Implementing and Proving the TLS 1.3 Record Layer

Karthikeyan Bhargavan[*]   Antoine Delignat-Lavaud[†]   Cédric Fournet[†]

Markulf Kohlweiss[†]   Jianyang Pan[*]   Jonathan Protzenko[†]   Aseem Rastogi[†]

Nikhil Swamy[†]   Santiago Zanella-Béguelin[†]   Jean Karim Zinzindohoué[*]

December 30, 2016

Abstract

The record layer is the main bridge between TLS applications and internal sub-protocols. Its core functionality is an elaborate authenticated encryption: streams of messages for each sub-protocol (handshake, alert, and application data) are fragmented, multiplexed, and encrypted with optional padding to hide their lengths. Conversely, the sub-protocols may provide fresh keys or signal stream termination to the record layer.

Compared to prior versions, TLS 1.3 discards obsolete schemes in favor of a common construction for Authenticated Encryption with Associated Data (AEAD), instantiated with algorithms such as AES-GCM and ChaCha20-Poly1305. It differs from TLS 1.2 in its use of padding, associated data and nonces. It encrypts the content-type used to multiplex between sub-protocols. New protocol features such as early application data (0-RTT and 0.5-RTT) and late handshake messages require additional keys and a more general model of stateful encryption.

We build and verify a reference implementation of the TLS record layer and its cryptographic algorithms in F⋆, a dependently typed language where security and functional guarantees can be specified as pre- and post-conditions. We reduce the high-level security of the record layer to cryptographic assumptions on its ciphers. Each step in the reduction is verified by typing an F⋆ module; when the step incurs a security loss, this module precisely captures the corresponding game-based security assumption.

We first verify the functional correctness and injectivity properties of our implementations of one-time MAC algorithms (Poly1305 and GHASH) and provide a generic proof of their security given these properties. We show the security of AEAD given any secure one-time MAC and PRF. We extend AEAD, first to stream encryption, then to length-hiding, multiplexed encryption. Finally, we build a security model of the record layer against an adversary that controls the TLS sub-protocols. We compute concrete security bounds for the AES-GCM and ChaCha20-Poly1305 ciphersuites, and derive recommended limits on sent data before re-keying. Combining our functional correctness and security results, we obtain the first verified implementations of the main TLS 1.3 record ciphers.

We plug our implementation of the record layer into an existing TLS library and confirm that the combination interoperates with Chrome and Firefox, and thus that experimentally the new TLS record layer (as described in RFCs and cryptographic standards) is provably secure.

---

[*]INRIA Paris-Rocquencourt. E-mail: `karthikeyan.bhargavan@inria.fr`, `panyang314@gmail.com`, `jean-karim.zinzin-dohoue@inria.fr`

[†]Microsoft Research. E-mail: `{antdl,fournet,markulf,protz,aseemr,nswamy,santiago}@microsoft.com`

1 Introduction

Transport Layer Security (TLS) is the main protocol for secure communications over the Internet. With the fast growth of TLS traffic (now most of the Web [51]), numerous concerns have been raised about its security, privacy, and performance. These concerns are justified by a history of attacks against deployed versions of TLS, often originating in the record layer.

History and Attacks Wagner and Schneier [52] report many weaknesses in SSL 2.0. The MAC construction offers very weak security regardless of the encryption strength. The padding length is unauthenticated, allowing an attacker to truncate fragments. Stream closure is also unauthenticated; although an end-of-stream alert was added in SSL 3.0, truncation attacks persist in newer TLS versions [13, 47].

The MAC-pad-encrypt mode is not generically secure [33], and is brittle in practice, despite encouraging formal results [2, 12, 42]. Many padding oracle attacks have surfaced over the years, ranging from attacks exploiting straightforward issues (e.g. implementations sending padding error alerts) to more advanced attacks using side channels (such as Lucky13 [1] or POODLE [40]). Even though padding oracle attacks are well known, they remain difficult to prevent in TLS implementations and new variants tend to appear regularly [48]. The CBC mode of operation is also not secure against chosen-plaintext attacks when the IV is predictable (as in TLS 1.0), which is exploited in the BEAST attack [22]. Random explicit IVs [16] and CBC for 64-bit block ciphers mode [11] are also vulnerable to birthday attacks. Finally, fragment compression can be exploited in adaptive chosen plaintext attacks to recover secrets [43].

Even with provably-secure algorithms, functional correctness and memory safety are essential to preserve security guarantees that implementation bugs can easily nullify. For instance, the OpenSSL implementation of Chacha20-Poly1305 has been found to contain arithmetic flaws [15] and more recently, a high severity buffer overflow vulnerability [50].

Changes in TLS 1.3 The IETF aims to robustly fix the weaknesses of the record layer by adopting a single AEAD mode for all ciphersuites, thus deprecating all legacy modes (MAC-only, MAC-pad-encrypt, RFC 7366 [28] encrypt-then-MAC, compress-then-encrypt). The new AEAD mode is designed to be provably-secure and modular, supporting algorithms such as AES-GCM, AES-CCM, and ChaCha20-Poly1305 within the same framework. The usage of AEAD has also been improved: authentication no longer relies on associated data, whereas implicit nonces derived from initialization vectors (IV) and sequence numbers yield better security and performance.

What is the Record Layer? In the key exchange literature, a common viewpoint is to treat each key generated in the key schedule as belonging to a specific, independent application. Under this model, the handshake encryption key is only used by the handshake to encrypt its own messages, and must be separate from the application data key used only to encrypt application data fragments. This model does not fit the actual use of keys in any version of TLS: it fails to capture TLS 1.2 renegotiation (where handshake messages are interleaved with the application data stream), TLS 1.3 post-handshake authentication and re-keying, or even alerts in any TLS version. In our modularization of TLS, following Bhargavan et al. [12], we consider that each sub-protocol of TLS—handshake, change cipher spec (CCS), alert and application data (AppData)—defines its own data stream. The role of the record is to multiplex all of these streams into one, corresponding to network messages after fragmentation, formatting, padding, and optional record-layer encryption. Under this model, the record layer is the exclusive user for all non-exported keys generated by the key schedule, and there is no need to assign keys to any given sub-protocol stream.

Figure 1 illustrates the stream multiplexing for a TLS 1.3 handshake with 0-RTT data and one re-keying from the point of view of the client. Separate channels are used for writing and reading. Within each channel, a band in the figure represents a stream, and arrows represent message fragments (incoming
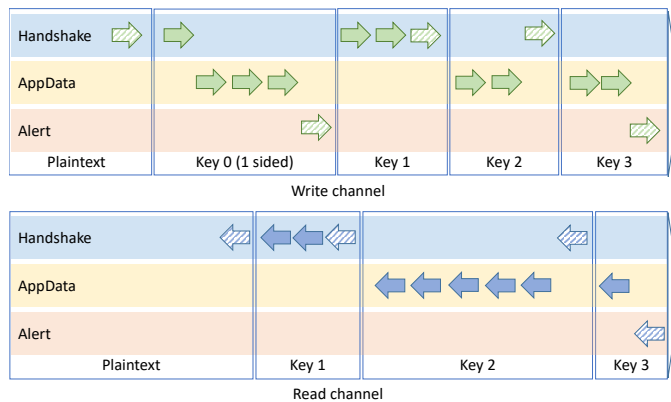
Figure 1: Multiplexing of sub-protocol streams by the record layer, depicting a TLS 1.3 0-RTT handshake with re-keying.

for left arrows, outgoing for right arrows) over time (flowing from left to right). Dashed arrows represent fragments used to signal key changes to the record layer. In TLS 1.2, CCS messages signal key changes; in TLS 1.3 this function is taken over by handshake and alert messages.

Related Work Since the first draft of TLS 1.3 in April 2014, the handshake and key schedule have undergone significant analysis efforts [18, 20, 25, 30, 31, 35] as it evolved over 18 iterations (at the time of writing). In contrast, few authors have analyzed changes to the record layer: Fischlin et al. [24] and Badertscher et al. [3] analyze an early draft that did not feature many of the current changes (for instance, it still relied on associated data to authenticate record meta-data), and Bellare and Tackmann [7] specifically focus on the way nonces are derived from IVs. This discrepancy may be explained by the difficulty of analyzing the record independently of the handshake protocol, and more generally, of defining the precise scope of its functionality.

As noted by several authors [21, 29], TLS occasionally uses the same keys to encrypt handshake messages and application data, e.g., for finished and post-handshake messages. This, unless carefully modeled in a custom security model [12] breaks key-indistinguishability. Authenticated and Confidential Channel Establishment (ACCE) [29, 36] is a game-based model that combines the handshake and the record protocols to preempt these composition problems. While ACCE models capture complex features of TLS such as renegotiation [27], its focus is primarily on the handshake, and it is unclear how to capture features such as post-handshake authentication [21]. Other limits of ACCE models are discussed in [3]. A recent paper by Krawczyk [34] proposes a new security definitions based on a key usability [19] rather than key indistinguishability to address this challenge.

Our contributions We contribute a reference implementation of the TLS record layer and its underlying cryptographic algorithms. We define its security as an indistinguishability game and we show a reduction with concrete bounds (Table 1) for any distinguisher to standard, low-level cryptographic assumptions. Our proof follows the structure depicted in Figure 2; from the bottom up:

1. We build a generic library for one-time message authentication codes (MAC) based on the Wegman-Carter-Shoup construction. We implement GHASH and Poly1305, and prove the functional correctness (w.r.t. mathematical specifications), memory safety, and encoding injectivity of the corresponding low-level implementations (§3). Similarly, we build a library for pseudo-random functions (PRF), and
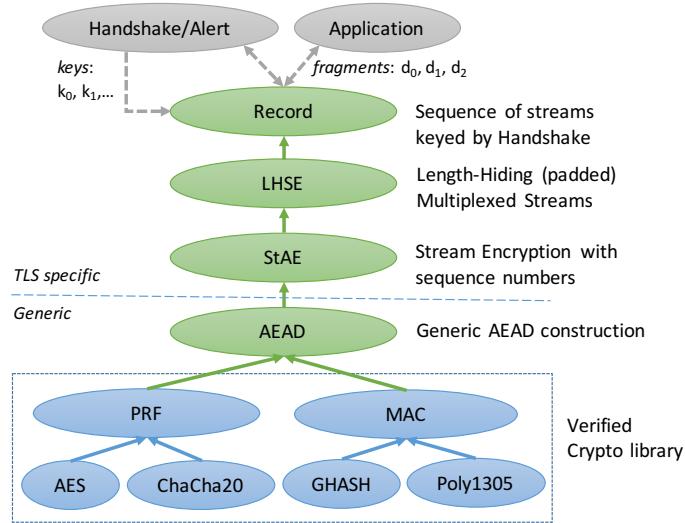
3

Figure 2: Modular structure of our proof. Green arrows denote security reductions proved by typing.

provide memory-safe implementations for AES and Chacha20 (§4). We express the security guarantees of both libraries as idealizations that we justify cryptographically.

2. We describe a generic AEAD construction that captures both RFC 5288 [44] for AES-GCM (as described in NIST SP800-38D [23]) and RFC7539 [41] for ChaCha20-Poly1305 through an interface compatible with RFC5116 [39]. We show that this construction satisfies a standard notion of AEAD security (§5) that combines indistinguishability under chosen-plaintext attacks (IND-CPA) with ciphertext integrity (INT-CTXT). Our proof applies to our F* implementation, and is verified by typing from these libraries upward.

3. From AEAD, we build and verify stream encryption, which uses AEAD nonces and record sequence numbers according to the TLS version-specific format (§6).

4. From stream encryption, we build a length-hiding encryption interface by adding padding, the TLS-specific content type multiplexing, and version-specific associated data (§7).

5. From length-hiding stream encryption with multiplexing, we implement the TLS record layer by adding interfaces to the handshake and alert sub-protocols that extend streams to sequences of streams by installing and enabling keys (§8). This captures novel protocol features of TLS 1.3 such as early application data (0-RTT and 0.5 RTT), late handshake messages, and re-keying. Based on our security bound, we propose a re-keying strategy that compensates for potential weaknesses in AES-GCM.

6. We evaluate our implementation of the TLS record layer (§9) by linking our AES-GCM and ChaCha20-Poly1305 ciphersuites to the handshake implementation of miTLS [12]. We confirm network interoperability with other TLS libraries both for TLS 1.2 and draft 18 of TLS 1.3. Our code and formal development maximize reuse between TLS 1.2 and 1.3.

Additional Materials This work is part of a larger project (https://project-everest.github.io/) which aims to build and deploy a verified secure HTTPS stack. The code presented in the paper is split

between several open-source projects, listed below. For convenience, a script available from `git clone https://github.com/project-everest/everest --branch record-layer-tr` installs these projects and points to the relevant code version.

- The F⋆ programming language: https://github.com/FStarLang/FStar/ including libraries and sample code, notably our verified code for the AEAD algorithms and constructions. We also rely on a new compiler from F⋆ to C code, available at https://github.com/FStarLang/kremlin.

- A new version of miTLS, extended to TLS 1.3, implemented in F⋆: https://github.com/mitls/mitls-fstar/. While verification of the full codebase is in progress, our repository contains verified protocol-specific code for the record layer fragment reported in this paper.

- The HACL⋆ cryptographic library: https://github.com/mitls/hacl-star/ including verified implementations for additional algorithms, as well as standalone security applications.

2 Compositional Verification by Typing

To implement and verify the record layer, we adopt a compositional approach to functional correctness and cryptographic security based on F⋆ [49], a dependently-typed programming language. This section explains our approach on two examples: arithmetic in a prime field used in Poly1305, and basic authenticated encryption. We refer the reader to [26] for a general presentation of this approach and [4] for a probabilistic semantics of F⋆ and additional cryptographic examples.

We use F⋆ not only to implement cryptographic constructions, but also as the formal syntax for their game-based security specifications. This is akin to the approach taken by Bhargavan et al. [14] in their proof of a TLS 1.2 handshake implementation using F7, an ancestor of F⋆. In contrast to F7, F⋆ supports an effectful programming style that is more efficient and closer to the cryptographic pseudo-code of code-based games [5]. For most of the presentation we use such pseudo-code instead of more precise and verbose F⋆ code. We do not assume familiarity with F⋆ and we use a simplified syntax that elides many details, especially in type annotations that are not relevant for the developments in this paper.

Functional Correctness of Poly1305 In F⋆, we specify arithmetic in the field $GF(2^{130} - 5)$ for the Poly1305 MAC algorithm as follows:

```
val p = 2^130 − 5 (∗ the prime order of the field ∗)
type elem = n:nat {n < p} (∗ abstract field element ∗)
let x +@ y : Tot elem = (x + y) % p (∗ field addition ∗)
let x *@ y : Tot elem = (x * y) % p (∗ field multiplication ∗)
```

This code uses F⋆ infinite-precision mathematical integers to define the prime order p of the field and the type of field elements. (The formula $\{n < p\}$ states that this type is inhabited by natural numbers n smaller than p.) It also defines two infix operators for addition and multiplication in the field in terms of arithmetic on infinite-precision integers. Their result is annotated with Tot elem, to indicate that these operations are pure total functions that return field elements. The F⋆ typechecker automatically checks that the result is in the field; it would trigger an error if e.g. we omitted the reduction modulo p. These operations are convenient to specify polynomial computations (see §3.2) but highly inefficient.

Instead, typical 32-bit implementations of Poly1305 represent field elements as mutable arrays of 5 unsigned 32-bit integers, each holding 26 bits. This representation evenly spreads out the bits across the integers, so that carry-overs during arithmetic operations can be delayed. It also enables an efficient modulo

operation for p. We show below an excerpt of the interface of our lower-level verified implementation, relying on the definitions above to specify their correctness.

```
abstract type repr = buffer UInt32.t 5 (* 5-limb representation *)
val select: memory → r:repr → Tot elem (* current value held in r *)

val multiply: e0:repr → e1:repr → ST unit
  (requires live e0 ∧ live e1 ∧ disjoint e0 e1)
  (modifies e0)
  (ensures select e0' = select e0 *@ select e1)
```

The type repr defines the representation of field elements as buffers (mutable arrays) of 5 32-bit integers. It is marked as abstract, to protect the representation invariant from the rest of the code. Functions are declared with a series of argument types (separated by →) ending with a return type and an effect (e.g. Tot or ST). Functions may have logical pre- and post-conditions that refer to their arguments, their result, and their effects on the memory. If they access buffers, they typically have a pre-condition requiring their caller to prove that the buffers are 'live' in the current memory. They also explicitly state which buffers they modify.

The total function select is used only in specifications; it reads the value of an element from the program memory. We use it, for example, in the stateful specification of multiply. In the types above, we keep the memory argument implicit, writing select e and select e' for the values of e in initial and final memories, respectively. (In real F* code, pre- and post-conditions take these memories as explicit arguments.)

The multiply function is marked as ST, to indicate a stateful computation that may use temporary stack-based allocations. It requires that its arguments e0 and e1 be live and disjoint; it computes the product of its two arguments and overwrites e0 with the result. Its post-condition specifies the result in terms of the abstract field multiplication of the arguments.

Implementing and proving that multiply meets its mathematical specification involves hundreds of lines of source code, including a custom Bignum library with lemmas on integer representations and field arithmetic (see §9). Such code is easy to get wrong, but once F* typechecks it, we are guaranteed that our low-level code is safe (e.g. it never accesses buffers out of bound, or de-allocated buffers) and functionally correct (since their results are fully specified). All F* types and specifications are then erased, hence the compiled code only performs efficient low-level operations.

Authenticated Encryption: Real Interface Let us consider a simplified version of the authenticated encryption (AE) functionality at the core of the TLS record layer. In F*, we may write an AE module with the following interface:

```
val ℓ_p: nat
val ℓ_c: nat
type lbytes (ℓ:nat) = b:bytes{length b = ℓ}
type bbytes (ℓ:nat) = b:bytes{length b ≤ ℓ}
type plain = lbytes ℓ_p
type cipher = lbytes ℓ_c
abstract type key
val keygen: unit → ST key
val decrypt: key → cipher → Tot (option plain)
val encrypt: k:key → p:plain → ST (c:cipher{decrypt k c = Some p})
```

Plaintexts and ciphertexts are represented as immutable bytestrings of fixed lengths ℓ_p and ℓ_c. We frequently rely on type abbreviations to statically enforce length checks for fixed-length bytestrings using

lbytes ℓ, and for bounded-length bytestrings using bbytes ℓ. (Our presentation uses immutable bytestrings for simplicity, whereas our record-layer implementation also uses mutable buffers of bytes.)

The interface defines an abstract type key; values of this type can only be generated via keygen and accessed via encrypt and decrypt. The internal representation of keys is hidden from all other modules to protect their integrity and secrecy.

The function keygen needs to generate randomness by calling an effectful external function; so we give this function the ST effect to indicate that the computation is impure and stateful (even though it does not explicitly modify the memory.) In particular, two calls to keygen may result in different results. The function encrypt would typically generate a nonce for use in the underlying AE construction, and hence is also marked as stateful. In contrast, decrypt is deterministic, so is marked with the Tot effect. Its result is an optional plain: either Some p if decryption succeeds, or None otherwise.

Our interface does not express any security guarantees yet, but it does require a functional correctness guarantee, namely that decryption undoes encryption. Besides, the F* type system implicitly checks for memory safety and integer overflows.

