:** Messages m sent in a session between C and S are secret as long as the private keys of C and S are not revealed before the end of the session, and the server chooses a $mode_S$ with a strong Diffie-Hellman group, a strong hash function, and a strong authenticated encryption algorithm.

If we further assume that TLS 1.3 clients and servers only support strong algorithms, we can simplify the above query to show that all messages sent between uncompromised principals are kept secret. In the rest of this paper, we assume that TLS 1.3 only enables strong algorithms, but that earlier versions of the protocol may continue to support weak algorithms.

Messages m_1 from the client to the server enjoy strong authentication and protection from replays:

- **1-RTT Authentication (and Replay Prevention):** If a message m is accepted by S over a session with an honest

C , then this message corresponds to a unique message sent by the C over a matching session.

However the authentication guarantee for messages $m_{0.5}, m_1$ received by the client is weaker. Since the client does not know whether the server sent this data before or after receiving the client’s second flight, the client and server sessions may disagree about the client’s identity. Hence, for these messages, we can only verify a weaker property:

- **0.5-RTT Weak Authentication (and Replay Prevention):** If a message m is accepted by C over a session with an honest S , then this message corresponds to a unique message sent by S over a server session that matches all values in the client session except (possibly) the client’s public key pk_C , the resumption master secret psk' , and the channel identifier cid .

We note that by allowing the server to send 0.5-RTT data, Draft-18 has weakened the authentication guarantees for all data received by an authenticated client. For example, if a client requests personal data from the server over a client-authenticated 1-RTT session, a network attacker could delay the client’s second flight (`Certificate-Finished`) so that when the client receives the server’s 0.5-RTT data, it thinks that it contains personal data, but the server actually sent data intended for an anonymous client.

Verifying TLS 1.3 1-RTT composed with TLS 1.2. We combine our model with the TLS 1.2 model described at the end of §II so that each client and server supports both versions. We then ask the same queries as above, but only for sessions where the server chooses TLS 1.3 as the version in $mode_S$. Surprisingly, ProVerif finds two counterexamples.

First, if a server supports `WeakRSADecryption` with RSA key transport in TLS 1.2, then the attacker can use the RSA decryption oracle to forge TLS 1.3 server signatures and hence break our secrecy and authentication goals. This attack found by ProVerif directly corresponds to the cross-protocol Bleichenbacher attacks described in [47], [7]. It shows that removing RSA key transport from TLS 1.3 is not enough, one must disable the use of TLS 1.2 RSA mode on any server whose certificate may be accepted by a TLS 1.3 client.

Second, if a client or server supports a weak hash function for signatures in TLS 1.2, then ProVerif shows how the attacker can exploit this weakness to forge TLS 1.3 signatures in our model, hence breaking our security goals. This attack corresponds to the SLOTH transcript collision attack on TLS 1.3 signatures described in [23]. To avoid this attack, TLS 1.3 implementations must disable weak hash functions in all supported versions, not just TLS 1.3.

After disabling these weak algorithms in TLS 1.2, we can indeed prove all our expected security goals about Draft-18 1-RTT, even when it is composed with TLS 1.2.

We may also ask whether TLS 1.3 clients and servers can be downgraded to TLS 1.2. If such a version downgrade takes place, we would end up with a TLS 1.2 session, so we need to state the query in terms of sessions where $mode_S$ contains TLS 1.2. ProVerif finds a version downgrade attack on a TLS 1.3 session, if the client and server support weak Diffie-Hellman groups in TLS 1.2. This attack closely mirrors the flaw described in [14]. Draft-13 introduced a

countermeasure in response to this attack, and we verify that by adding it to the model, the downgrade attack disappears.

Although our models of TLS 1.3 and 1.2 are individually verified in a few seconds each, their composition takes several minutes to analyze. As we add more features and modes to the protocol, ProVerif takes longer and requires more memory. Our final composite model for all modes of TLS 1.3+1.2 takes hours on a powerful workstation.

IV. 0-RTT WITH SEMI-STATIC DIFFIE-HELLMAN

In earlier versions of TLS, the client would have to wait for two round-trips of handshake messages before sending its request. 1-RTT in TLS 1.3 brings this down to one round trip, but protocols like QUIC use a "zero-round-trip" (0-RTT) mode, by relying on a *semi-static* (long-term) Diffie-Hellman key. This design was adapted for TLS in the OPTLS proposal [52] and incorporated in Draft-7 (along with a fix we proposed, as described below).

Protocol Flow. The protocol is depicted in Figure 3. Each server maintains a Diffie-Hellman key pair (s, g^s) and publishes a signed server configuration containing g^s . As usual, a client initiates a connection with a `ClientHello` containing its ephemeral key g^x . If a client has already obtained and cached the server's certificate and signed configuration (in a prior exchange for example), then the client computes a shared secret $g^{x \cdot s}$ and uses it to derive an initial set of shared keys which can then immediately be used to send encrypted data. To authenticate its 0-RTT data, the client may optionally send a certificate and a signature over the client's first flight.

The server then responds with a `ServerHello` message that contains a fresh ephemeral public key g^y . Now, the client and server can continue with a regular 1-RTT handshake using the new shared secret $g^{x \cdot y}$ in addition to $g^{x \cdot s}$.

The 0-RTT protocol continued to evolve from Draft-7 to Draft-12, but in Draft-13, it was removed in favor of a PSK-based 0-RTT mode. Even though Diffie-Hellman-based 0-RTT no longer exists in Draft-18, we analyze its security in this section, both for posterity and to warn protocol designers about the problems they should watch out for if they decide to reintroduce DH-based 0-RTT in a future version of TLS.

Verification with ProVerif. We modeled the protocol in ProVerif and wrote queries to check whether the 0-RTT data m_0 is (forward) secret and authentic. ProVerif is able to prove secrecy but finds that m_0 is not forward secret if the semi-static key s is compromised once the session is over. ProVerif also finds a Key Compromise Impersonation attack on authentication: if g^s is compromised, then an attacker can forge 0-RTT messages from C to S . Furthermore, the 0-RTT flight can be replayed by an attacker and the server will process it multiple times, thinking that the client has initiated a new connection each time. In addition to these three concerns, which were documented in Draft-7, ProVerif also finds a new attack, explained below, that breaks 0-RTT authentication if the server's certificate is not included in the 0-RTT client signature.

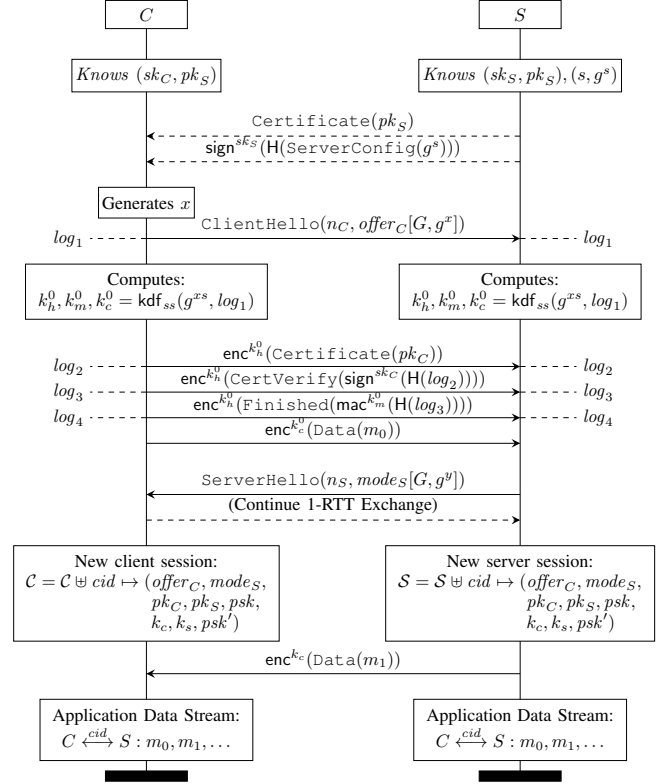


Figure 3: DH-based 0-RTT in TLS 1.3 Draft-12, inspired by QUIC and OPTLS.

Unknown Key Share Attack on DH-based 0-RTT in QUIC, OPTLS, and TLS 1.3. We observe that in the 0-RTT protocol, the client starts using g^s without having any proof that the server knows s . So a dishonest server M can claim to have the same semi-static key as S by signing g^s under its own key sk_M . Now, suppose a client connects to M and sends its client hello and 0-RTT data; M can simply forward this whole flight to S , which may accept it, because the semi-static keys match. This is an *unknown key share* (UKS) attack where C thinks it is talking to M but it is, in fact, connected to S .

In itself, the UKS attack is difficult to exploit, since M does not know $g^{x \cdot s}$ and hence cannot decrypt or tamper with messages between C and S . However, if the client authenticates its 0-RTT flight with a certificate, then M can forward C 's certificate (and C 's signature) to S , resulting in a *credential forwarding* attack, which is much more serious. Suppose C is a browser that has a page open at website M ; from this page M can trigger any authenticated 0-RTT HTTPS request m_0 to its own server, which then uses the credential forwarding attack to forward the request to S , who will process m_0 as if it came from C . For example, M may send a POST request that modifies C 's account details at S .

