

A preliminary version of this paper appeared in the *24th European Symposium on Research in Computer Security, ESORICS 2019*. This is the full version.

# Secure Communication Channel Establishment: TLS 1.3 (over TCP Fast Open) vs. QUIC

Shan Chen<sup>1</sup>, Samuel Jero<sup>2</sup>, Matthew Jagielski<sup>3</sup>, Alexandra Boldyreva<sup>1</sup>, and Cristina Nita-Rotaru<sup>3</sup>

<sup>1</sup>Georgia Institute of Technology  
dragoncs16@gmail.com sasha@gatech.edu

<sup>2</sup>Purdue University  
sjero@sjero.net

<sup>3</sup>Northeastern University  
jagielski.m@husky.neu.edu c.nitarotaru@neu.edu

November 13, 2020

## Abstract

Secure channel establishment protocols such as Transport Layer Security (TLS) are some of the most important cryptographic protocols, enabling the encryption of Internet traffic. Reducing latency (the number of interactions between parties before encrypted data can be transmitted) in such protocols has become an important design goal to improve user experience. The most important protocols addressing this goal are TLS 1.3, the latest TLS version standardized in 2018 to replace the widely deployed TLS 1.2, and Quick UDP Internet Connections (QUIC), a secure transport protocol from Google that is implemented in the Chrome browser. There have been a number of formal security analyses for TLS 1.3 and QUIC, but their security, when layered with their underlying transport protocols, cannot be easily compared.

Our work is the first to thoroughly compare the security and availability properties of these protocols. Towards this goal, we develop novel security models that permit “layered” security analysis. In addition to the standard goals of server authentication and data confidentiality and integrity, we consider the goals of IP spoofing prevention, key exchange packet integrity, secure channel header integrity, and reset authentication, which capture a range of practical threats not usually taken into account by existing security models that focus mainly on the cryptographic cores of the protocols. Equipped with our new models we provide a detailed comparison of three low-latency layered protocols: TLS 1.3 over TCP Fast Open (TFO), QUIC over UDP, and QUIC[TLS] (a new design for QUIC that uses TLS 1.3 key exchange) over UDP. In particular, we show that TFO’s cookie mechanism does probably achieve the security goal of IP spoofing prevention. Additionally, we find several new availability attacks that manipulate the early key exchange packets without being detected by the communicating parties. By including packet-level attacks in our analysis, our results shed light on how the reliability, flow control, and congestion control of the above layered protocols compare, in adversarial settings. We hope that our models will help protocol designers in their future protocol analyses and that our results will help practitioners better understand the advantages and limitations of secure channel establishment protocols.

**Keywords:** applied cryptography, provable security, TLS, QUIC, secure channel, availability, network protocols

# 1 Introduction

**Motivation.** Nowadays, more than half of all Internet traffic is encrypted according to a 2017 EFF report [31], with Google reporting that 95% of its traffic is encrypted as of October 2020 [1]. This trend has also been facilitated by efforts like the free digital certificate issuer Let’s Encrypt servicing 87 million active (unexpired) certificates and 150 million unique domains at the end of 2018 [2].

This widespread Internet traffic encryption is enabled by protocols that allow two parties (where one or both parties have a public key certificate) to establish a secure communication channel over the insecure Internet. Typically, the parties first authenticate all parties holding a public key certificate and agree on a session key — the key exchange phase. Then, this session key is used to encrypt the communication during the session — the secure channel phase. We will refer to such protocols as secure channel establishment protocols.

The main secure channel establishment protocol in use today is Transport Layer Security (TLS). The session key establishment of the widely deployed standard TLS 1.2 [61] requires 2 round-trip times (RTTs) of end-to-end communication for a full connection and 1 RTT for resumption (which saves 1 RTT), before sending encrypted data. TLS 1.2 typically runs on top of TCP [56] for reliable transport, which adds another 1 RTT of establishing a TCP connection before the TLS connection. Further, this TCP cost is paid every time the two parties communicate with each other, even if the connection is interrupted and then immediately resumed. Given that most encrypted traffic is web traffic, this cost represents a significant performance bottleneck, a nuisance to users, and financial loss to companies. For instance, back in 2006 Amazon found that every 100ms of latency cost them 1% in sales [48], while a typical RTT on a connection from New York to London is 70ms [70].

Not surprisingly, many efforts in recent years have focused on reducing latency in secure channel establishment protocols. The focus has been on reducing the number of interactions (or RTTs) during session establishment and resumption without sacrificing much security. The most important protocols addressing this goal are TLS 1.3 [60], Google’s Quick UDP Internet Connections (QUIC) [63, 21, 46], and the new design for QUIC [36] (which we refer to as QUIC[TLS] [69] to indicate that it borrows the key exchange from TLS 1.3).

With TLS 1.3, it is possible to achieve 0-RTT resumption, i.e., the client is able to send encrypted 0-RTT data on the first flight to the server with session resumption. This is achieved by utilizing a session ticket that was saved during a previous communication and multiple keys (which we call stage keys) that can be set within one session, of which some keys are set faster (with slightly less security) so that data can be encrypted earlier. However, for the most common case where TLS 1.3 runs over TCP, the client still has to wait 1 RTT before transmitting any encrypted data due to the aforementioned TCP connection. One optimization for TCP, called TCP Fast Open (TFO) [59, 18], extends TCP to allow for 0-RTT resumption connections, so that the client may begin data transmission immediately. The mechanism underlying this optimization is a cookie saved from previous communication, similar to the ticket used by TLS 1.3. As a result, TLS 1.3 over TFO offers 0-RTT resumption.

Like TLS 1.3, QUIC uses weaker initial keys, under which data can be encrypted earlier, and a token saved from previous communication between the parties. However, unlike TLS, QUIC operates over the very simple UDP [55] rather than TCP. Instead of relying on TCP for reliability, flow control, and congestion control, QUIC implements its own data transmission functionality, integrating connection establishment with key exchange. These features allow QUIC to have 1-RTT full connections and 0-RTT resumption connections.

In Table 1 we show the cost of establishing full and resumption connections for several layered protocol options achieving end-to-end security. These include TLS 1.2 over TCP, TLS 1.3 over TCP, TLS 1.3 over TFO, QUIC over UDP, and QUIC[TLS] over UDP. It is clear that the last three win in terms of the number of interactions. *But how does their security compare?*

**Related Work.** At first glance, the question is easy to answer. Recent works have done formal security analyses of TLS 1.3 [42, 12, 25, 19, 43, 26, 47, 28, 22, 11, 20, 15, 27] and Google’s QUIC [30, 49]. Most works confirm that (the cryptographic cores of) both protocols are provably secure under

Table 1: Latency comparison of layered protocols

Layered Protocol	Full Connection	Resumption Connection
TCP+TLS 1.2	3-RTT	2-RTT
TCP+TLS 1.3	2-RTT	1-RTT
<b>TFO+TLS 1.3</b>	2-RTT	<b>0-RTT</b>
<b>UDP+QUIC</b>	1-RTT	<b>0-RTT</b>
<b>UDP+QUIC[TLS]</b>	1-RTT	<b>0-RTT</b>

reasonable computational assumptions. Moreover, as shown in [49, 28], their 0-RTT data transmission designs cannot achieve the same strong security guaranteed by classical key exchange protocols with at least one RTT. In particular, for TLS 1.3 and QUIC, their 0-RTT keys do not provide forward secrecy and the 0-RTT data suffers from replay attacks.<sup>1</sup> Overall, it might seem that all three layered protocols mentioned above are equally secure.

However, a closer look reveals that the answer is not that simple. First, all aforementioned formal security analyses, except for [49] analyzing QUIC’s IP-spoofing protection (also known as address validation), did not consider packet-level availability attacks, i.e., those targeting availability properties like reliable delivery of messages (via packets). Therefore, it is not clear at the packet level what security can be achieved and what attacks can be prevented by these protocols. In other words, we have no formal understanding of what security can be obtained when layering protocols. Also, TFO uses some cryptographic primitives, such as a cookie, to prevent IP spoofing, but no formal analysis has been done. Furthermore, the security of QUIC[TLS] has not been formally analyzed (although some security aspects can be reduced to those of Google’s QUIC and TLS 1.3).

**Our Contributions.** To compare security, we first need to define a general protocol syntax for secure channel establishment and fix a security model for it. Since the only provable security analysis that studies security related to packet-level availability attacks is [49], we take their *quick connections (QC)* protocol syntax and the associated *quick authenticated and confidential channel establishment (QACCE)* security model as our starting point.

**PROTOCOL DEFINITION.** To accommodate protocol syntaxes of TLS 1.3 and QUIC[TLS], we extend the QC protocol to a more general *multi-stage authenticated and confidential channel establishment (msACCE)* protocol, which allows more keys to be established during each session. Note that QC is a two-stage extension (i.e., two keys established within one session) to the ACCE protocol syntax proposed by Jager *et al.* [37], while our msACCE further extends QC to capture multiple keys.

**SECURITY DEFINITIONS.** We extend QC’s associated QACCE security model [49] to two msACCE security models — msACCE-std and msACCE-pauth — that are general enough for all layered secure channel establishment protocols listed in Table 1. The former is fairly standard and is for core cryptographic security notions such as Server Authentication and Channel Security. The latter security model is novel and is for packet-level security beyond those captured by the standard notions.

For the second model that deals with packet-level availability attacks, we first follow QACCE [49] to consider IP-spoofing security and further extend it to our stronger IP-Spoofing Prevention notion that additionally captures IP-spoofing attacks in the full connections. Then, we design several novel notions for packet-level authentication as follows.

First, we define Header Integrity to capture the integrity of the whole unencrypted packet header. (Note that previous models like QACCE only cover the header integrity implied by the authenticity security of the underlying authenticated encryption scheme.) To enable fine-grained security analyses and comparisons, we split the above notion into two related ones, Key Exchange (KE) Header Integrity and Secure Channel (SC) Header Integrity, which capture header integrity during the key exchange

<sup>1</sup>Note that it is possible to construct key exchange protocols to provide both forward secrecy for 0-RTT keys and replay resistance for 0-RTT data, but the existing constructions [34, 24, 5] require either large server-side storage or computationally expensive public-key operations.

phase and secure channel phase respectively. Furthermore, we define the notion of KE Payload Integrity to cover availability attacks that modify the payloads of packets sent during key exchange. We note that unlike the availability attacks shown in [49], successful attacks under our new notions do not affect the client’s session key establishment and therefore are harder or impossible to detect by the client. Finally, we formalize the new goal of Reset Authentication to deal with attacks forging a reset packet to abruptly terminate (without completing) an honest party’s session, i.e., it ensures that a session should be reset only by the intended party who lost its state (e.g., due to server reboots).<sup>2</sup> SECURITY ANALYSES. Equipped with our new models, we study the security and availability properties provided by TFO+TLS 1.3, UDP+QUIC, and UDP+QUIC[TLS], with results summarized below.

