

Proving the TLS Handshake Secure (as it is)

Karthikeyan Bhargavan¹, Cédric Fournet², Markulf Kohlweiss², Alfredo Pironti¹, Pierre-Yves Strub³, and Santiago Zanella-Béguelin^{1,2}

¹ INRIA {`firstname.name`}@inria.fr

² Microsoft Research {`fournet,markulf`}@microsoft.com

³ IMDEA Software Institute `pierre-yves@strub.nu`

Abstract. The TLS Internet Standard features a mixed bag of cryptographic algorithms and constructions, letting clients and servers negotiate their use for each run of the handshake. Although many ciphersuites are now well-understood in isolation, their composition remains problematic, and yet it is critical to obtain practical security guarantees for TLS, as all mainstream implementations support multiple related runs of the handshake and share keys between algorithms.

We study the provable security of the TLS handshake, as it is implemented and deployed. To capture the details of the standard and its main extensions, we rely on mTLS, a verified reference implementation of the protocol. We propose new agile security definitions and assumptions for the signatures, key encapsulation mechanisms (KEM), and key derivation algorithms used by the TLS handshake. To validate our model of key encapsulation, we prove that both RSA and Diffie-Hellman ciphersuites satisfy our definition for the KEM. In particular, we formalize the use of PKCS#1v1.5 and build a 3,000-line EASYCRYPT proof of the security of the resulting KEM against replayable chosen-ciphertext attacks under the assumption that ciphertexts are hard to re-randomize.

Based on our new agile definitions, we construct a modular proof of security for the mTLS reference implementation of the handshake, including ciphersuite negotiation, key exchange, renegotiation, and resumption, treated as a detailed 3,600-line executable model. We present our main definitions, constructions, and proofs for an abstract model of the protocol, featuring series of related runs of the handshake with different ciphersuites. We also describe its refinement to account for the whole reference implementation, based on automated verification tools.

1 Introduction

TLS is the most widely deployed protocol for securing communications and yet, after two decades of attacks, patches and extensions, its practical security remains unresolved. One of the most troublesome aspects of the protocol is its handling of a large number of cryptographic algorithms and constructions. New extensions are added to the protocol and its implementations, while older features remain for backward compatibility. Thus, TLS clients and servers offer many choices, and each run of the handshake involves a negotiation of the best protocol version, ciphersuite, and extensions available at both ends. Such a trade-off between flexibility and security creates several problems:

- (1) It makes the security of TLS depend on its correct configuration, inasmuch as some versions (e.g. SSL2) and algorithms (e.g. MD5 and RC4) are much weaker than others, and may also suffer from different implementation flaws. In theory, only very restrictive configurations have been proved secure. In practice, dangerous mis-configurations of TLS are commonplace.
- (2) It complicates the protocol logic, as the integrity of the negotiation itself relies on algorithms being negotiated; this is a persistent source of attacks, from protocol regression in SSL2 [27] to version fallback in current browsers [18].
- (3) It demands stronger security assumptions, to reflect the fact that honest parties may use the same key materials with different algorithms. Intuitively, TLS *on its own* enables a range of chosen-protocol attacks whereby a weak algorithm (chosen by the attacker) may compromise the security of stronger algorithms (chosen by honest parties). We detail below several constructions of TLS that demand joint assumptions on collections of algorithms. Surprisingly, prior work on the provable security of TLS failed to make this observation or left it implicit.

Besides interference between multiple algorithms, TLS features dependencies between multiple runs of the handshake. For instance, a client connection may first run an RSA-based session to establish a master secret and keys for the record layer, then run a second session on the same connection, possibly with different algorithms and certificates. Using a parallel connection, the client may run a third *resumption* handshake, re-using the master secret of a prior session to derive new keys. At that point, the security of those keys depends on algorithms and constructions used in three runs of the handshake. (See for instance [5] for recent attacks involving three related handshakes.) This is in sharp contrast with prior work on the provable security of TLS [13, 16, 17], which focus on a fixed run of the protocol, for a fixed choice of algorithms.

1.1 Cryptographic Agility. *Agile security* considers families of schemes or protocols, all serving the same purpose, when the same keys are shared across members of the family. Acar et al. [1] propose agile definitions for pseudo-random functions (PRF) and encryption schemes, and advocate agility as a major practical concern for protocols like TLS. Instead, *combined*, or *joint security* [12] studies the sharing of keys between constructions serving different purposes, e.g. encryption and signing. TLS requires both agile and joint security; in the remainder we let the term *agility* encompass both concepts.

The agility mechanisms of TLS are primarily driven by *ciphersuites* of the form `TLS_e_s_WITH_r`, which indicates a key encapsulation mechanism (KEM) e and signature scheme s for the handshake, and an authenticated encryption scheme r for the record layer. For instance, the commonly-used ciphersuite `TLS_RSA_WITH_AES_256_CBC_SHA` indicates an RSA handshake: the client sends a fresh premaster secret encrypted under the server public key; both parties use it to extract a master secret, used in turn as the seed of a SHA1-based PRF to derive 4 keys for SHA1-based MACs and AES encryption in CBC mode. TLS 1.2 currently has 314 registered ciphersuites. More precisely, the choice of algorithms depends on additional data exchanged during the handshake (hence subject

to active attacks), including protocol versions, certificate requests, certificate chains, and various extensions in the first two messages of the handshake (e.g. for choosing hash functions and elliptic curves). Still, because of key reuse across algorithms, we stress that the security of TLS does not reduce to the security of a few thousand fixed-algorithm variants of the handshake.

1.2 Empirical Study of Web Servers and Browsers. Using an online analyzer [24], we gathered extended information on server configurations for 215 of the top 500 domains,⁴ including the TLS versions, ciphersuites, certificates, and extensions they offer. These servers accept 64 ciphersuites, with an average of 12 and standard deviation of 6. They still widely deploy weak algorithms: 70% accept at least one ciphersuite with MD5 and 90% at least one with RC4. All servers but one offer several versions; 37% offer only SSL3 and TLS 1.0; 56% offer all 4 versions from SSL3 to TLS 1.2. Although now forbidden by the standard, 3% still accept SSL2.

We also tested 12 TLS clients, including major web browsers (Chrome, Firefox, Internet Explorer, Safari) and libraries (NSS, OpenSSL, SChannel, Secure Transport). These clients similarly propose a large number of ciphersuites, ranging from 19 to 36; they all propose weak hash (MD5) or encryption methods (RC4, or even no encryption).

1.3 Cross-Ciphersuite Attacks. As a first example, most TLS servers are configured to use the same RSA certificate both for signing handshake messages and for decrypting premaster secrets. Experimentally, 69% of the servers we tested propose at least one ciphersuite using RSA for encryption and one using it for signing, and *all* 138 of those use the same key for both purposes.

As a second example, Mavrogiannopoulos et al. [20] report a cross-protocol attack between plain Diffie-Hellman (DH) and Elliptic-Curve Diffie-Hellman (ECDH) ciphersuites, due to a mis-interpretation of the signed group description sent by the server. Each family of ciphersuites is (a priori) secure in isolation, but configurations enabling a DH client and an ECDH server are subject to their attack.