The unknown key share attack described above applies to both QUIC and OPTLS, but remained undiscovered despite several security analyses of these protocols [42], [56], [52],

because these works did not consider client authentication, and hence did not formulate an authentication goal that exposed the flaw. We informed the authors of QUIC and they acknowledged our attack. They now recommend that users who need client authentication should not use QUIC, and should instead move over to TLS 1.3. We also informed the authors of the TLS 1.3 standard, and on our suggestion, Draft-7 of TLS 1.3 included a countermeasure for this attack: the client signature and 0-RTT key derivation include not just the handshake log but also the cached server certificate. With this countermeasure in place, ProVerif proves authentication for 0-RTT data.

V. PRE-SHARED KEYS FOR RESUMPTION AND 0-RTT




Figure 4: TLS 1.3 Draft-18 PSK-based Resumption and 0-RTT.

Aside from the number of round-trips, the main cryptographic cost of a TLS handshake is the use of public-key algorithms for signatures and Diffie-Hellman, which are still significantly slower than symmetric encryption and MACs. So, once a session has already been established between a client and server, it is tempting to reuse the symmetric session key established in this session as a pre-shared symmetric key in new connections. This mechanism is called *session resumption* in TLS 1.2 and is widely used in HTTPS where a single browser typically has many parallel and sequential connections to the same website. In TLS 1.2, pre-shared keys (PSKs) are also used instead of certificates by resource-constrained devices that cannot afford public-key encryption. TLS 1.3 combines both these use-cases in a single PSK-based handshake mode that combines resumption, PSK-only handshakes, and 0-RTT.

Protocol Flow. Figure 4 shows how this mode extends the regular 1-RTT handshake; in our analysis, we only consider

PSKs that are established within TLS handshakes, but similar arguments apply to PSKs that are shared out-of-band. We assume that the client and server have established a pre-shared key psk in some earlier session. The client has cached psk , but in order to remain state-less, the server has given the client a ticket containing psk encrypted under an encryption key k_t . As usual, the client sends a `ClientHello` with its ephemeral key share g^x and indicates that it prefers to use the shared PSK psk . To prove its knowledge of psk and to avoid certain attacks (described below), it also MACs the `ClientHello` with a *binder* key k^b derived from the psk . The client can then use psk to already derive an encryption key for 0-RTT data m_0 and start sending data without waiting for the server’s response. When the server receives the client’s flight, it can choose to accept or reject the offered psk . Even if it accepts the psk , the server may choose to reject the 0-RTT data, it may choose to skip certificate-based authentication, and (if it does not care about forward secrecy) it may choose to skip the Diffie-Hellman exchange altogether. The recommended mode is PSK-DHE, where psk and g^{xy} are both mixed into the session keys. The server then sends back a `ServerHello` with its choice and the protocol proceeds with the appropriate 1-RTT handshake and completes the session.

Verifying PSK-based Resumption. We first model the PSK-DHE 1-RTT handshake (without certificate authentication) and verify that it still meets our usual security goals:

- **PSK-DHE 1-RTT (Forward) Secrecy** Any message m sent over a PSK-DHE session in 1-RTT is secret as long as the PSK psk and the ticket encryption key k_t are not compromised until the end of the session.
- **PSK-DHE 1-RTT Authentication and Replay Protection** Any message m received over a PSK-DHE session in 1-RTT corresponds to a unique message sent by a peer over a matching session (notably with the same psk) unless psk or k_t are compromised during the session.
- **PSK-DHE 1-RTT Unique Channel Identifier** The values psk' , ems , and $H(log_7)$ generated in a DHE or PSK-DHE session are all unique channel identifiers.

Notably, data sent over PSK-DHE is forward secret even if the server’s long term ticket encryption key k_t is compromised after the session. In contrast, pure PSK handshakes do not provide this forward secrecy.

The authentication guarantee requires that the client and server must agree on the value of the PSK psk , and if this PSK was established in a prior session, then the unique channel identifier property says that the client and server must transitively agree on the prior session as well. An earlier analysis of Draft-10 in Tamarin [35] found a violation of the authentication goal because the 1-RTT client signature in Draft-10 did not include the server’s `Finished` or any other value that was bound to the PSK. This flaw was fixed in Draft-11 and hence we are able to prove authentication for Draft-18.

Verifying PSK-based 0-RTT. We extend our model with the 0-RTT exchange and verify that m_0 is authentic and secret. The strongest queries that ProVerif can prove are the

following:

- **PSK-based 0-RTT (Forward) Secrecy** A message m_0 sent from C to S in a 0-RTT flight is secret as long as psk and k_t are never compromised.
- **PSK-based 0-RTT Authentication** A message m_0 received by S from C in a 0-RTT flight corresponds to some message sent by C with a matching `ClientHello` and matching psk , unless the psk or k_t are compromised.

In other words, PSK-based 0-RTT data is not forward secret and is vulnerable to replay attacks. As can be expected, it provides a symmetric authentication property: since both C and S know the psk , if either of them is compromised, the attacker can forge 0-RTT messages.

An Attack on 0-RTT Client Authentication. Up to Draft-12, the client could authenticate its 0-RTT data with a client certificate in addition to the PSK. This served the following use case: suppose a client and server establish an initial 1-RTT session (that outputs psk') where the client is unauthenticated. Some time later, the server asks the client to authenticate itself, and so they perform a PSK-DHE handshake (using psk') with client authentication. The use of psk' ensures continuity between the two sessions. In the new session, the client wants to start sending messages immediately, and so it would like to use client authentication in 0-RTT.

To be consistent with Draft-12, suppose we remove the outer binder MAC (using k^b) on the `ClientHello` in Figure 4, and we allow client authentication in 0-RTT. Then, if we model this protocol in ProVerif and ask the 0-RTT authentication query again, ProVerif finds a credential forwarding attack, explained next.

Suppose a client C shares psk with a malicious server M , and M shares a different psk' with an honest server S . If C sends an authenticated 0-RTT flight (certificate, signature, data m_0) to M , M can decrypt this flight using psk , re-encrypt it using psk' , and forward the flight to S . S will accept the authenticated data m_0 from C as intended for itself, whereas C intended to send it only to M . In many HTTPS scenarios, as discussed in §IV, M may be able to control the contents of this data, so this attack allows M to send arbitrary requests authenticated by C to S .

This attack was not discovered in previous analyses of TLS 1.3 since many of them did not consider client authentication; the prior Tamarin analysis [35] found a similar attack on 1-RTT client authentication but did not consider 0-RTT client authentication. The attacks described here and in [35] belong to a general class of *compound authentication* vulnerabilities that appear in protocols that compose multiple authentication credentials [17]. In this case, the composition of interest is between PSK and certificate-based authentication. We found a similar attack on 1-RTT server authentication in pure PSK handshakes.

In response to our attack, Draft-13 included a `resumption_context` value derived from the psk in the handshake hash, to ensure that the client's signature over the hash cannot be forwarded on another connection (with a different psk'). This countermeasure has since evolved to

the MAC-based design showed in Figure 4, which has now been verified in this paper.

The Impact of Replay on 0-RTT and 0.5-RTT. It is now widely accepted that asynchronous messaging protocols like 0-RTT cannot be easily protected from replay, since the recipient has no chance to provide a random nonce that can ensure freshness. QUIC attempted to standardize a replay-prevention mechanism but it has since abandoned this mechanism, since it cannot prevent attackers from forcing the client to resend 0-RTT data over 1-RTT [66].

Instead of preventing replays, TLS 1.3 Draft-18 advises applications that they should only send non-forward-secret and idempotent data over 0-RTT. This recommendation is hard to systematically enforce in flexible protocols like HTTPS, where all requests have secret cookies attached, and even GET requests routinely change state.

We argue that replays offer an important attack vector for 0-RTT and 0.5-RTT data. If the client authenticates its 0-RTT flight, then an attacker can replay the entire flight to mount *authenticated replay* attacks. Suppose the (client-authenticated) 0-RTT data asks the server to send a client's bank statement, and the server sends this data in a 0.5-RTT response. An attacker who observes the 0-RTT request once, can replay it any number of times to the server from anywhere in the world and the server will send it the user's (encrypted) bank statement. Although the attacker cannot complete the 1-RTT handshake or read this 0.5-RTT response, it may be able to learn a lot from this exchange, such as the length of the bank statement, and whether the client is logged in.

In response to these concerns, client authentication has now been removed from 0-RTT. However, we note that similar replay attacks apply to 0-RTT data that contains an authentication cookie or OAuth token. We highly recommend that TLS 1.3 servers should implement a replay cache (based on the client nonce n_C and the ticket age) to detect and reject replayed 0-RTT data. This is less practical in server farms, where time-based replay mitigation may be the only alternative.