Table 2: Security comparison

	TLS 1.3 +TFO	QUIC +UDP	QUIC[TLS] +UDP
0-RTT Key Forward Secrecy [49, 28]	<b>X</b>	<b>X</b>	<b>X</b>
0-RTT Data Anti-Replay [49, 28]	<b>X</b>	<b>X</b>	<b>X</b>
Server Authentication	✓	✓	✓
Channel Security	✓	✓	✓
IP-Spoofing Prevention	✓	✓	✓
KE Header Integrity	<b>X</b>	<b>X</b>	<b>X</b>
KE Payload Integrity	✓	<b>X</b>	<b>X</b>
SC Header Integrity	<b>X</b>	✓	✓
Reset Authentication	<b>X</b>	<b>X</b>	✓

We first confirm that all protocols provably satisfy the standard security notions of Server Authentication and Channel Security given that their building blocks are secure. The results mostly follow from prior works and we just have to argue that they still hold for the extended model. Similarly, prior results showed that QUIC is secure against IP spoofing, which can be extended to our stronger notion for both UDP+QUIC and UDP+QUIC[TLS]. As for TFO+TLS 1.3, its IP-Spoofing Prevention security relies on TCP sequence number randomization and TFO’s cookie mechanism (but no prior former analysis confirmed its security). We prove that TFO+TLS 1.3 does satisfy this security assuming that the underlying TFO cookie generation function is a pseudorandom function.

Regarding SC Header Integrity, we show that while UDP+QUIC and UDP+QUIC[TLS] are secure, TFO+TLS 1.3, on the other hand, is insecure because it allows header-only packets to be sent in the secure channel phases and does not authenticate the TCP headers of encrypted packets. This theoretical result captures practical availability attacks that the networking community has been slowly uncovering via manual investigation over the last 30 years [64, 40, 4, 33, 45, 44, 38, 58, 17, 32, 57, 52, 71, 39], such as TCP flow control manipulation, TCP acknowledgment injection, etc.

We next show that none of the three protocols satisfies KE Header Integrity. For TFO+TLS 1.3 this result leads to a TFO cookie removal attack that we discover, which allows the attacker to undermine the whole benefit of TFO without being detected by the communicating parties. Then, we show that UDP+QUIC is not secure in the sense of KE Payload Integrity. This leads to a new availability attack that we call ServerReject Triggering. Note that unlike the QUIC attacks (e.g., server config replay attack, connection ID manipulation attack, etc.) discovered in [49], ServerReject Triggering is harder to detect. Similar attacks also work for UDP+QUIC[TLS]. We show that TFO+TLS 1.3, on the other hand, achieves KE Payload Integrity.

<sup>2</sup>Our Reset Authentication notion for stateless reset security may be confused with *secure termination* proposed by [13]. However, secure termination targets a very different goal: it ensures that upon receipt of a termination message the receiver has received all messages that were sent and will ever be sent by the sender, i.e., it ensures a session is completed as expected. As modeled in [13], this property can be reduced to the security of the underlying channel protocol, which in our case is guaranteed by Channel Security (and hence not captured by our second model). Note that all termination message(s) are encrypted and authenticated in the established secure channel, while a reset packet used for a stateless reset typically cannot be protected because the issuing party has already lost all session states including the keying material.

We further show that neither TFO+TLS 1.3 nor UDP+QUIC provides Reset Authentication, justifying the TCP Reset attack [71] relevant for TFO+TLS 1.3 and the PublicReset attack for UDP+QUIC. On the other hand, we prove that UDP+QUIC[TLS] does guarantee Reset Authentication with its new stateless reset feature, by relying on the unpredictability of the reset tokens.

We note that the new UDP+QUIC[TLS] protocol achieves the strongest overall security of the three designs. While formally it does not provide KE Payload Integrity, the related attacks can also happen in TFO+TLS 1.3 in a similar way; the latter satisfies KE Payload Integrity mainly because its availability functionalities are all carried in its protocol headers rather than payloads. More importantly, UDP+QUIC[TLS] is the only protocol that guarantees Reset Authentication.

Even though QUIC[TLS] may not be able to sustain the competition in the long run despite stronger security, we hope our models will help protocol designers and practitioners better understand the important security aspects of existing and new secure channel establishment protocols.

**Paper Organization.** The rest is organized as follows. We provide an overview of relevant design information for TFO, TLS 1.3, QUIC, and QUIC[TLS] in Section 2. Section 3 and Appendix A specify our notations and preliminaries. Section 4 formally defines our msACCE protocol and its security. Section 5 provides the details of our security analyses and Section 6 concludes our paper.

## 2 Background

Network protocols are designed and implemented following a layered network stack model where each layer has its own functionality, defines an interface for use by higher layers, and relies only on the properties of lower layers. In this work, we are concerned with three layers: network, represented by the IP protocol; transport, represented by UDP, TCP, or TFO; and application, represented by TLS, QUIC, or QUIC[TLS].

### 2.1 TLS 1.3 over TFO

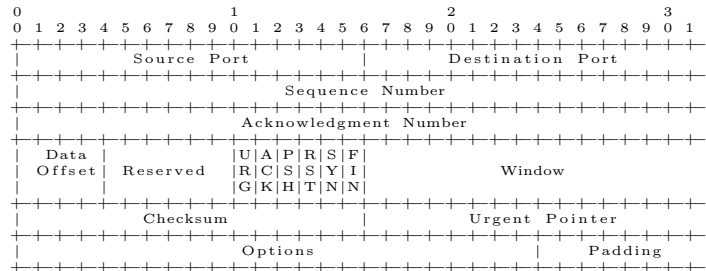


Figure 1: TCP header [56]

**TCP Fast Open.** TFO is an optimization for the TCP protocol. TCP itself provides the following services to an application (or higher protocol): (1) reliability, (2) ordered delivery, (3) flow control, and (4) congestion control. It is connection-oriented and consists of three phases: connection establishment, data transfer, and connection tear-down. TCP relies on control information from its header to implement this functionality. For example, as shown in Figure 1, control bits specify what type of packets are sent over the network, which determines whether the packets are establishing a new connection, sending data, acknowledging data, or tearing down the connection.

The disadvantage of layering protocols is that higher level protocols have no control over the internal mechanics of lower level protocols and can interact with them only through defined interfaces. A protocol using standard TCP for transferring data needs to wait for connection establishment at

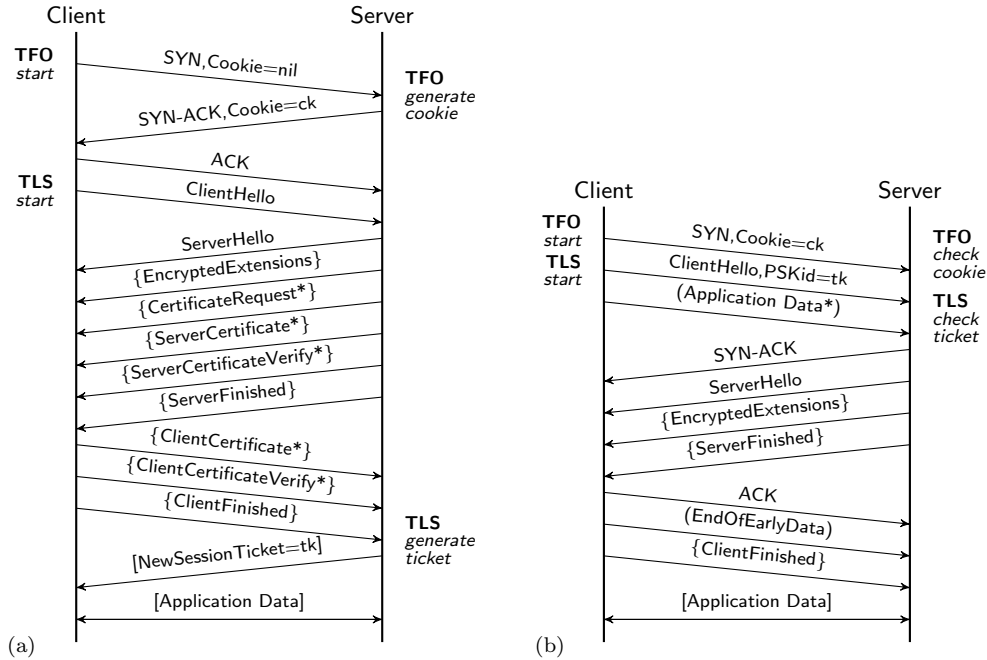


Figure 2: TFO+TLS 1.3 (EC)DHE 2-RTT full handshake (a) and PSK-(EC)DHE 0-RTT resumption handshake (b). \* indicates optional messages. () indicates messages protected using the 0-RTT keys derived from a pre-shared key (PSK). {} and [] indicate messages protected with initial and final keys.

the TCP layer to complete before it receives notification of a new connection and can begin its own processing and data transfer.

The TFO optimization introduces a simple modification to the TCP connection establishment handshake to reduce the 1-RTT connection establishment latency of TCP and allow for 0-RTT handshakes, so that data transmission may begin immediately. TFO fulfills the same design goals mentioned for TCP above, assuming the connection is established correctly.

The mechanism through which 0-RTT is achieved is a cookie that is obtained by the client the first time it communicates with a server and cached for later uses. This cookie is intended to prevent replay attacks while avoiding the need for servers to keep expensive state. It is generated by the server, authenticates client IP address, and has a limited lifetime. Generation and verification have low overhead.

Cookies are sent in the TFO option field in SYN packets. The first two message exchanges in Figure 2 (a), show how a cookie is obtained. The client requests a cookie by using the TFO option in the SYN with the cookie field set to empty (nil), indicating that it would like to use TFO. The server generates an appropriate cookie and places it in the TFO option field of the SYN-ACK. The client caches this cookie for subsequent connections to this server. If a cookie was not provided, the client instead caches the negative response, indicating that TFO connections should not be tried to this server, for some time.

In subsequent connections to this server (first message in Figure 2 (b)), the client places its cached TFO cookie in the TFO option in the SYN packet. The client is also allowed to send 0-RTT data in the remainder of the SYN packet. This might be an HTTP GET request or a TLS ClientHello message. When the server receives the SYN, it will validate the cookie. If the cookie is valid, it responds with a SYN-ACK acknowledging the 0-RTT data and a response to the 0-RTT data. If the cookie is invalid (expired or otherwise), a full handshake is required and any initial data ignored.

**TLS 1.3.** TLS 1.3 provides confidentiality, authentication, and integrity of communication over a

secure channel between a client and a server. This is accomplished in two phases – the handshake protocol and the record protocol. The handshake sets up appropriate parameters for the record protocol to achieve these three goals. These include parameters like the cipher suite to use and the shared secret key. Unfortunately, the handshake in TLS 1.2 takes 2-RTTs to complete. Additionally, the naive layering of TLS 1.2 over TCP, as traditionally used for HTTPS, would require a full 3-RTTs before the HTTP request could be sent. Fortunately, the recently standardized TLS 1.3 [60] provides many improvements over TLS 1.2. Most relevant for our purposes, it enables 0-RTT handshakes at the TLS level.