Authenticated Encryption: Security Given an implementation of AE, one usually measures its concrete security as the advantage of an adversary $\mathcal{A}$ that attempts to guess the value of b in the following game:

$$
\begin{array}{l}
\textbf{Game } \mathsf{Ae}(\mathcal{A}, \mathsf{AE}) \\
\hline
b \xleftarrow{\$} \{0,1\};\ L \leftarrow \varnothing;\ k \xleftarrow{\$} \mathsf{AE.keygen}() \\
b' \leftarrow \mathcal{A}^{\mathsf{Encrypt},\mathsf{Decrypt}}();\ \textbf{return } (b \overset{?}{=} b')
\end{array}
$$

$$
\begin{array}{l}
\textbf{Oracle } \mathsf{Encrypt}(p) \\
\hline
\textbf{if } b \textbf{ then } c \xleftarrow{\$} \mathsf{byte}^{\ell_c};\ L[c] \leftarrow p \\
\textbf{else } c \leftarrow \mathsf{AE.encrypt}\ k\ p \\
\textbf{return } c
\end{array}
\qquad
\begin{array}{l}
\textbf{Oracle } \mathsf{Decrypt}(c) \\
\hline
\textbf{if } b \textbf{ then } p \leftarrow L[c] \\
\textbf{else } p \leftarrow \mathsf{AE.decrypt}\ k\ c \\
\textbf{return } p
\end{array}
$$

The adversary $\mathcal{A}$ is a program that can call the two oracle functions to encrypt and decrypt using a secret key k. In the real case ($b = 0$) they just call the real AE implementation. In the ideal case ($b = 1$), Encrypt replaces the ciphertext with a fresh random bytestring of the same length and logs the encryption in L, while Decrypt performs decryption by a lookup in the log, returning either a plaintext recorded earlier or $\perp$ when lookup fails. (In devising the Decrypt oracle we had a choice: we could instead have returned $\perp$ when queried on a challenge ciphertext already recorded in L regardless of b. This strategy of rejecting challenge ciphertexts is idiomatic in cryptographic definitions in which Decrypt always returns $\perp$ when $b = 1$. We decided to allow decryption of challenge ciphertexts when $b = 0$ and implement the ideal behavior when $b = 1$ using table lookup. This allows us to more closely align game-based definitions with idealized cryptographic libraries.) Ideal AE is perfectly secure, inasmuch as the ciphertext does not depend on the plaintext. Thus, we define AE security by saying that the attacker cannot easily distinguish between the ideal and real cases.

For this game, we define $\mathcal{A}$'s advantage probabilistically as $|2 \Pr[\mathsf{Ae}(\mathcal{A}, \mathsf{AE})] - 1|$, e.g. an adversary flipping a coin to guess b will succeed with probability $\frac{1}{2}$ and has 0 advantage.

We adopt a more flexible notation for indistinguishability games: we keep the sampling of b and the call to the adversary implicit, and instead indicate the oracles available to this adversary. Hence, we write the game above (with the same oracles) equivalently as

$$
\begin{array}{l}
\textbf{Game } \mathsf{Ae}^b(\mathsf{AE}) \\
\hline
L \leftarrow \varnothing\ ;\ k \xleftarrow{\$} \mathsf{AE.keygen}();\ \textbf{return } \{\mathsf{Encrypt}, \mathsf{Decrypt}\}
\end{array}
$$

This notation facilitates the re-use of oracles for building other games, much like $F^\star$ modules. In general, we write G^b to refer to an indistinguishability game G where the adversary $\mathcal{A}$ tries to guess the value of the random bit b by calling the oracles returned by G. For all such games, we equivalently define the advantage as $\left| \Pr[\mathcal{A}^{G^1} = 1] - \Pr[\mathcal{A}^{G^0} = 1] \right|$.

Embedding games in $F^\star$ modules Although we wrote the game Ae^b in pseudo-code, each game in this paper reflects a verified $F^\star$ module, also written e.g. AE^b, that uses a boolean flag b to select between real and ideal implementations of the underlying cryptographic module AE. For example, AE^b may define the key type and encrypt function as

```
abstract type key = {key: AE.key; log: encryption_log}
let encrypt (k:key) (p:plain) =
  if b then
    let c = random_bytes ℓ_c in
    k.log ← k.log ++ (c,p);
    c
  else AE.encrypt k.key p
```

where the (private) key representation now includes both the real key and the ideal encryption log. The encrypt function uses k.log to access the current log, and ++ to append a new entry in the log, much as the Encrypt oracle.

Idealization Interfaces The idealized module AE^b can be shown to implement the following typed interface that reflects the security guarantee of the Ae^b game:

```
abstract type key
val log: memory → key → Spec (seq (cipher × plain)) (* reads k.log *)
val keygen: unit → ST k:key
  (ensures b ⇒  log k' = ∅)
val encrypt: k:key → p:plain → ST (c:cipher)
  (ensures b ⇒  log k' = log k ++ (c,p))
val decrypt: k:key → c:cipher  → ST (o:option plain)
  (ensures b ⇒  o = lookup c (log k))
```

The interface declares keys as abstract, hiding both the real key value and the ideal log, and relies on the log to specify the effects of encryption and decryption. To this end, it provides a log function that reads the current content of the log—a sequence of ciphertexts and plaintexts. This function is marked as Spec, indicating that it may be used *only in specification* and will be discarded by the compiler after typechecking.

Each of the 3 ensures clauses above uses this ghost function to specify the state of the log before (log k) and after the call (log k'). Hence, the interface states that, in the ideal case, the function keygen creates a key with an empty log; encrypt k p returns a ciphertext c and extends the log for k with an entry mapping c to p; and decrypt k c returns exactly the same result as a lookup for c in the current k.log. This post-condition formally guarantees that decrypt succeeds if and only if it is passed a ciphertext that was generated by encrypt; in other words it guarantees both functional correctness and authentication (a notion similar to INT-CTXT).

AE^b is also parameterized by a module $Plain^b$ that defines abstract plaintexts, with an interface that allows access to their concrete byte representation only when $b = 0$ (for real encryption). By typing AE^b, we verify that our idealized functionality is independent (information-theoretically) from the actual values of the plaintexts it processes.

From the viewpoint of the application, the plaintext abstraction guarantees that AE^1 preserves the confidentiality of encrypted data (as in classic information flow type systems), and its security guarantees can be used to build higher-level authentication guarantees. For instance, the application may prove, as an invariant, that only well-formed messages are encrypted using a given key, and thus that parsing and processing of any decrypted message always succeeds.

Probabilistic Semantics In $\mathrm{F}^\star$, we model randomness generation (e.g. random_bytes) using primitive probabilistic sampling functions, returning e.g. true or false with probability $\frac{1}{2}$. Two Boolean terminating $\mathrm{F}^\star$ expressions A^0 and A^1 are equivalent, written $A^0 \approx A^1$, when they return true with the same probability. They are ϵ-equivalent when $\epsilon = |\Pr[A^1 \Downarrow \mathsf{true}] - \Pr[A^0 \Downarrow \mathsf{true}]|$ where $A \Downarrow v$ denotes that program A evaluates to value v in the language semantics of $\mathrm{F}^\star$. These definitions extend to program evaluation contexts, written $A^b[\_]$, in which case ϵ depends on the program plugged into the context, which intuitively stands for the adversary. Equipped with these definitions, we can develop code-based game-playing proofs following the well-established approach of Bellare and Rogaway [5] directly applied to $\mathrm{F}^\star$ programs rather than pseudo-code.

For example, we can reformulate AE security as $\mathsf{AE}^1[\mathcal{A}] \approx_\epsilon \mathsf{AE}^0[\mathcal{A}]$, where $\mathcal{A}$ now ranges over well-typed Boolean programs parameterized by the two functions encrypt and decrypt defined by AE^b. Our definition of ϵ-equivalence between real and ideal implementations of AE^b closely matches the definition of $\mathcal{A}$'s advantage in the Ae^b game.

Concrete security definitions and reductions As illustrated for AE below, our security definitions consist of a game and a notation for the adversary advantage, parameterized by some measures of their use of the oracles (e.g. how many times the oracle is called). We intend to provide concrete bounds on those advantages, as a function of their parameters. To this end, our reduction theorems will relate this advantage for a given construction to the advantages of its building blocks.

Definition 1 (AE-security). *Given* AE, *let* $\epsilon_{\mathsf{Ae}}(\mathcal{A}[q_e, q_d])$ *be the advantage of an adversary* $\mathcal{A}$ *that makes* q_e *queries to* Encrypt *and* q_d *queries to* Decrypt *in the* $\mathsf{Ae}^b(\mathsf{AE})$ *game.*

We can either assume that this definition holds for our real AE module with an ϵ that is small for realistic adversaries (possibly relying on functional correctness and some prior proof of security), or we can prove that our AES-GCM module (say) achieves some bound on ϵ, by reduction to a simpler assumptions on the AES cipher. In later sections, we will show how we can precisely compute the adversary $\mathcal{A}$'s advantage in the game above from a related adversary $\mathcal{B}$'s advantage in winning the PRF game on the underlying cipher (e.g. AES). The proof relies on standard cryptographic game transformations that are applied manually at the level of $\mathrm{F}^\star$ code, combined with functional correctness proofs about the real and ideal code, verified automatically by $\mathrm{F}^\star$.

Games vs Idealized Modules There are several differences between the games we present and the actual modules of our implementation. Standard-compliant modules include many details elided in informal games; they also use lower-level representations to yield more efficient code, and require additional type annotations to keep track of memory management.

These modules are part of a general-purpose verified cryptographic libraries, providing real functionality (when idealizations flags are off), so they always support multiple instances of their functionality. Here, AE^b has a function to generate keys, passed as parameters to the encrypt function, whereas the game oracle uses a single, implicit key. (This difference can usually be handled by a standard hybrid-argument reduction.)

Modules preferably rely on the $\mathrm{F}^\star$ type system to enforce the rules of the games. Hence, dynamic checks in games (say, to test whether a nonce has already been used) are replaced with static pre-conditions on

typed adversaries. Similarly, types enforce many important but trivial conditions, such as the length of oracle arguments, often kept implicit in the paper.

3 One-Time MACs

We begin with constructions for MACs. As detailed in §5, AEAD uses fresh key materials for each message authentication, so we consider their security when MACing just once.

We treat the two main constructions, GHASH and Poly1305, using the same definitions, code, and proofs, inasmuch as possible. We initially suppose that the whole key is freshly generated for each MAC (as in ChaCha20-Poly1305), before presenting the general case where a part of the key is shared between multiple MACs (as in AES-GCM).

3.1 One-time MAC functionality and security

We outline below our interface for message authentication code (MAC), omitting its functional specification (see §9).

> val ℓ_{k_0}: nat   *(∗ static key length, may be 0 ∗)*
> val ℓ_k: n:nat $\{\ell_{k_0} \leq \ell_k\}$ *(∗ complete key length ∗)*
> val ℓ_t: nat   *(∗ tag length ∗)*
> val ℓ_m: nat   *(∗ maximal message length ∗)*
> type key_0 = lbytes ℓ_{k_0} *(∗ static key shared between MACs ∗)*
> type key = lbytes ℓ_k *(∗ one−time key (including static key) ∗)*
> type tag = lbytes ℓ_t *(∗ authentication tag ∗)*
> type message = b:bbytes ℓ_b {wellformed b}
> val keygen$_0$: unit → ST key_0
> val keygen: key_0 → ST key
> val verify: key → message → tag → Tot bool
> val mac: k:key → m:message → Tot (t:tag{verify k m t})

This interface defines concrete byte formats for keys, tags, and messages. authenticated messages are strings of at most ℓ_m bytes that comply with an implementation-specific well-formedness condition. (We need such a condition for GHASH.) We let m range over well-formed messages.

Key-generation functions are marked as stateful (ST) to reflect their use of random sampling. Static keys (of type key_0) may be used to generate multiple one-time keys (of type key), defining keygen as e.g. $k \xleftarrow{\$} k_0 \| \mathsf{byte}^{\ell_k - \ell_{k_0}}$. To begin with, we assume $\ell_{k_0} = 0$ so that k_0 is the empty string ε.

The two main functions produce and verify MACs. Their correctness is captured in the verify postcondition of mac: verification succeeds at least on the tags correctly produced using mac with matching key and message.

One-Time Security MAC security is usually defined using computational unforgeability, as in the following game:

Game UF-1CMA$(\mathcal{A}, \mathsf{MAC})$	**Oracle** Mac(m)
$k \xleftarrow{\$} \mathsf{MAC.keygen}(\varepsilon);\ log \leftarrow \bot$	**if** $log \neq \bot$ **return** $\bot$
$(m^\star, t^\star) \leftarrow \mathcal{A}^{\mathsf{Mac}}$	$t \leftarrow \mathsf{MAC.mac}(k, m)$
return $\mathsf{MAC.verify}(k, m^\star, t^\star)$	$log \leftarrow (m, t)$
$\qquad\qquad \wedge\ log \neq (m^\star, t^\star)$	**return** t

The Mac oracle permits the adversary a single chosen-message query (recorded in *log*) before he has to produce a forgery. The advantage of $\mathcal{A}$ playing the UF-1CMA game is defined as $\epsilon_{\mathsf{UF\text{-}1CMA}}(\mathcal{A}[\ell_m]) \triangleq \Pr[\mathsf{UF\text{-}1CMA}(\mathcal{A}, \mathsf{MAC}) = 1]$.

We seek a stronger property for AEAD—the whole ciphertext must be indistinguishable from random bytes—and we need a decisional game for type-based composition, so we introduce a variant of unforgeability that captures indistinguishability from a random tag (when r is set).

Definition 2 (IND-UF-1CMA). *Let $\epsilon_{\mathsf{Mac1}}(\mathcal{A}[\ell_m, q_v])$ be the advantage of an adversary $\mathcal{A}$ that makes q_v Verify queries on messages of length at most ℓ_m in the following game:*

Game $\mathsf{Mac1}^b(\mathsf{MAC})$	**Oracle** $\mathsf{Mac}(m)$
$k \xleftarrow{\$} \mathsf{MAC.keygen}(\varepsilon);\ log \leftarrow \bot$	**if** $log \neq \bot$ **return** $\bot$
return $\{\mathsf{Mac}, \mathsf{Verify}\}$	$t \leftarrow \mathsf{MAC.mac}(k, m)$
	if $b \wedge r$
Oracle $\mathsf{Verify}(m^\star, t^\star)$	$\quad t \xleftarrow{\$} \mathsf{byte}^{\mathsf{MAC}.\ell_t}$
if b **return** $log = (m^\star, t^\star)$	$log \leftarrow (m, t)$
return $\mathsf{MAC.verify}(k, m^\star, t^\star)$	**return** t

In this game, the MAC oracle is called at most once, on some chosen message m; it returns a tag t and logs (m, t). Conversely, Verify is called q_v times before and after calling MAC. When b is set, the game idealizes MAC in two ways: verification is replaced by a comparison with the log; and (when r is set) the tag is replaced with random bytes.

We show (in the appendix) that our definition implies UF-1CMA when $q_v \geq 1$ and that random tags are neither necessary nor sufficient for unforgeability. We are not aware of much prior work on Mac1 with r set. A pairwise independent hash function would satisfy our IND-UF-1CMA definition but may require longer keys [45].

Multi-Instance Security with a Shared Key In the AEAD construction, we instantiate a one-time MAC for every encryption and decryption. AES-GCM uses a static MAC key computed from the AEAD key, that is shared by all MAC instances. This state sharing is not captured by the games above. To this end, we extend the $\mathsf{Mac1}^b$ game into a multi-instance version $\mathsf{MMac1}^b$ with a setup that invokes the keygen_0 function to generate any key materials reused across instances.

In the multi-instance case it is convenient to support two kinds of instances: honest instances are created with Keygen and idealized as in $\mathsf{Mac1}^b$; dishonest instances are created with Coerce and use the real implementation regardless of b. (Formally Coerce does not make the model stronger, as an adversary can run all algorithms himself. The finer model is however useful in hybrid arguments, and for composition with a PRF in §4.)

Definition 3 (m-IND-UF-1CMA). *Let $\epsilon_{\mathsf{MMac1}}(\mathcal{A}[\ell_m, q_v, q_i])$ be the advantage of an adversary $\mathcal{A}$ that creates q_i instances and makes at most q_v Verify queries overall on messages of length at most ℓ_m in the game:*

Game $\mathsf{MMac1}^b(\mathsf{MAC})$	**Oracle** $\mathsf{Keygen}(n)$
$log \leftarrow \varnothing;\ k \leftarrow \varnothing;\ H \leftarrow \varnothing$	**if** $k[n] \neq \perp$ **return** $\perp$
$k_0 \stackrel{\$}{\leftarrow} \mathsf{MAC.keygen}_0()$	$k[n] \leftarrow \mathsf{MAC.keygen}(k_0)$
return $\{\mathsf{MAC}, \mathsf{Verify},$	$H \leftarrow H \cup n$
$\quad \mathsf{Coerce}, \mathsf{Keygen}\}$	**Oracle** $\mathsf{Coerce}(n, k)$
	if $k[n] \neq \perp$ **return** $\perp$
Oracle $\mathsf{Mac}(n, m)$	$k[n] \leftarrow k$
if $k[n] = \perp$ **return** $\perp$	
if $log[n] \neq \perp$ **return** $\perp$	**Oracle** $\mathsf{Verify}(n, m, t)$
$t \leftarrow \mathsf{MAC.mac}(k[n], m)$	**if** $k[n] = \perp$ **return** $\perp$
if $b \wedge n \in H$	$v \leftarrow \mathsf{MAC.verify}(k[n], m, t)$
$\quad t \stackrel{\$}{\leftarrow} \mathsf{byte}^{\mathsf{MAC}.\ell_t}$	**if** $b \wedge n \in H$
$log[n] \leftarrow (m, t)$	$\quad v \leftarrow log[n] = (m, t)$
return t	**return** v

We confirm that $\mathsf{Mac1}$ is a special case of $\mathsf{MMac1}$ security and that, even with a static key, it suffices to consider a single verification query. (The proofs are in the appendix.)

Lemma 1 (MMac1 reduces to Mac1). *Given $\mathcal{A}$ against $\mathsf{MMac1}^b$, when $\ell_{k_0} = 0$, we construct $\mathcal{B}$ (linearly in q_i) against $\mathsf{Mac1}^b$ such that:*

$$\epsilon_{\mathsf{MMac1}}(\mathcal{A}[\ell_m, q_v, q_i]) \leq q_i\, \epsilon_{\mathsf{Mac1}}(\mathcal{B}[\ell_m, q_v]).$$

Lemma 2. *Given $\mathcal{A}$ against $\mathsf{MMac1}$ we construct $\mathcal{B}$ such that:*

$$\epsilon_{\mathsf{MMac1}}(\mathcal{A}[\ell_m, q_v, q_i]) \leq q_v\, \epsilon_{\mathsf{MMac1}}(\mathcal{B}[\ell_m, 1, q_i]).$$

Verified Implementation m-IND-UF-1CMA security reflects the type-based security specification of our idealized module $\mathsf{MMac1}^b$, which has an interface of the form

```
val log: memory → key → Spec (option (message × tag))
val mac: k:key → m:message → ST (t:tag)
  (requires log k = None)
  (ensures log k′ = Some(m,t))
val verify: k:key → m:message → t:tag → ST (v:bool)
  (ensures b ⇒ v = (log k′ = Some(m,t)))
```

The types of mac and verify express the gist of our security property: the specification function log gives access to the current content of the log associated with a one-time key; mac requires that the log be empty (None in $\mathrm{F}^\star$) thereby enforcing our one-time MAC discipline; verify ensures that, when b is set, verification succeeds if and only if mac logged exactly the same message and tag. Their implementation is automatically verified by typing $\mathsf{MMac1}^b$. However, recall that typing says nothing about the security loss incurred by switching b—this is the subject of the next subsection.