VI. COMPUTATIONAL ANALYSIS OF TLS 1.3 DRAFT-18

Our ProVerif analysis of TLS 1.3 Draft-18 identifies the necessary conditions under which the symbolic security guarantees of the protocol hold. We now use the tool CryptoVerif [24] to see whether these conditions are sufficient to obtain cryptographic security proofs for the protocol in a more precise computational model. In particular, under the assumption that the algorithms used in TLS 1.3 Draft-18 satisfy certain strong cryptographic assumptions, we prove that the protocol meets our security goals.

Proofs in the computational model are hard to mechanize, and CryptoVerif offers less flexibility and automation than ProVerif. To obtain manageable proofs, we focus only on TLS 1.3 (we do not consider TLS 1.2) and we ignore downgrade attacks. We split the protocol into three pieces and prove them separately using CryptoVerif, before composing them manually to obtain a proof for the full protocol.

A. Cryptographic Assumptions

We make the following assumptions about the cryptographic algorithms supported by TLS 1.3 clients and servers.

Diffie-Hellman. We assume that the Diffie-Hellman groups used in TLS 1.3 satisfy the gap Diffie-Hellman (GDH) assumption [61]. This assumption means that given g , g^a , and g^b for random a, b , the adversary has a negligible probability to compute g^{ab} , even when the adversary has access to a decisional Diffie-Hellman oracle, which tells him given G, X, Y, Z whether there exist x, y such that $X = G^x$, $Y = G^y$, and $Z = G^{xy}$.

In our proof, we require GDH rather than the weaker decisional Diffie-Hellman (DDH) assumption, in order to prove secrecy of keys on the server side as soon as the server sends its `Finished` message: at this point, if the adversary controls a certificate accepted by the client, he can send its own key share y' to the client to learn information on $g^{x'y'}$, and that would be forbidden under DDH. We also require that $x^y = x'^y$ implies $x = x'$ and that $x^y = x^y'$ implies $y = y'$, which holds when the considered Diffie-Hellman group is of prime order. This is true for all groups currently specified in TLS 1.3, and our proof requires it for all groups included in the future.

We also assume that all Diffie-Hellman group elements have a binary representation different from $0^{len_H()}$. This assumption simplifies the proof by avoiding a possible confusion between handshakes with and without Diffie-Hellman exchange. Curve25519 does have a 32-byte zero element, but excluding zero Diffie-Hellman shared values is already recommended to avoid points of small order [54].

Finally, we assume that all Diffie-Hellman group elements have a binary representation different from $len_H() || \text{"TLS 1.3,"} || z || h || 0x01$. This helps ease our proofs by avoiding a collision between $hkdf\text{-extract}(es, e)$ and $derive\text{-secret}(es, pbk, \text{""})$ or $derive\text{-secret}(es, ets_c, log_1)$. This assumption holds with the currently specified groups and labels, since group elements have a different length than the bitstring above. The technical problem identified by our assumption was independently discovered and discussed on the TLS mailing list [67], and has led to a change in Draft-19 which will make this assumption unnecessary.

Signatures. We assume that the function `sign` is unforgeable under chosen-message attacks (UF-CMA) [43]. This means that an adversary with access to a signature oracle has a negligible probability of forging a signature for a message not signed by the signature oracle. Only the oracle has access to the signing key; the adversary has the public key.

Hash Functions. We assume that the function `H` is collision-resistant [36]: the adversary has a negligible probability of finding two different messages with the same hash.

HMAC. We need two assumptions on HMAC-H:

We require that the functions $x \mapsto \text{HMAC-H}^{0^{len_H()}}(x)$ and $x \mapsto \text{HMAC-H}^{kdf_0}(x)$ are independent random oracles, in order to justify the use of HMAC-H as a randomness extractor in the HKDF construct. This assumption can itself be justified as follows. Assuming that the compression func-

tion underlying the hash function is a random oracle, Theorem 4.4 in [38] shows that HMAC is indistinguishable [33] from a random oracle, provided the MAC keys are less than the block size of the hash function minus one, which is true for HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512. It is then easy to show that $x \mapsto \text{HMAC-H}^{0^{len_H()}}(x)$ and $x \mapsto \text{HMAC-H}^{kdf_0}(x)$ are indistinguishable from independent random oracles in this case.

We assume that HMAC-H is a pseudo-random function (PRF) [9], that is, HMAC-H is indistinguishable from a random function provided its key is random and used only in HMAC-H, when the key is different from $0^{len_H()}$ and kdf_0 . We avoid these two keys to avoid confusion with the two random oracles above. Since keys are chosen randomly with uniform probability from a set `key` (with cardinality $|\text{key}|$), the only consequence of avoiding these keys is that $\frac{2}{|\text{key}|}$ is added to the probability of breaking the PRF assumption.

Authenticated Encryption. The authenticated encryption scheme is IND-CPA (indistinguishable under chosen plaintext attacks) and INT-CTXT (ciphertext integrity) [11], provided the same nonce is never used twice with the same key. IND-CPA means that the adversary has a negligible probability of distinguishing encryptions of two distinct messages of the same length that it has chosen. INT-CTXT means that an adversary with access to encryption and decryption oracles has a negligible probability of forging a ciphertext that decrypts successfully and has not been returned by the encryption oracle.

B. Verifying 1-RTT Handshakes without Pre-Shared Keys

To prove the security of TLS 1.3 in `CryptoVerif`, we first establish some lemmas about the primitives, as detailed in Appendix A. Then, we split the protocol into three parts, as shown in Figure 5, and verify them in sequence, before composing them by hand into a proof for the full protocol. This modular hybrid approach allows us to have proofs of manageable complexity, and to obtain results even when keys are reused many times, such as when several PSK-based resumptions are performed, which would otherwise be out of scope of `CryptoVerif`.

We first consider the initial 1-RTT handshake shown in Figure 2, until the new client and server session boxes. We model a honest client and a honest server, which are willing to interact with each other, but also with dishonest clients and servers included in the adversary. We do not consider details of the negotiation (or the `RetryRequest` message). We give the handshake keys (k_c^h and k_s^h) to the adversary, and let it encrypt and decrypt the handshake messages, so our security proof does not rely on the encryption of the handshake.

We assume that the server is always authenticated and consider both the handshake with and without client authentication. The honest client and server may be compromised at any time: the secret key of the compromised participant is then sent to the adversary, and the compromise is recorded by defining a variable `corruptedClient` or `corruptedServer`.

The outputs of this protocol are the application traffic secrets ats_c and ats_s (the derivation of the keys k_c and k_s from these secrets is left for the record protocol), the exporter master secret ems , and the resumption master secret psk' (later used as pre-shared key). CryptoVerif proves the following properties:

- **Key Authentication:** If the client terminates a session with the server and the server is not compromised, then the server has accepted a session with the client, and they share the same parameters: the keys ats_c , ats_s , and ems and all messages sent in the protocol until the server `Finished` message. (We can make no claim on the client `Finished` message because it has not been received by the server at this point, nor on psk' because it depends on the client `Finished` message.)

In our CryptoVerif model, we formalize this property by adding an event `ClientTerm(...)` in the client, executed when the client terminates a session (that is, sends his `Finished` message) with an honest server (that is, `corruptedServer` is not defined). We similarly define an event `ServerAccept(...)` at the server, executed when the server accepts a session (that is, sends his `Finished` message). The arguments of these events include the session keys and all the messages sent in the protocol until the server `Finished` message. We then ask CryptoVerif or prove an authentication query that states that, with overwhelming probability, each execution of event `ClientTerm` corresponds to a distinct execution of event `ServerAccept` with the same arguments.

Conversely, if a server terminates a session with an honest client, and either the client is authenticated and not compromised, or the client key share $g^{x'}$ accepted by the server was generated by the client, then the client must have accepted a session with the server, and they must agree on the established keys and on all messages sent in the protocol. We state this property as a CryptoVerif query and verify it.

- **Replay Prevention:** The authentication properties stated above are already injective, that is, they guarantee that each session of the client (resp. server) corresponds to a distinct session of the server (resp. client), and consequently, they forbid replay attacks.
- **(Forward) Secrecy of Keys:** The keys ats_c , ats_s , ems , and psk' exchanged in several protocol sessions are indistinguishable from independent fresh random values. This property means for instance that the keys psk' remains secret (indistinguishable from independent fresh random values) even if ats_c , ats_s , ems are given to the adversary, and similarly for the other keys. Secrecy holds on the client side when the server is not compromised before the end of the session. It holds on the server side when the client is authenticated and not compromised before the end of the session or when the key share $g^{x'}$ used by the server comes from the client. We prove secrecy of ats_c , ats_s , and ems on the server side when the key share $g^{x'}$ comes from the client as soon as the server sends its `Finished` message. This property allows us to prove security of 0.5-RTT messages by composition with the

record protocol.

- **Unique Channel Identifier:** When cid is psk' or $H(\log_7)$, we do not use CryptoVerif as the result is immediate: if a client session and a server session have the same cid , then these sessions have the same \log_7 by collision-resistance of H (which implies collision-resistance of HMAC-H), so all their parameters are equal.