In a TLS 1.3 full connection (see Figure 2 (a) starting from the 4th message), the client begins by sending a `ClientHello` message containing a list of ciphersuites the client is willing to use with key shares for each and optional extensions. The server responds with a `ServerHello` message containing the ciphersuite to use and its key share. At this point, an initial encryption key is derived and all future messages are encrypted. The server also sends an `EncryptedExtensions` message containing any extension data, a `CertificateRequest` message if doing client authentication, a `ServerCertificate` message containing the server’s certificate, a `ServerCertificateVerify` message containing a signature over the handshake with the private key corresponding to the server’s certificate, and a `ServerFinished` message containing an HMAC of all messages in the handshake. The client receives these messages, verifies their contents, and responds with `ClientCertificate` and `ClientCertificateVerify` messages if doing client authentication before finishing with a `ClientFinished` message containing an HMAC of all messages in the handshake. At this point, a final encryption key is derived and used for encrypting all future messages. If the server supports 0-RTT connections, one final handshake message, the `NewSessionTicket` message, will be sent by the server to provide the client with an opaque session ticket to be used in a resumption session.

In later TLS 1.3 resumption connections to this server, the client uses the session ticket established in the prior full connection to do a 0-RTT connection. In this case, the client sends a `ClientHello` message indicating a pre-shared-key ciphersuite, a ciphersuite to use for the final key, and the cached session ticket. The client can then derive an encryption key and begin sending 0-RTT data. The server will verify the session ticket, use it to establish the same encryption key, and send a `ServerHello` message containing the ciphersuite to use and its final key share. At this point, an initial encryption key is derived and all future messages are encrypted. The server also sends an `EncryptedExtensions` message containing any extension data and a `ServerFinished` message containing an HMAC of all messages in the handshake. The client receives these messages, verifies their contents, and responds with an `EndOfEarlyData` message and a `ClientFinished` message containing an HMAC of all messages in the handshake. At this point, a final encryption key is derived and used for encrypting all future messages.

**TLS 1.3 over TFO.** TLS assumes that lower layers provide reliable, in-order delivery of TLS messages. As a result, TLS is usually layered on top of TCP, which provides these properties. This usually results in a delay for the TCP handshake followed by a delay for the TLS handshake. This is obviously undesirable. However, the combination of TLS 1.3 and TFO enables true 0-RTT connections.

In a full connection to a TFO+TLS 1.3 server, the client requests a TFO cookie in the TCP SYN and then does a full TLS 1.3 handshake once the TCP connection completes. This takes 3-RTTs (see Figure 2 (a)), but provides a cached TFO cookie and cached TLS session ticket.

In subsequent resumption connections to this server, the client can use the TFO cookie to establish a 0-RTT TCP resumption connection and include the TLS 1.3 `ClientHello` message in the SYN packet. This `ClientHello` message can use the cached TLS session ticket to start a 0-RTT resumption handshake. Thus, the TCP and TLS 1.3 connections are established at the same time, as shown in Figure 2 (b).

## 2.2 QUIC over UDP

**UDP.** UDP [55] is an extremely simple transport protocol providing unreliable datagram delivery, the




Figure 3: UDP header [55]

ability to multiplex data between multiple applications, and an optional checksum. A UDP sender simply wraps the message to be sent with a UDP header (see Figure 3) and the receiver unwraps the message and delivers it to the application, after possibly verifying the checksum. No other processing is performed.

UDP has been typically used for applications where low latency is crucial, like video gaming and real-time streaming video. As a result, it can traverse NAT devices and firewalls that often block unknown or rare protocols.

**QUIC.** QUIC is a transport protocol developed by Google and implemented by Chrome and Google servers since 2013 [63, 21, 46]. It now provides service for the majority of requests by Chrome to Google properties [68]. QUIC’s goal was to provide secure communication comparable with TLS while achieving reduced connection setup latency compared to traditional TCP+TLS 1.2. To do so, it provides the following services to applications: (1) reliability, (2) in order delivery, (3) flow control, (4) congestion control, (5) data confidentiality, and (6) data authenticity. For repeated connections to the same server it also provides (7) 0-RTT connections, enabling useful data to be sent in the first round trip. In short, QUIC provides a very similar set of services to TFO+TLS 1.3.

Instead of modifying TCP to enable 0-RTT connection establishment, QUIC replaces TCP entirely, using UDP to provide application multiplexing and enabling it to traverse the widest possible swath of the Internet. QUIC then provides all other guarantees itself.

QUIC packets contain a public header and a set of frames that are encrypted and authenticated after initial connection setup. The header contains a set of public flags, a unique 64-bit connection identifier, a variable-length packet number, and optional version and nonce fields. All other protocol information is carried in control and stream (data) frames that are encrypted and authenticated.

To provide 0-RTT, QUIC caches important information about the server that will enable the client to determine the encryption key to be used for each new connection. As shown in Figure 4 (a), the first time a client contacts a given server it has no cached information, so it sends an empty (Inchoate) `ClientHello` message. The server responds with a `ServerReject` message containing the server’s certificate and three pieces of information for the client to cache. The first of these is an object called an `scfg`, or server config. The `scfg` contains a variety of information about the server, including a Diffie-Hellman share from the server, supported encryption and signing algorithms, and flow control parameters. This `scfg` has a defined lifetime and is signed by the server’s private key to enable authentication using the server’s certificate. Along with the `scfg`, the server sends the client a source-address token or `stk`. The `stk` is used to prevent IP spoofing. It contains an encrypted version of the client’s IP address and a timestamp.

With this cached information, a client can establish an encrypted connection with the server. It first ensures that the `scfg` is correctly signed by the server’s certificate and is valid, then sends a `ClientHello` indicating the `scfg` it is using, the `stk` value it has cached, a Diffie-Hellman share for the client, and a client nonce. After sending the `ClientHello`, the client can create an initial encryption key and send additional encrypted `Application Data` packets. In fact, to take advantage of the 0-RTT connection establishment it must do so. When the server receives the `ClientHello` message, it validates the `stk` and client nonce parameters and creates the same encryption key using the server share from the `scfg` and the client’s share from the `ClientHello` message.

At this point, both client and server have established the connection and setup encryption keys and all further communication between the parties is encrypted. However, the connection is not forward



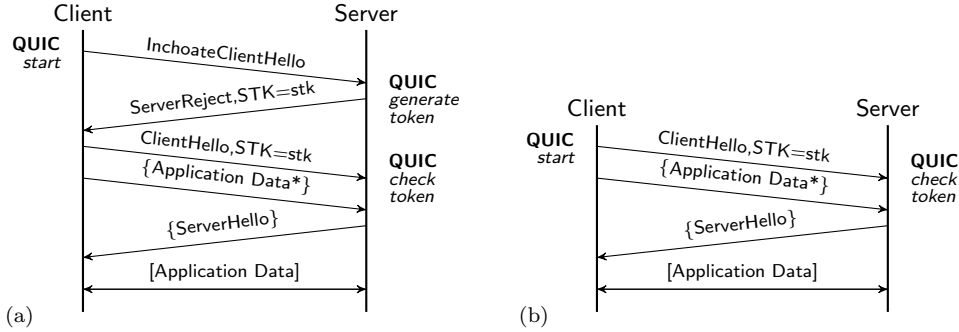


Figure 4: UDP+QUIC 1-RTT full handshake (a) and 0-RTT resumption handshake (b). \* indicates optional messages. {} and [] indicate messages protected with initial and final keys.

secret yet, meaning that compromising the server would compromise all previous communication because the server’s Diffie-Hellman share is the same for all connections using the same `scfg`. To provide forward secrecy for all data after the first RTT, the server sends a `ServerHello` message after receiving `ClientHello` which contains the client’s newly generated Diffie-Hellman share. Once the client receives `ServerHello`, client and server derive and begin using the new forward-secret encryption key.

If a client has previously connected to a server, it can instead initiate a resumption connection to the same server. This consists of only the last two steps of a full connection, sending the `ClientHello` and `ServerHello` messages as shown in Figure 4 (b).

### 2.3 QUIC with TLS 1.3 Key Exchange

A new version of QUIC [36], which also supports 0-RTT, describes several improvements of the previous design. The most important change is replacing QUIC’s key exchange with the one from TLS 1.3, as specified in the latest Internet draft [69]. We provide more details (e.g., the new stateless reset feature) of this new QUIC (denoted by QUIC[TLS]) in Section 5.

## 3 Preliminaries

**Notations.** Let  $\{0, 1\}^*$  denote the set of all finite-length binary strings (including the empty string  $\varepsilon$ ) and  $\{0, 1\}^n$  denote the set of  $n$ -bit binary strings.  $[n]$  denotes the set of integers  $\{1, 2, \dots, n\}$ . For a finite set  $\mathcal{R}$ , let  $|\mathcal{R}|$  denote its size and  $r \stackrel{\$}{\leftarrow} \mathcal{R}$  denote sampling  $r$  uniformly at random from  $\mathcal{R}$ . For a binary string  $s$ , let  $|s|$  denote its length in bits.  $y \leftarrow F(x)$  (resp.  $y \stackrel{\$}{\leftarrow} F(x)$ ) denotes  $y$  being the output of the deterministic (resp. probabilistic) function  $F$  with input  $x$ . Let  $x \leftarrow a$  denote assigning value  $a$  to variable  $x$ . We use the wildcard  $\cdot$  to indicate any valid input of a function.

**Public Key Infrastructure.** For simplicity, we do not consider certificates or certificate checks but assume any public key associated with a party is supported by a *public key infrastructure (PKI)* and hence certified and bound to the party’s identity. That is, we omit PKI details but simply assume that the binding of public keys to party identities is publicly known by default.

**PRF and AEAD.** In Appendix A we recall the security definitions of a *pseudorandom function (PRF)*  $F$  and a stateful *authenticated encryption with associated data (AEAD)* scheme `sAEAD` with authentication level  $al \in [4]$  (i.e., protecting against the first  $al$  types of the following attacks: forgeries, replays, reordering, or dropping) [41, 14]. Accordingly, there we provide the definitions for the corresponding advantages:  $\text{Adv}_F^{\text{prf}}(A)$ ,  $\text{Adv}_{\text{sAEAD}}^{\text{aead-al}}(A)$ . We also refer to [62] for the syntax and security definitions of a nonce-based AEAD scheme.

## 4 Multi-Stage Authenticated and Confidential Channel Establishment

In this section, we define the syntax and two security models for *multi-stage authenticated and confidential channel establishment (msACCE)* protocols.

### 4.1 Protocol Syntax

Our msACCE protocol is an extension to the *quick connection (QC)* protocol proposed by Lychev *et al.* [49] and the *multi-stage key exchange (MSKE)* protocol proposed by Fischlin and Günther [30] (and further developed by [25, 26, 47, 28]). Even though the authors of [49] claimed their QC protocol syntax to be general, TLS 1.3 does not fit it well because TLS 1.3 has two initial keys and one final key in 0-RTT resumption while QC captures only one initial key. On the other hand, the MSKE protocol and its extensions focus only on the key exchange phases.

Our msACCE protocol syntax inherits many parts of the QC protocol syntax but extends it to a multi-stage structure and additionally covers session resumptions (formally, unlike QC), session resets, and header-only packets exchanged in secure channel phases. msACCE also employs a more general stateful authenticated encryption syntax (rather than a nonce-based one) to capture the encryption schemes of both TLS 1.3 and QUIC.