Our third example concerns the record algorithms (the r in `TLS_e_s_WITH_r`). Recall that both parties derive keys for r immediately after the KEM phase, and start using them before verifying the Finished messages that confirm the integrity of the handshake. As an optimization, the optional False Start TLS extension [19] lets clients send private application data before key confirmation. Depending on r , the *same* key materials are split into IVs, MAC keys, and encryption keys of various lengths. Hence, the client and the server may start using the same bits with different algorithms r_C and r_S , for instance as an IV at the client and as a MAC key at the server. To our knowledge, we are the first to report this cross-algorithm attack against [19]. We do not have an exploit based on two standard record algorithms (r_C, r_S) but one can easily design a pair of

⁴ <http://www.alexa.com/topsites/global>, as of January 2014, excluding domains with no valid HTTPS certificate.

schemes strong in isolation and subject to the attack, and key recovery attacks against any standard algorithm r_C could be used to attack strong r_S algorithms.

1.4 Multiple Sessions and Connections. Following the standard, we recall TLS terminology for multiple related handshakes; this differs from the key-exchange model of Bellare & Rogaway [3] with only one kind of sessions and no shared state between sessions. Local instances of the protocol provide a *connection* (concretely, taking ownership of a TCP connection), either as client or as server. Each connection goes through a sequence of *epochs*, each epoch running one *handshake*. For a given connection, we refer to additional handshakes in the sequence as *renegotiations*. We refer to epochs performing full handshakes as *sessions*, and to epochs performing abbreviated handshakes as *resumptions*. We have a transition from the current epoch to the next each time a handshake *completes* by successfully processing the last message of the handshake. Abstractly, the local instance never stops; it is then ready to send (or receive) the first message of the next handshake.

Sessions intend to establish a fresh *master secret*, associated with data extracted from the handshake messages that record its origin and purpose, and used to derive fresh keys for the record layer. *Resumptions* instead rely on a prior complete session to save the cost of public-key cryptography and directly derive fresh keys using the algorithms and master secret of the original session. For each epoch, the handshake consists of a series of messages exchanged using the current record-layer protection mechanisms, initially in the clear, then typically using authenticated encryption.

1.5 Proving the TLS Handshake Secure. The scope of this paper is the TLS handshake, as it is specified in the Internet Standard and (to a lesser extent) as it is commonly used. We model multiple, related sessions and connections, and the agility issues caused by multiple ciphersuites featuring RSA and DHE key exchanges. We also model unilateral and mutual authentication, based on RSA and (EC)DSA signatures. On the other hand, we do not cover static DH, PSK, and ECDHE key exchanges, and we do not investigate the joint usage of keys for signing and encryption. (See the full paper for their discussion.)

Our main result is provable security for a standard-compliant, reference implementation of the handshake, seen as a detailed cryptographic model of the protocol. Our provably-secure handshake code consists of 3,600 lines of F#. Its security relies on new agile assumptions, notably for its KEMs. We reduce them to lower-level assumptions on RSA encryption and Diffie-Hellman exchange, using a 3,000-line EASYCRYPT [2] proof. Working with a reference implementation, and testing it against mainstream implementations, forces us to handle the details of multiple handshakes and algorithms. Proving it secure requires both modularity and automation.

A feature of TLS that traditionally resists abstraction is that the handshake releases algorithms and derived keys to the record layer *before* the handshake completes, so that its last messages can be exchanged as TLS fragments protected by the new keys. We revisit the cryptographic folklore that the handshake can only be proved secure by including these encrypted messages. The kernel of the

lore is that it cannot be proved using a Bellare & Rogaway-style key-exchange definition. To achieve modularity, we separate record-key generation from handshake completion: our main definition releases the record keys in the middle of the handshake, before signaling its completion a few messages later. Since the handshake does *not* rely on record-layer protection, we can safely let the handshake adversary control both the network and the record layer. Completion is still necessary to confirm that the record keys are secure before encrypting any application data—but not for encrypting handshake Finished messages.

We stress that this paper establishes the security of the *handshake*, seen as a component of TLS, not the full communications protocol. Our main construction provides key indistinguishability, and ensures agreement on parameters for the record layer. Our results complement those of Bhargavan et al. [4], who describe mTLS, an implementation of TLS verified in the computational model of cryptography; they focus on the main TLS API and application security, but rely on stronger, ad hoc assumptions for RSA and Diffie-Hellman ciphersuites. Our handshake is integrated with mTLS, which provides additional definitions and verified code for the record layer and the protocol logic. (Their security model ensures in particular that the record keys are used for protecting application data only after handshake completion [4].) By composing our results with theirs, we obtain security for a reference implementation of the TLS standard and the sample applications built and verified on top of mTLS.

1.6 Overview of the Paper. We see the use of a verified reference implementation and automated tools as essential to precisely account for multiple related epochs and algorithms in TLS; §6 briefly describes our use of high-level programming, type systems, and provers to carry out modular cryptographic verification at this scale. To present our result and explain its proof structure, however, we rely on more succinct definitions and constructions, given in §2–5 and outlined below. This more abstract treatment suffices to convey the main ideas, but it necessarily omits many aspects of the handshake, such as its message formats. We refer to the standard [9] or the implementation for the details. Also, for simplicity, we do not model forward secrecy and state reveal e.g. for master secrets, and we consider only static compromise for long-term keys.

Signatures (§2) and certificates. We begin with a relatively simple agile definition. TLS supports three core signature algorithms, $s \in \{RSA, DSA, ECDSA\}$, used with a range of algorithms h to hash the text before signing. The hash algorithm depends on protocol versions, ciphersuites and extensions. TLS does not enforce any key-based hash algorithm policy, so we need a notion of security that tolerates *some* weak algorithms in the standard. For instance, a verifier tricked into using MD5 may remain secure, provided the signer only uses SHA1, and vice-versa. For each core algorithm s , we define h^* - H -security against an adversary that must forge a valid signature for algorithms (s, h^*) , given access to signing oracles for any algorithms (s, h) with $h \in H$. We show that a family of secure schemes may not be jointly secure, but we leave open its concrete analysis for the range of algorithms used in TLS.

Our model excludes any validation rules for certificates and their PKI, an important problem outside the scope of the TLS standard. Our constructions simply authenticate the exchanged certificate chains, and use a specification function to extract from them the public keys used in the handshake.

Master secrets, key encapsulation, and key derivation (§3). Following Krawczyk et al. [17], we use KEMs [8] to model key-exchange; this allows us to unify RSA and Diffie-Hellman within the same formalism. Instead of treating the whole handshake as a KEM, however, following Morrissey et al. [21], we decompose it into *premaster secret*, *master secret*, and *record-key derivation* phases; this yields the modularity we need e.g. for modeling the re-use of master secrets between handshakes. We show how to securely construct a master secret KEM from a premaster secret KEM for RSA and Diffie-Hellman ciphersuites (Theorem 1) and, independently, how to derive record keys and Finished messages from master secrets (see the full paper). We formalize the proof of Theorem 1 in EASYCRYPT. For RSA, this involves showing that countermeasures to Bleichenbacher’s attacks [6, 15] provide enough protection against chosen-ciphertext attacks. We rely on the assumption that PKCS#1v1.5 ciphertexts are hard to re-randomize; we leave open the problem of further reducing this conjecture to standard RSA assumptions. Our result does not directly compare to the one of Krawczyk et al. as their KEM also includes key derivation and Finished messages, whereas we rely on this new assumption. To comply with the standard, we also support agility in the algorithm used to extract master secrets from a premaster secrets. As for agile signatures, we arrive at a definition parameterized by an algorithm for the encryptor and a set of algorithms for the decryptor.

Once established, the master secret is used to key a pseudo-random function (PRF) for multiple epochs for two purposes: (1) to derive the record-layer key materials for the epoch; and (2) to compute the MACs of all messages exchanged in an epoch to verify its integrity. Our corresponding security definition (in the full paper) requires that adversaries commit to a record-layer algorithm r before key derivation. This let us support the negotiation of r without having to make agile assumptions for the record layer, as discussed in §1.3.