When cid is ems , collision-resistance just yields that the client and server sessions have the same \log_4 . CryptoVerif proves that, if a client session and a server session both terminate successfully with the same \log_4 , then they have the same \log_7 and the same keys, so all their parameters are equal.

We need to guide CryptoVerif in order to prove these properties, with the following main steps. We first apply the security of the signature under the server key sk_s . We introduce tests to distinguish cases, depending on whether the Diffie-Hellman share received by the server is a share $g^{x'}$ from the client, and whether the Diffie-Hellman share received by the client is the share g^y generated by the server upon receipt of $g^{x'}$. Then we apply the random oracle assumption on $x \mapsto \text{HMAC-H}^{\text{kdf}_0}(x)$, replace variables that contain $g^{x'y}$ with their values to make equality tests $m = g^{x'y}$ appear, and apply the gap Diffie-Hellman assumption. At this point, the handshake secret hs is a fresh random value. We use the properties on the key schedule established in Appendix A to show that the other keys are fresh random values, and apply the security of the MAC and of the signature under the client key sk_c .

C. Verifying Handshakes with Pre-Shared Keys

We now analyze the handshake protocol in Figure 4, up until the new client and server sessions are established. The protocol begins with 0-RTT and continues on to 1-RTT. We consider both variants of PSK-based 1-RTT, with and without Diffie-Hellman exchange.

We ignore the ticket $\text{enc}^{k_t}(psk)$ and consider a honest client and a honest server that initially share the pre-shared key psk . Dishonest clients and servers may be included in the adversary. As in the previous section, we give the handshake keys (k_c^h and k_s^h) to the adversary and ignore handshake encryption. Certificates for the client and server are optional, since they are already authenticated via the psk ; we do not rely on authentication in our proofs and consider that the adversary performs the signature and verification operations on certificates if they occur.

The outputs of this protocol are the client early traffic secret ets_c (the derivation of the key k_c from ets_c is left for the record protocol), the application traffic secrets ats_c and ats_s , the exporter master secret ems , and the resumption master secret psk' . We run CryptoVerif on our model to obtain the following verification results:

- **Key Authentication:** CryptoVerif shows the same authentication properties as for the handshake without pre-shared key, assuming that both participants are uncompromised. Notably, however, CryptoVerif cannot prove authentication of ets_c . While the binder $\text{mac}^{k_b}(\dots)$ authenticates most of the client `ClientHello` message, the client

may offer several pre-shared keys and send a binder for each of these keys. Only the binder for the pre-shared key selected by the server is checked. Hence the adversary may alter another of the proposed binders, yielding a different log_1 and a different ets_c on the server side. This is not a serious attack, as the record protocol will fail if ets_c does not match on the client and server sides.

- **Replay Prevention:** CryptoVerif proves that all the authentication properties shown above are injective, thus showing the absence of replays for ats_c , ats_s , and ems . However, CryptoVerif cannot prove replay protection for the 0-RTT session key ets_c , and indeed the client ClientHello message can be replayed, yielding the same key ets_c for several sessions of the server even though there is a single session of the client.
- **Secrecy of Keys:** The keys ets_c , ats_c , ats_s , ems , and psk' exchanged in several protocol sessions are indistinguishable from independent fresh random values. Secrecy holds both on the client side and on the server side except that, on the server side, the keys ets_c are not independent of each other since an adversary may force the server to accept several times the same key ets_c by replaying the client ClientHello message. We prove the secrecy of ats_c , ats_s , and ems on the server side as soon as the server sends its Finished message.
- **Forward Secrecy:** CryptoVerif is unable to prove secrecy of the keys when psk is compromised after the end of the session, even assuming that hkdf-extract is a random oracle. Secrecy obviously does not hold in this case for the handshake without Diffie-Hellman exchange. We believe that it still holds for the handshake with Diffie-Hellman exchange; our failure to prove it in this case is due to the current limitations of CryptoVerif.
- **Unique Channel Identifier:** We proceed as in the handshake without pre-shared key. We additionally notice that, if a client session and a server session have the same log_7 , then they have the same psk . Indeed, by collision-resistance of $mac = HMAC-H$, they have the same k^b , so the same es , so the same psk .

D. Verifying the Record Protocol

The third component of TLS 1.3 is the record protocol that encrypts and decrypts messages after the new client and server sessions have been established in Figures 2 and 4.

In our model, we assume that the client and server share a fresh random traffic secret. We generate an encryption key and an initialization vector (IV), and send and receive encrypted messages using those key and IV, and a counter that is distinct for each message. (Our model is more detailed than the symbolic presentation given in the figures as we consider the IV and the counter.) We also generate a new traffic secret as specified in the key update mechanism of TLS 1.3 Draft-18 (Section 7.2). CryptoVerif proves the following properties automatically:

- **Key and Message Secrecy:** CryptoVerif proves that the updated traffic secret is indistinguishable from a fresh random value. It also proves that, when the adversary provides two sets of plaintexts m_i and m'_i of the same

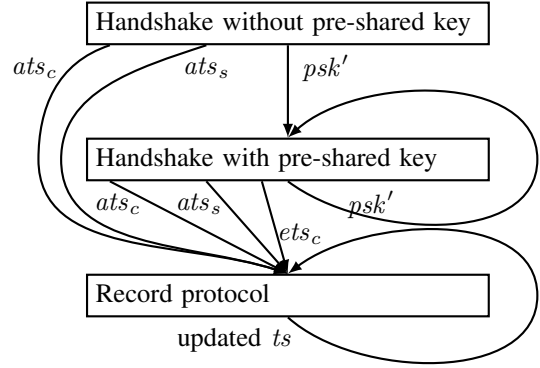


Figure 5: Structure of the CryptoVerif proof

padded length, it is unable to determine which of two sets is encrypted, even when the updated traffic secret is leaked.

- **Message Authentication:** CryptoVerif proves that, if a message m is decrypted by the receiver with a counter c , then the message m has been encrypted and sent by an honest sender with the same counter c .
- **Replay Prevention:** The authentication property above is injective, that is, any sent application data may be accepted at most once by the receiver.

E. A Composite Proof for TLS 1.3 Draft-18

We compose these results using a hybrid argument (as in [40]). Figure 5 summarizes the structure of the composition; more details are given in the full version [13].

First, we use the secrecy property of the initial handshake to replace all session keys with independent fresh random values. We rely on authentication and replay prevention to show that the same replacement is performed in matching sessions of the client and server.

Then, we use the security properties of the record protocol using ats_c and ats_s as traffic secrets, to obtain secrecy, forward secrecy (with respect to the compromise of sk_S and sk_C), authentication, and replay prevention for application messages in both directions. The security of the record protocol also shows that the updated traffic secrets generated during subsequent key updates preserve these properties.

Using the key psk' provided by the initial handshake, we then apply the security of the PSK-based handshake, to obtain that the keys ets_c , ats_c , ats_s , and psk' provided by this handshake are independent fresh random values. (The forward secrecy property of the initial handshake allows us to leak the keys sk_S and sk_C , so that the adversary can indeed perform the signature operations related to certificates, as we assumed in our model of handshakes with pre-shared keys.) We then apply the security of the record protocol to ats_c and ats_s , as above, for 1-RTT messages. We also apply it to ets_c for 0-RTT messages, but since the handshake does not prevent replays for this key, the composition will not prevent replays for messages sent under this key.

Finally, we apply the security of the PSK-based handshake again to the newly obtained psk' , hence obtaining composite security for arbitrary sequences of PSK-based resumptions.

VII. REF-TLS: A REFERENCE TLS 1.3 IMPLEMENTATION WITH A VERIFIED PROTOCOL CORE

In today’s web ecosystem, TLS is used by wide variety of client and server applications to establish secure channels across the Internet. For example, Node.js servers are written in JavaScript and can accept HTTPS connections using a Node’s builtin `https` module that calls OpenSSL. Popular desktop applications, such as WhatsApp messenger, are also written in JavaScript using the Electron framework (which combines Node.js with the Chromium rendering engine); they connect to servers using the same `https` module.

Our goal is to develop a high-assurance reference implementation of TLS 1.3, called RefTLS, that can be seamlessly used by Electron apps and Node.js servers. We want our implementation to be small, easy to read and analyze, and effective as an early experimental version of TLS 1.3 that real-world applications can use to help them transition to TLS 1.3, before it becomes available in mainstream libraries like OpenSSL. Crucially, we want to be able to verify the security of the core protocol code in RefTLS, and show that it avoids both protocol-level attacks as well as implementation bugs in its protocol state machine [12].

In this section, we describe RefTLS and evaluate its progress towards these goals. RefTLS has been used as a prototype implementation of TLS Draft-13 to Draft-18, interoperating with other early TLS 1.3 libraries. Its protocol core has been symbolically analyzed with ProVerif, and it has been successfully integrated into Electron applications.