The detailed protocol syntax is defined below.

A msACCE protocol is an interactive protocol between a client and a server. The protocol consists of one or more stages and each stage consists of two phases (that are formally defined later): key exchange and secure channel. The client and server establish keys in key exchange phases and exchange messages encrypted and decrypted with these keys in secure channel phases. Messages are exchanged via *packets*. A packet consists of source and destination IP addresses<sup>3</sup>  $IP_s, IP_d \in \{0, 1\}^{32} \cup \{0, 1\}^{128}$ , a header, and a payload. Each party  $P$  has a unique IP address  $IP_P$ . Both parties can keep static state (shared among several sessions) and volatile state (used only within one session). All states are initialized to the empty string  $\varepsilon$ .

The protocol is associated with the security parameter  $\lambda \in \mathbb{N}_+$ , a key generation algorithm  $\text{Kg}$  that takes as input  $1^\lambda$  and outputs a public and secret key pair, a header space<sup>4</sup> (for transport and application layers)  $\mathcal{H} \subseteq \{0, 1\}^*$ , a payload space  $\mathcal{PD} \subseteq \{0, 1\}^*$ , header and payload spaces  $\mathcal{H}_{\text{rst}} \subseteq \mathcal{H}$ ,  $\mathcal{PD}_{\text{rst}} \subseteq \mathcal{PD}$  for reset packets (described later), a resumption state space  $\mathcal{RS} \subseteq \{0, 1\}^*$ , a stateful AEAD scheme  $\text{sAEAD} = (\text{sG}, \text{sl}, \text{sE}, \text{sD})$  (with a message space  $\mathcal{M} \subseteq \{0, 1\}^*$ , an associated data space  $\mathcal{AD} \subseteq \{0, 1\}^*$ , and a state space  $\mathcal{ST} \subseteq \{0, 1\}^*$ ), *disjoint*<sup>5</sup> message spaces  $\mathcal{M}_{\text{KE}}, \mathcal{M}_{\text{SC}}, \mathcal{M}_{\text{PRST}} \subseteq \mathcal{M}$  with  $\mathcal{M}_{\text{KE}}, \mathcal{M}_{\text{SC}}$  for messages encrypted during key exchange and secure channel phases respectively and  $\mathcal{M}_{\text{PRST}}$  for pre-reset messages (described later) encrypted in a secure channel phase. During the protocol’s execution, the server keeps a (static) local time state  $\text{loct}$  that represents its current time period, which is initialized to 0 and gets incremented (by one) periodically. The server may keep a (static) configuration state  $\text{scfg}$  that is updated whenever  $\text{loct}$  is incremented (e.g., for security reasons); if so, the protocol is further associated with a server configuration generation function  $\text{scfg\_gen}$ , which takes as input  $1^\lambda$ , the server secret key (generated by  $\text{Kg}$ ), and a timestamp (which is the content of  $\text{loct}$ ), then outputs a new server configuration state.<sup>6</sup>

<sup>3</sup>For the network-layer protocols, we only consider the Internet Protocol and its IP address header fields because our model mainly focuses on the application and transport layers and additionally only captures the IP-spoofing attack.

<sup>4</sup>Some protocol header fields (e.g., port numbers, checksums, etc.) can be excluded if they are not the focus of the security analysis.

<sup>5</sup>Disjointness is a reasonable assumption as practical protocols (such as those in Table 1) enforce different leading bits for different types of messages.

<sup>6</sup>In the QC protocol syntax [49]  $\text{scfg}$  is also updated at the beginning of each time period, but there the notion of time was modeled rather informally. We refer to [6, 65] for “timed” models proposed to analyze authentication and key exchange protocols whose security relies crucially on time synchronization, which is not the case for the protocols analyzed in this work (for which timestamps, if any, are essentially only used by servers to expire resumption tokens). In order to not further complicate our already complex model, we choose to simplify our time modeling by replacing

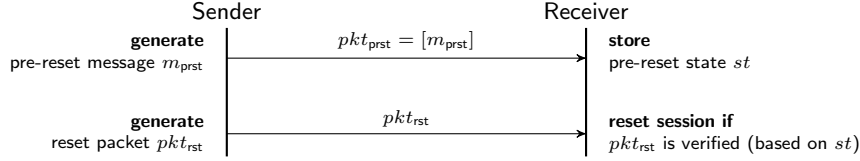


Figure 5: Stateless reset.  $\square$  indicates messages protected in a secure channel.

A *reset* packet enables the sender, who lost its volatile session state due to some error condition (e.g., server reboots, denial-of-service attacks, etc.) and hence lost its previous connection with the receiver, to reset (i.e., terminate without completing) their session. A *pre-reset* message (e.g., a reset token in UDP+QUIC[TLS]) may have been sent to the receiver in a secure channel phase<sup>7</sup> as a *pre-reset* packet before the sender lost its state, in order to let the receiver verify the sender’s reset packet, as shown in Figure 5. For simplicity, we assume each party generates *at most one* pre-reset message for each session. A *non-reset* packet is not a reset packet. A *header-only* packet has no payload.

We say a party *rejects* a packet if its processing the packet leads to an error (defined according to the protocol), and *accepts* it otherwise.

The protocol has two modes, *full* and *resumption*. Its corresponding executions are referred to as the full and resumption sessions. Each resumption session is associated with a *unique* previous full session and we say the resumption session *resumes* its associated full session.

In the beginning of a full or resumption session, each party takes as input a list of messages<sup>8</sup>  $\mathcal{M}^{\text{snd}} = (M_1, \dots, M_l), M_i \in \mathcal{M}_{\text{SC}}, l \in \mathbb{N}$  (where the total message length  $|\mathcal{M}^{\text{snd}}|$  is polynomial in  $\lambda$  and  $\mathcal{M}^{\text{snd}}$  can be empty) as well as the other party’s IP address. For a full session, the server takes as input its associated public and secret key pair (generated by  $\text{Kg}$ ) and the client takes the server’s public key as input<sup>9</sup>. For a resumption session, each party additionally takes as input its own (static) resumption state  $rs \in \mathcal{RS}$  (set in the associated full session). In either case, the client sends the first packet to start a session and each party may set its static state if it is not set yet. In particular, if  $\text{scfg\_gen}$  is defined and the server’s configuration state  $\text{scfg}$  has not been set, the server sets  $\text{loct}$  to 0 and runs  $\text{scfg\_gen}$  to set its  $\text{scfg}$ , which may be further updated during the protocol’s execution.

A  $D$ -stage (for an arbitrary  $D \in \mathbb{N}_+$ ) msACCE protocol consists of  $D$  successive stages and each stage, e.g., the  $d$ -th ( $d \in [D]$ ) stage, consists of one or two phases described as follows:

**1) Key Exchange.** At the end of this phase each party sets its  $d$ -th stage key  $k^d = (k_c^d, k_s^d)$ . At most one of  $k_c^d$  and  $k_s^d$  could be  $\perp$ , i.e., unused.<sup>10</sup> If this is the final stage in a full session, each party can send additional messages<sup>11</sup> in  $\mathcal{M}_{\text{KE}}$  encrypted with  $k^d$  and by the end of this phase each party sets its own resumption state.

**2) Secure Channel.** This phase is mandatory for the final stage but optional for other stages. In this phase, the parties can exchange messages from their input lists as well as pre-reset messages, encrypted and decrypted using the associated stateful AEAD scheme with  $k^d$ .<sup>12</sup> The client uses  $k_c^d$

timestamps with local counters. The advancement of (server-side) time periods is triggered by the NextTP oracle as described later in our security model.

<sup>7</sup>A pre-reset message can also be carried within an *encrypted* key exchange packet. We consider it encrypted as a separate secure channel packet to get a clean packet-authentication security model described later.

<sup>8</sup>Note that in reality the exchanged messages may depend on each other, but modeling that complicates our already complex protocol syntax and it has no real implications on our security analysis. Furthermore, for simplicity we consider transportation of *atomic* messages rather than a data *stream*, where the latter was modeled in [29] and later extended to capture multiplexing [54].

<sup>9</sup>Since the client knows the identifier of its intended communicating peer when the protocol starts, our protocol is analyzed in the *pre-specified peer model* [51].

<sup>10</sup>This captures the case where a 0-RTT key only consists of a client encryption key while the server encryption key does not exist.

<sup>11</sup>This captures the post-handshake key exchange messages that are used for session resumption, post-handshake authentication, key update, etc.

<sup>12</sup>Note that our work does not capture key updates inside the secure channel, which was modeled in [35].

to encrypt and the server uses it to decrypt, whereas the server uses  $k_s^d$  to encrypt and the client uses it to decrypt. They may also send reset or header-only packets. At the end of this phase, each party outputs a list of received messages (which may be empty).

Each message exchanged between the parties must belong to some unique phase at some unique stage. One stage’s second phase and the next stage’s first phase may overlap, and the two phases in the final stage may also overlap. We call the final stage key the *session* key and the other stage keys the *interim* keys.

**Correctness.** Consider a client and a server running a  $D$ -stage msACCE protocol in either mode without sending any reset packet. Each party’s input message list  $\mathcal{M}^{\text{snd}}$ , in which the messages are sent among  $D$  stages according to any possible partitioning  $\mathcal{M}^{\text{snd}} = \mathcal{M}_1^{\text{snd}}, \dots, \mathcal{M}_D^{\text{snd}}$ , is equal to the other party’s total output message list  $\mathcal{M}^{\text{rcv}} = \mathcal{M}_1^{\text{rcv}}, \dots, \mathcal{M}_D^{\text{rcv}}$ , in which the message order is preserved. Each party terminates its session upon receiving the other party’s reset packet.

**Relations to Prior ACCE-type Protocols.** With our more general protocol syntax, the ACCE [37] and QC [49] protocols can be classified into 1-stage and 2-stage msACCE protocols respectively.

## 4.2 Security Models

We propose two security models respectively for basic authenticated and confidential channel security and novel packet authentication. Our models do not consider the key exchange and secure channel phases independently, as was the case in the MSKE model (and its extensions) used by some previous QUIC and TLS 1.3 security analyses [30, 25, 26, 47, 28, 15, 27], because QUIC’s key exchange and secure channel phases are inherently inseparable and the TLS 1.3 full handshake does not fit into a composability framework, as discussed in [49, 26].

### 1) msACCE Standard Security Model:

In this msACCE standard (msACCE-std) security model, we consider the standard security goals for msACCE protocols: server authentication<sup>13</sup> and channel security (which captures data confidentiality and integrity).

Our msACCE-std model is based on the standard security portion of the QACCE model [49], but extends it to capture four authentication levels (following those of stateful authenticated encryption schemes) and any number of stages. Recall that QACCE considers only two stages within a session and hence cannot be used to analyze the TLS 1.3 0-RTT resumption. Inspired by the “timed” models proposed by [6, 65], our model introduces a new NextTP oracle (defined later) to formally capture the advancement of time periods on the server side, which was not modeled properly in QACCE. Finally, like QACCE and other previous models, we consider a very powerful adversary who can control communications between honest parties, can adaptively learn their stage keys, and can adaptively corrupt servers to learn their long-term keys and secret states.