Our implementation of $\mathsf{MMac1}^b$ supports the two constructions described next, including code and functional correctness proofs for their algorithms. It also provides a more efficient interface for computing MACs incrementally. Instead of actually concatenating all authenticated materials in a message, the user creates a stateful hash, then repeatedly appends 16-byte words to the hash, and finally calls mac or verify on this hash, with a type that binds the message to the final hash contents in their security specifications. Our code further relies on indexed abstract types to separate keys and hashes for different instances of the functionality, and to support static key compromise.

3.2 Wegman-Carter-Shoup (WCS) Constructions

Next, we set up notations so that our presentation applies to multiple constructions, including GHASH and Poly1305; we factor out the encodings to have a core security assumption on sequences of field elements; we verify their injectivity; we finally prove a concrete bound in general, and in particular for GHASH and Poly1305.

From bytes to polynomials and back   In addition to fixed lengths for keys and tags, the construction is parameterized by

- a field $\mathbb{F}$;

- an encoding function $\overline{\cdot}$ from messages to polynomials in $\mathbb{F}$, represented as sequences of coefficients $\overline{m} \in \mathbb{F}^*$.

- a truncation function from $e \in \mathbb{F}$ to $\mathsf{tag}(e) \in \mathsf{byte}^{\ell_t}$;

The key consists of two parts: an element $r \in \mathbb{F}$ and a one-time pad $s \in \mathsf{byte}^{\ell_t}$. We assume that r and s are sampled uniformly at random, from some $R \subseteq \mathbb{F}$ and from byte^{ℓ_t}, respectively. We write $r\|s \leftarrow k$ for the parsing of key materials into r and s, including the encoding of r into R.

Generic Construction   Given a message m encoded into the sequence of d coefficients $\overline{m}_0, \dots, \overline{m}_{d-1}$ of a polynomial $\overline{m}(x) = \sum_{i=1..d} \overline{m}_{d-i} x^i$ in $\mathbb{F}$, the tag is computed as:

$$\mathsf{hash}_r(m) \leftarrow \mathsf{tag}(\overline{m}(r)) \qquad \text{in } \mathbb{F} \text{ (before truncation)}$$
$$\mathsf{mac}(r\|s, m) \leftarrow \mathsf{hash}_r(m) \boxplus s \quad \text{in } \mathsf{byte}^{\ell_t}$$

where the blinding operation $\boxplus$ is related to addition in $\mathbb{F}$ (see specific details below). We refer to $\mathsf{hash}_r(m)$, the part of the construction without blinding, as the hash.

We describe the two WCS instantiations employed in TLS.

GHASH [23] uses the Galois field $GF(2^{128})$, defined as the extension $GF(2)[x]/x^{128} + x^7 + x^2 + x + 1$, that is, the field of polynomials with Boolean coefficients modulo the irreducible polynomial $x^{128} + x^7 + x^2 + x + 1$. Such polynomials are represented as 128-bit vectors. Conveniently, polynomial addition, the blinding operation $\boxplus$, and its inverse $\boxminus$ simply correspond to 128-bit XOR. Polynomial multiplication is also efficiently supported on modern processors. The message encoding $\overline{\cdot}$ simply splits the input message into 16-byte words, seen as integers in $0..2^{128} - 1$; and the tag truncation is the identity. For AES-GCM, GHASH has a $\mathsf{keygen0}$ function that samples a single random $r \in GF(2^{128})$ shared across MAC instances.

Poly1305 [9] uses the prime field $GF(p)$ for $p = 2^{130} - 5$, that is, the field of integer addition and multiplication modulo p, whose elements can all be represented as 130-bits integers. Its message encoding $\overline{\cdot}$ similarly splits the input message into 16-byte words, seen as integers in $0..2^{128} - 1$, then adds 2^{ℓ} to each of these integers, where ℓ is the word length in bits. (Hence, the encoding precisely keeps track of the length of the last word; this feature is unused for AEAD, which applies its own padding to ensure $\ell = 128$.) The truncation function is $\mathsf{tag}(e) = e \mod 2^{128}$. The blinding operation $\boxplus$ and its inverse $\boxminus$ are addition and subtraction modulo 2^{128}. For ChaCha20-Poly1305, both r and s are single-use ($\ell_{k_0} = 0$) but our proof also applies to the original Poly1305-AES construction [9] where r is shared.

Injectivity Properties   We intend to authenticate messages, not just polynomial coefficients. To this end, we instantiate our $\mathsf{wellformed}$ predicate on messages and we show (in $\mathrm{F}^\star$) that

$\forall\ (m0:\ \mathsf{bytes})\ (m1:\ \mathsf{bytes}).$
$\mathsf{wellformed}\ m0\ \wedge\ \mathsf{wellformed}\ m1\ \wedge\ \mathsf{Poly.equals}\ \overline{m0}\ \overline{m1} \Rightarrow\ m0 = m1$

where Poly.equal specifies polynomial equality by comparing two sequences of coefficients, extending the shorter sequence with zero coefficients if necessary. This enables the (conditional) composition of MACs with suitable well-formedness predicates for AEAD in TLS. This is required for GHASH as it is otherwise subject to 0-message truncations.

We verify that the property above suffices to prove that both encodings are secure, and also that it holds in particular once we define wellformed as the range of formatted messages for AEAD (which are 16-byte aligned and embed their own lengths; see §5). We also confirm by typing that, with Poly1305, there is no need to restrict messages: its encoding is injective for all bytestrings [9, Theorem 3.2].

Security We give a theorem similar to those in prior work [9, 32, 46] but parameterized by the underlying field $\mathbb{F}$, encoding $\overline{}$, truncation tag, and blinding operation $\boxplus$. It covers all uses of AES-GCM and ChaCha20-Poly1305 in TLS.

Consider the MMac1 definition covering both shared and fresh values for r. Let q_v be the number of oracle calls to Verify (for which $log[n] \neq (m^\star, t^\star)$) and d a bound on the size (expressed in number of field elements) of the messages in calls to Mac and Verify.

Theorem 1. *The Wegman-Carter-Shoup construction for messages in $\mathbb{F}^{d-1}$ is m-IND-UF-1CMA secure with concrete bound $\epsilon_{\mathsf{MMac1}}(\mathcal{A}[\ell_m, q_v, q_i]) = \frac{d.\tau.q_v}{|R|}$ with $d = \ell_m/16$, and $\tau = 1$ for GHASH and $\tau = 8$ for Poly1305.*

The proof (in the appendix) uses Lemma 2 then establishes a bound $\frac{d.\tau}{|R|}$ for an adversary that makes a single Verify query. This bound follows from an $\frac{d.\tau}{|R|}$-almost-$\boxminus$-universal property, which has been separately proved for GHASH [38] and Poly1305 [9]; the appendix also includes its proof for all instantiations of hash_r for TLS.

Concrete bounds for GHASH The range size for r is 2^{128} and there is no tag truncation, hence by Lemma 2 we get a straight $\epsilon = \frac{d.q_v}{2^{128}}$, so for TLS the main risk is a failure of our PRF assumption on AES. We come back to this in §7.

Concrete bound for Poly1305 The effective range R of r is reduced, first by uniformly sampling in $0..2^{128} - 1$, then by clamping 22 bits, to uniformly sampling one element out of $|\ R\ |= 2^{106}$ potential values. We lose another 3 bit of security from the truncation of $\mathbb{F}$ to byte^{ℓ_t} and by applying Lemma 2 we arrive at $\epsilon = \frac{d.q_v}{2^{103}}$.

4 Pseudo-Random Functions for AEAD

We now consider the use of symmetric ciphers in counter mode, both for keying one-time MACs and for generating one-time pads for encryption. We model ciphers as PRFs. For TLS, we will use AES or ChaCha20, and discuss PRF/PRP issues in §7. A pseudo-random function family PRF implements the following interface:

```
type key
val keygen: unit → ST key
val ℓ_d : nat (* fixed domain length *)
```

```
val ℓ_b : nat (* fixed block length *)
type domain = lbytes ℓ_d
type block = lbytes ℓ_b
val eval: key → domain → Tot block (* functional specification *)
```

This interface specifies an abstract type for keys and a key-generation algorithm. (Type abstraction ensures that these keys are used only for PRF computations.) It also specifies concrete, fixed-length bytestrings for the domain and range of the PRF, and a function to compute the PRF. We refer to the PRF outputs as blocks. As usual, we define security as indistinguishability from a random function with lazy sampling.

Definition 4 (PRF security). *Let $\epsilon_{\mathsf{Prf}}(\mathcal{A}[q_b])$ be the advantage of an adversary $\mathcal{A}$ that makes q_b Eval queries in the game:*

Game $\mathsf{Prf}^b(\mathsf{PRF})$	**Oracle** $\mathsf{Eval}(m)$
$T \leftarrow \varnothing$	**if** $T[m] = \bot$
$k \xleftarrow{\$} \mathsf{PRF.keygen}()$	$\quad$ **if** b **then** $T[m] \xleftarrow{\$} \mathsf{byte}^{\ell_b}$
return $\{\mathsf{Eval}\}$	$\quad$ **else** $T[m] \leftarrow \mathsf{PRF.eval}(k, m)$
	return $T[m]$

The AEAD constructions we consider use PRFs both to generate keys for the one-time MAC used to authenticate the ciphertext and to generate a one-time pad for encryption and decryption. Accordingly, we partition the domain and use a specialized security definition, with a separate eval function for each usage of the PRF. (This will enable us to give more precise types to each of these functions.)

We assume the PRF domain consists of concatenations of a fixed-sized counter j and a nonce n, written $j\|n$. This notation hides minor differences between AEAD algorithm specifications, e.g. AES-GCM uses $n\|j$ instead $j\|n$. Our implementation handles these details, and verifies that $j\|n$ is injective for all admissible values of j and n.

For key generation, AES-GCM uses the PRF to derive both a static MAC key k_0 generated from the PRF (with nonce and counter 0) and a 1-time MAC key for each nonce (with counter 1), whereas Poly1305 uses a pure 1-time MAC key for each nonce (with counter 0). To handle both cases uniformly, we introduce a parameter $j_0 \in \{0, 1\}$ to shift the counter before concatenation with the nonce. In the following, we assume a *compatible* MAC, meaning that either $j_0 = 0 \wedge \ell_{k_0} = 0 \wedge \ell_k \leq \ell_b$ or $j_0 = 1 \wedge \ell_{k_0} \leq \ell_b \wedge \ell_k - \ell_{k_0} \leq \ell_b$.

For pad generation, counter mode encrypts plaintext blocks as $p \oplus \mathsf{eval}(j\|n)$ and decrypts by applying the same pad to the ciphertext. In the PrfCtr game below, we separate encryption and decryption, and we fuse the block generation and the XOR, so that we can give separate types to plaintexts and ciphertexts. (We truncate the block in case it is smaller than the input, as required for the last block in counter-mode.)

Definition 5 (PrfCtr security). *Given PRF and MAC, let $\epsilon_{\mathsf{PrfCtr}}(\mathcal{A}[q_b, q_g])$ be the advantage of an adversary $\mathcal{A}$ that makes q_b queries to either EvalEnx or EvalDex and q_g queries to EvalKey in the following game:*

Game $\mathsf{PrfCtr}^b(\mathsf{PRF}, \mathsf{MAC})$

$T \leftarrow \varnothing;\ R \leftarrow \varnothing$
$k \xleftarrow{\$} \mathsf{PRF.keygen}()$
$k_0 \xleftarrow{\$} \mathsf{MAC.keygen0}()$
if $j_0 \wedge \neg b$
 $o \leftarrow \mathsf{PRF.eval}(k, 0^{\ell_b})$
 $k_0 \leftarrow \mathsf{truncate}(o, \mathsf{MAC}.\ell_{k_0})$
return $\{\mathsf{EvalKey}, \mathsf{EvalEnx},$
     $\mathsf{EvalDex}\}$

Oracle $\mathsf{EvalKey}(j\|n)$

if $j \neq j_0$ **return** $\perp$
if $T[j\|n] = \perp$
 if b
  $k_m \xleftarrow{\$} \mathsf{MAC.keygen}(k_0)$
 else
  $o \leftarrow \mathsf{PRF.eval}(k, j\|n)$
  $k_m \leftarrow \mathsf{truncate}(k_0\|o, \ell_k)$
 $T[j\|n] \leftarrow k_m$
return $T[j\|n]$

Oracle $\mathsf{EvalEnx}(j\|n, p)$

if $j \leq j_0$ **return** $\perp$
$o \xleftarrow{\$} \mathsf{Eval}(j\|n)$
$c \leftarrow p \oplus \mathsf{truncate}(o, |p|)$
return c

Oracle $\mathsf{EvalDex}(j\|n, c)$

if $j \leq j_0$ **return** $\perp$
$o \xleftarrow{\$} \mathsf{Eval}(j\|n)$
$p \leftarrow c \oplus \mathsf{truncate}(o, |c|)$
return p

Lemma 3 (PrfCtr^b reduces to Prf^b). *Given* PRF, MAC, *and* $\mathcal{A}$ *against* $\mathsf{PrfCtr}^b(\mathsf{PRF}, \mathsf{MAC})$, *we construct* $\mathcal{B}$ *against* $\mathsf{Prf}^b(\mathsf{PRF})$ *such that:*

$$\epsilon_{\mathsf{PrfCtr}}(\mathcal{A}[q_b, q_g]) = \epsilon_{\mathsf{Prf}}(\mathcal{B}[q_b + q_g + j_0]).$$

The proof is in the appendix. Intuitively, we have a perfect reduction because, in all cases, the specialized game still samples a single fresh block for each $j\|n$ for a single purpose, and returns a value computed from that block.

In the next section, once b holds and the MAC has been idealized, we will use two oracles that further idealize encryption and decryption:

Oracle $\mathsf{EvalEnx}'(j\|n, p)$

if $j \leq j_0$ **return** $\perp$
if $T[j\|n] \neq \perp$ **return** $\perp$
if b' $c \xleftarrow{\$} \mathsf{byte}^{|p|}$
else $c \xleftarrow{\$} \mathsf{EvalEnx}(j\|n, p)$
$T[j\|n] \leftarrow (p, c)$
return c

Oracle $\mathsf{EvalDex}'(j\|n, c)$

if $j \leq j_0$ **return** $\perp$
if $T[j\|n] = (p, c)$ for some p
 return p
else return $\perp$

When b' holds, encryption samples c instead of $o = p \oplus c$, and records the pair (p, c) instead of just $p \oplus c$; and decryption simply performs a table lookup. This step is valid provided the block at $j\|n$ is used for encrypting a single p and decrypting the resulting c. The oracles enforce this restriction dynamically (on their second lines) whereas our code enforces it statically, using type-based preconditions on $\mathsf{EvalEnx}$ or $\mathsf{EvalDex}$ implied by the AEAD invariant of §5.

Verified Implementation Lemma 3 and the subsequent step are not currently verified by typing. (Still, note that the sampling of c instead of o is justified by F*'s probabilistic semantic and could be verified using the relational typing rule for sample in RF* [4])

We use an idealized PRF module with two idealization flags (for b and for b') that directly corresponds to the specialized game $\mathsf{PrfCtr}^{b,b'}$, parameterized by a Cipher module that implements real AES128, AES256, and Chacha20 (depending on an algorithmic parameter alg) and by a MAC module. The separation of the PRF domain is enforced by typing: depending on alg, j_0, j, b, and b', its range includes keys, blocks, and pairs (p, c).

5 From MAC and PRF to AEAD

We implement the two main AEAD constructions used by TLS 1.3 and modern ciphersuites of TLS 1.2. We show that their composition of a PRF and a one-time MAC yields a standard notion of AEAD security. Our proof is generic and carefully designed to be modular and TLS-agnostic: we share our AEAD code between TLS 1.2 and 1.3, and plan to generalize it for other protocols such as QUIC.

AEAD functionality Our authenticated encryption with associated data (AEAD) has a real interface of the form

```
val ℓ_n: nat              (* fixed nonce length *)
val ℓ_a: n:nat{n < 2^32} (* maximal AD length *)
val ℓ_p: n:nat{n < 2^32} (* maximal plaintext length *)
val cipherlen: n:nat{n ≤ ℓ_p} → Tot nat
type nonce = lbytes ℓ_n
type ad = bbytes ℓ_a
type plain = bbytes ℓ_p
type cipher = bytes


val decrypt: key → nonce → ad → c:cipher →
  ST (option (p:plain{length c = cipherlen (length p)}))
val encrypt: k:key → n:nonce → a:ad → p:plain →
  ST (c:cipher{length c = cipherlen (length p))
```

with two main functions to encrypt and decrypt messages with associated data of variable lengths, and types that specify the cipher length as a function of the plain length. (We omit declarations for keys, similar to those for PRFs §4.)

Definition 6 (Aead security). *Let $\epsilon_{\mathsf{Aead}}(\mathcal{A}[q_e, q_d, \ell_p, \ell_a])$ be the advantage of an adversary that makes at most q_e Encrypt and q_d Decrypt queries on messages and associated data of lengths at most ℓ_p and ℓ_a in the game:*

Game $\mathsf{Aead}^b(\mathsf{AEAD})$	**Oracle** $\mathsf{Decrypt}(n, a, c)$
$C \leftarrow \varnothing$	**if** b
$k \xleftarrow{\$} \mathsf{AEAD.KeyGen}()$	**if** $C[n] = (a, p, c)$ *for some* p
return $\{\mathsf{Encrypt}, \mathsf{Decrypt}\}$	**return** p
	return $\perp$
Oracle $\mathsf{Encrypt}(n, a, p)$	**else**
if $C[n] \neq \perp$ **return** $\perp$	$p \leftarrow \mathsf{AEAD.decrypt}(k, n, a, c)$
if b $\;c \xleftarrow{\$} \mathsf{byte}^{\mathsf{cipherlen}(|p|)}$	**return** p
else $c \leftarrow \mathsf{AEAD.encrypt}(k, n, a, p)$	
$C[n] \leftarrow (a, p, c)$	
return c	

Our definition generalizes AE in §2; it has a richer domain with plaintext and associated data of variable lengths; a function cipherlen from plaintext lengths to ciphertext lengths; and nonces n. It similarly maintains a log of encryptions, indexed by nonces. Crucially, Encrypt uses the log to ensure that each nonce is used at most once for encryption.

Generic AEAD Construction Given a PRF and a compatible MAC, AEAD splits plaintexts into blocks which are then blinded by pseudo-random one-time pads generated by calling PRF on increasing counter values, as shown in §4. (Blocks for MAC keys and the last mask may require truncation.)