Flow and ProScript. RefTLS is written in Flow [32], a typed variant of JavaScript. Static typing in Flow guarantees the absence of a large class of classic JavaScript bugs, such as reading a missing field in an object. Consequently, our code looks very much like a program in a typed functional language like OCaml or F#. We would like to verify the security of all our Flow code, but since Flow is a fully-fledged programming language, it has loops, mutable state, and many other features that are hard to automatically verify.

In earlier work, we developed a typed subset of JavaScript called ProScript [48] that was designed for writing cryptographic protocol code that could be compiled automatically to ProVerif. ProScript is also a subset of Flow and so we can reuse its ProVerif compiler to extract symbolic models from the core protocol code in RefTLS, if we write it carefully.

ProScript code is written defensively, in that it cannot, even accidentally, access external libraries or extensible JavaScript functionalities such as object instantiation, or redefinable properties such as `Array.split`. These restrictions are necessary in JavaScript where external functions can completely redefine the behavior of all libraries and object prototypes. The resulting style enforces syntactic scoping and strict type checking for all variables and functions, and disallows implicit coercions, object prototype access, and dynamic extensions of arrays and objects.

For ease of analysis, ProScript disallows loops, recursion, and only allows access mutable state through a well defined `table` interface. These are significant restrictions, but as we show, the resulting language is still expressive enough to write the core composite protocol code for TLS 1.0-1.3.




Figure 6: RefTLS Architecture. The library is written in Flow, a typed subset of JavaScript. The protocol core is verified by translation to ProVerif. The cryptographic library, message formatting and parsing, and the runtime framework are trusted. The application and parts of the RefTLS library are untrusted (assumed to be adversarial in our model).

Implementation Structure. Figure 6 depicts the architecture of RefTLS and shows how it can be safely integrated into larger, unverified and untrusted applications. At the top, we have Node.js and Electron applications written in JavaScript. RefTLS exposes an interface to these applications that exactly matches that of the default Node.js `https` module (which uses OpenSSL), allowing these applications to transparently use RefTLS instead of OpenSSL.

The RefTLS code itself is divided into untrusted Flow code that handles network connections and implements the API, a verified protocol module, written in ProScript, and some trusted but unverified Flow code for parsing and serializing TLS messages. All this code is statically typechecked in Flow. The core protocol module, called RefTLS-CORE, implements all the cryptographic operations of the protocol. It exposes an interface that allows RefTLS to drive the protocol, but hides all keying material and sensitive session state within the core module. This isolation is currently implemented via the Node module system; but we can also exploit Electron’s multi-threading feature in order to provide thread-based isolation to the RefTLS-CORE module, allowing it to only be accessed through a pre-defined RPC interface. Strong isolation for RefTLS-CORE allows us to verify it without relying on the correctness of the rest of the RefTLS codebase.

However, RefTLS still relies on the security and correctness of the crypto library and the underlying Electron, Node.js, and JavaScript runtimes. In the future, we may be able to reduce this trusted computing base by relying on verified crypto [73], verified JavaScript interpreters [29], and least-privilege architectures, such as ESpectro [69], which can control access to dangerous libraries from JavaScript.

A Verified Protocol Core. In RefTLS-CORE, we develop, implement and verify (for the first time) a composite state machine for TLS 1.2 and 1.3 (shown in Appendix B). Each state transition is implemented by a ProScript function that processes a flight of incoming messages, changes the session state, and produces a flight of outgoing messages. For

TLS 1.3 clients, these functions are `get_client_hello`, `put_server_hello`, and `put_server_finished`; servers use the functions `put_client_hello`, `get_server_finished`, and `put_client_finished`.

We then use the ProScript compiler to translate this module into a ProVerif script that looks much like the protocol models described in earlier sections of this paper. (See [48] for details of the translation.) Each pure function in ProScript translates to a ProVerif function; functions that modify mutable state are translated to ProVerif processes that read and write from tables. The interface of the module is compiled to a top-level process that exposes a subset of the protocol functions to the adversary over a public channel.

The adversary can call these functions in any order and any number of times, to initiate connections in parallel, to provide incoming flights of messages, and to obtain outgoing flights of messages. The ProVerif model uses internal tables, not accessible to the attacker, to manage state updates between flights and preserve state invariants through the protocol execution.

Our approach allows us to quickly obtain verifiable ProVerif models from running RefTLS code. For example, we were able to rapidly prototype changes to the TLS 1.3 specification between Draft-13 and Draft-18, while testing for interoperability and analyzing the core protocol at the same time. In particular, we extracted a model from our Draft-18 implementation, and verified our security goals from §III and §V with ProVerif.

We engineered the ProScript compiler to generate readable ProVerif models that can be modified by a protocol analyst to experiment with different threat models. We are working towards applying the same automated translation approach towards CryptoVerif models. CryptoVerif syntax differs slightly from the ProVerif syntax, yet there is ongoing work in the CryptoVerif team to have it accept the same source syntax as ProVerif. However, the kind of models that are easy to verify using CryptoVerif differ from the models that ProVerif can automatically verify, and the assumptions on cryptographic primitives will always remain different. Therefore, even if the source syntax is the same, we may need to adapt our compiler to generate different models for ProVerif and CryptoVerif.

Evaluation: Verification, Interoperability, Performance.

The full RefTLS codebase consists of about 6500 lines of Flow code, including 3000 lines of trusted libraries (mostly message parsing), 2500 lines of untrusted application code, and 1000 lines of verified protocol core. From the core, we extracted an 800 line protocol model in ProVerif and composed it with our generic library from §II. Verifying this model took several hours on a powerful workstation.

RefTLS implements TLS 1.0-1.3, and interoperates with all major TLS libraries for TLS 1.0-1.2. Fewer libraries currently implement TLS 1.3, but RefTLS participated in the IETF Hackathon and achieved interoperability with other implementations of Draft-14. It now interoperates with NSS (Firefox) and BoringSSL (Chrome) for Draft-18.

By implementing Node’s `https` interface, we are able to naturally integrate RefTLS within any Node or Electron

application. We demonstrate the utility of this approach by integrating RefTLS into the Brave web browser, which is written in Electron. We are able to intercept all of Brave’s HTTPS requests and reliably fulfill them through RefTLS.

We benchmarked RefTLS against Node.js’s default OpenSSL-based HTTPS stack when run against an OpenSSL peer over TLS 1.2. In terms of computational overhead, RefTLS is two times slower than Node’s native library, which is not surprising since RefTLS is written in JavaScript, whereas OpenSSL is written in C. In exchange for speed, RefTLS offers an early implementation of TLS 1.3 and a verified protocol core. Furthermore, in many application scenarios, network latency dominates over crypto, so the performance penalty of RefTLS may not be that noticeable.

VIII. DISCUSSION AND RELATED WORK

Symbolic Analysis of TLS 1.3. We symbolically analyzed a composite model of TLS 1.3 Draft-18 with optional client authentication, PSK-based resumption, and PSK-based 0-RTT, running alongside TLS 1.2 against a rich threat model, and we established a series of security goals. In summary, 1-RTT provides forward secrecy, authentication and unique channel identifiers, 0.5-RTT offers weaker authentication, and 0-RTT lacks forward secrecy and replay protection.

We discovered potential vulnerabilities in 0-RTT client authentication in earlier draft versions. These attacks were presented at the TLS Ready-Or-Not (TRON) workshop and contributed to the removal of certificate-based 0-RTT client authentication from TLS 1.3. The current design of PSK binders in Draft-18 is also partly inspired by these kinds of authentication attacks.

TLS 1.3 has been symbolically analyzed before, using the Tamarin prover [35]. ProVerif and Tamarin are both state-of-the-art protocol analyzers with different strengths. Tamarin can verify arbitrary compositions of protocols by relying on user-provided lemmas, whereas ProVerif is less expressive but offers more automation. In terms of protocol features, the Tamarin analysis covered PSK and ECDHE handshakes for 0-RTT and 1-RTT in Draft-10, but did not consider 0-RTT client certificate authentication or 0.5-RTT data. On the other hand, they do consider delayed (post-handshake) authentication, which we did not consider here.

The main qualitative improvement in our verification results over theirs is that we consider a richer threat model that allows for downgrade attacks, and that we analyze TLS 1.3 in composition with previous versions of the protocol, whereas they verify TLS 1.3 in isolation.

Our full ProVerif development consists of 1030 lines of ProVerif; including a generic library incorporating our threat model (400 lines), processes for TLS 1.2 (200 lines) and TLS 1.3 (250 lines), and security queries for TLS 1.2 (50 lines) and TLS 1.3 (180 lines). All proofs complete in about 70 minutes on a powerful workstation. In terms of manual effort, these models took about 3 person-weeks for a ProVerif expert.