Agile security model (§4) and proof (§5) for sequences of handshakes.

The main two goals of the handshake are to establish shared keys for the record layer, and to agree on many parameters, including those used in the handshake itself. To this end, we propose a new security definition that covers multiple epochs on different connections, related by resumptions and renegotiations. We equip our adversary (informally including the rest of TLS, the application, and the network) with oracles to create honest connections and long-term keys for clients and servers, to control their usage, and to exchange handshake messages. Each honest instance of the protocol represents a connection, and logs a sequence of *local assignments*, recording its view on the successive epochs of the connection. This enables us to capture TLS assignments in a generic manner. Our main integrity result is that, when a handshake completes, and under suitable conditions on algorithms and keys, honest clients and servers agree on all assignments for all epochs on the connection. More explicitly, for new sessions,

both parties agree on a unique label; the negotiation algorithms, parameters, and key-exchange values; and the optional certificate chains for the client and the server. For resumptions, both parties agree on the label of the session being resumed, as well as a fresh unique label for key derivation.

We also provide secure key derivation, depending on distinguished exchange-value assignments for each ciphersuite. They are somewhat similar to session identifiers in Bellare-Rogaway models but are used to define both *safety*, akin to freshness, and partnering. A session is *safe* when honest client and server agree on these assignments, under suitable conditions on algorithms and long-term keys. As discussed above, our definition immediately releases all connection keys. We guarantee that the keys of safe sessions are indistinguishable from fresh random keys; this accounts for selective session key reveal and test queries in Bellare-Rogaway models. Additionally, we provide *verified safety*, that is, sufficient conditions on the recorded long-term keys that enable honest parties to infer that their session is safe.

Our main result (§5, Theorem 2) reduces the concrete security of the TLS handshake to agile assumptions on the constructions used for signatures, KEMs, and PRFs. Each epoch assigns a distinguished agility-parameter a , selecting all algorithms for the epoch. The theorem statement is parameterized by a predicate α on a that holds whenever all algorithms selected by a are (assumed to be) secure. Thus, it provides meaningful security only for epochs where $\alpha(a)$ holds, despite any other epochs. If α is always false, there is nothing to prove. If we care specifically about one ciphersuite, say `TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA`, we may apply our theorem with α set to true only when a selects that ciphersuite. This already improves on non-agile results for TLS that assume all honest parties agree *in advance* on a ciphersuite and reject any others.

Our model accounts for agility with respect to record algorithms, and yields channel security for mITLS without agile assumptions on the algorithms r used in the record layer. We thus validate the use of stateful LHAE [23] for clients and servers that negotiate r . We require, however, that no application data be sent before the Finished messages are verified. For implementations that violate this requirement [19], stronger agile assumptions seem unavoidable.

Code-Based Verified Implementation (§6). We finally present the reference implementation of the handshake we integrated into mITLS, and its verification against our security definition, based on the same modular proof structure but at a greater level of detail, relying on type-based verification for scalability. Our code supports the standard and commonly-used extensions; we tested it against various mainstream TLS clients and servers, using 4 versions ranging from SSL3 to TLS 1.2, 12 ciphersuites, and various subsets of extensions. It improves on the original mITLS code [4], which supported less features, and whose security relied on monolithic, TLS-specific assumptions for RSA and DH ciphersuites. The full paper reports experimental results showing that our code runs handshakes with reasonable performance. To enable its automated verification, our code is structured into small, independent modules (that is, program libraries) parameterized by algorithm descriptors. For instance, our library code for the

HMAC-based PRF used in TLS implements agility before calling selected core algorithms, e.g. SHA1. In contrast, the code that implements SHA1 is outside the scope of our verification effort—we document our agile cryptographic assumption on it, and call a standard library. Each cryptographic construction used in the handshake corresponds to a separate library in the code. We define the security of libraries for multiple keys and multiple algorithms; the corresponding definitions and reductions to single-key security of individual algorithms appear in the full paper.

In summary, our work sheds light on important design and implementation issues of TLS. To our knowledge, we provide the first provable-security results for TLS that account for algorithm agility. We are also the first to give an abstract security model for handshakes related by resumption and renegotiation.

Further reading. Our website <http://www.mitls.org> provides additional materials: the MITLS source code; the EASYCRYPT proof of Theorem 1; and a companion paper with empirical data on TLS handshakes, auxiliary definitions, constructions, and proofs, and extended discussions of attacks and related work.

2 Agile Signatures

An *agile signature scheme* consists of three algorithms: **KeyGen** is a standard key generation algorithm, while **Sign** and **Verify** take an extra agility parameter. For instance, given a core signature scheme $s = (\text{keygen}, \text{sign}, \text{verify})$, the hash-then-sign scheme $S_s = (\text{KeyGen}, \text{Sign}, \text{Verify})$ of TLS is defined as follows: **KeyGen** \triangleq **keygen** generates a key pair for algorithm s ; **Sign** (h, sk, m) \triangleq **sign** $(sk, h(m))$ computes a signature using the core scheme s and hash algorithm h ; and **Verify** (h, pk, m, σ) \triangleq **verify** $(pk, h(m), \sigma)$ verifies a purported signature σ for message m hashed with algorithm h . We define existential unforgeability under chosen-message attacks (EUF-CMA) for agile signatures.

Definition 1 (EUF-CMA). *Let $(\text{KeyGen}, \text{Sign}, \text{Verify})$ be an agile signature scheme, p^* a parameter, and P a set of parameters, and consider the following forgery game:*

<p>Game EUF \triangleq $pk, sk \leftarrow \text{KeyGen}(); M := \emptyset$ $m', \sigma \leftarrow \mathcal{A}^{\text{SIGN}}(pk)$ return $m' \notin M \wedge \text{Verify}(p^*, pk, m', \sigma)$</p>	<p>Oracle SIGN(p, m) \triangleq if $p \notin P$ then return \perp $M := M \cup \{m\}$ return Sign(p, sk, m)</p>
--	--

The scheme is (ϵ, t, p^, P) -secure against EUF-CMA if, for any \mathcal{A} that runs in time t , the EUF game returns **true** with probability at most ϵ .*

This definition generalizes plain EUF-CMA security; the two coincide for a scheme with fixed hash algorithm h , i.e. $(p^*, P) = (h, \{h\})$. We do not require $p^* \in P$; for instance, one may pragmatically assume that forging an MD5-based signature is hard when given only SHA1-based signatures. Indeed, the attacks of Stevens et al. [26] rule out $(\text{MD5}, \{\text{MD5}, \dots\})$ -security, but $(\text{MD5},$

$\{\text{SHA1}\}$)-security may still hold. On the other hand, non-agile security does not imply agile security. Consider for instance the scenario where the pre-image security of MD5 is broken. Then the attacks described by Naccache and Shparlinski [22] are likely to break $(\text{SHA256}, \{\text{MD5}, \text{SHA256}\})$ -security, even though $(\text{SHA256}, \{\text{SHA256}\})$ -security would still hold.

The TLS standard features the following hash-then-sign schemes: prior to version 1.2, RSA PKCS#1v1.5 signatures use the concatenation of MD5 and SHA1 hashes and (EC)DSA signatures use SHA1. TLS 1.2 introduces additional agility to facilitate migration from MD5 and SHA1 to stronger algorithms. Designers are aware of agility problems, and prescribe ad hoc countermeasures [9, §7.4.3]. The standard still requires that (EC)DSA use SHA1, delaying the migration to stronger algorithms. It also adds an encoding of the hash algorithm identifier to guarantee that all hash algorithms have disjoint range.