17

To authenticate the ciphertext and associated data, the construction formats them into a single 16-byte-aligned buffer (ready to be hashed as polynomial coefficients as described in §3) using an encoding function declared as val encode: bbytes ℓ_p × bbytes ℓ_a → Tot bbytes $(\ell_p + \ell_a + 46)$ and implemented (in pseudo-code) as

Function encode(c, a)	**Function** pad$_{16}(b)$
return pad$_{16}(a) \,|\, pad_{16}(c)$	$r, b_1, \ldots, b_r \leftarrow$ split$_{16}(b)$
$|\, $length$_8(a) \,|\, $length$_8(c)$	**return** $b \,|\, $zeros$(16 - |b_r|)$

where the auxiliary function split$_\ell(b)$ splits the bytestring b into a sequence of r non-empty bytestrings, all of size ℓ, except for the last one which may be shorter. (that is, if $r, b_1, \ldots b_r \leftarrow$ split$_b(\ell)$, then $b = b_1 \,\|\, \cdots \,\|\, b_r.$); where zeros$(\ell)$ is the bytestring of ℓ zero bytes; and where length$_8(n)$ is the 8-byte representation of the length of n. Thus, our encoding adds minimal zero-padding to a and c, so that they are both 16-bytes aligned, and appends a final 16-byte encoding of their lengths.

Recall that the domain of MAC messages is restricted by the wellformed predicate. We now define wellformed b = ∃ (c:cipher) (a:ad). b = encode c a and typecheck the property listed in §3 that ensures injectivity of the polynomial encoding.

The rest of the AEAD construction is defined below, using an operator $otp \overline{\oplus} p$ that abbreviates the expression truncate$(otp, |p|) \oplus p$, and a function untag$_{16}$ that separates the ciphertext from the tag.

The main result of this section is that it is Aead-secure when PRF is Prf-secure and MAC is MMac1-secure:

Function keygen()

$k \xleftarrow{\$} $ PRF.keygen()$; \ k_0 \leftarrow \varepsilon$
if j_0
    $o \leftarrow$ PRF.eval$(k, 0^{\ell_b})$
    $k_0 \leftarrow$ truncate$(o, MAC.\ell_{k_0})$
return $k_0 \| k$

Function encrypt(K, n, a, p)	**Function** decrypt(K, n, a, c)
$(k_0, k) \leftarrow$ split$_{\ell_{k_0}}(K); \ c \leftarrow \varepsilon$	$(k_0, k) \leftarrow$ split$_{\ell_{k_0}}(K); \ p \leftarrow \varepsilon$
$k_1 \leftarrow$ PRF.eval$(k, j_0 \| n)$	$k_1 \leftarrow$ PRF.eval$(k, j_0 \| n)$
$k_m \leftarrow$ truncate$(k_0 | k_1, MAC.\ell_k)$	$k_m \leftarrow$ truncate$(k_0 | k_1, MAC.\ell_k)$
$r, p_1, \ldots, p_r \leftarrow$ split$_{\ell_b}(p);$	$(c, t) \leftarrow$ untag$_{16}(c)$
for $j = 1..r$	$m \leftarrow$ encode(c, a)
$\quad otp \leftarrow$ PRF.eval$(k, j_0 + j \| n)$	**if** ¬MAC.verify(k_m, m, t)
$\quad c \leftarrow c \,\|\, (otp \overline{\oplus} p_j)$	$\quad$ **return** $\perp$
$t \leftarrow$ MAC.mac$(k_m, $encode$(c, a))$	$r, c_1, \ldots, c_r \leftarrow$ split$_{\ell_b}(c);$
return $c \| t$	**for** $j = 1..r$
	$\quad otp \leftarrow$ PRF.eval$(k, j_0 + j \| n)$
	$\quad p \leftarrow p \,\|\, (otp \overline{\oplus} c_j)$
	return p

Theorem 2 (AEAD construction). *Given* $\mathcal{A}$ *against* Aead*, we construct* $\mathcal{B}$ *against* Prf *and* $\mathcal{C}$ *against* MMac1*, with:*

$$\epsilon_{\mathsf{Aead(AEAD)}}(\mathcal{A}[q_e, q_d, \ell_p, \ell_a]) \leq \epsilon_{\mathsf{Prf(PRF)}}(\mathcal{B}[q_b])$$
$$+ \epsilon_{\mathsf{MMac1(MAC)}}(\mathcal{C}[\ell_p + \ell_a + 46, q_d, q_e + q_d])$$

where q_b *(the number of distinct queries to the PRF) satisfies:*

$$q_b \leq j_0 + q_e \left(1 + \left\lceil \frac{\ell_p}{\ell_b} \right\rceil \right) + q_d$$

18

Proof sketch. The proof relies on the $\mathsf{PrfCtr}^{b,b'}$ and $\mathsf{MMac1}^b$ idealizations. It involves a sequence of transformations from Aead^0 to Aead^1 that inline successively more idealizations. Therefore, we introduce a parametric game $\mathsf{AeadCtr}(X)$ for any game X that returns $\mathsf{EvalKey}$, $\mathsf{EvalEnx}$, $\mathsf{EvalDex}$, Mac, and Verify oracles:

Game $\mathsf{AeadCtr}(X)$
$\overline{(\mathsf{EvalKey}, \mathsf{EvalEnx}, \mathsf{EvalDex}, \mathsf{Mac}, \mathsf{Verify}) \leftarrow \mathsf{X}()}$
return $\{\mathsf{Encrypt}, \mathsf{Decrypt}\}$

Oracle $\mathsf{Encrypt}(n,a,p)$	**Oracle** $\mathsf{Decrypt}(n,a,c)$
if $C[n] \neq \bot$ **return** $\bot$	$c,t \leftarrow \mathsf{untag}_{16}(c)$
$\mathsf{EvalKey}(n); \ c \leftarrow \varepsilon$	$\mathsf{EvalKey}(n)$
$r, p_1, \ldots, p_r \leftarrow \mathsf{split}_{\ell_b}(p)$	**if** $\neg\mathsf{Verify}(n, \mathsf{encode}(c,a), t)$
for $j = 1..r$	$\quad$**return** $\bot$
$\quad c \leftarrow c \,\|\, \mathsf{EvalEnx}(j_0 + j\|n, p_j)$	$r, c_1, \ldots, c_r \leftarrow \mathsf{split}_{\ell_b}(c); \ p \leftarrow \varepsilon$
$c \leftarrow c \,\|\, \mathsf{Mac}(n, \mathsf{encode}(c,a))$	**for** $j = 1 \ldots r$
$C[n] \leftarrow (a, p, c)$	$\quad p \leftarrow p \,\|\, \mathsf{EvalDex}(j_0 + j\|n, c_j)$
return c	**return** p

When X is obtained from PrfCtr^0 and $\mathsf{MMac1}^0$ we have a game that is equivalent to Aead^0. We first switch to PrfCtr^1 to get random MAC keys and then idealize $\mathsf{MMac1}^1$. When X is obtained from PrfCtr^1 and $\mathsf{MMac1}^1$ ciphertexts are authenticated and we can switch to $\mathsf{PrfCtr}^{1,0}$ and then to $\mathsf{PrfCtr}^{1,1}$. At this stage the PRF table contains randomly sampled ciphertext blocks and decryption corresponds to table lookup in this table. This is ensured on the code by our AEAD invariant. The full proof is in the appendix. $\qquad\square$

Verified Implementation We outline below the idealized interface of our main AEAD^b module built on top of (the idealized interfaces of) $\mathsf{PrfCtr}^{b,b'}$ and $\mathsf{MMac1}^b$, both taken as cryptographic assumption, and documented by the games with the same name on paper. We focus on types for encryption and decryption:

```
abstract type key (* stateful key, now containing the log C *)
val log: memory → key → Spec (seq (nonce × ad × cipher × plain)
val keygen : unit → ST (k:key)
  (ensures b ⇒ log k = ∅)
val encrypt: k:key → n:nonce → a:ad → p:plain → ST (c:cipher)
  (requires b ⇒ lookup_nonce n (log k) = None)
  (ensures (b ⇒ log k' = log k ++ (n,a,c,p)))
val decrypt: k:key → n:nonce → a:ad → c:cipher → ST (o:option plain)
  (ensures b ⇒ o = lookup (n,a,c) (log k))
```

As in §2, we have a multi-instance idealization, with a log for each instance stored within an abstract, stateful key; and we provide a ghost function log to access its current contents in logical specifications. Hence, key generation allocates an empty log for the instance; encryption requires that the nonce be fresh and records its results; and decryption behaves exactly as a table lookup, returning a plaintext if, and only if, it was previously stored in the log by calling encryption with the same nonce and additional data.

This step of the construction is entirely verifiable by typing. To this end, we supplement its implementation with a precise invariant that relates the AEAD log to the underlying PRF table and MAC logs. For each entry in the log, we specify the corresponding entries in the PRF table (one for the one-time MAC key, and one for each block required for encryption) and, for each one-time MAC key entry, the contents of the MAC log (an encoded message and the tag at the end of the ciphertext in the AEAD log entry). By typing the AEAD code that implements the construction, we verify that the invariant is preserved as it completes its series of calls to the PRF and MAC idealized interfaces. Hence, although our code for decryption does

not actually decrypt by a log lookup, we prove that (when b holds) its results always matches the result of a lookup on the current log. As usual, by setting all idealization flags to false, the verified code yields our concrete TLS implementation.

Security bounds Theorem 2 can be specialized to provide precise security bounds for the various AEAD ciphersuites:

Construction	$\epsilon_{\mathsf{Aead}}(\mathcal{A}[q_e, q_d, \ell_p, \ell_a]) \leq$
AES-GCM	$\epsilon_{\mathsf{Prf}}\left(\mathcal{B}\left[q_e\left(1 + \frac{\ell_p}{16}\right) + q_d + 1\right]\right)$ $+ \frac{q_d}{2^{128}} \cdot \left(\frac{\ell_p + \ell_a + 46}{16}\right)$
ChaCha20-Poly1305	$\epsilon_{\mathsf{Prf}}\left(\mathcal{B}\left[q_e\left(1 + \frac{\ell_p}{64}\right) + q_d\right]\right)$ $+ \frac{q_d}{2^{103}} \cdot \left(\frac{\ell_p + \ell_a + 46}{16}\right)$

In this paper, we only consider MAC schemes that provide information theoretic security although we do not rely on this fact in the proof of AEAD.

6 From AEAD to Stream Encryption (StAE)

TLS requires stream encryption: message fragments must be received and processed in the order they were sent, thereby defeating attempts to delete or re-order network traffic. To this end, encryption and decryption use a local sequence number to generate distinct, ordered nonces for AEAD.

In practice, it is difficult to prevent multiple honest servers from decrypting and processing the same 0-RTT encrypted stream. Since decryption is now stateful, we must generalize our model to support multiple parallel decryptors for each encryptor. In our security definitions, we thus add a `genD` oracle to generate new decryptors (with local sequence numbers set to zero) from a given encryptor.

Otherwise, the stateful construction is quite simple: TLS 1.3 combines the sequence number with a static, random 'initialization vector' (IV) in the key materials to generate pairwise-distinct nonces for encrypting fragments using AEAD. In contrast, TLS 1.2 nonces concatenate the static IV with a per-fragment explicit IV that is sent alongside the ciphertext on the network (except for ciphersuites based on ChaCha20-Poly1305 which follow the TLS 1.3 nonce format). Some TLS 1.2 implementations incorrectly use uniformly random explicit IVs [16]. This is much inferior to using the sequence number because of the high collision risk on 64 bits. Therefore, in our implementation, we use the following nonce construction:

$$n = \begin{cases} bigEndian_8(seqn) \parallel iv_4 & \text{for AES-GCM in TLS 1.2} \\ bigEndian_{12}(seqn) \oplus iv_{12} & \text{otherwise} \end{cases}$$

where the indices indicate lengths in bytes. We discuss at the end of the section the impact of random explicit IVs in TLS 1.2 on the concrete security bound. The use of longer static IVs in TLS 1.3 is a practical improvement, as (informally) it acts as an auxiliary secret input to the PRF and may improve multi-user security [7]. This is particularly clear for ChaCha20, where the key, nonce, and counter are just copied side by side to the initial cipher state.

We easily verify (by typing) that both constructions are injective for $0 \leq seqn < 2^{64}$, which is required (also by typing) to meet the 'fresh nonce' pre-condition for calling AEAD encryption. Formally, the state

invariant for StAE encryption is that $0 \leq seqn < 2^{64}$ and the underlying AEAD log has an entry for every nonce n computed from a sequence number smaller than $seqn$.

StAE functionality A stream authenticated encryption functionality StAE implements the following interface:

```
type seqn_t = UInt64.t
val q_e: seqn_t (* maximal number of encryptions *)
val cipherlen: n:nat{ n ≤ ℓ_p } → Tot nat (* e.g. ℓ_p + MAC.ℓ_t *)

type role = E | D
abstract type state (r:role)
val seqn: mem → state r → Spec seqn_t
val gen: unit → ST (s:state E) (ensures seqn s' = 0)
val genD: state E → ST (s:state D) (ensures seqn s' = 0)
val encrypt: s:state E → ad → p:plain →
  ST (c:cipher{length c = cipherlen (length p))
    (requires seqn s < q_e) (ensures seqn s' = seqn s + 1)
val decrypt: s:state D → ad → c:cipher →
  ST (o:option (p:plain{length c = cipherlen (length p)}))
    (requires seqn s < q_e)
    (ensures seqn s' = if o = None then seqn s else seqn s + 1)
```

We omit type declarations for plain, cipher and ad as they are similar to AEAD. For TLS, the length of additional data ℓ_a can be 0 (TLS 1.3) or 13 (TLS 1.2) and the length of IVs ℓ_{iv} is 12. Compared to previous functionalities, the main change is that keys are replaced by states that include a 64-bit sequence number. Accordingly, in this section we assume that at most 2^{64} fragments are encrypted. The stateful function gen initializes the encryptor state used by the encryption algorithm, while genD initializes a decryptor state used by the decryption algorithm. The stateful encrypt and decrypt functions require that the sequence number in the key state does not overflow (seqn $s < q_e$) and ensure that it is incremented (only on success in the case of decryption). In pseudo-code, authenticated stream encryption is constructed as follows:

Function gen()
$k \xleftarrow{\$} \text{AEAD.keygen}()$
$iv \xleftarrow{\$} \text{byte}^{\ell_{iv}}$
return $\{k \leftarrow k;$
$\quad iv \leftarrow iv; \ seqn \leftarrow 0\}$

Function genD(s)
return $\{k \leftarrow s.k;$
$\quad iv \leftarrow s.iv; \ seqn \leftarrow 0\}$

Function encrypt(s, a, p)
$n \leftarrow \text{nonce}(s.iv, s.seqn)$
$c \xleftarrow{\$} \text{AEAD.encrypt}(s.k, n, a, p)$
$s.seqn \leftarrow s.seqn + 1$
return c

Function decrypt(s, a, c)
$n \leftarrow \text{nonce}(s.iv, s.seqn)$
$p \leftarrow \text{AEAD.decrypt}(s.k, n, a, c)$
if $(p = \bot)$ **return** $\bot$
$s.seqn \leftarrow s.seqn + 1$
return p

Definition 7 (Stae). *Let $\epsilon_{\text{Stae}}(\mathcal{A}[q_e, q_d, \ell_p, \ell_a])$ be the advantage of an adversary $\mathcal{A}$ that makes q_e encryption*

queries and q_d decryption queries in the game below.

Game $\mathsf{Stae}^b(\mathsf{StAE})$
***
$s \xleftarrow{\$} \mathsf{StAE.gen}()$
$D \leftarrow \varnothing \quad E \leftarrow \varnothing$
return $\{\mathsf{GenD}, \mathsf{Encrypt},$
$\qquad\qquad \mathsf{Decrypt}\}$

Oracle $\mathsf{GenD}(d)$
***
if $(D[d] \neq \bot)$ **return** $\bot$
$D[d] \leftarrow \mathsf{StAE.genD}(s)$

Oracle $\mathsf{Encrypt}(a, p)$
***
if b
$\quad c \leftarrow \mathsf{byte}^{\mathsf{cipherlen}(|p|)}$
else
$\quad c \leftarrow \mathsf{StAE.encrypt}(s, a, p)$
$E[s.seqn - 1, a, c] \leftarrow p$
return c

Oracle $\mathsf{Decrypt}(d, a, c)$
***
if $(D[d] = \bot)$ **return** $\bot$
if b
$\quad p \leftarrow E[D[d].seqn, a, c]$
$\quad$ **if** $(p \neq \bot)$
$\qquad D[d].seqn \leftarrow D[d].seqn + 1$
else
$\quad p \leftarrow \mathsf{StAE.decrypt}(D[d], a, c)$
return p

The game involves a single encryptor, a table of decryptors D, and a log of encryptions E. For brevity, it relies on the stateful encryptor and decryptors specified above, e.g. encrypts increments $s.seqn$ and $\mathsf{Encrypt}$ records the encryption with sequence number $s.seqn - 1$. (Equivalently, it could keep its own shadow copies of the sequence numbers.) In contrast with AEAD, decryption only succeeds for the current sequence number of the decryptor.

Our definition corresponds most closely to level-4 (stateful) LHAE of [17]. In both definitions the requirement is that decrypt only successfully decrypted a prefix of what was sent. We discuss minor differences: in their game the adversary needs to distinguish m_0 from m_1, while in our game he has to distinguish ciphertexts from random. Also, instead of rejecting all decryptions when $b = 0$ and forbidding in-sync challenge ciphertexts when $b = 1$ we return in-sync challenge plaintexts in both cases. These changes are superficial; the rational behind them is to keep our game close to our idealizing modules in the implementation. Another difference is that we do not require real decryption to continue rejecting ciphertexts upon decryption failure. We also leave length-hiding and stream termination to §7. So our definition compares to their definitions with the maximal plaintext length set to $|p|$.

Theorem 3 (Stae perfectly reduces to Aead)**.** *Given $\mathcal{A}$ against Stae, we construct $\mathcal{B}$ against Aead with*

$$\epsilon_{\mathsf{Stae}}(\mathcal{A}[q_e, q_d, \ell_p, \ell_a]) = \epsilon_{\mathsf{Aead}}(\mathcal{B}[q_e, q_d, \ell_p, \ell_a]).$$

Proof. We build $\mathcal{B}$ by composing $\mathcal{A}$ with a variant of the construction given above that calls Aead oracles instead of AEAD functions. $\qquad\square$

In TLS 1.2 with AES-GCM, there is a probability $p(q_e) \approx \frac{q_e^2}{2^{65}}$ of collision when the IVs are randomly sampled. This probably already exceeds 2^{-32} after less than 2^{17} fragments. This is more than sufficient to mount practical attacks against HTTPS websites, as demonstrated in [16]. Therefore, random explicit IVs must not be used in TLS 1.2.

7 TLS Content Protection: Length-Hiding Stream Encryption

We are now ready to use stream encryption for protecting TLS 1.3 traffic, which consists of a sequence of protocol-message fragments, each tagged with their content type, while hiding their content, their type,
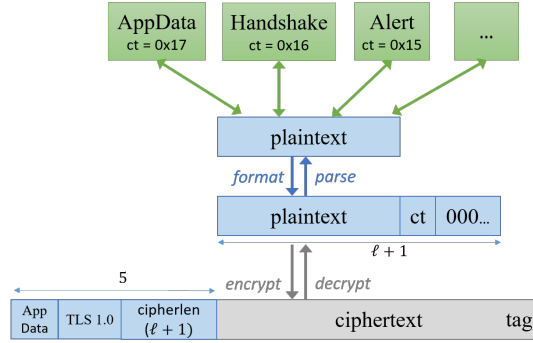
Figure 3: Constructing a TLS 1.3 record fragment

and their actual size before padding. The steps taken by the record layer to construct encrypted fragments are depicted in Figure 3, with apparent size ℓ after padding. The last line adds the (unprotected) record header; for backward compatibility, it pretends to be a TLS 1.0 AppData record irrespective of its actual encrypted content type. On the other hand, TLS does not attempt to hide the record boundaries (as e.g. SSH) so we do not expect indistinguishability from random for the resulting record.