The detailed security model is defined below.

**Protocol Entities.** The set of parties  $\mathcal{P}$  consists of two disjoint type of parties: clients  $\mathcal{C}$  and servers  $\mathcal{S}$ , i.e.,  $|\mathcal{P}| = |\mathcal{C}| + |\mathcal{S}|$ .

**Session Oracles.** To capture multiple sequential and parallel protocol executions, each party  $P \in \mathcal{P}$  is associated with a set of session oracles  $\pi_P^1, \pi_P^2, \dots$ , where  $\pi_P^i$  models  $P$  executing a protocol instance in session  $i \in \mathbb{N}_+$ .

**Session Identifiers.** As part of the security model, *session identifiers* are used to model entity authentication, session key confirmation, and handshake integrity. Compared to the general definition of matching conversations [10, 37] used for modeling entity authentication, a session identifier `sid` is defined according to the protocol specifications and security goals, often as a *subset* of the entire communication transcript. For instance, QUIC’s `sid` in QACCE [49] is defined as the second-round key exchange messages, i.e., `ClientHello` and `ServerHello`, while the first-round messages are excluded to

<sup>13</sup>Our msACCE-std model focuses on the most common unilateral server authentication, but can be extended to mutual authentication, e.g., as described in [42].

allow for valid but different source-address tokens or signatures. Similarly, TLS 1.2’s `sid` in ACCE [42] is defined as the first three key exchange messages, while the rest are excluded to allow for valid but different encrypted `Finished` messages. A msACCE protocol may have different session identifiers in full and resumption modes, but for simplicity we use the same notation `sid`.

**Peers and Partners.** We say a client oracle and a server oracle are each other’s *peer* if they share the same session identifier `sid`; we say they are each other’s *partner* if they share the same first-stage session identifier `sid1` (i.e., `sid` restricted to the first stage), which intuitively means that they set the first stage key with each other. Note that a client oracle may have more than one partner if `sid1` consists of only message(s) sent from the client oracle, which can be replayed to the same<sup>14</sup> server to establish multiple (identical) first-stage keys. Therefore, a session oracle’s partner may not be its final unique communication peer.

**Security Experiment.** In order to define security, we consider the following security experiment run between a challenger and an adversary  $\mathcal{A}$ . In the beginning of the experiment, the challenger initializes the states of all parties and then runs `Kg` for all servers to generate their public-secret key pairs. The static states of all parties are properly set and in particular `scfg_gen` (if defined) is executed for each server (where `loct`  $\leftarrow$  0) to set its `scfg` state. All other states are initialized to the empty string  $\varepsilon$ . The experiment is associated with an authentication level  $al \in [4]$  for channel security. Each oracle  $\pi_P^i$  is associated with a random bit  $b_P^i \xleftarrow{\$} \{0,1\}$ . Let  $N \in \mathbb{N}_+$  denote the maximum number of msACCE protocol instances for each party and  $D \in \mathbb{N}_+$  denote the maximum number of stages in each session. The adversary  $\mathcal{A}$  is given all public keys and IP addresses associated with all parties and then interacts with session oracles and parties via the following queries:

- **Connect**( $\pi_C^i, S$ ), for  $C \in \mathcal{C}, S \in \mathcal{S}, i \in [N]$ .

This query allows the adversary to ask a specified client oracle to start a full session with a specified server.  $\pi_C^i$  outputs the first packet it would send to  $S$  in a full session according to the protocol. This output packet is returned to  $\mathcal{A}$  instead of  $S$ . After this query, we say  $S$  is the *intended server* of  $\pi_C^i$ . The session oracle created in this query (i.e.,  $\pi_C^i$ ) must have never been used before.

- **Resume**( $\pi_C^i, S, i'$ ), for  $C \in \mathcal{C}, S \in \mathcal{S}, i, i' \in [N], i' < i$ .

This query allows the adversary to ask a specified client oracle to start a resumption session with a specified server to resume a specified full session between the two parties, if the associated previous client oracle has set its resumption state.  $\pi_C^i$  inputs the resumption state set by  $\pi_C^{i'}$  and outputs the first packet it would send to  $S$  in a resumption session according to the protocol, where the output packet is returned. This query is only allowed if  $\pi_C^{i'}$  has set its resumption state. The session oracle created in this query (i.e.,  $\pi_C^i$ ) must have never been used before.

- **Send**( $\pi_P^i, pkt$ ), for  $P \in \mathcal{P}, i \in [N], pkt \in \{0,1\}^*$ .

This query allows the adversary to send any packet to a specified session oracle and get its response in a key exchange phase. If  $\pi_P^i$  is in a key exchange phase,  $pkt$  is sent to  $\pi_P^i$  and the response is returned, otherwise, returns  $\perp$ .

- **Reveal**( $\pi_P^i, d$ ), for  $P \in \mathcal{P}, i \in [N], d \in [D]$ .

This query allows the adversary to learn any stage key of a specified session oracle. The contents of  $\pi_P^i$ ’s  $d$ -th stage key state  $k^d$  is returned. After this query, we say  $k^d$  was *revealed*.

- **Corrupt**( $S$ ), for  $S \in \mathcal{S}$ .

This query allows the adversary to learn the long-term secret and static state of a specified server. The long-term secret key, configuration state (if any), resumption states, and other static state of  $S$  are returned. After this query, we say  $S$  was *corrupted*.

- **NextTP**( $S$ ), for  $S \in \mathcal{S}$ .

This query allows the adversary to advance the time period of a specified server. If `scfg` is defined, it further allows the adversary to reconfigure the specified server. The local time state `loct` of  $S$  is incremented by one and `scfg_gen` (if defined) is executed to update the configuration state `scfg`; the value of `loct` is returned. If `scfg_gen` is executed, we say  $S$  was *reconfigured* (at time period `loct`).

<sup>14</sup>In practice, 0-RTT replay attacks can be mounted to *different* servers with the same public-secret key pair. However, 0-RTT key exchange message(s) replayed to other servers with different public-secret key pairs should be rejected.

- $\text{Encrypt}(\pi_P^i, d, ad, m_0, m_1)$ , for  $P \in \mathcal{P}, i \in [N], d \in [D], ad \in \mathcal{AD}, m_0, m_1 \in \mathcal{M}_{\text{SC}} \cup \mathcal{M}_{\text{pRST}} \cup \{\text{rst}\}$ . This query allows the adversary to specify any associated data and any two secure channel or pre-reset messages of the same length, then get the ciphertext of one message determined by  $b_P^i$ , the random bit associated with a specified session oracle at a specified stage. This query also keeps a state  $m_{\text{prst}}$  (initialized to  $\varepsilon$ ) for  $\pi_P^i$  to store one of the specified pre-reset messages determined by  $b_P^i$  and allows the adversary to retrieve it on request. (Recall that each oracle has at most one pre-reset message, so this query stores only the first pre-reset message and rejects others.)

The detailed procedures are listed as follows:

- 1: if  $m_0, m_1$  are of different types (i.e.,  $\mathcal{M}_{\text{SC}}$  or  $\mathcal{M}_{\text{pRST}}$  or  $\{\text{rst}\}$ ) or  $|m_0| \neq |m_1|$  or  $\pi_P^i$  is not in its  $d$ -th secure channel phase or  $k_p^d = \perp$  (where  $p = c$  if  $P \in \mathcal{C}$  and  $p = s$  if  $P \in \mathcal{S}$ ), return  $\perp$
- 2: if  $m_0, m_1 \in \mathcal{M}_{\text{pRST}}$ , set  $m_{\text{prst}} \leftarrow m_{b_P^i}$  if  $m_{\text{prst}} = \varepsilon$  or return  $\perp$  otherwise
- 3: if  $m_0 = m_1 = \text{rst}$ , return  $m_{\text{prst}}$
- 4: (upon setting each encryption stage key, initialize  $st_e$  with  $\text{sAEAD.sl}$  and set  $u \leftarrow 0, sent \leftarrow \varepsilon$ )
- 5:  $u \leftarrow u + 1, (sent.ct_u, st'_e) \xleftarrow{\$} \text{sAEAD.sE}(k_p^d, ad, m_{b_P^i}, st_e)$
- 6:  $(sent.ad_u, st_e) \leftarrow (ad, st'_e)$
- 7: return  $sent.ct_u$

- $\text{Decrypt}(\pi_P^i, d, ad, ct)$ , for  $P \in \mathcal{P}, i \in [N], d \in [D], ad \in \mathcal{AD}, ct \in \{0, 1\}^*$ .

This query allows the adversary to specify any associated data and any ciphertext to be decrypted by the *partner(s)* of a specified session oracle at a specified stage, then get the secret bit  $b_P^i$  if the decrypted message is a valid secure channel or pre-reset message and this query is “out-of-sync” (defined below), otherwise, it gets  $\perp$ .

The detailed procedures are listed as follows:

- 1: if  $\pi_P^i$  is not in its  $d$ -th secure channel phase or  $k_p^d = \perp$  (where  $p = c$  if  $P \in \mathcal{C}$  and  $p = s$  if  $P \in \mathcal{S}$ ), return  $\perp$
- 2: (upon setting each decryption stage key, initialize  $st_d$  with  $\text{sAEAD.sl}$  and set  $v \leftarrow 0, rcvd \leftarrow \varepsilon, \text{outofsync} \leftarrow 0$ )
- 3:  $v \leftarrow v + 1, rcvd.ct_v \leftarrow ct, (m, st'_d) \leftarrow \text{sAEAD.sD}(k_p^d, ad, ct, st_d)$
- 4:  $(rcvd.ad_v, st_d) \leftarrow (ad, st'_d)$
- 5: if  $m \notin \mathcal{M}_{\text{SC}} \cup \mathcal{M}_{\text{pRST}}$ , set  $m \leftarrow \perp$
- 6: if  $(al = 4) \wedge \text{cond}_4$  or  $(al \leq 3) \wedge (m \neq \perp) \wedge \text{cond}_{al}$ , set  $\text{outofsync} \leftarrow 1$
- 7: if  $(\text{outofsync} = 1) \wedge (m \neq \perp)$ , return  $b_P^i$ , otherwise, return  $\perp$

The above “out-of-sync” condition (see procedure 6) varies with the associated authentication level  $al \in [4]$ . We recall the four authentication conditions  $\text{cond}_{al}$  defined in [14] below:

$$\text{cond}_1: (\nexists y : (ct = sent.ct_y) \wedge (ad = sent.ad_y))$$

$$\text{cond}_2: (\nexists y : (ct = sent.ct_y) \wedge (ad = sent.ad_y)) \vee (\exists w < v : ct = rcvd.ct_w)$$

$$\text{cond}_3: (\nexists y : (ct = sent.ct_y) \wedge (ad = sent.ad_y)) \vee (\exists w, x, y : (w < v) \wedge (sent.ct_x = rcvd.ct_w) \wedge (sent.ct_y = ct) \wedge (x \geq y))$$

$$\text{cond}_4: (u < v) \vee (ct \neq sent.ct_v) \vee (ad \neq sent.ad_v)$$

Note that  $\text{cond}_{al}$  constrains the adversary’s ability to make an “out-of-sync”  $\text{Decrypt}$  query (which is necessary to reveal the secret bit). That is, the more easily  $\text{cond}_{al}$  can be satisfied, the more powerful the adversary is, and hence the higher authentication level the security experiment captures. In particular,  $\text{cond}_1$  corresponds to the lowest authentication level that only guarantees no forgeries;  $\text{cond}_2$  adds a clause to additionally capture replays;  $\text{cond}_3$  extends  $\text{cond}_2$  to further capture reordering (note that  $\text{cond}_3$  constrained on  $x = y$  is equal to  $\text{cond}_2$ ); finally,  $\text{cond}_4$  corresponds to the highest authentication level that prevents forgeries, replays, reordering, and dropping.