Computational Proofs for TLS 1.3. We presented the first mechanically-checked cryptographic proof for TLS 1.3, developed using the CryptoVerif prover. We prove secrecy,

forward secrecy with respect to the compromise of long-term keys, authentication, replay prevention (except for 0-RTT data), and existence of a unique channel identifier for TLS 1.3 draft-18. Our analysis considers PSK modes with and without DHE key exchange, with and without client authentication. It includes 0-RTT and 0.5-RTT data, as well as key updates, but not post-handshake authentication.

Unlike the ProVerif analysis, our CryptoVerif model does not consider compositions of client certificates and pre-shared keys in the same handshake. It also does not account for version or ciphersuite negotiation; instead, we assume that the client and server only support TLS 1.3 with strong cryptographic algorithms. The reason we limit the model in this way is to make the proofs more tractable, since CryptoVerif is not fully automated and requires significant input from the user. With future improvements in the tool, we may be able to remove some of these restrictions.

CryptoVerif is better suited to proofs than finding attacks. Sometimes, proof failures in CryptoVerif might lead us towards computational attacks that do not appear at the symbolic level, but we did not find such attacks in our model of TLS 1.3. We failed to prove forward secrecy for handshakes that use both pre-shared keys and Diffie-Hellman, but this failure is due to limitations in our tool, not due to an attack. Our proofs required some unusual assumptions on public values in Diffie-Hellman groups to avoid confusions between different key exchange modes; these ambiguities are inherent in Draft-18 but have been fixed in Draft-19, making some of our assumptions unnecessary.

In comparison with previous cryptographic proofs of draft versions of TLS 1.3 [40], [52], [55], our cryptographic assumptions and proof structure is similar. The main difference in this work is that our proof is mechanized, so we can easily adapt and recheck our proofs as the protocol evolves.

Our full CryptoVerif development consists of 1895 lines, including new definitions and lemmas for the key schedule (570 lines), a model of the initial handshake (550 lines), a model of PSK-based handshakes (625 lines), and a model of the record protocol (150 lines). For different proofs, we sometimes wrote small variations of these files, and we do not count all those variations here. All proofs completed in about 6 minutes. The total verification effort took about 5 person-weeks for a CryptoVerif expert.

Verifying TLS Implementations. Specifications for protocols like TLS are primarily focused on interoperability; the RFC standard precisely defines message formats, cryptographic computations, and expected message sequences. However, it says little about what state machine these protocol implementations should use, or what APIs they should offer to their applications. This specification ambiguity is arguably the culprit for many implementation bugs [12] and protocol flaws [15] in TLS.

In the absence of a more explicit specification, we advocate the need for verified reference implementations of TLS that can provide exemplary code and design patterns on how to deploy the protocol securely. We proposed one such implementation, RefTLS, for use in JavaScript applications. The core protocol code in RefTLS implements both TLS

1.2 and 1.3 and has been verified using ProVerif. However, RefTLS is a work-in-progress and many of its trusted components remain to be verified. For example, we did not verify our message parsing code or cryptographic libraries, and our verification results rely on the correctness of the unverified ProScript-to-ProVerif compiler [48].

The symbolic security guarantees of RefTLS are weaker than those of computationally-verified implementations like miTLS [21]. However, unlike miTLS, our analysis is fully automated and it can quickly find attacks. The type-based technique of miTLS requires significant user intervention and is better suited to building proofs than finding attacks.

Other Verification Approaches. In addition to ProVerif and CryptoVerif, there are many symbolic and computational analysis tools that have been used to verify cryptographic protocols like TLS. As discussed above, Tamarin [68] was used to symbolically analyze TLS 1.3 Draft-10 [35]. EasyCrypt [8] has been used to develop cryptographic proofs for various components used in TLS, including the MAC-Encode-Encrypt construction used in the record layer [5].

Our ProScript-to-ProVerif compiler is inspired by previous works on deriving ProVerif models from F# [20], Java [6], and JavaScript [16]. Such translations have been used to symbolically and computationally analyze TLS implementations [18]. An alternative to model extraction is to synthesize a verified implementation from a verified model; [30] shows how to compile CryptoVerif models to OCaml and uses it to derive a verified SSH implementation.

The most advanced case studies for verified protocol implementations use dependent type systems, because they scale well to large codebases. Refinement types for F# have been used to prove both symbolic [19] and cryptographic security properties, with applications to TLS [21]. The F* programming language [70] has been used to verify small protocols and cryptographic libraries [73]. Similar techniques have been applied to the cryptographic verification of Java programs [53].

IX. CONCLUSION AND FUTURE WORK

TLS 1.3 is a social and technical experiment in the collaborative design of a practical protocol with regular input and review from the academic research community. It seeks to reverse the traditional pattern where security analyses are performed several years after standardization, when it may be too late to change how implementations work. This paper describes our contribution to this standardization effort.

We present verification results for symbolic models in ProVerif, computational models in CryptoVerif, and a reference implementation in JavaScript of TLS 1.3 Draft-18. There are still many features and aspects of the emerging protocol standard that remain to be analyzed. Furthermore, the formal connections between our ProVerif models, CryptoVerif proofs, and JavaScript code are not as strong as could be desired. We have focused on proof automation and readable models as a pragmatic first step, but we are working on formal proofs of correctness for our translations from Flow to ProVerif and CryptoVerif, so that we can obtain strong guarantees for our protocol source code.