Formatting: Content Type and Length Hiding Encryption and decryption rely on formatting and padding functions over a type fragment indexed by a length ℓ indicating the public *maximal length* of its content, which are specified as follows:

> type len = n:nat $\{n \leq 2^{14}\}$ *(∗ valid record length in TLS ∗)*
> type fragment (ℓ:len) = {ct:byte; data:bbytes ℓ}
> val parse: ℓ:len $\to$ lbytes (ℓ+1) $\to$ Tot (option (fragment ℓ))
> val format: ℓ:len $\to$ f:fragment ℓ $\to$ Tot (p:lbytes (ℓ+1))
>   (ensures parse ℓ p = Some f)

These functions must be carefully implemented to prevent any side channel. We also construct and parse records into headers and payloads using functions

> val parse_record: r:record $\to$ Tot (option (n:nat $\times$ c:lbytes n))
> val format_record: n:nat $\to$ c:lbytes n $\to$ Tot (r:record)
>   (ensures parse_record r = Some (n,c))

These function specifications suffice to establish our theorems below. We now give the concrete format function for TLS 1.3:

$$\textbf{Function } \mathsf{format}(\ell : \mathsf{len}, f : \mathsf{fragment}\ \ell)$$
$$\overline{f.data \,\|\, [f.ct] \,\|\, pad_0(\ell - |f.data|)}$$

where $\mathsf{pad}_0\ n$ is the string of n 0x00 bytes. We verify the post-condition of format by typing. We omit the corresponding parse function and the code for processing headers.

The implementation of parse and format, and the converse function for parsing a bytestring into a fragment value, require precautions to avoid leaking the actual contents length using side-channels. The code for processing headers does *not* depend on the fragment, only on its length after padding.

Stream Closure As explained in §6, stream integrity ensures that decrypted traffic is a prefix of encrypted traffic. Complementarily, the TLS record layer relies on well-defined *final fragments*, specified as a predicate val final: fragment ℓ $\to$ Tot bool, to ensure that no further encryptions are performed on a stream after sending such a fragment.

23

For LHAE, we extend the stateful key of StAE to record the termination status of the stream, which can be queried with the predicate val closed: mem → state r → Spec bool. Furthermore, we extend the post-condition of encryption to ensure that the state s′ after encrypting fragment f satisfies closed s′ = final f. Therefore, upon receiving a final fragment, the decryptor is guaranteed to have received the whole data stream. This control-flow mechanism only depends on stream integrity; it does not involve additional cryptography, but provides a useful additional guarantee to the application, and may help achieve forward secrecy by erasing the stream key immediately after closure.

LHSE Construction and Game The LHSE construction is:

Function encrypt(s, ℓ, f)	**Function** decrypt(s, r)
if closed(s) **return** $\perp$	**if** closed(s) **return** $\perp$
$p \leftarrow$ format(ℓ, f)	$\ell, c \leftarrow$ parse_record(v)
$c \leftarrow$ StAE.encrypt$(s, [], p)$	$p \leftarrow$ StAE.decrypt$(s, [], c)$
if (final f) $s \leftarrow closed$	$f \leftarrow$ parse(ℓ, p)
return format_record(ℓ, c)	**if** $(f \neq \perp \wedge$ final $f)$ $s \leftarrow closed$
	return f

with the same state as StAE—we omit the unmodified functions for generating encryptors and decryptors. When a final fragment is sent or received, we erase the StAE state.

The TLS 1.3 construction uses empty associated data, relying on implicit authentication of the underlying key and sequence number. (Our code also supports the older TLS 1.2 construction, which uses 13 bytes of associated data in total, obtained by appending the protocol version and the content type to the sequence number of stream encryption.)

Definition 8 (Lhse). *Given* LHSE*, let* $\epsilon_{\mathsf{Lhse}}(\mathcal{A}[q_e, q_d])$ *be the advantage of an adversary* $\mathcal{A}$ *that makes* q_e *encryption queries and* q_d *decryption queries in the game below.*

Game Lhse^b(LHSE)	**Oracle** GenD(d)
$s \xleftarrow{\$}$ Lhse.gen()	**if** $(D[d] = \perp)$
$D \leftarrow \varnothing;\ F \leftarrow \varnothing$	$\quad D[d] \leftarrow$ LHSE.genD(s)
return $\{$GenD, Encrypt, Decrypt$\}$	

Oracle Encrypt(ℓ, f)	**Oracle** Decrypt(d, v)
if b	**if** $(D[d] = \perp)$ **return** $\perp$
$\quad r \leftarrow$ LHSE.encrypt$(s, \ell, f_{\mathsf{final}(f)})$	$s_d \leftarrow D[d]$
else	**if** b
$\quad r \leftarrow$ LHSE.encrypt(s, ℓ, f)	$\quad$ **if** closed(s_d) **return** $\perp$
$F[s.seqn - 1, r] \leftarrow f$	$\quad f \leftarrow F[s_d.seqn, r]$
return v	$\quad$ **if** $(f \neq \perp)$ $s_d.seqn{+}{+}$
	$\quad$ **if** $(f \neq \perp \wedge$ final $f)$ $s_d \leftarrow closed$
	else
	$\quad f \leftarrow$ LHSE.decrypt(s_d, v)
	return f

where f_0 *(respectively,* f_1*) is a fragment (respectively, a final fragment), with fixed content type and data* 0^ℓ*.*

The game logs the encryption stream in F, indexed by fragment sequence numbers and ciphertexts. It has an oracle for creating decryptors; it stores their state in a table D, indexed by some abstract d chosen by the adversary. It does not model stream termination, enforced purely by typing the stream content.

Theorem 4 (Lhse *perfectly reduces to* Stae).
Given $\mathcal{A}$ *against* Stae*, we construct* $\mathcal{B}$ *against* Aead *with*

$$\epsilon_{\mathsf{Lhse}}(\mathcal{A}[q_e, q_d]) = \epsilon_{\mathsf{Stae}}(\mathcal{B}[q_e, q_d, 2^{14} + 1, \ell_a])$$

where ℓ_a *is* 0 *for TLS 1.3 and* 13 *for TLS 1.2.*

Proof. We build $\mathcal{B}$ by composing $\mathcal{A}$ with a variant of the construction given above that calls Stae oracles instead of its functions. The 1.3 record does not use authenticated data, $\ell_a = 0$, (while in TLS 1.2 it was $\ell_a = 13$). The maximum fragment size is $\ell_p = 2^{14} + 1$. We encrypt one additional byte for the content type. □

Multi-Stream LHSE In the next section (as in our interface above), we use a multi-instance Lhse game, defined below.

> **Game** Multi(Lhse^b)
> ―――――――――――――――――――――――――
> $E \leftarrow \varnothing$; **return** $\{\mathsf{Gen}, \mathsf{GenD}, \mathsf{Encrypt}, \mathsf{Decrypt}\}$
>
> **Oracle** Gen(i) | **Oracle** GenD(i, d)
> ――――――――――――――― | ―――――――――――――――
> **if** $(E[i] = \bot)\ E[i] \xleftarrow{\$} \mathsf{Lhse}^b()$ | **if** $(E[i] = \bot)\ E[i] \xleftarrow{\$} \mathsf{Lhse}^b()$
> | $E[i].\mathsf{GenD}(d)$
>
> **Oracle** Encrypt(i, ℓ, f) | **Oracle** Decrypt(i, d, v)
> ――――――――――――――― | ―――――――――――――――
> **if** $(E[i] = \bot)$ **return** $\bot$ | **if** $(E[i] = \bot)$ **return** $\bot$
> **return** $E[i].\mathsf{Encrypt}(\ell, f)$ | **return** $E[i].\mathsf{Decrypt}(d, v)$

For every fresh index i passed to Gen, we spawn an instance of Lhse and we record its state and oracle in table E. In all other cases, the oracles above now look up the shared instance at i and forward the call to the instance oracle.

Security bounds for TLS Table 1 gives the concrete bounds by ciphersuites, setting ℓ_p to $2^{14} + 1$ and ℓ_a to 0 (or 13 for TLS 1.2).

Chacha20 uses a Davies-Meyer construction and is considered a good PRF. For AES-GCM ciphersuites, blocks are relatively small (16 bytes) so we incur a loss of $\frac{q_b^2}{2^{129}}$ by the PRP/PRF switching lemma [6], intuitively accounting for the probability of observing collisions on ciphertext blocks and inferring the corresponding plaintext blocks are different. As observed e.g. by Luykx and Paterson [37], this factor limits the amount of data that can be sent securely using AES-GCM.

Based on their suggestion to send at most $2^{24.5}$ fragments with the same key (itself based on a proof by [10] for the UF-1CMA security of GHASH that avoids the full PRF-PRP switching loss), our implementation may automatically trigger TLS re-keying after sending $2^{24.5}$ fragments. This strategy results in the bound in the last row, which no longer depends quadratically on q_e and thus achieves a poor man's form of beyond birthday bound security.

The Luykx and Paterson bound shows that when $q_e \leq 2^{50}$ and $q_d \leq 2^{60}$ then the success probability of any attacker breaking the integrity of AES-GCM is bounded by 2^{-56}. Conditioned on all decryption queries having failed, they can then assume $q_d = 0$ when considering the bound for confidentiality.

To tolerate an attacker success probability of at most ϵ, they compute the maximum number of encrypted fragments q_e such that $q_e \leq \frac{\sqrt{2^{129}\epsilon}-1}{(1+(2^{14}+1)/16)}$. For instance as long as at most $2^{24.5}$ records are encrypted they bound the success probability by $\epsilon = 2^{-60}$.

Ciphersuite	$\epsilon_{\mathsf{Lhse}}(\mathcal{A}[q_e, q_d]) \leq$
General bound	$\epsilon_{\mathsf{Prf}}(\mathcal{B}[q_e(1 + \lceil (2^{14}+1)/\ell_b \rceil) + q_d + j_0])$ $+ \ \epsilon_{\mathsf{MMac1}}(\mathcal{C}[2^{14}+1+46, q_d, q_e + q_d])$
ChaCha20-Poly1305	$\epsilon_{\mathsf{Prf}}(\mathcal{B}\left[q_e\left(1 + \left\lceil \frac{(2^{14}+1)}{64} \right\rceil \right) + q_d\right]) + \frac{q_d}{2^{93}}$
AES128-GCM AES256-GCM	$\epsilon_{\mathsf{Prp}}(\mathcal{B}[q_b]) + \frac{q_b^2}{2^{129}} + \frac{q_d}{2^{118}}$ where $q_b = q_e(1+\lceil(2^{14}+1)/16\rceil)+q_d+1$
AES128-GCM AES128-GCM	$\frac{q_e}{2^{24.5}}\left(\epsilon_{\mathsf{Prp}}(\mathcal{B}[2^{34.5}]) + \frac{1}{2^{60}} + \frac{1}{2^{56}}\right)$ with re-keying every $2^{24.5}$ records (counting q_e for all streams, and $q_d \leq 2^{60}$ per stream)

Table 1: Summary of security bounds for the AEAD ciphersuites in TLS.

Let us now consider the bound for simultaneously achieving confidentiality and integrity. By solving for q_e we have: $q_e \leq \frac{\sqrt{2^{129}\epsilon - 2^{11}q_d - q_d - 1}}{1 + \lceil 2^{14}+1\rceil/16}$. For $q_d = 2^{34}$ we have $q_e \leq 2^{22.72}$, while for $q_d = 2^{35}$ we no longer get a meaningful security bound. For $q_d = 2^{32}$ one can send a more satisfying $q_e \leq 2^{24.21}$ records. If as in 1RTT only a single failed decryption query is accepted, then $q_d = 1$ and we get $q_e = 2^{24.49}$, essentially the same bound as Luykx and Paterson [37].

Exhaustive key search To better understand the terms ϵ_{Prp} and ϵ_{Prf} above, let us recall an exhaustive, ideal-cipher attack. Assume the adversary can search for q_K keys (each time encrypting a known plaintext block and comparing it with an oracle encryption of that block). The attack succeeds with probability $\frac{q_K}{2^{\ell_K}}$. With AES128 exhaustive key search is the best attack when this term dominates the quadratic term $(q_K \geq 2q_b^2)$, and the third term $(q_K \geq 2^{11}q_d)$.

8 The TLS 1.3 Record Protocol

Figure 4 presents the TLS 1.3 protocol from draft 18, focusing only on how it drives the record layer. In particular, this presentation ignores most of the details of the handshake.

The client sends the `ClientHello` message in cleartext, and may then immediately install a fresh 0-RTT key k_0^c in the record layer and use it to encrypt a stream of *early data*.

The server receives this message and, if it accepts 0-RTT, also installs the 0-RTT key k_0^c and decrypts the client's data. Otherwise, it discards this first stream. In parallel, the server sends a `ServerHello` message that allows both parties to derive encryption keys k_h^c and k_h^s for the handshake messages, and k_1^c and k_1^s for application data. The server installs k_h^s in the record and uses it to encrypt a first stream of handshake messages, ending with a finished message that triggers the installation of key k_1^s. If the server supports 0.5-RTT, it may immediately start using k_1^s for sending application data.

Once 0-RTT stream is complete (signaled by an end-of-early-data alert) and after processing the `ServerHello`, the client installs the handshake keys k_h^c and k_h^s for encryption and decryption. It completes the handshake by sending its own encrypted stream, ending with a finished message, and installs the of application traffic keys k_1^c and k_1^s.

Upon completing the handshake, the server also installs k_1^c for decryption. After this point, the con-

nection is fully established and both parties use the installed application traffic keys for all content types: AppData, Alert, and even Handshake messages (such as `KeyUpdate`).

Later, the client (or the server) can terminate their current output stream by sending either a `KeyUpdate` handshake message or a close-notify alert message. In the first case, it installs and starts using the next application traffic key k_2^c (then k_3^c, etc). The server (or the client) responds accordingly, with a `KeyUpdate` or a close-notify. In the first case, it installs and starts using the next traffic keys k_2^c and k_2^s. In the second case, the connection is closed.

In all cases, each party uses a single stream at a time in each direction, for sending and receiving all content types, and each stream ends with a distinguished message that clearly indicates its final fragment. 0-RTT data ends with an end-of-early-data alert; encrypted handshake messages in both directions end with the finished message; 0.5 and 1-RTT data streams end with a key update or a close-notify alert. This precaution ensures that any truncations at the end of a stream will cause a connection failure, rather than continuing with the next stream.

Performance/Security Trade-Offs. 0-RTT and 0.5-RTT significantly decrease communications latency, but they yield weaker application security. 0-RTT traffic has weaker forward secrecy and is subject to replays: if multiple servers may accept the connection and (as usual) do not share an anti-replay cache, then they may all receive and process (prefixes of) the same early traffic data. This motivates our model with multiple decryptors, and also requires the server application to defer some effects of early-data processing till handshake completion. Also, since data is sent before ciphersuite negotiation, the client may use relatively weak algorithms (or broken implementations) that the server would otherwise have a chance to refuse.

0.5-RTT incurs similar, lesser risks as the server sends data before the handshake completes. The server is subject to 0-RTT truncation attacks if it starts sending data before receiving the client's end of early data. Also, if the server relies on a client signature, it should not send sensitive data before handshake completion. In contrast with 0-RTT, sending 0.5-RTT traffic is a local configuration issue for the server; the client receives 0.5-RTT data after completing the handshake and does not distinguish it from 1-RTT data.

TLS 1.2 is routinely deployed with 'FalseStart', which is similar to 0.5-RTT but in the other direction, the client may locally decide to start sending encrypted application data as soon as it can compute the keys, before handshake completion. This places additional trust in the client's ciphersuite whitelist, inasmuch as sensitive data may be sent before confirming their correct negotiation with the server.

A Minimal Record Game We present a simplified, more liberal model of the Record that seeks to abstract away from the details of how the connection evolves. This facilitates the statement of our theorem, but as usual our results apply to the full $F^\star$ implementation, which does carefully keep track of the sequence of keys, as outlined at the end of this section.

We abstract the state of the connection by a context bitstring; as the handshake progresses, we concatenate more relevant handshake parameters to the context. For instance, after `ClientHello`, the context consists of the client's nonce n_C and its proposed ciphersuites and key exchange values; after `ServerHello`, it additionally contains the server nonce n_S, algorithm choice, key exchange value, etc.

Instead of modeling duplex channels between clients and servers, we consider separate sequences of streams in each direction. Our game (Figure 6) models re-keying and context extension for a sequence of streams (all in the same direction), covering 0RTT, 0.5RTT, and 1RTT traffic, relying on the multi-instance game $\mathsf{SE} = \mathsf{Multi}(\mathsf{Lhse})$ (see §7).

The game has oracles Init and InitD for generating multi-stream encryptors and decryptors in their initial state, indexed by n and m, respectively. We assume that their arguments determine the record algorithm. Their state consist of a current context, a current stream number j, and a local map I from stream numbers
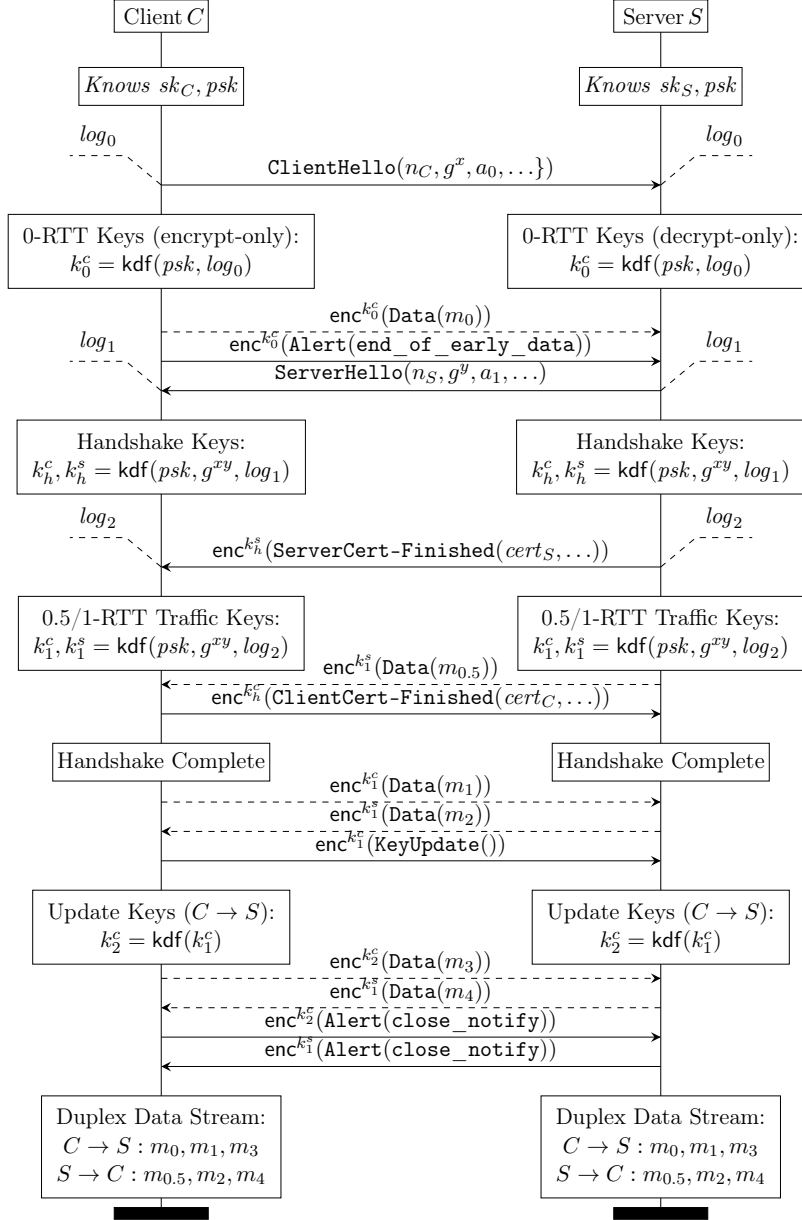
Figure 4: TLS 1.3 Draft-18 Protocol, seen from the viewpoint of the Record Protocol. Dotted arrows represent zero or more (encrypted) application data fragments. Each key stands for an instance of LHSE.

to the value of the context when they were installed. We use variables ctx, j, and I to refer to the fields of $E[m]$ and $D[n]$, respectively.