**Advantage Measures.** An adversary  $\mathcal{A}$  against a msACCE protocol  $\Pi$  in msACCE-std has the following advantage measures (that each takes the security parameter as an implicit input).

*Server Authentication.* We define  $\text{Adv}_{\Pi}^{\text{s-auth}}(\mathcal{A})$  as the probability that there exist a client oracle  $\pi_C^i$  and its intended server  $S$  such that the following conditions hold:

1.  $\pi_C^i$  has set its session key;
2.  $S$  was not corrupted before  $\pi_C^i$  set its session key;
3. No interim keys of  $\pi_C^i$  or its partner(s) were revealed;<sup>15</sup>
4. There is no unique server oracle  $\pi_S^j$  that is  $\pi_C^i$ 's peer.

The above captures the attacks in which the adversary impersonates a server to make the client mistakenly believe that it shares the session key with the server.

(*level-al*)<sup>16</sup> *Channel Security*. We define  $\mathbf{Adv}_{\Pi}^{\text{cs-al}}(\mathcal{A})$  as  $|2\Pr[b_P^i = b'] - 1|$ , where  $al \in [4]$  is the specified authentication level and  $(P, i, b')$  is output by  $\mathcal{A}$ , such that the following conditions hold:

1. If  $P = S \in \mathcal{S}$ ,  $\pi_S^i$  has a peer  $\pi_C^i$ ; if  $P = C \in \mathcal{C}$ , denote  $S$  as  $\pi_C^i$ 's intended server;
2. If  $S$  was corrupted, then 1)  $S$  was corrupted after  $\pi_P^i$  set its last stage key and 2) for any ( $d$ -th) stage key of  $\pi_P^i$  that is not required to provide *forward secrecy*, no  $\mathbf{Encrypt}(\pi_P^i, d, \cdot, \cdot, \cdot)$  queries were made before  $S$  was (if ever) reconfigured after  $\pi_P^i$  set its  $d$ -th stage key.
3. No stage keys of  $\pi_P^i$  or its partner(s) were revealed;
4. If an  $\mathbf{Encrypt}(\pi_P^i, \cdot, \cdot, \cdot, \cdot)$  query was made on two distinct pre-reset messages, then later no  $\mathbf{Encrypt}(\pi_P^i, \cdot, \cdot, \mathbf{rst}, \mathbf{rst})$  queries were made.

The above captures the attacks in which the adversary compromises the confidentiality or integrity of secure channel messages without revealing stage keys or the hidden pre-reset message or corrupting the server before the client sets its last stage key (which may not be the session key). If some stage key is not supposed to provide forward secrecy, the adversary is further prevented from accessing the communication of that stage if the server was not reconfigured before it was corrupted.

## 2) msACCE Packet-Authentication Security Model:

In this msACCE packet-authentication (msACCE-pauth) security model, we consider security goals related to packet authentication beyond those captured by the msACCE-std model. Note that msACCE-std essentially focuses only on the packet fields in the application layer, while msACCE-pauth further covers transport-layer headers and IP addresses.

First, we consider IP spoofing prevention as with the QACCE model, but, as illustrated later, generalize one of the QACCE queries to additionally capture IP spoofing attacks in the full sessions. Then, more importantly, we define four novel packet-level security notions (elaborated later): *KE Header Integrity*, *KE Payload Integrity*, *SC Header Integrity*, and *Reset Authentication*, which enable a comprehensive and fine-grained security analysis of layered protocols.

In particular, KE Header and Payload Integrity respectively capture the header and payload integrity of key exchange packets. Such security issues have not been investigated before and, as we show later, lead to new availability attacks for both TFO+TLS 1.3 and UDP+QUIC. Furthermore, we employ SC Header Integrity to capture the header integrity of non-reset packets in secure channel phases. Note that, unlike the availability attacks shown in [49], successful attacks breaking our security notions are *harder or impossible to detect* by the client as they do not affect the client's session key establishment, so these attacks could be more harmful in this sense. Finally, our model captures malicious *undetectable* session resets in a secure channel phase with Reset Authentication.

<sup>15</sup>More precisely, we allow revealing other stage keys (if any), except the first stage key, of a  $\pi_C^i$ 's partner which observes the same session identifier at only the first stage but not the next one, because such a partner's key exchange message is never received by  $\pi_C^i$ . Similar condition relaxation also holds for our other security notions.

<sup>16</sup>Note that Server Authentication does not depend on  $al$ . To see this, consider the weaker Server Authentication notion defined in the same way except its security experiment has no  $\mathbf{Decrypt}$  query and the  $\mathbf{Encrypt}$  query takes only one message as input. One can easily reduce our Server Authentication security with an arbitrary authentication level  $al$  to the above weaker Server Authentication security. Therefore, those notions are all equivalent.

As with the msACCE-std model, msACCE-pauth captures multiple stages and considers a very powerful adversary. It also inherits the same definitions of protocol entities, session oracles, session identifiers, peers, and partners. The security experiment and advantage measures are defined below.

**Security Experiment.** Consider the same experiment setups as in msACCE-std, except that no random bit  $b_P^i$  is needed. The adversary  $\mathcal{A}$  is given all the public parameters and interacts with the session oracles via the same **Connect**, **Resume**, **Send**, **Reveal**, **Corrupt**, **NextTP** queries as in the msACCE-std model,<sup>17</sup> as well as the following:

- **Connprivate**( $\pi_C^i, \pi_S^j, \text{cmp}$ ), for  $C \in \mathcal{C}, S \in \mathcal{S}, i, j \in [N], \text{cmp} \in \{0, 1\}$ .

This query allows the adversary to run a complete or partial full session between any specified client and server oracles without observing their communication. This query always returns  $\perp$ . If  $\text{cmp} = 1$ ,  $\pi_C^i$  and  $\pi_S^j$  run a *complete* full session privately without showing their communication to  $\mathcal{A}$ . If  $\text{cmp} = 0$ ,  $\pi_C^i$  and  $\pi_S^j$  run a *partial* full session privately such that the last packet sent from  $\pi_C^i$  right before  $\pi_S^j$  sets its first stage key is blocked. Note that this query extends the QACCE **Connprivate** query [49] to model IP-spoofing attacks targeting both *full* and resumption sessions, with help of an input flag  $\text{cmp}$ .

- **Pack**( $\pi_P^i, ad, m$ ), for  $P \in \mathcal{P}, i \in [N], ad \in \mathcal{AD}, m \in \mathcal{M}_{\text{SC}} \cup \mathcal{M}_{\text{pRST}} \cup \{\text{prst}, \text{rst}\}$ .

This query allows the adversary to specify any associated data and any message in a secure channel phase, then get the packet output by a specified session oracle. It also allows the adversary to ask a specified session oracle to output its pre-reset packet, which is either for a specified pre-reset message (if  $m \in \mathcal{M}_{\text{pRST}}$ ) or for a real one (hidden from the adversary) generated according to the protocol (if  $m = \text{prst}$ ), or output its reset packet (if  $m = \text{rst}$ ).

The detailed procedures are listed as follows:

- 1: if  $\pi_P^i$  is not in a secure channel phase, return  $\perp$
- 2: if  $m \in \mathcal{M}_{\text{pRST}} \cup \{\text{prst}\}$  and  $\pi_P^i$  has output its pre-reset packet before, return  $\perp$
- 3: if  $m = \text{rst}$ ,  $\pi_P^i$  outputs its reset packet, which is returned
- 4: if  $m \neq \text{rst}$ ,  $\pi_P^i$  outputs the (encrypted) packet that it would send to its partner(s) according to the protocol, for the specified associated data  $ad$  and message  $m$  if  $m \in \mathcal{M}_{\text{SC}}$  or for the pre-reset message if  $m \in \mathcal{M}_{\text{pRST}} \cup \{\text{prst}\}$ , then this packet is returned

- **Deliver**( $\pi_P^i, pkt$ ), for  $P \in \mathcal{P}, i \in [N], pkt \in \{0, 1\}^*$ .

This query allows the adversary to deliver any packet to a specified session oracle and get its response in a secure channel phase. If  $\pi_P^i$  is in a secure channel phase,  $pkt$  is delivered to  $\pi_P^i$  and its response (if any) is returned, otherwise, returns  $\perp$ .

**Advantage Measures.** An adversary  $\mathcal{A}$  against a msACCE protocol  $\Pi$  in msACCE-pauth has the following associated advantage measures (that each takes the security parameter as an implicit input).

*IP-Spoofing Prevention.* We define  $\text{Adv}_{\Pi}^{\text{ipsp}}(\mathcal{A})$  as the probability that there exist a client oracle  $\pi_C^i$  and a server oracle  $\pi_S^j$  such that the following conditions hold:

1.  $\pi_S^j$  has set its first stage key right after a **Send**( $\pi_S^j, (\text{IP}_C, \text{IP}_S, \cdot, \cdot)$ ) query;
2.  $S$  was not corrupted before  $\pi_S^j$  set its first stage key;
3. For any session oracle  $\pi_S^y$  of  $S$  that shares the same time period (i.e., **loct** value) with  $\pi_S^j$ , the only queries made by  $\mathcal{A}$  about  $C$  (before  $\pi_S^j$  set its first stage key) were:
  - **Connprivate**( $\pi_C^x, \pi_S^y, \cdot$ ) for some  $x \in [N]$ , or
  - **Send**( $\pi_S^y, (\text{IP}_C, \text{IP}_S, \cdot, \cdot)$ ), where  $(\text{IP}_C, \text{IP}_S, \cdot, \cdot)$  was the last packet received by  $\pi_S^y$  right before it set its first stage key.

The above captures the attacks in which the adversary fools a server into accepting a spurious connection request seemingly from an impersonated client, without observing any previous communication between the client and server in the same time period associated with the accepting server.

<sup>17</sup>Note that **Encrypt** and **Decrypt** queries are not needed because msACCE-pauth does not consider data confidentiality and integrity that are already captured by msACCE-std.



Such attacks can lead to exhaustion of server resources, i.e., denial of service, by triggering excessive stage key derivation and probably 0-RTT data processing.