Given algorithms h and h' with disjoint ranges, if the core signature scheme itself is (ϵ, t) -EUF-CMA secure on their joint range, then we have $(\epsilon', t', h, \{h, h'\})$ -security for the corresponding agile hash-then-sign signature scheme, where the difference between ϵ, t and ϵ', t' depends on the reduction to the collision resistance of h . Sadly, the core signature schemes used in TLS are not EUF-CMA secure. The best we can do, for now, is thus to assume that the hash-then-sign signature scheme that uses them meets Definition 1.

3 Master Secrets & Key Encapsulation

Following [14, 17], we model the basic key-exchange functionality of TLS as different variations on KEMs. However, we separate the derivation of the master secret from the derivation of keys for the record-layer. We model the premaster secret phase for RSA and Diffie-Hellman exchanges as agile KEMs (`keygen`, `!enc`, `dec`) parameterized by a 2-byte protocol version string.

RSA. `keygen` generates a fresh RSA key pair (pk, sk) ; `enc` (pv, pk) appends a randomly chosen 46-byte string to pv to obtain the premaster secret pms , and returns it with the ciphertext c resulting from its PKCS#1v1.5 encryption under pk ; `dec` (pv, sk, c) decrypts c with sk . If the padding is correct and the decrypted pms is exactly 48 bytes long, it returns pms with the first 2 bytes replaced by pv , otherwise it returns \perp ; such errors are handled in our ms -KEM below.

Diffie-Hellman. `keygen` selects group parameters pp , generates a fresh pair of DH values (g^x, x) , and returns $pk = (pp, g^x)$ and $sk = (pk, x)$ as public and private KEM keys; `enc` $(pv, (pp, g^x))$ samples y and returns $pms = g^{xy}$ and $c = g^y$; `dec` $(pv, (pk, x), c)$ returns $c^x = g^{xy}$. The ciphertext space guarantees that c is in a large prime-order subgroup specified by pk . In contrast to the RSA pms -KEM, neither `enc` nor `dec` depend on pv .

On their own, these two premaster secret KEMs are *not* secure under any indistinguishability notion, even under relatively weak active attacks such as, for instance, plaintext-checking attacks (PCA): recall the Bleichenbacher attack, and the lack of active security for basic Diffie-Hellman (e.g., querying a plaintext-checking oracle on c^r and pms^r for any $r \neq 1$, suffices to distinguish a random

pms from the one encapsulated in c). Rather than using pms as a key, TLS feeds it through an agile *key extraction function* (KEF) parameterized by a hash algorithm, to compute the master secret ms .

We model this phase of the handshake as an *agile labeled KEM*, extending the labeled KEMs of [14, 17] with an agility parameter. Given an agile (unlabeled) pms -KEM $e = (\text{keygen}, \text{enc}, \text{dec})$ and an agile key extraction function family KEF, the master secret KEM $E_e = (\text{KeyGen}, \text{Enc}, \text{Dec})$ of TLS is defined as follows:

- $\text{KeyGen}() \triangleq \text{keygen}()$;
- $\text{Enc}(pv, h, pk, \ell) \triangleq pms, c \leftarrow \text{enc}(pv, pk)$; $ms \leftarrow \text{KEF}(pv, h, pms, \ell)$;
return ms, c
generates a premaster secret pms and a ciphertext c using e , then derives a master secret ms for ℓ using KEF.
- $\text{Dec}(pv, h, sk, \ell, c) \triangleq pms \leftarrow \text{dec}(pv, sk, c)$; **if** $pms = \perp$ **then** $pms \leftarrow pv \parallel \$$;
return $\text{KEF}(pv, h, pms, \ell)$
decrypts the ciphertext c to obtain pms . If decryption fails, it computes a fake pms by appending a random 46-byte string to pv (this is never the case for DH). It returns the value obtained from pms and ℓ using the agile KEF.

We assume sufficient checks to ensure that all arguments are well-formed before calling the master secret KEM algorithms; e.g., for Diffie-Hellman, our code validates group parameters and checks that pk and c belong to a large prime-order subgroup before calling Dec .

We define security for agile labeled KEMs as indistinguishability under replayable chosen-ciphertext attacks (IND-RCCA), a relaxation of CCA security, first introduced for public-key encryption by Canetti et al. [7].

Definition 2 (IND-RCCA). *Let $(\text{KeyGen}, \text{Enc}, \text{Dec})$ be an agile labeled KEM, p^* a parameter, P a set of parameters; and consider the following game:*

Game $\text{RCCA} \triangleq$	Oracle $\text{ENC}(\ell) \triangleq$	Oracle $\text{DEC}(p, \ell, c) \triangleq$
$pk, sk \leftarrow \text{KeyGen}()$	if $\ell \in L$ then return \perp	if $\ell \in L \vee p \notin P$ then return \perp
$K, L := \emptyset$	$k_0, c \leftarrow \text{Enc}(p^*, pk, \ell)$	$L := L \cup \{\ell\}$
$b \leftarrow \{0, 1\}$	$k_1 \leftarrow \$$	$k \leftarrow \text{Dec}(p, sk, \ell, c)$
$b' \leftarrow \mathcal{A}^{\text{ENC}, \text{DEC}}(pk)$	$K(\ell) := K(\ell) \cup \{k_0, k_1\}$	if $k \in K(\ell)$ then return \perp
return $(b' = b)$	return k_b, c	return k

The *RCCA advantage* of \mathcal{A} , $\text{Adv}_{p^*, P}^{\text{RCCA}}(\mathcal{A})$ is defined as $2 \Pr[\text{RCCA} : b' = b] - 1$. The scheme is (ϵ, t, p^*, P) -secure against IND-RCCA- n when the advantage of any adversary \mathcal{A} running in time t and making at most n queries to ENC is at most ϵ . We write *IND-RCCA* instead of *IND-RCCA-1*.

The check $\ell \in L$ in the decryption oracle reflects a property of TLS: honest servers decrypt at most once for each nonce. The check $\ell \in L$ in the encryption oracle is analogous to the restriction of Krawczyk et al. [17] to define IND-CCCA security for non-agile KEMs.

The lemma below (proved by a standard hybrid argument in the full paper) enables us to prove security for a single query, then use the multi-query variant for reasoning about TLS in our main theorem.

Lemma 1. *If a KEM (KeyGen, Enc, Dec) is $(\epsilon/n, t', p^*, P)$ -secure against IND-RCCA, then it is (ϵ, t, p^*, P) -secure against IND-RCCA- n , where $t' = t + O(n \cdot t_{\text{Enc}})$ and t_{Enc} is the worst-case cost of algorithm Enc.*

Next, we define the assumptions for our main theorem on the TLS master secret KEM: *non-randomizability under plaintext-checking attacks* (NR-PCA) and *one-wayness under plaintext-checking attacks* (OW-PCA).