REFERENCES

- [1] D. Adrian, K. Bhargavan, Z. Durumeric, P. Gaudry, M. Green, J. A. Halderman, N. Heninger, D. Springall, E. Thomé, L. Valenta *et al.*, “Imperfect forward secrecy: How Diffie-Hellman fails in practice,” in *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2015, pp. 5–17.
- [2] M. R. Albrecht and K. G. Paterson, “Lucky microseconds: A timing attack on Amazon’s S2N implementation of TLS,” in *EUROCRYPT*, 2016, pp. 622–643.
- [3] N. AlFardan, D. J. Bernstein, K. G. Paterson, B. Poettering, and J. C. Schuldt, “On the security of RC4 in TLS,” in *USENIX Security Symposium*, 2013, pp. 305–320.
- [4] N. J. AlFardan and K. G. Paterson, “Lucky thirteen: Breaking the TLS and DTLS record protocols,” in *2013 IEEE Symposium on Security and Privacy (SP 2013)*, 2013, pp. 526–540.
- [5] J. B. Almeida, M. Barbosa, G. Barthe, and F. Dupressoir, “Verifiable Side-Channel Security of Cryptographic Implementations: Constant-Time MEE-CBC,” in *Fast Software Encryption (FSE)*, 2016, pp. 163–184.
- [6] M. Avalle, A. Pironti, R. Sisto, and D. Pozza, “The Java SPI framework for security protocol implementation,” in *Availability, Reliability and Security (ARES), 2011 Sixth International Conference on*, Aug 2011, pp. 746–751.
- [7] N. Aviram, S. Schinzel, J. Somorovsky, N. Heninger, M. Dankel, J. Steube, L. Valenta, D. Adrian, J. A. Halderman, V. Dukhovni, E. Kasper, S. Cohny, S. Engels, C. Paar, and Y. Shavitt, “DROWN: breaking TLS using SSLv2,” in *USENIX Security Symposium*, 2016, pp. 689–706.
- [8] G. Barthe, F. Dupressoir, B. Grégoire, C. Kunz, B. Schmidt, and P.-Y. Strub, “EasyCrypt: A tutorial,” in *Foundations of Security Analysis and Design VII (FOSAD)*, ser. Lecture Notes in Computer Science. Springer, 2014, vol. 8604, pp. 146–166.
- [9] M. Bellare, “New proofs for NMAC and HMAC: Security without collision-resistance,” in *Advances in Cryptology (CRYPTO)*, 2006, pp. 602–619.
- [10] M. Bellare, J. Kilian, and P. Rogaway, “The security of the cipher block chaining message authentication code,” *Journal of Computer and System Sciences*, vol. 61, no. 3, pp. 362–399, Dec. 2000.
- [11] M. Bellare and C. Namprempe, “Authenticated encryption: Relations among notions and analysis of the generic composition paradigm,” in *Advances in Cryptology – ASIACRYPT’00*, 2000, pp. 531–545.
- [12] B. Beurdouche, K. Bhargavan, A. Delignat-Lavaud, C. Fournet, M. Kohlweiss, A. Pironti, P.-Y. Strub, and J. K. Zinzindohoue, “A messy state of the union: taming the composite state machines of TLS,” in *IEEE Symposium on Security & Privacy (Oakland)*, 2015.
- [13] K. Bhargavan, B. Blanchet, and N. Kobeissi, “Verified models and reference implementations for the TLS 1.3 standard candidate,” Inria, Research report RR-9040, 2017.
- [14] K. Bhargavan, C. Brzuska, C. Fournet, M. Green, M. Kohlweiss, and S. Z. Béguelin, “Downgrade resilience in key-exchange protocols,” in *IEEE Symposium on Security and Privacy (Oakland)*, 2016, pp. 506–525.
- [15] K. Bhargavan, A. Delignat-Lavaud, C. Fournet, A. Pironti, and P.-Y. Strub, “Triple handshakes and cookie cutters: Breaking and fixing authentication over TLS,” in *IEEE Symposium on Security & Privacy (Oakland)*, 2014, pp. 98–113.
- [16] K. Bhargavan, A. Delignat-Lavaud, and S. Maffei, “Language-based defenses against untrusted browser origins,” in *USENIX Security Symposium*, 2013, pp. 653–670.
- [17] K. Bhargavan, A. Delignat-Lavaud, and A. Pironti, “Verified contributive channel bindings for compound authentication,” in *Network and Distributed System Security Symposium (NDSS ’15)*, 2015.
- [18] K. Bhargavan, C. Fournet, R. Corin, and E. Zălinescu, “Verified cryptographic implementations for TLS,” *ACM TOPLAS*, vol. 15, no. 1, pp. 3:1–3:32, 2012.
- [19] K. Bhargavan, C. Fournet, and A. D. Gordon, “Modular verification of security protocol code by typing,” in *ACM Symposium on Principles of Programming Languages (POPL)*, 2010, pp. 445–456.
- [20] K. Bhargavan, C. Fournet, A. D. Gordon, and S. Tse, “Verified interoperable implementations of security protocols,” *ACM Transactions on Programming Languages and Systems*, vol. 31, no. 1, 2008.
- [21] K. Bhargavan, C. Fournet, M. Kohlweiss, A. Pironti, and P.-Y. Strub, “Implementing TLS with verified cryptographic security,” in *IEEE Symposium on Security & Privacy (Oakland)*, 2013. [Online]. Available: pubs/implementing-tls-with-verified-cryptographic-security-sp13.pdf
- [22] K. Bhargavan and G. Leurent, “On the practical (in-)security of 64-bit block ciphers: Collision attacks on HTTP over TLS and OpenVPN,” in *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2016, pp. 456–467.
- [23] —, “Transcript collision attacks: Breaking authentication in TLS, IKE, and SSH,” in *ISOC Network and Distributed System Security Symposium (NDSS)*, 2016.
- [24] B. Blanchet, “A computationally sound mechanized prover for security protocols,” *IEEE Transactions on Dependable and Secure Computing*, vol. 5, no. 4, pp. 193–207, 2008.
- [25] —, “Automatic verification of correspondences for security protocols,” *Journal of Computer Security*, vol. 17, no. 4, pp. 363–434, 2009.
- [26] —, “Security protocol verification: Symbolic and computational models,” in *Principles of Security and Trust (POST)*, 2012, pp. 3–29.
- [27] —, “Modeling and verifying security protocols with the applied pi calculus and ProVerif,” *Foundations and Trends in Privacy and Security*, vol. 1, no. 1–2, pp. 1–135, Oct. 2016.
- [28] D. Bleichenbacher, “Chosen ciphertext attacks against protocols based on the RSA encryption standard PKCS# 1,” in *Annual International Cryptology Conference*, ser. Lecture Notes in Computer Science, vol. 1462. Springer, 1998, pp. 1–12.
- [29] M. Bodin, A. Charguéraud, D. Filaretti, P. Gardner, S. Maffei, D. Naudziuniene, A. Schmitt, and G. Smith, “A trusted mechanised javascript specification,” in *ACM Symposium on the Principles of Programming Languages (POPL)*, 2014, pp. 87–100.
- [30] D. Cadé and B. Blanchet, “Proved generation of implementations from computationally secure protocol specifications,” *Journal of Computer Security*, vol. 23, no. 3, pp. 331–402, 2015.
- [31] S. Chaki and A. Datta, “Aspicer: An automated framework for verifying security protocol implementations,” in *2009 22nd IEEE Computer Security Foundations Symposium*. IEEE, 2009, pp. 172–185.
- [32] A. Chaudhuri, “Flow: Abstract interpretation of javascript for type checking and beyond,” in *ACM Workshop on Programming Languages and Analysis for Security (PLAS)*, 2016.
- [33] J.-S. Coron, Y. Dodis, C. Malinaud, and P. Puniya, “Merkle-Damgård revisited: How to construct a hash function,” in *Advances in Cryptology (CRYPTO)*, 2005, pp. 430–448.
- [34] V. Cortier, S. Kremer, and B. Warinschi, “A survey of symbolic methods in computational analysis of cryptographic systems,” *Journal of Automated Reasoning*, vol. 46, no. 3–4, pp. 225–259, 2011.
- [35] C. Cremers, M. Horvat, S. Scott, and T. van der Merwe, “Automated analysis and verification of TLS 1.3: 0-RTT, resumption and delayed authentication,” in *IEEE Symposium on Security and Privacy (Oakland)*, 2016, pp. 470–485.
- [36] I. B. Damgård, “A design principle for hash functions,” in *Advances in Cryptology—CRYPTO89*, 1989, pp. 416–427.
- [37] T. Dierks and E. Rescorla, “The Transport Layer Security (TLS) Protocol Version 1.2,” IETF RFC 5246, 2008.
- [38] Y. Dodis, T. Ristenpart, J. Steinberger, and S. Tessaro, “To hash or not to hash again? (In)differentiability results for H^2 and HMAC,” in *Advances in Cryptology (Crypto)*, 2012, pp. 348–366.
- [39] D. Dolev and A. C. Yao, “On the security of public key protocols,” *IEEE Transactions on Information Theory*, vol. 29, no. 2, pp. 198–207, 1983.
- [40] B. Dowling, M. Fischlin, F. Günther, and D. Stebila, “A cryptographic analysis of the TLS 1.3 handshake protocol candidates,” in *ACM Conference on Computer and Communications Security (CCS)*, 2015, pp. 1197–1210.
- [41] M. Fischlin, F. Günther, B. Schmidt, and B. Warinschi, “Key confirmation in key exchange: A formal treatment and implications for TLS 1.3,” in *IEEE Symposium on Security and Privacy (Oakland)*, 2016, pp. 452–469.
- [42] M. Fischlin and F. Günther, “Multi-stage key exchange and the case of Google’s QUIC protocol,” in *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2014, pp. 1193–1204.