*KE Header Integrity.* We define  $\mathbf{Adv}_{\Pi}^{\text{int-keh}}(\mathcal{A})$  as the probability that there exist a client oracle  $\pi_C^i$  and a server oracle  $\pi_S^j$  such that the following conditions hold:

1.  $\pi_C^i$  has set its session key and has a peer  $\pi_S^j$ ;
2.  $S$  was not corrupted before  $\pi_C^i$  set its session key;
3. No interim keys of  $\pi_C^i$  or its partner(s) were revealed;
4. In a key exchange phase before  $\pi_C^i$  set its session key,  $\pi_C^i$  (resp.  $\pi_S^j$ ) accepted a packet with a header that was not output by  $\pi_S^j$  (resp.  $\pi_C^i$ ).

The above captures the attacks in which the adversary modifies the protocol header of a key exchange packet of the communicating parties without affecting the client setting its session key. In the above definition, we assume that a client sets its session key *immediately* after sending its last key exchange packet(s) (if any). Then, a forged packet that leads to a successful attack cannot be any of these last packet(s), which have not yet been sent to the server. The same assumption is made for KE Payload Integrity defined below.

*KE Payload Integrity.* We define  $\mathbf{Adv}_{\Pi}^{\text{int-kep}}(\mathcal{A})$  as the probability that there exist a client oracle  $\pi_C^i$  and a server oracle  $\pi_S^j$  such that the same (1)~(3) conditions as in the above KE Header Integrity notion hold and the following holds:

4. In a key exchange phase before  $\pi_C^i$  set its session key,  $\pi_C^i$  (resp.  $\pi_S^j$ ) accepted a packet with a payload that was not output by  $\pi_S^j$  (resp.  $\pi_C^i$ ).

The above captures the attacks in which the adversary modifies the payload of a key exchange packet of the communicating parties without affecting the client setting its session key. Note that the above notion can be trivially achieved if the session identifier includes the entire key exchange transcript (excluding some headers), which is the case for TLS 1.3. However, some other protocols like QUIC only have partial transcripts as their session identifiers.

*SC Header Integrity.* We define  $\mathbf{Adv}_{\Pi}^{\text{int-h}}(\mathcal{A})$  as the probability that  $\mathcal{A}$  outputs  $(P, i)$  such that the same (1)~(3) conditions as in the Channel Security notion hold and the following holds:

4. In a secure channel phase, a partner of  $\pi_P^i$  accepted a non-reset packet with a header that was not output by  $\pi_P^i$  (via Pack queries), or a partner of  $\pi_P^i$  accepted a non-reset header-only packet.

The above captures the attacks in which the adversary creates a valid non-reset secure channel packet by forging the protocol header without breaking any Channel Security conditions. Note that in the above security notion an invalid header forgery is detected *immediately* after the malicious packet is received and processed, while the detection of invalid packet forgeries in a key exchange phase (e.g., for plaintext packets) can be *delayed* to the point when the client sets its session key, according to the definitions of KE Header and Payload Integrity.

*Reset Authentication.* We define  $\mathbf{Adv}_{\Pi}^{\text{rst-auth}}(\mathcal{A})$  as the probability that  $\mathcal{A}$  outputs  $(P, i)$  such that the same (1)~(3) conditions as in the Channel Security notion hold and the following holds:

4. In a secure channel phase, a partner of  $\pi_P^i$  accepted a reset packet and 1) if  $P \in \mathcal{S}$  then  $P$  was not corrupted before  $\pi_P^i$  accepted that reset packet and 2) the  $\text{Pack}(\pi_P^i, \cdot, m)$  queries made by  $\mathcal{A}$  were only for  $m \in \mathcal{M}_{\text{SC}} \cup \{\text{prst}\}$  (i.e.,  $m \notin \mathcal{M}_{\text{pRST}} \cup \{\text{rst}\}$ ).

The above captures the attacks in which the adversary forges a valid reset packet without breaking any Channel Security conditions or corrupting the secret static state used to generate the reset packet. Note that such attacks are *undetectable* by the accepting party that resets the session, as opposed to a network attacker that simply drops packets which may eventually lead to a session reset.

REMARK ABOUT msACCE SECURITY MODEL COMPLETENESS AND LOW-LAYER INTEGRITY. Note that the payload integrity in secure channels is captured by Channel Security. Our msACCE-std and msACCE-pauth models *completely* capture the authentication (or integrity) of all packet fields in the transport and application layers. Furthermore, msACCE-pauth captures (network-layer) IP-Spoofing Prevention against weaker off-path attackers (i.e., those can only inject packets without observing the communication), but leaves other integrity attacks on low layers (e.g., network, link, and physical layers) uncovered. Such attacks may affect packet forwarding, node-to-node data transfer, or raw data transmission, which are outside the scope of our work.

## 5 Provable Security Analysis

Equipped with msACCE security models, we now analyze and compare the security of TFO+TLS 1.3, UDP+QUIC, and UDP+QUIC[TLS]. The security results are summarized in Table 2 (recalled below). As mentioned in the Introduction, by [49, 28] results, none of the above protocols achieves forward secrecy for 0-RTT keys or protects against 0-RTT data replays (which contribute to the first two rows in the table). We now move to the detailed analyses and start with TFO+TLS 1.3.

Table 2: Security comparison (reminder)

	TLS 1.3 +TFO	QUIC +UDP	QUIC[TLS] +UDP
0-RTT Key Forward Secrecy [49, 28]	<b>X</b>	<b>X</b>	<b>X</b>
0-RTT Data Anti-Replay [49, 28]	<b>X</b>	<b>X</b>	<b>X</b>
Server Authentication	✓	✓	✓
Channel Security	✓	✓	✓
IP-Spoofing Prevention	✓	✓	✓
KE Header Integrity	<b>X</b>	<b>X</b>	<b>X</b>
KE Payload Integrity	✓	<b>X</b>	<b>X</b>
SC Header Integrity	<b>X</b>	✓	✓
Reset Authentication	<b>X</b>	<b>X</b>	✓

### 5.1 TLS 1.3 over TFO

#### 1) Protocol Description:

Referring to msACCE protocol syntax, the TFO+TLS 1.3 2-RTT full handshake (see Figure 2 (a)) is a 2-stage msACCE protocol in the full mode and the TFO+TLS 1.3 0-RTT resumption handshake (see Figure 2 (b)) is a 3-stage msACCE protocol in the resumption mode. Note that we focus only on the main components of the handshakes and omit more advanced features such as 0.5-RTT data, client authentication, and post-handshake messages (except `NewSessionTicket`). In a full handshake, the initial keys are set after sending or receiving `ServerHello` and the final keys (i.e., session keys) are set after sending or receiving `ClientFinished` (but only handshake messages up to `ServerFinished` are used for final key generation). In a 0-RTT resumption handshake, the parties set 0-RTT keys to encrypt or decrypt 0-RTT data, after sending or receiving `ClientHello`.

According to the TFO and TLS 1.3 specifications [18, 60], the TFO+TLS 1.3 header contains the TCP header (see Figure 1). We ignore some uninteresting header fields such as port numbers and the checksum because modifying them only leads to redirected or dropped packets. Such adversarial capabilities are already considered in the msACCE security models. We thus define the header space  $\mathcal{H}$  as containing the following TCP header fields: a 32-bit sequence number `sqn`, a 32-bit acknowledgment number `ack`, a 4-bit data offset `off`, a 6-bit reserved field `resv`, a 6-bit control bits field `ctrl`, a 16-bit window `window`, a 16-bit urgent pointer `urqp`, a variable-length ( $\leq 320$ -bit) padded options `opt`. For encrypted packets,  $\mathcal{H}$  additionally contains TLS 1.3 record header fields: an 8-bit type `type`, a 16-bit

version `ver`, and a 16-bit length `len`. We further define reset packets as those with the RST bit (i.e., the 4-th bit of `ctrl`) set to 1. `scfg_gen` erases the server’s resumption states (if any) and refreshes the (static) key  $k_{\text{ck}} \xleftarrow{\$} \{0, 1\}^{128}$  for the TFO cookie generation function  $F$ , which for example can be an AES-128 block cipher with output truncated to 64 bits as suggested in [59, 18]. If the session ticket of TLS 1.3 is implemented as a self-encrypted and self-authenticated value, `scfg_gen` also refreshes the server’s (static) session ticket encryption key  $stek \xleftarrow{\$} \{0, 1\}^{128}$  for a secure authenticated encryption scheme (e.g., AES-GCM [50]).

TLS 1.3 enforces different content types for encrypted key exchange and secure channel messages. For simplicity, we define  $\mathcal{M}_{\text{KE}}$  and  $\mathcal{M}_{\text{SC}}$  as consisting of bit strings differing in their first bits. Note that  $\mathcal{M}_{\text{PRST}} = \emptyset$ . In Appendix B, we formalize TLS 1.3’s stateful AEAD scheme  $\text{sAEAD}_{\text{TLS}}$  based on its underlying nonce-based AEAD scheme (instantiated with AES-GCM or other schemes as documented in [60]). The associated data space  $\mathcal{AD}$  contains the TLS 1.3 record header.

We refer to Figure 6 in Appendix D for the remaining details of TFO and refer to [27] for the detailed descriptions of TLS 1.3 handshake messages and key generations as well as [26, 60] for details about TLS 1.3’s handling of session tickets.

## 2) Security Results:

TFO+TLS 1.3’s session identifier  $\text{sid}_{\text{TLS}}$  is defined as all key exchange messages (that could be encrypted) from `ClientHello` to `ServerFinished`, excluding TCP headers and IP addresses. The  $\text{msACCE-std}$  security of TFO+TLS 1.3 is by definition independent of TCP headers and is hence provided by the TLS 1.3 component. Previous works [47, 27] only proved TLS 1.3’s authenticated key exchange security, i.e., the stage keys are authenticated and indistinguishable from random ones under reasonable computational assumptions. In Appendix C, we outline how their security results can be adapted to prove TLS 1.3’s Server Authentication and level-4 Channel Security in our  $\text{msACCE-std}$  model, by additionally relying on the level-4 AEAD security of  $\text{sAEAD}_{\text{TLS}}$  (which can be reduced to security of the underlying nonce-based AEAD as shown in [22]).

The  $\text{msACCE-pauth}$  security analyses are shown as follows.

**IP-Spoofing Prevention.** This security of TFO+TLS 1.3 is provided by the TFO component through TCP sequence number randomization and TFO cookies. By modeling the cookie generation function  $F$  as a PRF, we have the following theorem with the proof in Appendix D:

**Theorem 1** *For any efficient adversary  $\mathcal{A}$  making at most  $q_{\text{T}}$ ,  $q_{\text{S}}$  queries respectively to `NextTP`, `Send`, there exists an efficient adversary  $\mathcal{B}$  such that:*

$$\text{Adv}_{\text{TFO+TLS 1.3}}^{\text{ipsp}}(\mathcal{A}) \leq (|\mathcal{S}| + q_{\text{T}}) \text{Adv}_F^{\text{prf}}(\mathcal{B}) + \frac{q_{\text{S}}(|\mathcal{S}| + q_{\text{T}})}{2^{|\text{sqn}|}},$$

where  $|\mathcal{S}|$  is the number of servers and  $|\text{sqn}|$  is the bit length of a TCP sequence number.