Definition 3 (NR-PCA, OW-PCA). *Let (keygen, enc, dec) be an agile (un-labeled) KEM, p^* a parameter, and P a set of parameters. Consider the two games:*

Game OW-PCA \triangleq $pk, sk \leftarrow \text{keygen}()$ $k^*, c^* \leftarrow \text{enc}(p^*, pk)$ $k \leftarrow \mathcal{A}^{\text{PCO}}(pk, c^*)$ return $(k = k^*)$	Game NR-PCA \triangleq $pk, sk \leftarrow \text{keygen}()$ $k^*, c^* \leftarrow \text{enc}(p^*, pk)$ $c \leftarrow \mathcal{A}^{\text{PCO}}(pk, c^*)$ return $(c \neq c^* \wedge k^* = \text{dec}(p^*, sk, c))$	Oracle PCO (p, k, c) \triangleq if $p \notin P \vee k = \perp$ then return \perp $k' \leftarrow \text{Dec}(p, sk, c)$ return $(k' = k)$
--	---	--

The NR-PCA advantage of \mathcal{A} , $\text{Adv}_{p^*, P}^{\text{NR-PCA}}(\mathcal{A})$ is the probability that the NR-PCA game returns true. The KEM is (ϵ, t, p^*, P) -secure against NR-PCA if the advantage of any adversary \mathcal{A} running in time t is at most ϵ . OW-PCA advantage and security are defined analogously.

The full paper gives preliminary theorems and conjectures on these assumptions, and relates our agile IND-RCCA KEMs to prior work and more standard assumptions. We hope this will stimulate further cryptanalytic work on TLS.

Our main result on KEMs is that the generic ms -KEM E_e of TLS is IND-RCCA secure if the underlying pms -KEM e is both NR-PCA and OW-PCA secure. The proof (in the full paper) has been formalized using EASYCRYPT. The proof is in the random oracle model for the agile KEF. As explained above, we consider the single challenge case.

Theorem 1 (RCCA from NR-PCA and OW-PCA). *Let \mathcal{A} be a (p^*, P) -RCCA adversary for E_e running in time $t_{\mathcal{A}}$ and making at most q_{KEF} and q_{DEC} queries to the random and decryption oracle, respectively. Let $p^* = (pv^*, h^*)$ and $P' \triangleq \{pv \mid (pv, h) \in P\}$. There exist an OW-PCA adversary \mathcal{B} and an NR-PCA adversary \mathcal{C} against e , both running in time $t_{\mathcal{A}} + O(q_{\text{DEC}} \cdot q_{\text{KEF}})$, such that*

$$\text{Adv}_{p^*, P}^{\text{RCCA}}(\mathcal{A}) \leq 2 \left(\text{Adv}_{pv^*, P'}^{\text{NR-PCA}}(\mathcal{B}) + \text{Adv}_{pv^*, P'}^{\text{OW-PCA}}(\mathcal{C}) + 2^{|pv| - |pms|} (q_{\text{KEF}} + q_{\text{DEC}}) \right).$$

The factor $2^{|pv| - |pms|}$ is the entropy of the value $pv \parallel \$$ used to derive the master secret when RSA decryption fails, as recommended by TLS 1.2 to mitigate Bleichenbacher attacks. With the DH pms -KEM, decryption never fails (as the ciphertext validation is done beforehand) so the last term above can be omitted.

4 Defining Agile Security for Sequences of Handshakes

Our security definition for handshakes is general enough to apply to TLS, as specified in the standard and coded in mITLS, while hiding implementation details like message formats and specific cryptographic constructions. The adversary creates and interacts with multiple instances i of a handshake protocol Π by calling Π 's oracles, detailed below. Each instance has a fixed role \mathcal{R} , either \mathcal{C} for Client or \mathcal{S} for Server, and models a connection endpoint.

- $\text{KeyGen}(v)$ creates and stores a new honest keypair for the long-term public-key algorithm v (in TLS, ranging over s for signing and e for key encapsulation) and returns the associated public key. Similarly, $\text{KeyInject}(v, pk, sk)$ stores a dishonest keypair (assuming pk is not yet in the store).
- $\text{Init}(\mathcal{R}, \text{cfg}_{\mathcal{R}})$ creates an instance with role \mathcal{R} and local configuration $\text{cfg}_{\mathcal{R}}$; it returns a fresh handle i .
- $\text{Send}_i(\text{frag})$ lets an existing instance i process a fragment, depending on its current state. As a result, the instance may update its state, assign local variables, and return a response. (In TLS, responses range over sequences of handshake and CCS message fragments, intended to be sent to the peer, as well as error messages.)
- $\text{Control}_i(\text{env})$ changes the global, internal state of the handshake, e.g., enabling the adversary to control access to stored sessions and private keys by the protocol the next time Send will be called, or to trigger a renegotiation request. This single oracle accounts for many control functions in the mITLS handshake implementation. For example, Control provides the environment with means to reject certificates that it deems invalid.

Each instance maintains its private local state (e.g. using local variables). Each instance can go through a sequence of epochs (e.g. recording the number of cycles in the state machine). For each epoch, it records a sequence of *variable assignments*, extended as the result of calls to Send and Control . Each variable is assigned at most once in every epoch. The selection and ordering of assignments within an epoch depends on the protocol; for instance, a client epoch may assign its client-certificate variable, then send a message to the server, causing the server epoch to record the same assignment later in the protocol.

Our definition is based on local variable assignments, which summarize the view of clients and servers so far about each epoch. This is adequate to model the handshake as a component within TLS, but this differs from models based on *matching conversations* [3] that compare the (unparsed) messages they have sent and received so far. We use assignments to express the main goals of the protocol, for instance assigning a fresh random value to the record key variable k ; and agreeing on all assignments as a session completes. We list below the main variables used in our presentation, but our definition can account for a more detailed model of the TLS handshake.

Unless explicitly mentioned for key-exchange materials, these variables are public: the adversary can read them, but not change them; the protocol can

ℓ	epoch identifier; in TLS, the concatenation of the client and server random values.
$\ell_{session}$	resumption identifier; in TLS, the identifier of the epoch that completed the session being resumed. (The mTLS code also assigns the TLS <i>sessionId</i> , chosen by the server, but we do <i>not</i> use it as an identifier as it is not necessarily unique.)
ac, as	client and server negotiation parameters; in TLS, they consist of protocol versions, ciphersuites, and extension messages.
a	agility parameter; in TLS, the protocol version, the negotiated ciphersuite, and data extracted from the first flight of messages sent by the server.
$cert_C, cert_S$	client and server certificate chains. In TLS, these certificates are optional; e.g. the assignment $cert_C := \perp$ denotes the absence of client certificate.
ex_C, ex_S	client and server exchange variables, possibly secret, used to specify safety.
k	record key for the epoch; in TLS, depending on a , this key is usually split into 4 keys for MAC & encrypt.
$complete$	successful completion flag, marking the end of the handshake for this epoch.

write them once in every epoch, but not read them. (This restriction matters only for the record key, as we replace it with a random value.) The agility-parameter variable a determines the algorithms and constructions used by the handshake. Our security properties are conditioned by a strength predicate $\alpha(a)$ that indicates whether those algorithms are strong enough to secure the epoch. When the role of an epoch is clear from the context, the *peer* refers to the opposite role, and the *peer-exchange variable* refers to the exchange variable of the opposite role (e.g. ex_C when \mathcal{R} is \mathcal{S}).

We deliberately avoid modeling certificate validation. For the handshake, certificate chains are authenticated, uninterpreted bitstrings. We leave as future work supplementing our model with an application-level certificate infrastructure above the mTLS API. We assume given a public specification function $\text{pk}(cert)$ that returns either the public key associated with a certificate chain, or \perp . The session state does not need to explicitly mention public keys, but public keys can appear in exchange variables.

A security model for a protocol describes how queries are answered and how session variables are assigned. Next, we define properties of these models as they interact with an adversary.

Definition 4 (Honesty, Safety, Matching Algorithms and Completion).