- [43] S. Goldwasser, S. Micali, and R. Rivest, “A digital signature scheme secure against adaptive chosen-message attacks,” *SIAM Journal of Computing*, vol. 17, no. 2, pp. 281–308, April 1988.
- [44] R. Hamilton, J. Iyengar, I. Swett, and A. Wilk, “QUIC: A UDP-based multiplexed and secure transport,” 2016, IETF Internet Draft.
- [45] K. E. Hickman, “The SSL protocol,” 1995, IETF Internet Draft, <https://tools.ietf.org/html/draft-hickman-netscape-ssl-00>.
- [46] T. Jager, F. Kohlar, S. Schäge, and J. Schwenk, “On the security of TLS-DHE in the standard model,” in *CRYPTO 2012*, 2012, pp. 273–293.
- [47] T. Jager, J. Schwenk, and J. Somorovsky, “On the security of TLS 1.3 and QUIC against weaknesses in PKCS#1 v1.5 encryption,” in *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2015, pp. 1185–1196.
- [48] N. Kobeissi, K. Bhargavan, and B. Blanchet, “Automated verification for secure messaging protocols and their implementations: A symbolic and computational approach,” in *IEEE European Symposium on Security and Privacy (EuroS&P)*, 2017.
- [49] H. Krawczyk, “Cryptographic extraction and key derivation: The HKDF scheme,” in *Advances in Cryptology (CRYPTO)*, 2010, pp. 631–648.
- [50] —, “A unilateral-to-mutual authentication compiler for key exchange (with applications to client authentication in tls 1.3),” in *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2016, pp. 1438–1450.
- [51] H. Krawczyk, K. G. Paterson, and H. Wee, “On the security of the TLS protocol: A systematic analysis,” in *CRYPTO 2013*, 2013, pp. 429–448.
- [52] H. Krawczyk and H. Wee, “The OPTLS protocol and TLS 1.3,” in *IEEE European Symposium on Security & Privacy (Euro S&P)*, 2016, cryptography ePrint Archive, Report 2015/978.
- [53] R. Küsters, T. Truderung, and J. Graf, “A framework for the cryptographic verification of Java-like programs,” in *IEEE Computer Security Foundations Symposium (CSF)*, 2012, pp. 198–212.
- [54] A. Langley, M. Hamburg, and S. Turner, “Elliptic curves for security,” IRTF RFC 7748 <https://tools.ietf.org/html/rfc7748>, Jan. 2016.
- [55] X. Li, J. Xu, Z. Zhang, D. Feng, and H. Hu, “Multiple handshakes security of TLS 1.3 candidates,” in *IEEE Symposium on Security and Privacy (Oakland)*, 2016, pp. 486–505.
- [56] 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 & Privacy (Oakland)*, 2015, pp. 214–231.
- [57] U. Maurer and B. Tackmann, “On the soundness of authenticate-then-encrypt: formalizing the malleability of symmetric encryption,” in *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2010, pp. 505–515.
- [58] N. Mavrogiannopoulos, F. Vercauteren, V. Velichkov, and B. Preneel, “A cross-protocol attack on the TLS protocol,” in *ACM CCS*, 2012.
- [59] C. Meyer, J. Somorovsky, E. Weiss, J. Schwenk, S. Schinzel, and E. Tews, “Revisiting SSL/TLS implementations: New Bleichenbacher side channels and attacks,” in *23rd USENIX Security Symposium*. USENIX Association, 2014, pp. 733–748.
- [60] B. Möller, T. Duong, and K. Kotowicz, “This POODLE bites: exploiting the SSL 3.0 fallback,” <https://www.openssl.org/~bodo/ssl-poodle.pdf>, 2014.
- [61] T. Okamoto and D. Pointcheval, “The gap-problems: a new class of problems for the security of cryptographic schemes,” in *Practice and Theory in Public Key Cryptography (PKC)*, 2001, pp. 104–118.
- [62] K. G. Paterson, T. Ristenpart, and T. Shrimpton, “Tag size does matter: Attacks and proofs for the TLS record protocol,” in *ASIACRYPT*, 2011, pp. 372–389.
- [63] K. G. Paterson and T. van der Merwe, “Reactive and proactive standardisation of TLS,” in *Security Standardisation Research (SSR)*, 2016, pp. 160–186.
- [64] M. Ray, A. Pironti, A. Langley, K. Bhargavan, and A. Delignat-Lavaud, “Transport Layer Security (TLS) session hash and extended master secret extension,” 2015, IETF RFC 7627.
- [65] E. Rescorla, M. Ray, S. Dispensa, and N. Oskov, “TLS renegotiation indication extension,” IETF RFC 5746, 2010.
- [66] E. Rescorla, “0-RTT and Anti-Replay,” <https://www.ietf.org/mail-archive/web/tls/current/msg15594.html>, Mar. 2015.
- [67] —, “[TLS] PR#875: Additional Derive-Secret stage,” <https://www.ietf.org/mail-archive/web/tls/current/msg22373.html>, Feb. 2017.
- [68] B. Schmidt, S. Meier, C. Cremers, and D. Basin, “Automated analysis of Diffie-Hellman protocols and advanced security properties,” in *IEEE Computer Security Foundations Symposium (CSF)*, 2012, pp. 78–94.
- [69] D. Stefan, “Espectro project description,” 2016, <https://cseweb.ucsd.edu/~dstefan/#projects>.
- [70] N. Swamy, C. Hrițcu, C. Keller, A. Rastogi, A. Delignat-Lavaud, S. Forest, K. Bhargavan, C. Fournet, P.-Y. Strub, M. Kohlweiss, J.-K. Zinzindohoue, and S. Zanella-Béguélin, “Dependent types and multi-monadic effects in F*,” in *ACM Symposium on Principles of Programming Languages (POPL)*, 2016, pp. 256–270.
- [71] M. Vanhoef and F. Piessens, “All your biases belong to us: Breaking RC4 in WPA-TKIP and TLS,” in *USENIX Security Symposium*, 2015, pp. 97–112.
- [72] D. Wagner and B. Schneier, “Analysis of the SSL 3.0 protocol,” in *USENIX Electronic Commerce*, 1996.
- [73] J. K. Zinzindohoue, E. Bartzia, and K. Bhargavan, “A verified extensible library of elliptic curves,” in *IEEE Computer Security Foundations Symposium (CSF)*, 2016, pp. 296–309.

APPENDIX A.

LEMNAS ON PRIMITIVES AND ON THE KEY SCHEDULE

We show the following properties:

- $\text{mac}_H^k(m) = \text{mac}^k(H(m))$ is an SUF-CMA (strongly unforgeable under chosen message attacks) MAC. Indeed, since $\text{mac} = \text{HMAC-H}$ is a PRF, it is an SUF-CMA MAC as shown in [10], and this property is preserved by composition with a collision-resistant hash function.
- $\text{sign}_H^{sk}(m) = \text{sign}^{sk}(H(m))$ is an UF-CMA signature. Indeed, sign is an UF-CMA signature, and this property is preserved by composition with a collision-resistant hash function.

We also prove several lemmas on the key schedule of TLS 1.3, using CryptoVerif.

- When es is a fresh random value, $e \mapsto \text{hkdf-extract}(es, e)$ and $\log_1 \mapsto \text{derive-secret}(es, \text{ats}_c, \log_1)$ are indistinguishable from independent random functions, and $k^b = \text{derive-secret}(es, \text{pbk}, \text{“”})$ and $\text{hkdf-extract}(es, 0^{\text{len}_H})$ are indistinguishable from independent fresh random values independent from these random functions.
- When hs is a fresh random value, $\log_1 \mapsto \text{derive-secret}(hs, \text{hts}_c, \log_1) \parallel \text{derive-secret}(hs, \text{hts}_s, \log_1)$ is indistinguishable from a random function and $\text{hkdf-extract}(hs, 0^{\text{len}_H})$ is indistinguishable from a fresh random value independent from this random function.
- When ms is a fresh random value, the functions $\log_4 \mapsto \text{derive-secret}(ms, \text{ats}_c, \log_4) \parallel \text{derive-secret}(ms, \text{ats}_s, \log_4) \parallel \text{derive-secret}(ms, \text{ems}, \log_4)$ and $\log_7 \mapsto \text{derive-secret}(ms, \text{rms}, \log_7)$ are indistinguishable from independent random functions.
- When l_1, l_2, l_3 are pairwise distinct labels and s is a fresh random value, $\text{hkdf-expand-label}(s, l_i, \text{“”})$ for $i = 1, 2, 3$ are indistinguishable from independent fresh random values.

All random values considered above are uniformly distributed. We use these properties as assumptions in our proof of the protocol. This modular approach considerably reduces the complexity of the games that CryptoVerif has to consider.

These results suggest that the key schedule could be simplified by replacing groups of calls to derive-secret that




Figure 7: Client state machine

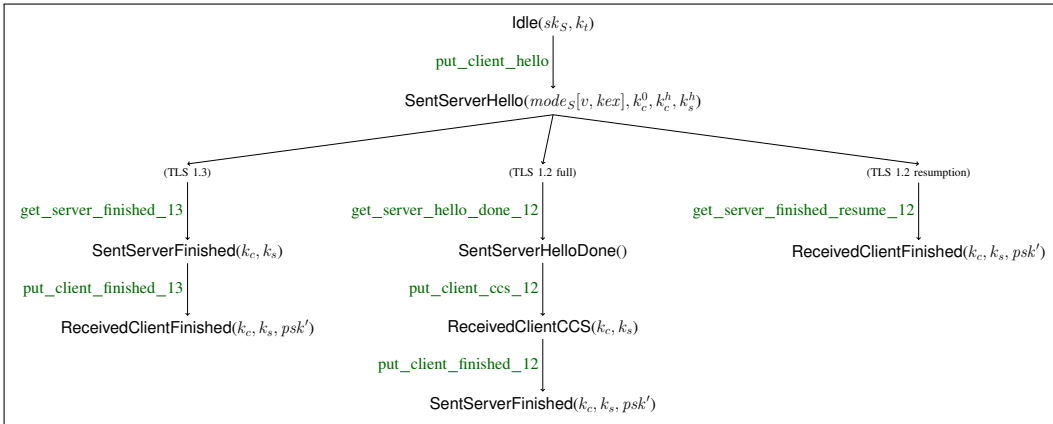


Figure 8: Server state machine

use the same key and log with a single call to `derive-secret` that would output the concatenation of several keys. The same remark also holds for calls to `hkdf-expand-label` that use the same key. This approach corresponds to the usage of expansion recommended in the formalization of HKDF [49], and would simplify the proof: some lemmas above would no longer be needed. We would also recommend replacing $ms = \text{hkdf-extract}(hs, 0^{\text{len}_H()})$ with $ms = \text{derive-secret}(hs, ms, \text{""})$: that would be more natural since we use the PRF property of HMAC-H for this computation and not the randomness extraction. If the argument $0^{\text{len}_H()}$ may change in the future, then we would support Krawczyk’s recommendation [67] of applying `hkdf-extract` to the result of `derive-secret` ($hs, ms, \text{""}$).

APPENDIX B. REFTLS PROTOCOL STATE MACHINES

Client. The RefTLS client implements the composite state machine shown in Figure 7 for TLS 1.3 and TLS 1.2. Each

state represents a point in the protocol where the client is either waiting for a flight of handshake messages from the server, or it has new session keys that it wishes to communicate to the record layer. Each arrow is annotated with the name of the function in RefTLS-CORE API that implements the corresponding state transition. Each transition may involve processing a flight of incoming messages, changing the session state, and producing a flight of outgoing messages.

Server. The RefTLS server implements a dual state machine for TLS 1.3 and TLS 1.2, as depicted in Figure 8. The server decides which protocol version and key exchange the handshake will use, and triggers the appropriate branch in the state machine by sending a `ServerHello`. Like the client, each of its state transition functions corresponds either to a flight of messages or to a change of keys.