Oracles Extend and ExtendD allow the local context to be extended (concatenated) with new information

```
type state role = {ctx:context; j:int; epochs: seq (LHSE.state role)}
type encryptor = state Encrypt
type decryptor = state Decrypt
val init: c:ctx → ST (e:encryptor)
    (ensures  e.current = −1 ∧ e.context = ctx)
val initD: c:ctx → ST (d:decryptor)
    (ensures  d.current = −1 ∧ d.context = ctx)
val extend: e:encryptor → c:ctx → ST unit
    (ensures  e'.context = e.context + ctx)
val extendD: d:decryptor → c:ctx → ST unit
    (ensures  d'.context = d.context + ctx)


val install: e:encryptor → ST unit
val installD: d:decryptor → ST unit
val encrypt: c:encryptor → ℓ:len → LHSE.fragment ℓ → ST record
val decrypt: c:decryptor → record → ST (ℓ, LHSE.fragment ℓ)
```

Figure 5: Simpler Record Interface Corresponding to Game

at any time.

Oracles Install and InstallD installs an LHSE instance (allocating it if it does not exist) for encryption and decryption, respectively. Recall that calls to SE.Gen are memoized, so that an encryptor and a decryptor share the same stream if and only if they agree on the stream sequence number and context.

Oracles Encrypt and Decrypt apply encryption and decryption to the currently-installed stream. Some fragments are final: they terminate the stream and signal the need to install a new stream before continuing.

Definition 9 (Record). *Let $\epsilon_{\mathsf{Record}}(\mathcal{A}[q_e, q_d, q_i])$ be the advantage of an adversary $\mathcal{A}$ that makes at most q_e encryption queries and q_d decryption queries for each of the q_i LHSE instances created using install queries in the game of Figure 6.*

Theorem 5 (Record reduces to Lhse).

$$\epsilon_{\mathsf{Record}}(\mathcal{A}[q_e, q_d, q_i]) \leq q_i \, \epsilon_{\mathsf{Lhse}}(\mathcal{B}[q_e, q_d]).$$

Our game complies with the idealized interface for LHSE and relies on its conditional idealization. If $b = 0$, then the oracles operate purely on local state, and simply implement a real sequence of encrypted streams, under the control of the record state machine. If $b = 1$, then we get perfect authentication of (a prefix of) the whole sequence of streams of fragments. (This property is verified by typing our idealized record implementation.) The *ctx* field of encryptors and decryptor represents their implicitly authenticated shared context: unless there is an encryptor with a matching context, the ideal encryption log is empty hence decryption will fail. In particular, as soon as the context includes `ServerCert-Finished`, and thus the fresh TLS nonces n_C and n_S, we know that there is at most one encryptor and one decryptor.

More precisely, consider encryptor and decryptor states $E[n]$ and $D[m]$. If $E[n].I[j] = D[m].I[j]$ then also $E[n].I[j'] = D[m].I[j']$ for any $j' < j$. Thus, for instance, when $D[m]$ receives a final fragment, we know that $E[n]$ and $D[m]$ agree on the whole sequence of communicated fragment for the first j streams. By Theorem 5 these guarantees also hold for the real record for any game adversary $\mathcal{A}$, except with probability $\epsilon_{\mathsf{Record}}(\mathcal{A})$.

29

Game $\mathsf{Record}^b(\mathsf{Lhse}^b(\mathsf{LHSE}))$

$E \leftarrow \varnothing \quad D \leftarrow \varnothing$
$\mathsf{SE.Gen}, \mathsf{SE.GenD}, \mathsf{SE.Encrypt}, \mathsf{SE.Decrypt} \xleftarrow{\$} \mathsf{Multi}(\mathsf{Lhse}^b)$
return $\{\mathsf{Gen}, \mathsf{Extend}, \mathsf{Install}, \mathsf{Encrypt},$
$\qquad \mathsf{GenD}, \mathsf{ExtendD}, \mathsf{InstallID}, \mathsf{Decrypt}\}$

Oracle $\mathsf{Init}(n)$	**Oracle** $\mathsf{InitD}(m, ctx_0)$
if $E[n] = \bot$	**if** $D[m] = \bot$
$\quad E[n] \leftarrow \{ctx \leftarrow n; j \leftarrow 0;$	$\quad D[m] \leftarrow \{ctx \leftarrow ctx_0;$
$\qquad\qquad I \leftarrow \varnothing\}$	$\qquad\qquad j \leftarrow 0; I \leftarrow \varnothing\}$

Oracle $\mathsf{Extend}(n, \delta)$	**Oracle** $\mathsf{ExtendD}(m, \delta)$
if $E[n]$ *exists*	**if** $D[m]$ *exists*
$\quad ctx \leftarrow ctx + \delta$	$\quad ctx \leftarrow ctx + \delta$

Oracle $\mathsf{Install}(n)$	**Oracle** $\mathsf{InstallD}(m)$
if $E[n]$ exists with $I[j] = \bot$	**if** $D[m]$ *exists with* $I[j] = \bot$
$\quad I[j] \leftarrow (ctx, j)$	$\quad I[j] \leftarrow (ctx, j)$
$\quad \mathsf{SE.Gen}(I[j])$	$\quad \mathsf{SE.GenD}(I[j])$

Oracle $\mathsf{Encrypt}(n, \ell, f)$	**Oracle** $\mathsf{Decrypt}(m, v)$
if $E[n]$ exists with $I[j] \neq \bot$	**if** $D[m]$ *exists with* $I[j] \neq \bot$
$\quad v \leftarrow \mathsf{SE.Encrypt}(I[j], \ell, f)$	$\quad f \leftarrow \mathsf{SE.Decrypt}(I[j], d, v)$
$\quad$ **if** (*final* f) $j \leftarrow j + 1$	$\quad$ **if** $f \neq \bot \wedge$ *final* f
$\quad$ **return** v	$\qquad j \leftarrow j + 1$
	$\quad$ **return** f

Figure 6: The TLS 1.3 Record Game

Application to 0RTT We briefly show how to control our game to model 0RTT and 0.5-RTT. For 0RTT, the client is the encryptor and the server is the decryptor. Both use the encryptor index n as initial context, representing the content of `ClientHello`, notably the fresh client random n_C. Conversely, the decryptor index m (including the fresh server random n_S) is *not* included in the initial context of InitD. As both parties install their first stream ($j = 0$) for 0RTT, this reflects that the underlying TLS key derivation (k_0^c in Figure 4) depends only on client-side information. Thus, although 0RTT traffic is protected, it may be decrypted by multiple server instances with different indexes m.

Calls to $\mathsf{ExtendD}$ and Extend reflect handshake communications in the other direction, as the `ServerCert-Finished` stream is sent and received, causing ctx to be extended with (at least) m. Afterwards, as the two parties successively install streams for the TLS keys k_h^c, k_1^c, k_2^c, ..., successful decryption guarantees agreement on a context that includes the pair n, m. Thus, in this usage of our Record game at most one server will successfully decrypt the first encrypted handshake fragment from the client, and from this point all streams are one-to-one.

Application to 0.5-RTT The server is the encryptor, the client the decryptor and, since they both have initial access to the first message exchange, we may select as index n that includes the client hello and server hello messages and implicitly authenticate the pair n_C, n_S. Thus, there is at most one honest client decryptor for 0.5-RTT and, from the client's viewpoint, successful decryption of the first handshake

fragment ensures agreement on this context. Still (at least from the record's viewpoint) the server is not guaranteed there is a matching decryptor until it receives `ClientCert-Finished` in the other direction and transitions to 1RTT.

Verified Implementation (Outline) Our TLS Record implementation supports sequences of streams for the full protocol described in Figure 4 and its TLS 1.2 counterpart.

Stream Sequences As described in the game above, it maintains a current stream for each direction, and it receives 'extend' and 'install' commands from the handshake protocol as the connection gets established. Its indexes (*ctx* in the game) consist of a summary of the handshake context available at the time of key derivation (always including the algorithms to use). In contrast with our game, which models all communications in a single direction, our code supports 'duplex' communications. This is necessary, for instance, for synchronizing key updates and connection closure between clients and servers. Our code also maintains a small (type-based) state machine that controls the availability of the current streams for sending application data.

Re-keying and Corruption The state machine enforces a limit on the maximum number of fragments that can be sent with one key to prevent sequence number overflows and account for the birthday bound weakness of AES-GCM. On key updates we delete old keys and model the corruption of individual streams using leak and coerce functions for keys. This is in keeping with the static corruption modeling of MITLS, e.g. to account for insecure handshake runs.

Fragment API Our code for the record is parameterized by a type of abstract application-data plaintexts, indexed by a unique stream identifier and an apparent fragment length. Type abstraction guarantees that, if the stream identifier is *safe* (a property that depends on the handshake and key derivation process), the idealized TLS implementation never actually accesses their actual length and contents, a strong and simple confidentiality property.

Our API has a configuration to control 0-RTT and 0.5-RTT as the connection is created. In particular, 0-RTT plaintexts have their own type, (indexed by an identifier that describe their 0-RTT context) which should help applications treat it securely. Conversely, 0.5-RTT is modeled simply by enabling earlier encryption of 1-RTT traffic.

Message API Our code also has a higher-level API with messages as (potentially large) bytestrings instead of individual fragments. As usual with TLS, message boundaries are application specific, whereas applications tend to ignore fragment boundaries. Nonetheless, our code preserves apparent message boundaries, never caches or fragments small messages, and supports message length-hiding by taking as inputs both the apparent (maximal) size ℓ_{max} of the message and its private (actual) size ℓ_m. It has a simple fragmentation loop, apparently sending up to 2^{14} bytes at each iteration, starting with $\ell_{max} - \ell_m$ bytes of padding follows by the actual data. (This ensures that an application that waits for the whole message never responds before receiving the last fragment.) We do not model de-fragmentation on the receiving end; our code delivers fragments as they arrive in a buffer provided by the application for reassembling its messages.

The correctness and security of this construction on top of the fragment API is verified by typing, essentially abstracting sequences of fragments into sequences of bytes for the benefit of the application, and turning close-notify alert fragments into end-of-files. (See also [24] for a cryptographic treatment of fragmentation issues.)

Module Name	Verification Goals	LoC	% annot	ML LoC	Time
StreamAE	Game StAE^b from §6	292	30%	339	519s
AEADProvider	Safety (dispatcher from OCaml to Crypto.AEAD or OpenSSL)	299	10%	583	349s
Crypto.AEAD	Proof of Theorem 2 from §5	2,806	87%	345	1267s
Crypto.Plain	Plaintext module for AEAD	130	20%	63	13s
Crypto.AEAD.Encoding	AEAD encode function from §5 and injectivity proof	395	30%	172	88s
Crypto.Symmetric.PRF	Game PrfCtr^b from §4	511	30%	489	155s
Crypto.Symmetric.Cipher	Agile PRF functionality	117	10%	131	9s
Crypto.Symmetric.AES	Safety (concrete reference implementation of	561	10%	1,024	62s
Crypto.Symmetric.Chacha20	block ciphers)	205	10%	172	38s
Crypto.Symmetric.UF1CMA	Game $\mathsf{MMac1}^b$ from §3	420	30%	303	54s
Crypto.Symmetric.MAC	Agile MAC functionality	312	10%	261	73s
Crypto.Symmetric.GF128	$GF(128)$ polynomial evaluation and GHASH encoding	264	20%	268	33s
Crypto.Symmetric.Poly1305	$GF(2^{130}-5)$ polynomial evaluation and Poly1305 encoding	1,297	60%	860	465s
Crypto.Symmetric.Poly1305.*	Bignum library and supporting lemmas for the functional correctness of field operations	4,264	80%	1310	450s
Crypto.Buffer.*	A verified model of mutable buffers (implemented natively)	1,870	N/A	0	625s
Total		13,743	56.5%	6,320	1h 10m

Table 2: Modules in our verified record layer implementation

9 Experimental evaluation

We evaluate our reference implementation of the TLS record layer both qualitatively (going over the verified goals of the various modules and how they relate to the games presented in the paper, and checking that our implementation interoperates with other TLS libraries) and quantitatively (measuring the verification and runtime performance).

Verification evaluation Table 2 summarizes the modular structure of our code. Most of the verification burden comes from the security proof of AEAD (totaling approximately 1900 lines of annotation out of a total of about 2500 lines of $F^\star$) and the functional correctness proof of the MAC implementations (totaling over 5000 lines of annotations and lemmas). For the latter, we implement a new big number library to verify the low-level implementations of various mathematical operations (such as point multiplication on elliptic curves or multiplication over finite fields) using machine-sized integers and buffers. We use it to prove the correctness of the polynomial computations for Poly1305 and GHASH.

Current limitations At the time of writing, our code for the record layer still includes a few assumptions that we are in the process of eliminating. Specifically, in Crypto.Symmetric.PRF, we admit a simple integer-normalization invariant (these integers are immutable, but not yet tracked by our invariant). We also have five local assumptions related to memory allocation and write effects, all due to a discrepancy in the style of specification used for low-level buffers in two modules of our development. Aside from these, our functional correctness proof of low-level arithmetic over GF(128) is not integrated with our main AEAD security development. Finally, in particular to experiment with other AEAD providers in the next section, the TLS-specific stateful encryption in StreamAE is verified on top of an idealized AEAD interface that slightly differs from the one exported from our idealized Crypto.AEAD construction. For instance, the

latter uses a lower-level memory model and a mutable representation of fragments. These limitations are minor with regards to the amount of code verified so far.

Besides, the record layer is part of a larger, partially-verified codebase—notably, we leave a complete verification of the TLS 1.3 handshake and its integration with our code as future work.

Interoperability Our record implementation supports both TLS 1.3 and 1.2 and exposes them through a common API. We have tested interoperability for our TLS 1.2 record layer with all major TLS implementations. For TLS 1.3 draft-14, we tested interoperability with multiple implementations, including BoringSSL, NSS, BoGo, and Mint, at the IETF96 Hackathon. At the time of writing, there are few implementations of TLS 1.3 draft-18, but we tested interoperability with Mint. In all cases, our clients were able to connect to interoperating servers using an ECDHE or PSK_ECDHE key exchange, then to exchange data with one of the following AEAD algorithms: AES256-GCM, AES128-GCM, and ChaCha20-Poly1305. Similarly, our servers were able to accept connections from interoperating clients that support the above ciphersuites.

Performance We evaluate the performance of our record layer implementation at two levels. First, we compare our implementation of AEAD encryption extracted to C using an experimental backend for $F^\star$ to OpenSSL 1.1.0 compiled with the no−asm option, disabling handwritten assembly optimizations. Our test encrypts a random payload of 2^{14} bytes with 13 bytes of constant associated data. We report averages over 1000 runs on an Intel Core i7 4600U CPU (2.1GHz) running Linux.

	Crypto.AEAD	OpenSSL
ChaCha20-Poly1305	17.95 cycles/byte	7.20 cycles/byte
AES256-GCM	454.16 cycles/byte	24.08 cycles/byte
AES128-GCM	357.23 cycles/byte	20.82 cycles/byte

Our implementation is 17-19 times slower than OpenSSL for AES-GCM and 2.5 times slower for ChaCha20-Poly1305. Note that the performance of custom assembly implementations can be signicantly better. OpenSSL with assembly can perform ChaCha20-Poly1305 in about one cycle per byte and can do AES-GCM in a small fraction of a cycle per byte.

Next, we measure the throughput of our record layer integrated into miTLS by downloading one gigabyte of random data from a TLS server on the local network. We compare two different integration methods: first, we extract the verified record layer in OCaml, and compile it alongside the OCaml-extracted miTLS. Then, we build an $F^\star$ interface to the C version of our record implementation and call it from miTLS. We compare these results with the default AEAD provider of miTLS (based on OpenSSL 1.1.0 with all optimizations), and curl (which uses OpenSSL for the full TLS protocol).

	OCaml	C	OpenSSL	curl
ChaCha20-Poly1305	74 KB/s	69 MB/s	354 MB/s	440 MB/s
AES256-GCM	68 KB/s	2.2 MB/s	398 MB/s	515 MB/s
AES128-GCM	89 KB/s	2.3 MB/s	406 MB/s	571 MB/s

We observe that miTLS is not a limiting factor in these benchmarks as its performance using the OpenSSL implementation of AEAD encryption is comparable to that of libcurl.

Unsurprisingly, the OCaml version of our verified implementation performs very poorly. This is due to the high overhead of both memory operations and arithmetic computations in the OCaml backend of $F^\star$ (which uses garbage-collected lists for buffers, and arbitrary-precision zarith integers). The C extracted version is over 30,000 times faster, but for AES, remains far slower than the Intel hardware-accelerated assembly version in OpenSSL. The comparison is fairer for ChaCha20, where we achieve 20% of the OpenSSL throughput.

Although our code is optimized for verification and modularity rather than performance, we do not believe that we can close the performance gap only by improving F* code. Instead, for performance-critical primitives such as AES, we intend to selectively link our F* code with assembly code proven to correctly implement a shared functional specification. We leave this line of research for future work.

References

[1] N. J. AlFardan and K. G. Paterson, "Lucky thirteen: Breaking the TLS and DTLS record protocols," in *2013 IEEE Symposium on Security and Privacy*, 2013, pp. 526–540.

[2] J. B. Almeida, M. Barbosa, G. Barthe, and F. Dupressoir, "Verifiable side-channel security of cryptographic implementations: Constant-time MEE-CBC," in *23rd International Conference on Fast Software Encryption, FSE 2016*, 2016, pp. 163–184.

[3] C. Badertscher, C. Matt, U. Maurer, P. Rogaway, and B. Tackmann, "Augmented secure channels and the goal of the TLS 1.3 record layer," in *9th International Conference on Provable Security, ProvSec 2015*, 2016, pp. 85–104.

[4] G. Barthe, C. Fournet, B. Grégoire, P. Strub, N. Swamy, and S. Zanella-Béguelin, "Probabilistic relational verification for cryptographic implementations," in *41st Annual ACM Symposium on Principles of Programming Languages, POPL 2014*, 2014, pp. 193–206.

[5] M. Bellare and P. Rogaway, "The security of triple encryption and a framework for code-based game-playing proofs," in *Advances in Cryptology – EUROCRYPT 2006*, 2006, pp. 409–426.