For a handshake protocol Π and a strength predicate $\alpha(\cdot)$, an adversary that calls Π 's oracles any number of times produces a trace of interleaved variable assignments for a series of epochs for each instance. In this trace:

- *As determined by its assigned agility parameter a : an epoch is either a session, with distinguished client- and server-exchange variables, or a resumption, with an $\ell_{session}$ variable; sessions (and their exchange variables) are either static or ephemeral; a static session has at least one static exchange variable; an ephemeral session has only ephemeral exchange variables.*
- *A (long-term) public key is honest for algorithm v if it was returned by a call to $\text{KeyGen}(v)$. A session's ephemeral server-exchange variable assignment is*

- honest if there is a server session with the same assignment to its server-exchange variable—and conversely for ephemeral client-exchange variables.
- A client session is safe if (i) $\alpha(a)$ holds; (ii) honest public keys for a 's algorithms are assigned to all static exchange variables; and (iii) there is a server session with the same assignment to the ephemeral server-exchange variable. A server session is safe if the converse holds.
 - A resumption is safe if $\alpha(a)$ holds and ℓ_{session} is the identifier of a safe and complete session.
 - An epoch has matching algorithm $r = \text{record}(a)$ when there is a peer epoch with the same identifier ℓ and algorithm r .
 - An epoch is complete when it includes the assignment $\text{complete} := 1$.

Anticipating on §5, for TLS we define the client exchange value ex_C to be the master secret ms together with the KEM public key pk , and the server-exchange variable ex_S to be the public key pk of the KEM. The latter is static for TLS-RSA, but ephemeral for TLS-DHE. Here ms is explicitly secret and ephemeral.

Definition 5 (Handshake Security). *Let Π be a handshake protocol, $\alpha(\cdot)$ a strength predicate, and \mathcal{A} an adversary that calls Π 's oracles any number of times. Consider the following properties:*

- (1) **Uniqueness:** *epoch identifiers are used at most once in each role.*
Let $\text{Adv}^U(\mathcal{A})$ be the probability that two different epochs with the same role assign the same value to ℓ when \mathcal{A} terminates.
- (2) **Verified Safety:** *if the peer of a session uses a strong signature algorithm to authenticate and the public-key for the peer signature is honest, then the peer-exchange variable assignment is honest.*
Let $\text{Adv}^S(\mathcal{A})$ be the probability that, when \mathcal{A} terminates, there is an epoch such that $\alpha(a)$ holds; the public key of the peer is honest; and the assignment to the peer exchange value is not honest (i.e. not assigned by any peer);
- (3) **Agile Key Derivation:** *depending on a random bit b , replace the record key assigned in safe epochs with matching algorithm r with a fresh $k \leftarrow \text{KeyGen}(r)$, assigning the same value to epochs that have the same identifier ℓ , algorithms $\text{kdf}(a)$ and exchange variables or resumption identifier.*
Let $\text{Adv}^K(\mathcal{A}) = 2p - 1$ where p is the probability that \mathcal{A} returns b .
- (4) **Agreement:** *for every safe and complete epoch, there is a safe epoch in the other role such that their two instances agree on all prior assignments.*
Let $\text{Adv}^I(\mathcal{A})$ be the probability that, when \mathcal{A} terminates: an instance created by $\text{Init}(\overline{\mathcal{R}}, \text{cfg})$ assigns $\text{complete} := 1$ in a safe epoch; and no instance created by $\text{Init}(\overline{\mathcal{R}}, \text{cfg}')$ begins with a series of epochs with the same assignments to all variables (up to, but possibly excluding $\text{complete} := 1$).

The handshake is (ϵ, t, α) -secure when for any adversary \mathcal{A} running in time t , we have $\text{Adv}^G(\mathcal{A}) \leq \epsilon$, for $G = U, S, K, I$.

Discussion. The properties above are given in chronological order: in TLS in particular, protocol instances first exchange fresh random values, then derive keys, and finally confirm the integrity of the session negotiation.

Property (1) simply ensures that ℓ provides a unique identifier, later authenticated using (4); we use these identifiers for matching client and server sessions.

Property (2) enables, for instance, a client that trusts both the negotiated algorithm and the server certificates to deduce that its server-exchange variable is honest, and conclude that its session is safe.

Property (3) idealizes the derived key; this is key indistinguishability. Recall that TLS uses the key before the two parties actually agree on the record algorithms. Conservatively, (3) idealizes the key only when the record algorithms match. As Krawczyk et al. [17], our model does not consider forward secrecy.

Property (4) guarantees agreement on all variable assignments at the client and server instances since their creation, not just the assignments of the current epoch. Hence, as soon as one epoch safely completes, the peers agree also on all prior epochs on that connection—even those that were not safe, or not verifiably safe. For TLS, this property holds only thanks to the (mandatory) secure renegotiation extension, which links each epoch to its predecessor. This property is closely related to the TLS renegotiation results of Giesen et al. [11]. They additionally propose an extension of TLS that would guarantee agreement on the full stream of application data, not just the handshake epochs. On the other hand, our model and security definition also cover resumptions and RSA ciphersuites, which are not covered by their results. Unlike previous analyses of TLS, our definition accounts for session resumptions. Property (4) guarantees agreement on the new epoch identifier ℓ and the identifier $\ell_{session}$ of the resumed session (and hence on the new record keys), as long as the original session is safe. The epochs of the original session may be on a different connection, between a different pair of instances; for those instances, safety for the original session independently guarantees agreement on all its original variable assignments.

TLS applications often group connections that use the same session or the same long-term key, allowing them to share resources and access rights. For example, web browsers allow all connections to the same server to share resources via the Same Origin Policy. It may seem desirable to guarantee a strong relationship between such connections, but our Property (4) guarantees agreement only for the sequence of epochs over a single connection. Indeed, the natural extension of this property to multiple connections does not hold for TLS, as shown by the triple handshake attack of Bhargavan et al. [5]. In this attack, an unsafe server-authenticated session is resumed on a new connection and then renegotiated with a new safe mutually-authenticated session. For the new safe epoch, Property (4) retroactively guarantees agreement on the prior resumption, but *not* on the original unsafe session that was resumed. Consequently, it is possible for a client and server instance to have a safe epoch but inconsistent variable assignments for the session associated with a prior resumed epoch; this leads to a variety of attacks, similar to the renegotiation attacks of Ray [25]. A stronger agreement can be achieved either at the application level, by checking agreement on prior connections, or by a protocol extension that includes a hash of the log of the original session in resumption handshakes [5]; we leave the modeling of this extension and its security for future work.

Compared with classic key exchange models [3] and the key exchange part of ACCE [13], our definition yields useful additional properties. Property (4) guarantees agreement on the negotiation parameters a_C and a_S for safe and complete epochs, thereby preventing version and ciphersuite rollback attacks.

Our definition also provides (some) security for anonymous connections, which can be composed with other authentication mechanisms to achieve application security. For example, renegotiation with client and server certificates may provide mutual authentication on top of an initial, safe, but anonymous handshake. Late application-level, client password authentication may also yield mutual authentication, as illustrated by mITLS [4].

5 Proving Agile Security for TLS Handshakes

We are now ready to reduce the security of TLS handshakes to the security of agile signatures, KEMs and PRFs. We structure the proof to apply simultaneously to the protocol, illustrated in Figures 1 and 2, and to its mITLS implementation.

Figure 1 shows the assignments performed by a client instance and a server instance that run two successive, matching handshakes on the same connection: for both instances, a static session, followed by a (renegotiated) resumption. Figure 2 similarly shows the assignments for an ephemeral session. The agility parameter a of the handshake indicates which algorithm to use for each underlying functionality. We write for instance $a := \text{alg}_C(\text{cfg}_C, a_S)$ to retrieve a from the client configuration and the negotiation parameter of the server; $e, p := \text{kem}(a)$ to retrieve the core algorithm e and public parameter of the master secret KEM from a ; and $E_e.\text{Enc}$ for encryption using the master-secret KEM for e .

Our second main theorem reduces the security of TLS handshakes to their underlying algorithms, depending on a *strength predicate* on their agility parameters. Its proof (in the full paper) relies on intermediate definitions for multi-key libraries and, as a first step, uses hybrid arguments to lift security from our agile definitions to the multi-key setting.

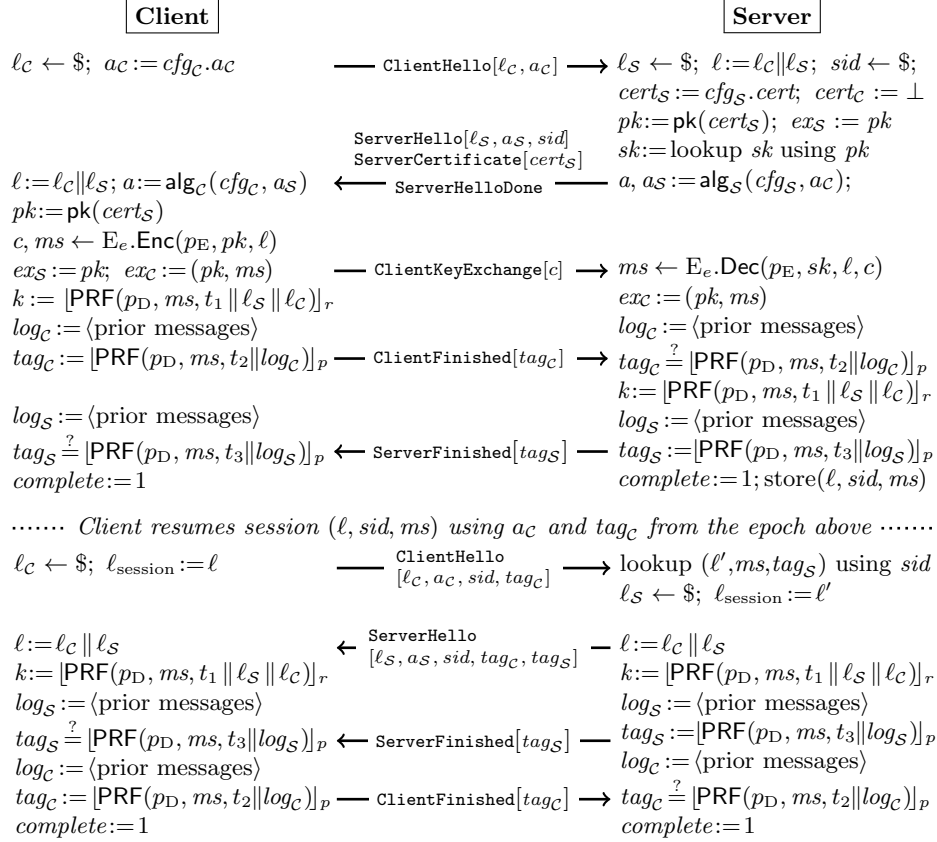
Theorem 2 (TLS Handshake). *Let a, a^* range over the agility parameters supported by TLS. Let $P_s = \{p^* \mid s, p^* := \text{sig}(a^*)\}$, $P_e = \{p^* \mid e, p^* := \text{kem}(a^*)\}$, and $P = \{p^* \mid p^* := \text{prf}(a^*)\}$. Let α be a strength predicate (Definition 4) such that the following assumptions hold:*

- (1) *If $\alpha(a)$ and $s, p := \text{sig}(a)$ then S_s is EUF-CMA $(\epsilon_{s,p}, t_{s,p}, p, P_s)$ -secure.*
- (2) *If $\alpha(a)$ and $e, p := \text{kem}(a)$ then E_e is IND-RCCA- $n_{ms}(\epsilon_{e,p}, t_{e,p}, p, P_e)$ -secure.*
- (3) *If $\alpha(a)$ and $p := \text{prf}(a)$ then PRF is an (ϵ_p, t_p, p, P) -secure PRF.*

Let n_s bound the number of calls to $S_s.\text{KeyGen}$. Let n and n_{ms} bound the number of epochs and sessions. Let n_e bound the number of calls to $E_e.\text{KeyGen}$, both for ephemeral and static KEMs. The TLS handshake is (ϵ, t, α) -secure, where

$$\epsilon = \sum_s \sum_p n_s \epsilon_{s,p} + \sum_e \sum_p n_e \epsilon_{e,p} + n_{ms} \sum_p \epsilon_p + n^2 (2^{-225} + 2^{-\min_p \lfloor \cdot \rfloor_p})$$

and where each t_ in the assumptions is at most t plus the cost of simulating Π in the reduction.*



Two epochs on the same connection: the first handshake establishes a session without client authentication using static keys; the second one resumes the session.

Conventions in the figure:

- (1) We use $\stackrel{?}{=}$ for checks; a failed check stops the instance.
- (2) We use $:=$ for assigning epoch variables; variables exchanged in messages are implicitly assigned, e.g. the server assigns ℓ_C and ac after parsing the first message.
- (3) We omit the extraction of the negotiated key exchange algorithm e and the parameters p_E, p_D from a ; for instance, we write p_D for $\text{prf}(a)$.
- (4) We omit **ChangeCipherSpec** messages: they are not part of the handshake protocol.
- (5) We write $\langle \text{prior messages} \rangle$ for the concatenation of all messages sent and received so far in the epoch, starting from the latest **ClientHello**.
- (6) We let $[\cdot]_r$ and $[\cdot]_p$ be functions that truncate to record-key and MAC sizes.
- (7) We let t_1, t_2, t_3 abbreviate the constant strings "derive key", "client finished", "server finished"; we write \parallel for bytestring concatenation.

Fig. 1. Abstract model of TLS handshake protocol (static handshake; resumption)

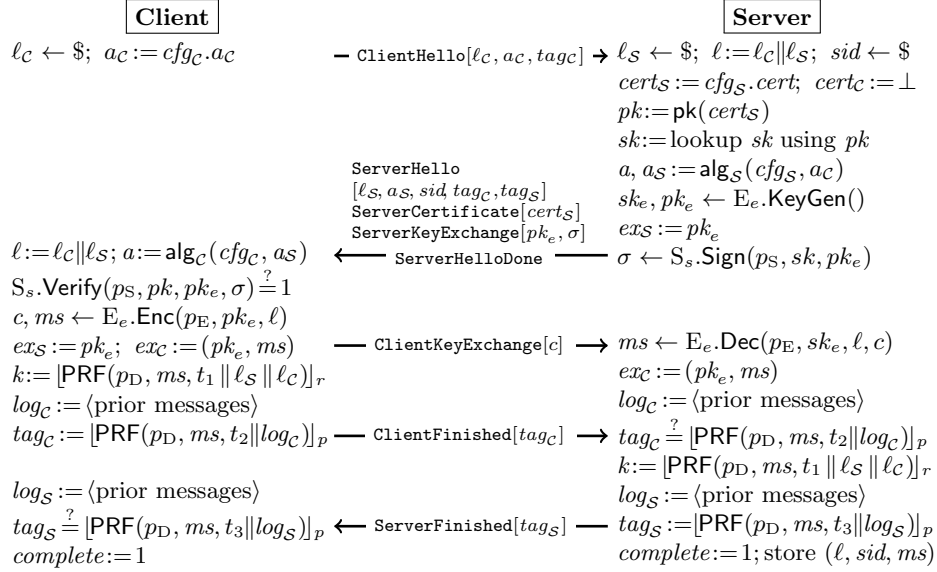


Fig. 2. Abstract model of TLS handshake protocol (ephemeral renegotiation)

Discussion. In the theorem, the sets P_s , P_e , and P represent the worst case. Indeed, signers may, for those keys that they consider honest, stop using signature algorithm s together with weak hash functions, like MD5, while TLS may still support verification using such hash algorithms for backward compatibility. To model such scenarios, one could instead add P_s , P_e , and P to the state of the experiment to record which hash algorithms have been used so far for signing, decrypting and deriving keys to obtain a more precise statement.