[6] ——, "Code-based game-playing proofs and the security of triple encryption," Cryptology ePrint Archive, Report 2004/331, 2004, http://eprint.iacr.org/2004/331.

[7] M. Bellare and B. Tackmann, "The multi-user security of authenticated encryption: AES-GCM in TLS 1.3," in *Advances in Cryptology – CRYPTO 2016*, 2016, pp. 247–276.

[8] M. Bellare, O. Goldreich, and A. Mityagin, "The power of verification queries in message authentication and authenticated encryption," *IACR Cryptology ePrint Archive*, vol. 2004, p. 309, 2004. [Online]. Available: http://eprint.iacr.org/2004/309

[9] D. J. Bernstein, "The Poly1305-AES message-authentication code," in *12th International Workshopo on Fast Software Encryption, FSE 2005*, 2005, pp. 32–49.

[10] ——, "Stronger security bounds for Wegman-Carter-Shoup authenticators," in *Advances in Cryptology – EUROCRYPT 2005*, 2005, pp. 164–180.

[11] K. Bhargavan and G. Leurent, "On the practical (in-)security of 64-bit block ciphers: Collision attacks on HTTP over TLS and OpenVPN," Cryptology ePrint Archive, Report 2016/798, 2016, http://eprint.iacr.org/2016/798.

[12] K. Bhargavan, C. Fournet, M. Kohlweiss, A. Pironti, and P. Strub, "Implementing TLS with verified cryptographic security," in *2013 IEEE Symposium on Security and Privacy*, 2013, pp. 445–459.

[13] 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 *2014 IEEE Symposium on Security and Privacy*, 2014, pp. 98–113.

[14] K. Bhargavan, C. Fournet, M. Kohlweiss, A. Pironti, P.-Y. Strub, and S. Zanella-Béguelin, "Proving the TLS handshake secure (as it is)," Cryptology ePrint Archive, Report 2014/182, 2014, http://eprint.iacr.org/2014/182/.

[15] H. Böck, "Wrong results with Poly1305 functions," https://mta.openssl.org/pipermail/openssl-dev/2016-March/006413, 2016.

[16] H. Böck, A. Zauner, S. Devlin, J. Somorovsky, and P. Jovanovic, "Nonce-disrespecting adversaries: Practical forgery attacks on GCM in TLS," Cryptology ePrint Archive, Report 2016/475, 2016, http://eprint.iacr.org/2016/475.

[17] C. Boyd, B. Hale, S. F. Mjølsnes, and D. Stebila, "From stateless to stateful: Generic authentication and authenticated encryption constructions with application to TLS," in *Topics in Cryptology – CT-RSA 2016*, 2016, pp. 55–71.

[18] 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 *2016 IEEE Symposium on Security and Privacy*, 2016, pp. 470–485.

[19] A. Datta, A. Derek, J. C. Mitchell, and B. Warinschi, "Computationally sound compositional logic for key exchange protocols," in *19th IEEE Computer Security Foundations Workshop, (CSFW-19 2006), 5-7 July 2006, Venice, Italy*, 2006, pp. 321–334.

[20] B. Dowling, M. Fischlin, F. Günther, and D. Stebila, "A cryptographic analysis of the TLS 1.3 handshake protocol candidates," in *22nd ACM Conference on Computer and Communications Security*, 2015, pp. 1197–1210.

[21] ——, "A cryptographic analysis of the TLS 1.3 draft-10 full and pre-shared key handshake protocol," http://eprint.iacr.org/2016/081, 2016.

[22] T. Duong and J. Rizzo, "Here come the $\oplus$ ninjas," Available at http://nerdoholic.org/uploads/dergln/beast_part2/ssl_jun21.pdf, May 2011.

[23] M. J. Dworkin, "Recommendation for block cipher modes of operation: Galois/Counter mode (GCM) and GMAC," National Institute of Standards & Technology, Tech. Rep. SP 800-38D, 2007.

[24] M. Fischlin, F. Günther, G. A. Marson, and K. G. Paterson, "Data is a stream: Security of stream-based channels," in *Advances in Cryptology - CRYPTO 2015*, 2015, pp. 545–564.

[25] M. Fischlin, F. Günther, B. Schmidt, and B. Warinschi, "Key confirmation in key exchange: A formal treatment and implications for TLS 1.3," in *2016 IEEE Symposium on Security and Privacy*, 2016, pp. 197–206.

[26] C. Fournet, M. Kohlweiss, and P. Strub, "Modular code-based cryptographic verification," in *18th ACM Conference on Computer and Communications Security, CCS 2011*, 2011, pp. 341–350.

[27] F. Giesen, F. Kohlar, and D. Stebila, "On the security of TLS renegotiation," in *2013 ACM Conference on Computer and Communications Security, CCS 2013*, 2013, pp. 387–398.

[28] P. Gutmann, "Encrypt-then-MAC for Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)," IETF RFC 7366, 2014.

[29] T. Jager, F. Kohlar, S. Schäge, and J. Schwenk, "On the security of TLS-DHE in the standard model," in *Advances in Cryptology – CRYPTO 2012*, 2012, pp. 273–293.

[30] 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 *22nd ACM Conference on Computer and Communications Security*, 2015, pp. 1185–1196.

[31] M. Kohlweiss, U. Maurer, C. Onete, B. Tackmann, and D. Venturi, "(de-) constructing TLS 1.3," in *Progress in Cryptology–INDOCRYPT 2015.* Springer, 2015, pp. 85–102.

[32] H. Krawczyk, "LFSR-based hashing and authentication," in *Advances in Cryptology – CRYPTO 1994*, 1994, pp. 129–139.

[33] ——, "The order of encryption and authentication for protecting communications (or: how secure is SSL?)," Cryptology ePrint Archive, Report 2001/045, 2001, http://eprint.iacr.org/2001/045.

[34] ——, "A unilateral-to-mutual authentication compiler for key exchange (with applications to client authentication in TLS 1.3)," in *23rd ACM Conference on Computer and Communications Security, CCS 2016*, 2016.

[35] H. Krawczyk and H. Wee, "The OPTLS protocol and TLS 1.3," Cryptology ePrint Archive, Report 2015/978, 2015, http://eprint.iacr.org/2015/978.

[36] H. Krawczyk, K. G. Paterson, and H. Wee, "On the security of the TLS protocol: A systematic analysis," in *Advances in Cryptology – CRYPTO 2013*, 2013, pp. 429–448.

[37] A. Luykx and K. G. Paterson, "Limits on authenticated encryption use in TLS," Personal webpage: http://www.isg.rhul.ac.uk/~kp/TLS-AEbounds.pdf, 2015.

[38] D. McGrew and J. Viega, "Flexible and efficient message authentication in hardware and software." Unpublished draft. Available online at http://www.cryptobarn.com/.

[39] D. McGrew, "An interface and algorithms for authenticated encryption," IETF RFC 5116, 2008.

[40] B. Möller, T. Duong, and K. Kotowicz, "This POODLE Bites: Exploiting The SSL 3.0 Fallback," Available at https://www.openssl.org/~bodo/ssl-poodle.pdf, 2014.

[41] Y. Nir and A. Langley, "ChaCha20 and Poly1305 for IETF protocols," IETF RFC 7539, 2015.

[42] K. G. Paterson, T. Ristenpart, and T. Shrimpton, "Tag size does matter: Attacks and proofs for the TLS record protocol," in *Advances in Cryptology – ASIACRYPT 2011*, 2011, pp. 372–389.

[43] J. Rizzo and T. Duong, "The CRIME Attack," September 2012.

[44] J. Salowey, A. Choudhury, and D. McGrew, "AES Galois Counter Mode (GCM) cipher suites for TLS," IETF RFC 5288, 2008.

[45] P. Sarkar, "A trade-off between collision probability and key size in universal hashing using polynomials," Cryptology ePrint Archive, Report 2009/048, 2009, http://eprint.iacr.org/2009/048.

[46] V. Shoup, "On fast and provably secure message authentication based on universal hashing," in *Advances in Cryptology – CRYPTO 1996*, 1996, pp. 313–328.

[47] B. Smyth and A. Pironti, "Truncating TLS connections to violate beliefs in web applications," Inria, Tech. Rep. hal-01102013, Oct. 2014. [Online]. Available: https://hal.inria.fr/hal-01102013

[48] J. Somorovsky, "Systematic fuzzing and testing of TLS libraries," in *23rd ACM Conference on Computer and Communications Security, CCS 2016*, 2016.

[49] N. Swamy, C. Hritcu, C. Keller, A. Rastogi, A. Delignat-Lavaud, S. Forest, K. Bhargavan, C. Fournet, P.-Y. Strub, M. Kohlweiss, J.-K. Zinzindohoue, and S. Zanella-Béguelin, "Dependent types and multi-monadic effects in F*," in *43nd ACM Symposium on Principles of Programming Languages, POPL 2016*, 2016, pp. 256–270.

[50] R. Święcki, "ChaCha20/Poly1305 heap-buffer-overflow," CVE-2016-7054, 2016.

[51] P. Swire, J. Hemmings, and A. Kirkland, "Online privacy and ISPs: ISP access to consumer data is limited and often less than access by others," Georgia Tech, Tech. Rep., 2016.

[52] D. Wagner and B. Schneier, "Analysis of the SSL 3.0 protocol," in *2nd USENIX Workshop on Electronic Commerce, WOEC 1996*, 1996, pp. 29–40.

Security proofs

Section 3

Claim 1. *IND-UF-1CMA implies UF-1CMA when $q_v \geq 1$ and random tags are neither necessary nor sufficient for unforgeability.*

(i) To see the latter, consider the $\mathsf{mac}(k, m)$ function that just returns the random key k as its tag for every message. Returning k for t^* break UF-1CMA but the scheme nevertheless satisfies the randomness requirement for $\mathcal{A}[\ell_m, 0]$ adversaries. (ii) To see the former, consider the $\mathsf{Mac1}^b$ game with $r = 0$ (denoted $\mathsf{Mac1}_0^b$). We can relate $\mathsf{Mac1}_0^b$ to the standard unforgeability definition as follows: clearly, an adversary who wins the UF-1CMA game can win the $\mathsf{Mac1}_0^b$ game by submitting his forgery (m, t) to Verify, which returns false if and only if $b = 1$. Conversely, an adversary that wins the $\mathsf{Mac1}_0^b$ game with advantage $\epsilon_{\mathsf{Mac1}_0}(\mathcal{A}[\ell_m, q_v])$ must query Verify on at least one forgery with the same probability, since the adversary's view is otherwise independent of b. In the reduction, $\mathcal{B}[\ell_m]$ guesses in which of the q_v Verify queries $\mathcal{A}$ submits his first forgery and return false for all prior queries.[1] It follows that:

$$\epsilon_{\mathsf{Mac1}_0}(\mathcal{A}[\ell_m, q_v]) \leq q_v \, \epsilon_{\mathsf{UF\text{-}1CMA}}(\mathcal{B}[\ell_m]).$$

Proof of Lemma 1. Let H_i, $i \in [0, q_i]$ be hybrids of $\mathsf{MMac1}^1$ that only idealize the first i instances: that is, the first i queries to $\mathsf{H}_i.\mathsf{Keygen}$ behave as $\mathsf{MMac}^1.\mathsf{Keygen}$ while later queries behave as $\mathsf{MMac}^1.\mathsf{Coerce}(n, k \xleftarrow{\$} \mathsf{byte}^{\ell_k})$. In particular, H_{q_i} is equivalent to MMac^1 as all keys are idealized, and H_0 is equivalent to MMac^0 as no key is idealized.

By the triangle inequality:

$$\begin{aligned}
\mathcal{A}[\ell_m, q_v, q_i]) \leq & \, |\Pr[\mathcal{A}(\mathsf{H}_{q_i}) = 1] - \Pr[\mathcal{A}(\mathsf{H}_{q_i-1}) = 1]| \\
& + |\Pr[\mathcal{A}(\mathsf{H}_{q_i-1}) = 1] - \cdots \\
& + |\Pr[\mathcal{A}(\mathsf{H}_1) = 1] - \Pr[\mathcal{A}(\mathsf{H}_0) = 1]|
\end{aligned}$$

In the reduction, $\mathcal{B}$ picks i at random between 1 and q_i and uses its oracles to $\mathsf{Mac1}^b$ such that $\mathcal{A}$ either interacts with H_i or H_{i-1} depending on b. $\mathcal{B}$ returns the output of $\mathcal{A}$.

---

[1]See [8] for a detailed discussion of this proof strategy, and the reason why strong unforgeability, i.e. $log \neq (m^\star, t^\star)$ and not just $\forall t. \; log \neq (m^\star, t)$, is necessary.

By definition, $\epsilon_{\mathsf{Mac1}}(\mathcal{B}[\ell_m, q_v])$ is:

$$\frac{1}{q_i}\left(\sum_{i=1}^{q_i} |\Pr[\mathcal{A}^{\mathsf{H}_i} = 1] - \Pr[\mathcal{A}^{\mathsf{H}_{i-1}} = 1]|\right)$$

and therefore, $\mathcal{A}[\ell_m, q_v, q_i]) \leq q_i\,\epsilon_{\mathsf{Mac1}}(\mathcal{B}[\ell_m, q_v])$.

Note that we could express a tighter bound in case we had an upper bound for the number of verification queries that $\mathcal{A}$ makes per instance, which could be as low as $\lceil q_v/q_i \rceil$. We directly prove $\mathsf{MMac1}$ security for GHASH and Poly1305, so this has no effect on our main result. $\qquad\square$

Proof of Lemma 2. In the reduction, $\mathcal{B}$ guesses in which Verify query $\mathcal{A}$ might make its first forgery (say, query number q). This guess is correct with probability $1/q_v$. Note that a successful $\mathcal{A}$ may distinguish purely based on the tag and never submit a forgery; in that case, the simulation provided by $\mathcal{B}$ is perfect until the end of the experiment. $\mathcal{B}$ simulates the q_v verification queries by returning 0 or 1 depending on whether the tag-message pair was generated using a query to Mac. If $\mathcal{B}$ guessed correctly, the simulation is perfect up to the submission of the forgery. Now $\mathcal{B}$ submits the message-tag pair of the query q to its single-call Verify oracle. Let b'' be the result of this query, and b' be the output of $\mathcal{A}$ after running to completion with all further Verify oracles simulated as before.

$\mathcal{B}$ output $b' \wedge \neg b''$ as its guess. If the q^{th} query result indicates a forgery, then clearly $b = 0$, otherwise we either guessed q wrongly or the adversary did not query on a forgery and so b' is the result of a perfect simulation.

In both cases, conditioned on having guessed correctly, the success probability of $\mathcal{B}$ is greater equal the success probability of $\mathcal{A}$ and thus $1/q_v \cdot \epsilon_{\mathsf{MMac1}}(\mathcal{A}[\ell_m, q_v, q_i]) \leq \epsilon_{\mathsf{MMac1}}(\mathcal{B}[\ell_m, 1, q_i])$. $\qquad\square$

Proof of Theorem 1. We make use of Lemma 2 and first prove the bound $\frac{d \cdot \tau}{|R|}$ for an adversary that makes only a single Verify query. We refer to the instance in which the Verify query is made as the target instance.

Game 0 For a given r and message m, we can compute s and t from one another. Thus, we first perform the following game transformation (with no security loss):

Oracle $\mathsf{Mac}(n, m)$
$\overline{}$
if $keys[n] = \perp$ **return** $\perp$
if $log[n] \neq \perp$ **return** $\perp$
$k \leftarrow keys[n]$
$r\|s \leftarrow k$
if $n \notin H$
$\quad t \leftarrow \mathsf{MAC.mac}(k, m)$
else if $b \vee vlog[n] = \perp$
$\quad t \xleftarrow{\$} \mathsf{byte}^{\ell_t}$
else
$\quad t \xleftarrow{\$} \mathsf{byte}^{\ell_t}$
$\quad (v^*, t^*, m^*) \leftarrow vlog[n]$
$\quad s' \leftarrow t \boxminus \mathsf{hash}_r(\overline{m})$
$\quad v \leftarrow t^\star = \mathsf{hash}_r(\overline{m}^\star) \boxplus s'$
$\quad$**if** $v^\star \neq v$
$\quad\quad t \leftarrow \mathsf{MAC.mac}(k, m)$
$log[n] \leftarrow (m, t)$
return t

Oracle $\mathsf{Verify}(n, m^\star, t^\star)$
$\overline{}$
if $keys[n] = \perp$ **return** $\perp$
$k \leftarrow keys[n]$
$r\|s \leftarrow k$
if $n \notin H$
$\quad v^\star \leftarrow \mathsf{MAC.verify}(k, m^\star, t^\star)$
else if b
$\quad v^\star \leftarrow log[n] = (m^\star, t^\star)$
else if $log[n] \neq \perp$
$\quad (m, t) = log[n]$
$\quad s' \leftarrow t \boxminus \mathsf{hash}_r(\overline{m})$
$\quad v^\star \leftarrow t^\star = \mathsf{hash}_r(\overline{m}^\star) \boxplus s'$
else
$\quad v^\star \leftarrow t^\star = \mathsf{hash}_r(\overline{m}^\star) \boxplus s$
$\quad vlog[n] \leftarrow (v^\star, t^*, m^*)$
return $v^\star$

When answering Mac before Verify, instead of sampling the one-time pad s and adding it to finalize the WCS MAC, we sample instead the tag t that will be released. For verification we compute s' from t, m and r.

When answering Verify before Mac, we verify as before, but store $(v^\star, t^\star, m^\star)$ in an additional table. For Mac we sample the tag t at random, and compute its s' from t, m and r. Only if the answer we would have given in verification differs between s and s', we recompute the tag using s.

Game 1 For the target instance, we can distinguish two cases: either the Mac oracle or the Verify oracle gets queried first. Consider the following game in which idealizations are marked by *(∗#∗)*:

Oracle $\mathsf{Mac}(n, m)$	**Oracle** $\mathsf{Verify}(n, m^\star, t^\star)$
if $keys[n] = \bot$ **return** $\bot$	**if** $keys[n] = \bot$ **return** $\bot$
if $log[n] \neq \bot$ **return** $\bot$	$k \leftarrow keys[n]$
$k \leftarrow keys[n]$	$r\|s \leftarrow k$
$r\|s \leftarrow k$	**if** $n \notin H$
if $n \notin H$	$\quad v^\star \leftarrow \mathsf{MAC.verify}(k, m^\star, t^\star)$
$\quad t \leftarrow \mathsf{MAC.mac}(k, m)$	**else if** b
else if $b \vee vlog[n] = \bot$	$\quad v^\star \leftarrow log[n] = (m^\star, t^\star)$
$\quad t \xleftarrow{\$} \mathsf{byte}^{\ell_t}$	**else if** $log[n] \neq \bot$
else	$\quad$ *(∗1∗)*
$\quad t \xleftarrow{\$} \mathsf{byte}^{\ell_t}$	$\quad v^\star \leftarrow log[n] = (m^\star, t^\star)$
$\quad$ *(∗3∗)*	**else**
$log[n] \leftarrow (m, t)$	$\quad$ *(∗2∗)*
return t	$\quad v^\star \leftarrow 0$
	$vlog[n] \leftarrow (v^\star, t^*, m^*)$
	return $v^\star$

We separately bound the probabilities for the following two events: *(∗1∗)* The oracle call to $\mathsf{Verify}(n, \overline{m}^\star, t^\star)$ accepts a forgery after calling $\mathsf{Mac}(n, \overline{m})$, *(∗2∗ or ∗3∗)* The oracle call to $\mathsf{Verify}(n, \overline{m}^*, t^\star)$ accepts a forgery before calling $\mathsf{Mac}(n, \overline{m})$ or the Mac query has to correct its tag t, because $v \neq v^\star$.