6 Verified Reference Implementation

We jointly programmed the TLS handshake and developed its proof. We finally outline our code, and explain how its structure and automated verification relate to the cryptographic models of §2–5; we provide additional details and performance results in the full paper. Our handshake implementation for MITLS consists of 3,600 lines of F# code plus 2,050 lines of F7 specifications; it supports four protocol versions, three key exchange mechanisms, two signature algorithms, and four hash functions. It deals mostly with the protocol aspects; indeed, our cryptographic proof for Theorem 1, conducted with EASYCRYPT, concerns less than 200 lines of F#. Conversely, Theorem 2 involves the full codebase and proving it requires a modular design and automated program verification techniques.

We adopt the type-based cryptographic verification method of Fournet et al. [10], previously applied to MITLS by Bhargavan et al. [4, §2]. The MITLS library consists of 45 modules, not counting application code or platform libraries. Each module implements a single cryptographic functionality or protocol component

and represents an abstraction boundary through its interface. A module is either trusted to be implemented correctly (e.g. the session database), or idealized under a cryptographic assumption (e.g. signatures) then verified, or perfectly verified (e.g. the protocol state machine). Each module interface specifies preconditions, postconditions, and type abstractions that govern the conditions under which secrets (keys, plaintexts, etc.) may be read or written by other modules.

We discuss the design of three important components that we modified during the course of this paper. *TLInfo* defines agility parameters and logical predicates (corresponding to α in Definition 4) that specify algorithmic strength, honesty for both long-term-keys and ephemeral secrets, matching record algorithms, and handshake completion events. This new logical model is more detailed than the original one [4]; furthermore, we extended the session structure and logical model to provide a general treatment of protocol extensions. *HandshakeMessages* implements message formatting and parsing; agreement (Definition 5(4)) depends on its details, since only formatted data is cryptographically authenticated. This code is complicated but not especially deep, and best handled using automated verification. *Handshake* implements the handshake state machine (*Send* in §5). Its code is not as simple as suggested by the KEMs of §3, since the TLS standard employs different sequences of messages for (say) RSA and DHE handshakes. Hence, we have similar but separate code for them, each of their interfaces complying with the KEM abstraction of §3. Also, our code handles errors and warnings, omitted in this presentation but also verified.

Our new results on the handshake, composed with prior results on mTLS [4] (the record layer, the top-level API, and various applications) yield agile, verified application security for TLS as it is.

References

1. T. Acar, M. Belenkiy, M. Bellare, and D. Cash. Cryptographic agility and its relation to circular encryption. In *EUROCRYPT 2010*, 2010.
2. G. Barthe, B. Grégoire, S. Héraud, and S. Zanella-Béguélin. Computer-aided security proofs for the working cryptographer. In *CRYPTO 2011*, 2011.
3. M. Bellare and P. Rogaway. Entity authentication and key distribution. In *CRYPTO'93*, 1993.
4. K. Bhargavan, C. Fournet, M. Kohlweiss, A. Pironti, and P.-Y. Strub. Implementing TLS with verified cryptographic security. In *IEEE Symposium on Security and Privacy*, 2013.
5. K. Bhargavan, A. Delignat-Lavaut, C. Fournet, A. Pironti, and P.-Y. Strub. Triple handshakes and cookie cutters: Breaking and fixing authentication over TLS. In *IEEE Symposium on Security and Privacy*, 2014.
6. D. Bleichenbacher. Chosen ciphertext attacks against protocols based on RSA encryption standard PKCS #1. In *CRYPTO'98*, 1998.
7. R. Canetti, H. Krawczyk, and J. B. Nielsen. Relaxing chosen-ciphertext security. In *CRYPTO 2003*, 2003.
8. R. Cramer and V. Shoup. Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack. *SIAM J. Computing*, 33(1):167–226, 2003.

9. T. Dierks and E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.2, 2008.
10. C. Fournet, M. Kohlweiss, and P.-Y. Strub. Modular code-based cryptographic verification. In *ACM CCS 11*, 2011.
11. F. Giesen, F. Kohlar, and D. Stebila. On the security of TLS renegotiation. In *ACM CCS 13*, 2013.
12. S. Haber and B. Pinkas. Securely combining public-key cryptosystems. In *ACM CCS 01*, 2001.
13. T. Jager, F. Kohlar, S. Schäge, and J. Schwenk. On the security of TLS-DHE in the standard model. In *CRYPTO 2012*, 2012.
14. J. Jonsson and B. S. Kaliski. On the security of RSA encryption in TLS. In *CRYPTO 2002*, 2002.
15. V. Klima, O. Pokorny, and T. Rosa. Attacking RSA-based sessions in SSL/TLS. In *CHESS 2003*, 2003.
16. F. Kohlar, S. Schge, and J. Schwenk. On the security of TLS-DH and TLS-RSA in the standard model. Cryptology ePrint Archive, Report 2013/367, 2013.
17. H. Krawczyk, K. G. Paterson, and H. Wee. On the security of the TLS protocol: A systematic analysis. In *CRYPTO 2013*, 2013.
18. A. Langley. Unfortunate current practices for HTTP over TLS, 2011. <http://www.ietf.org/mail-archive/web/tls/current/msg07281.html>.
19. N. M. Langley, A. and B. Moeller. Transport Layer Security (TLS) False Start. Internet Draft, 2010.
20. N. Mavrogiannopoulos, F. Vercauteren, V. Velichkov, and B. Preneel. A cross-protocol attack on the TLS protocol. In *ACM CCS 12*, 2012.
21. P. Morrissey, N. Smart, and B. Warinschi. A modular security analysis of the TLS handshake protocol. In *ASIACRYPT 2008*, 2008.
22. D. Naccache and I. E. Shparlinski. Divisibility, Smoothness and Cryptographic Applications. *ArXiv e-prints*, Oct. 2008.
23. K. G. Paterson, T. Ristenpart, and T. Shrimpton. Tag size does matter: Attacks and proofs for the TLS record protocol. In *ASIACRYPT 2011*, 2011.
24. Qualys SSL labs. SSL server test. <https://www.ssllabs.com/ssltest/analyze.html>.
25. M. Ray. Authentication gap in TLS renegotiation. http://extendedsubset.com/Renegotiating_TLS.pdf, 2009.
26. M. Stevens, A. Sotirov, J. Appelbaum, A. Lenstra, D. Molnar, D. A. Osvik, and B. de Weger. Short chosen-prefix collisions for MD5 and the creation of a rogue CA certificate. Cryptology ePrint Archive, Report 2009/111, 2009.
27. D. Wagner and B. Schneier. Analysis of the SSL 3.0 protocol. In *2nd USENIX Workshop on Electronic Commerce (WOEC'96)*, 1996.