This second event requires the adversary to guess either the value s or the value t, both sampled at random, to set $v^\star$ or v to 1 respectively. We bound this probability by $2^{-\ell_t+1}$.

The first event dominates the probability. The adversary wins if and only if $t^\star \boxminus t = \mathsf{hash}_r(\overline{m}^\star) \boxminus \mathsf{hash}_r(\overline{m})$ for $m \neq m^\star$. This probability is bound by the ϵ-almost-$\boxminus$-universal property of GHASH and Poly1305. □

Definition 10. *A family of function* $\{\mathsf{hash}_r\}_{r \in R}$ *is* ϵ*-almost-$\boxminus$-universal, if* $\forall x \neq x' \in \mathbb{F}^d, y \in \mathsf{byte}^{\ell_t}$:

$$\Pr_{r \leftarrow R}[\mathsf{hash}_r(x) \boxminus \mathsf{hash}_r(x') = y] \leq \epsilon$$

Theorem 6. *The family of function* $\{\mathsf{hash}_r\}_{r \in R}$ *with* $\mathsf{hash}_r = \mathsf{tag}(\sum_{i=1..d} \overline{m}_{d-i} r^i \text{ in } \mathbb{F})$ *is* $\frac{d \cdot \tau}{|R|}$*-almost-$\boxminus$-universal for*

$$\tau = \begin{cases} 8 & \textit{for } \bar{\cdot}, \mathsf{tag}, \mathbb{F} \textit{ and } R \textit{ of Poly1305} \\ 1 & \textit{for } \bar{\cdot}, \mathsf{tag}, \mathbb{F} \textit{ and } R \textit{ of GHASH} \end{cases}$$

Proof. We start with a lemma specific to the Poly1305 truncation, which leads to $\tau = 8$. For other variants of the construction, e.g. by further truncating the tag, any similar lemma would suffice in the proof below.

Lemma 4. *For all* $a, a^\star \in 0..2^{130} - 6$ *and* $y \in 0..2^{128} - 1$*, if* $a^\star - a \bmod 2^{128} = y$*, then* $a^\star - a = y + \delta \cdot 2^{128}$ *in* $\mathbb{F}$ *for some* $\delta \in -4..3$*.*

Proof. The hypothesis is equivalent to $(a^\star \bmod 2^{128}) - (a^\star \bmod 2^{128}) = y + \gamma \cdot 2^{128}$ for some $\gamma \in -1..0$, so we can use $\delta = \gamma + a^\star/2^{128} - a/2^{128}$. $\qquad\square$

We have that $\mathsf{hash}_r(x) \boxminus \mathsf{hash}_r(x') = y$ if and only if

$$\sum_{i=1..d}(x_{d-i} - x'_{d-i})r^i - y_\delta = 0 \quad \text{in } \mathbb{F} \tag{1}$$

for some $y_\delta = y + \delta \cdot 2^{128}$ with $\delta \in \begin{cases} -4..3 & \text{for Poly1305} \\ 0 & \text{for GHASH} \end{cases}$

Equation 1, defines a non-null polynomial of degree at most d and tests whether r is a root of that polynomial. Since each polynomial is not null, it has at most d roots, and the equality holds with probability at most $d/|R|$. We arrive at the bound stated in the theorem by multiplying by the number of possible polynomials in the construction-specific step. $\qquad\square$

Section 4

Proof of Lemma 3. When $\mathcal{A}$ queries EvalEnx or EvalDex (they are exactly the same except for parameter names), $\mathcal{B}$ calls Prf^b.Eval, XORs the value provided by $\mathcal{A}$, and truncates. To compute k_0 when $j_0 = 1$, $\mathcal{B}$ queries Prf.Eval on 0^{ℓ_b}. Whenever $\mathcal{A}$ queries PrfCtr.EvalKey on nonce $j_0\|n$, $\mathcal{B}$ queries Prf.Eval on the same input, prepends k_0, and truncates.

Analysis The requirement $j > j_0$ in EvalEnx and EvalDex, guarantees that EvalKey uses disjoint parts of the PRF's domain regardless of j_0. The use of $j_0\|n$ in EvalKey also guarantees that k_0 is computed using a disjoint element of the domain. As the simulation is perfect $\mathcal{B}$ can forward the output of $\mathcal{A}$ as its own guess, making in total $q_b + q_g + j_0$ queries to Prf. $\qquad\square$

Section 5

Proof of AEAD *construction.* As a first step of the proof of AEAD security we model the composition of cryptographic functionalities in our implementation using abstract keys with state.

Composing auxiliary games We introduce the following composed game $\mathsf{PrfCtr}^b(\mathsf{PRF}, \mathsf{MMac1}^b(\mathsf{MAC}))$ that restrict adversaries from accessing the real MAC keys. This will allow us to idealize both the PRF and the MAC functionalities in order to idealize the authentication of ciphertexts:

Game $\mathsf{PrfCtr}^b(\mathsf{PRF}, \mathsf{MMac1}^b)$	**Oracle** $\mathsf{EvalKey}(j\|n)$
$T \leftarrow \varnothing;\ M \leftarrow \varnothing;$	**if** $j \neq j_0$ **return** $\perp$
$(\mathsf{Mac}, \mathsf{Verify},$	**if** $T[j\|n] = \perp$
$\ \mathsf{Keygen}, \mathsf{Coerce}) \leftarrow \mathsf{MMac1}^b()$	$\quad$ **if** b
$k \xleftarrow{\$} \mathsf{PRF.keygen}()$	$\quad\quad \mathsf{Keygen}(n)$
if $j_0 \wedge \neg b$	$\quad$ **else**
$\quad o \leftarrow \mathsf{PRF.eval}(k, 0^{\ell_b})$	$\quad\quad o \leftarrow \mathsf{PRF.eval}(k, j\|n)$
$\quad k_0 \leftarrow \mathsf{truncate}(o, \ell_{k_0})$	$\quad\quad k_m \leftarrow \mathsf{truncate}(k_0\|o, \ell_k)$
$\quad \mathsf{MMac1}.k_0 \leftarrow k_0$	$\quad\quad \mathsf{Coerce}(n, k_m)$
return $\{\mathsf{EvalKey}, \mathsf{EvalEnx},$	$\quad T[j\|n] \leftarrow \top$
$\quad\quad\quad \mathsf{EvalDex}, \mathsf{Mac}, \mathsf{Verify}\}$	

with the same EvalEnx and EvalDex oracles as in §4. Thus, the adversary has access to the oracles EvalKey, EvalEnx and EvalDex of the PrfCtr^b game; in addition, for every nonce n queried to EvalKey, the adversary is given access to oracles Mac and Verify of the MMac1 game. In the composed game, the adversary is not given direct access to the MAC keys; instead the game calls MMac1.Coerce to turn PRF blocks into MAC keys.

When it is clear from the context we will omit the game parameters $\mathsf{PRF}, \mathsf{MAC}$ and instead write $\mathsf{PrfCtr}^b(\mathsf{MMac1}^b)$. Both the PrfCtr^b and the $\mathsf{MMac1}^b$ components of the game use the same random bit b to determine the winning condition. However, we commonly use composed games in reduction proofs where some component uses a fixed value of b, which we write as superscript. For instance, we may write $\mathsf{PrfCtr}^b(\mathsf{MMac1}^0)$ to compose PrfCtr with only the real variant of $\mathsf{MMac1}^b$.

Lemma 5 (PrfCtr(MMac1) reduces to MMac1 and Prf). *Given $\mathcal{A}$ against $\mathsf{PrfCtr}(\mathsf{MMac1})$, we construct $\mathcal{B}$ against Prf and $\mathcal{C}$ against $\mathsf{MMac1}$ such that:*

$$\epsilon_{\mathsf{PrfCtr}(\mathsf{MMac1})}(\mathcal{A}[q_b, q_i, \ell_m, q_v]) \leq \epsilon_{\mathsf{Prf}}(\mathcal{B}[q_b + q_i + j_0])$$
$$+ \epsilon_{\mathsf{MMac1}}(\mathcal{C}[\ell_m, q_v, q_i])$$

Proof. We write $\mathcal{A}^{\mathsf{PrfCtr}^b(\mathsf{MMac1}^b))}$ to refer to the adversary's output bit when interacting with such a game.
By definition, $\epsilon_{\mathsf{PrfCtr}(\mathsf{MMac1})}(\mathcal{A}[q_b, q_i, \ell_m, q_v])$ is:

$$| \Pr[\mathcal{A}^{\mathsf{PrfCtr}^1(\mathsf{MMac1}^1)} = 1] - \Pr[\mathcal{A}^{\mathsf{PrfCtr}^0(\mathsf{MMac1}^0)} = 1] |$$

By the triangle inequality:

$$| \Pr[\mathcal{A}^{\mathsf{PrfCtr}^1(\mathsf{MMac1}^1)} = 1] - \Pr[\mathcal{A}^{\mathsf{PrfCtr}^0(\mathsf{MMac1}^1)} = 1] |$$
$$\leq | \Pr[\mathcal{A}^{\mathsf{PrfCtr}^1(\mathsf{MMac1}^1)} = 1] - \Pr[\mathcal{A}^{\mathsf{PrfCtr}^1(\mathsf{MMac1}^0)} = 1] |$$
$$+ | \Pr[\mathcal{A}^{\mathsf{PrfCtr}^1(\mathsf{MMac1}^0)} = 1] - \Pr[\mathcal{A}^{\mathsf{PrfCtr}^0(\mathsf{MMac1}^0)} = 1] |$$

This justifies the game transformation that first switches from $\mathsf{PrfCtr}^b(\mathsf{MMac1}^0)$ to $\mathsf{PrfCtr}^b(\mathsf{MMac1}^1)$. We bound the difference:

$$| \Pr[\mathcal{A}^{\mathsf{PrfCtr}^1(\mathsf{MMac1}^1)} = 1] - \Pr[\mathcal{A}^{\mathsf{PrfCtr}^1(\mathsf{MMac1}^0)} = 1] |$$

by $\epsilon_{\mathsf{MMac1}}(\mathcal{C}[\ell_m, q_v, q_i])$, where $\mathcal{C}$ runs $\mathcal{A}$ against the game $\mathsf{PrfCtr}^1(\mathsf{MMac1}^b)$ by simulating PrfCtr^1 and returning the output of $\mathcal{A}$. Similarly, we bound the difference:

$$| \Pr[\mathcal{A}^{\mathsf{PrfCtr}^1(\mathsf{MMac1}^0)} = 1] - \Pr[\mathcal{A}^{\mathsf{PrfCtr}^0(\mathsf{MMac1}^0)} = 1] |$$

by $\epsilon_{\mathsf{PrfCtr}}(\mathcal{B}[q_b, q_i])$, where $\mathcal{B}$ runs $\mathcal{A}$ against the game $\mathsf{PrfCtr}^b(\mathsf{MMac1}^0)$ by simulating $\mathsf{MMac1}^0$ and returning the output of $\mathcal{A}$. We conclude by applying Lemma 3. $\qquad\square$

We can similarly define a composed game $\mathsf{PrfCtr}^{b,b'}(\mathsf{MMac1}^b)$ for the specialized game $\mathsf{PrfCtr}^{b,b'}$. Our proof uses the following lemma about $\mathsf{PrfCtr}^{1,b'}$, which can be lifted to $\mathsf{PrfCtr}^{1,b'}(\mathsf{MMac1}^1)$.

Lemma 6 ($\mathsf{PrfCtr}^{1,b'}$ has 0 advantage). *Let $\mathsf{PrfCtr}^{1,b'}$ be the PrfCtr game using instead oracles $\mathsf{EvalEnx}'$, $\mathsf{EvalDex}'$ with b set to 1 and $\mathcal{A}$ guessing instead b'.*
    For all $\mathcal{A}$ we have $\epsilon_{\mathsf{PrfCtr}'}(\mathcal{A}) = 0$.

Proof. Queries to $\mathsf{EvalEnx}'$ and $\mathsf{EvalDex}'$ are answered differently. The condition $T[j\|n] = \bot$ ensures that $\mathcal{A}$ can query $\mathsf{EvalEnx}'$ successfully only once. $\mathcal{B}$ still calls $\mathsf{Prf.Eval}$, XORs p, and truncates. But as b is set, both branches return random values independently of b'. $\mathsf{EvalDex}'$ is already independent of b'. $\qquad\square$

Main proof The proof proceeds in a sequence of game steps.

Game 0 is defined as $\mathsf{G0} = \mathsf{AeadCtr}(\mathsf{PrfCtr}^0(\mathsf{MMac1}^0))$: by inlining the real versions of $\mathsf{EvalPad}$, Mac and Verify from $\mathsf{PrfCtr}^0(\mathsf{MMac1}^0)$, one can easily see that this game is identical to $\mathsf{Aead}^0(\mathsf{AEAD})$, hence we have that $\Pr[\mathcal{A}^{\mathsf{Aead}^0} = 1] = \Pr[\mathcal{A}^{\mathsf{G0}} = 1]$.

Game 1 is defined as $\mathsf{G1} = \mathsf{AeadCtr}(\mathsf{PrfCtr}^1(\mathsf{MMac1}^1))$. Game 0 and Game 1 are the same game except for the bit b of PrfCtr, so a reduction $\mathcal{A}'$ can simply simulate the game and returns the output of $\mathcal{A}$ as its guess. Therefore:

$$
|Pr[\mathcal{A}^{\mathsf{G1}}[q_e, q_d, \ell_p, \ell_a] = 1] - \Pr[\mathcal{A}^{\mathsf{G0}}[q_e, q_d, \ell_p, \ell_a] = 1]|
$$
$$
= \epsilon_{\mathsf{PrfCtr(MMac1)}}(\mathcal{A}'[q_e\left(\left\lceil \frac{\ell_p}{\ell_b} \right\rceil\right), q_e + q_d, \ell_p + \ell_a + 46, q_d])
$$

The implementation inlined into Game 1 has random tags and one-time-pads and the MAC is perfectly unforgeable. We could now conclude that in this implementation, not only the pads, but the full ciphertexts are random. This, however, requires us to reason about the various encoding steps (splitting the plaintext and splicing the ciphertext taking into account the truncated block), to check that we never apply the same pad twice, and to show that decryption is not revealing information about the plaintext blocks. We implement this information-theoretic step of the proof in $\mathrm{F}^\star$, where we can get assurance that we capture the correct implementation details in the real implementation (by running it on known test vectors).

Game 2 is defined as $\mathsf{G2} = \mathsf{AeadCtr}(\mathsf{PrfCtr}^{1,0}(\mathsf{MMac1}^1))$. Game 1 and Game 2 are perfectly equivalent.

Proof. $\mathsf{EvalEnx}'$ calls $\mathsf{EvalEnx}$ after checking $T[j\|n] = \bot$. We will show that this restriction is already enforced by $\mathsf{AeadCtr}$.

Similarly for every call to $\mathsf{EvalDec}$ we will show that the restriction $T[j\|n] = (p, c)$ is already enforced by $\mathsf{AeadCtr}$ when the bit b of PrfCtr is set. In this case decryption using $\oplus$ in $\mathsf{G1}$ already leads to the same result as the lookup in T of $\mathsf{G2}$.

We prove both restrictions on $\mathsf{AeadCtr}$ using an $\mathrm{F}^\star$ invariant that relates the content of the table C of $\mathsf{AeadCtr}$ to the *log* of $\mathsf{MMac1}$:

$$
\forall n.(C[n] = \bot \implies \forall i \neq 0.T[j_0 + i\|n] = \bot \wedge log[n] = \bot) \wedge
$$
$$
(C[n] \neq \bot \implies \mathbf{let}\ (a, p, c\|t) = C[n]\ \mathbf{in}
$$
$$
p = \mathsf{fst}(T[j_0 + 1\|n]) \| \dots \| \mathsf{fst}(T[j_0 + \lceil |p|/\ell_b\rceil\|n]) \wedge
$$
$$
c = \mathsf{snd}(T[j_0 + 1\|n]) \| \dots \| \mathsf{snd}(T[j_0 + \lceil |p|/\ell_b\rceil\|n]) \wedge
$$
$$
log[n] = (\mathsf{encode}(c, a), t))
$$

For the first restriction, the $\mathsf{AeadCtr}$ game guarantees that $C[n] = \bot$ upon calling $\mathsf{Encrypt}$, consequently, $T[j\|n] = \bot$ upon calling $\mathsf{EvalEnx}$.

For the second restriction, the $\mathsf{MMac1}$ game guarantees that Verify only succeeds, if $\mathsf{encode}([c_{j\in[1,r]}], a)$ was recorded in *log* by a call to $\mathsf{Encrypt}$. We can thus conclude by the injectivity of encode that the restriction holds. $\qquad\square$

Game 3 is defined as $\mathsf{G3} = \mathsf{AeadCtr}(\mathsf{PrfCtr}^{1,1}(\mathsf{MMac1}^1))$. Game 2 and Game 3 are perfectly equivalent by Lemma 6. Together we have $Pr[\mathcal{A}^{\mathsf{G3}} = 1] - Pr[\mathcal{A}^{\mathsf{G1}} = 1] = 0$.

Note that the condition $T[j\|n] = (p, c)$ in $\mathsf{EvalDex}'(j\|n, c)$ restricts decryption to those cases where decryption using $\oplus$ when $b' = 0$ would already lead to the same result as the table lookup.

From the invariant we also infer that $\Pr[\mathcal{A}^{\mathsf{G3}} = 1] = \Pr[\mathcal{A}^{\mathsf{Aead}^1} = 1]$ and that the size of the PRF table T is $\sum_{i=1}^{|C|} C[n_i].|c| = |T|$. Ciphertexts thus consists of random PRF blocks that are never repeated. Consequently, they are themselves indistinguishable from random. Summarizing the game steps

$$
\begin{aligned}
\epsilon_{\mathsf{Aead}}(\mathcal{A}[q_e, q_d, \ell_p, \ell_a]) =&\ \Pr[\mathcal{A}^{\mathsf{Aead}^1} = 1] - \Pr[\mathcal{A}^{\mathsf{Aead}^0} = 1] \\
\leq&\ |\Pr[\mathcal{A}^{\mathsf{G3}} = 1] - \Pr[\mathcal{A}^{\mathsf{G1}} = 1]| + \\
&\ |\Pr[\mathcal{A}^{\mathsf{G1}} = 1] - \Pr[\mathcal{A}^{\mathsf{G0}} = 1]| \\
=&\ \epsilon_{\mathsf{PrfCtr(MMac1)}}(\mathcal{A}'[q_e \left\lceil \frac{\ell_p}{\ell_b} \right\rceil, q_e + q_d, \ell_p + \ell_a + 46, q_d]) \\
\leq&\ \epsilon_{\mathsf{Prf}}(\mathcal{B}[q_b]) + \\
&\ \epsilon_{\mathsf{MMac1}}(\mathcal{C}[\ell_p + \ell_a + 46, q_d, q_e + q_d])
\end{aligned}
$$

With $q_b = q_e \left(\left\lceil \frac{\ell_p}{\ell_b} \right\rceil \right) + q_e + q_d + j_0 = j_0 + q_e \left(1 + \left\lceil \frac{\ell_p}{\ell_b} \right\rceil \right) + q_d$ $\square$