REMARK. Note that the above second probability term is not very small since  $|\text{sqn}| = 32$ . Actually, an attacker can indeed successfully establish a TCP connection using a spoofed client IP address with roughly  $2^{32}$  random guesses. However, our security bound is still acceptable because each guess made by the attacker corresponds to an “online” `Send` query. Here “online” means that the attacker has to interact with the server every time it makes a guess, no matter how many resources it may consume offline. Such online attacks are easy to detect and suppress.

**KE Header Integrity.** TFO+TLS 1.3 does not achieve this security notion because TCP headers are never authenticated. We find a new practical attack below, where an efficient adversary  $\mathcal{A}$  can always get  $\text{Adv}_{\text{TFO+TLS 1.3}}^{\text{int-keh}}(\mathcal{A}) = 1$ :

*TFO Cookie Removal.*  $\mathcal{A}$  can first make  $\pi_C^{i'}$  complete a full handshake with  $\pi_S^{j'}$  (via `Connect`, `Send` queries), then query `Resume`( $\pi_C^i, S, i'$ ) ( $i' < i$ ) to get the output packet  $(\text{IP}_C, \text{IP}_S, H, pd)$ , which is a SYN packet with a TFO cookie.  $\mathcal{A}$  then modifies the `opt` field of  $H$  to get a new  $H' \neq H$  that contains no cookie. The resulting SYN packet will be accepted by a new server oracle  $\pi_S^j$ , which will then respond with a SYN-ACK packet that does not contain a TFO cookie, indicating a fallback to

the standard 3-way TCP. As a result, a 1-RTT handshake is needed to complete the connection and any 0-RTT data sent with SYN would be retransmitted. This eliminates the entire benefit of TFO without being detected, resulting in reduced performance and increased handshake latency. A similar attack is possible by removing the TFO cookie in a server’s SYN-ACK packet.

Interestingly, clients are supposed to cache negative TFO responses and avoid sending TFO connections again for a lengthy period of time. This is because the most likely explanation for this behavior is that the server does not support TFO, but only standard TCP [18]. As a result, performing this attack for a single connection prevents TFO from being used with this server for a lengthy time period (i.e., days or weeks).

**KE Payload Integrity.** TFO+TLS 1.3 is secure in this regard simply because  $\text{sid}_{\text{TLS}}$  consists of the payloads of all key exchange packets exchanged between the communicating parties before the client sets its session key. That is, for any client oracle that has a peer server oracle, by definition they observe the same  $\text{sid}_{\text{TLS}}$  and hence no key exchange packet payload can be modified, i.e.,  $\text{Adv}_{\text{TFO+TLS 1.3}}^{\text{int-kep}}(\mathcal{A}) = 0$  for any efficient adversary  $\mathcal{A}$ .

**SC Header Integrity.** TFO+TLS 1.3 does not achieve this security notion again because of the unauthenticated TCP headers. A efficient adversary  $\mathcal{A}$  can get  $\text{Adv}_{\text{TFO+TLS 1.3}}^{\text{int-h}}(\mathcal{A}) = 1$  by either modifying the TCP header of an encrypted packet (e.g., reducing the `window` value) or by forging a header-only packet (e.g., removing the payload of an encrypted packet and changing its `ack` value). Such packets are valid and will be accepted by the receiving session oracle.

The above fact exposes the adversary’s ability to arbitrarily modify or even entirely forge the information in the TCP header, which is being relied on to provide reliable delivery, in-order delivery, flow control, and congestion control for the targeted flow. This leads to a whole host of availability attacks that the networking community has been slowly uncovering via manual investigation over the last 30 years [64, 40, 4, 33, 45, 44, 38, 58, 17, 32, 57, 52, 71, 39]. Some of the practical attacks are described in Appendix E.

**Reset Authentication.** TFO+TLS 1.3 is insecure in this sense because its reset packet, TCP Reset, is an unauthenticated header-only packet. This leads to a practical attack below, where an efficient adversary  $\mathcal{A}$  always gets  $\text{Adv}_{\text{TFO+TLS 1.3}}^{\text{rst-auth}}(\mathcal{A}) = 1$ :

*TCP Reset Attack.*  $\mathcal{A}$  can first make two session oracles complete a handshake using `Connect`, `Send` queries, then use `Pack`, `Deliver` queries to let them exchange secure channel packets. By observing these packet headers,  $\mathcal{A}$  can easily forge a valid reset packet by setting its RST bit to 1 and the remaining header fields to reasonable values. This attack will cause TCP to tear down the connection immediately without waiting for all data to be delivered.

Note that even an off-path adversary who can only inject packets into the communication channel may be able to accomplish this attack. The injected TCP reset packet needs to be within the receive window for the client or server, but [71] demonstrated that a surprisingly small number of packets is needed to achieve this, thanks to the large receive windows typically used by implementations.

## 5.2 QUIC over UDP

### 1) Protocol Description:

Referring to msACCE protocol syntax, the UDP+QUIC 1-RTT full handshake (see Figure 4 (a)) is a 2-stage msACCE protocol in the full mode and the UDP+QUIC 0-RTT resumption handshake (see Figure 4 (b)) is a 2-stage msACCE protocol in the resumption mode. The initial keys are set after sending or receiving `ClientHello` and the final keys (i.e., session keys) are set after sending or receiving `ServerHello`.

According to the UDP and QUIC specifications [55, 21, 46], the UDP+QUIC header contains the UDP header (see Figure 3) and the QUIC header (described below). As with the TCP header, we ignore the port numbers and checksum in the UDP header. Similarly, we also ignore the UDP length field because it affects only the length of the QUIC header and payload, while the adversary is already

allowed to directly modify the packet length. We thus can completely omit the UDP header and define the header space  $\mathcal{H}$  as containing the following QUIC header fields: an 8-bit public flag `flag`, a 64-bit connection ID `cid`, a 64-bit packet number `pn`,<sup>18</sup> and other optional fields. We further define reset packets as those with the `PUBLIC_FLAG_RESET` bit (i.e., the 7-th bit of `flag`) set to 1. A reset packet header only contains `flag` and `cid`.

As with TLS 1.3, for UDP+QUIC we define  $\mathcal{M}_{\text{KE}}$  and  $\mathcal{M}_{\text{SC}}$  as consisting of bit strings differing in their first bits.  $\mathcal{M}_{\text{PRST}} = \emptyset$ . In Appendix B, we formalize QUIC’s stateful AEAD scheme `sAEADQUIC` based on its underlying nonce-based AEAD scheme (instantiated with AES-GCM [50]). The associated data space  $\mathcal{AD}$  contains the entire QUIC header.

We refer to [49] for the detailed descriptions of `scfg.gen` and QUIC handshake messages and key generations.

## 2) Security Results:

UDP+QUIC’s session identifier `sidQUIC` is defined as the `ClientHello` payload and `ServerHello`, excluding IP addresses. The `msACCE-std` security of UDP+QUIC follows from prior works as we discuss in Appendix F. Note that UDP+QUIC only achieves level-1 Channel Security (based on the level-1 AEAD security of `sAEADQUIC` proved in Appendix B), but, as discussed in [49], QUIC implicitly prevents packet dropping and reordering by authenticating `pn` in the packet header. It also prevents replays with frame sequence numbers encrypted in the payload. Therefore, UDP+QUIC essentially achieves level-4 authentication as TLS 1.3 does.

The `msACCE-pauth` security analyses are shown as follows.

**IP-Spoofing Prevention.** In [49], QUIC has been proven secure against IP spoofing based on the AEAD security. Their IP-spoofing security notion is the same as our IP-Spoofing Prevention notion for UDP+QUIC except that ours additionally captures attacks in full sessions. However, since source-address tokens are validated in both full and resumption sessions, their results can be trivially adapted to show that UDP+QUIC achieves IP-Spoofing Prevention.

**KE Header and Payload Integrity.** UDP+QUIC does not achieve these security notions because its first-round key exchange messages, i.e., `InchoateClientHello` and `ServerReject`, and any invalid `ClientHello` are not fully authenticated. Interestingly, a variety of existing attacks on QUIC’s availability discovered in [49] are all examples of key exchange packet manipulations (e.g., the server config replay attack, connection ID manipulation attack, etc.), but these attacks cause connection failure and hence are easy to detect. However, successful attacks breaking KE Header or Payload Integrity will be harder (if not impossible) to detect.

For KE Header Integrity, we do not find any harmful attacks but theoretical attacks exist. For instance, an efficient adversary  $\mathcal{A}$  can get  $\text{Adv}_{\text{UDP+QUIC}}^{\text{int-keh}}(\mathcal{A}) = 1$  as follows.  $\mathcal{A}$  can first query `Connect`( $\pi_C^i, S$ ) to get the output packet ( $\text{IP}_C, \text{IP}_S, H, pd$ ), then modify the `flag` and `pn` fields of  $H$  to get a new header  $H' \neq H$  that only changes `pn`’s length but not its underlying value. The resulting packet will be accepted by a new server oracle  $\pi_S^j$ . This attack has no practical impact on UDP+QUIC but it successfully modifies the protocol header without being detected.

For KE Payload Integrity, we find a new practical attack described below where an efficient adversary  $\mathcal{A}$  can get  $\text{Adv}_{\text{UDP+QUIC}}^{\text{int-kep}}(\mathcal{A}) \approx 1$ :

*ServerReject Triggering.*  $\mathcal{A}$  can first let  $\pi_C^{i'}$  complete a full handshake with  $\pi_S^{j'}$  with `Connect`, `Send` queries, then query `Resume`( $\pi_C^i, S, i'$ ) ( $i' < i$ ) to get the output `ClientHello` packet.  $\mathcal{A}$  then modifies its payload by replacing the source-address token `stk` with a random value, which with high probability is invalid. Sending this modified packet to a new server oracle  $\pi_S^j$  will trigger a `ServerReject` packet containing a new valid `stk`. This as a result downgrades the original 0-RTT resumption connection to a full 1-RTT connection, which causes increased latency and results in the retransmission of any

<sup>18</sup>Note that the real QUIC header only contains, say 32, lower bits of the sequence number, but this is enough to deduce the correct 64-bit value since the sequence numbers used by encryption are consecutive and start from 1. For simplicity, we consider a full 64-bit `pn` in the header as with the previous work [49].

0-RTT data. Note that this attack is hard to detect because  $\pi_C^i$  may think its original  $\text{stk}'$  has expired (although this does not happen frequently).

**SC Header Integrity.** UDP+QUIC is secure in this regard because it does not allow header-only packets to be sent in the secure channel phases and the *entire* protocol header is taken as the associated data authenticated by the underlying encryption scheme. Therefore, UDP+QUIC’s SC Header Integrity can be reduced to its level-1 Channel Security. Formally, for any efficient adversary  $\mathcal{A}$  there exists an efficient adversary  $\mathcal{B}$  such that  $\text{Adv}_{\text{UDP+QUIC}}^{\text{int-h}}(\mathcal{A}) \leq \text{Adv}_{\text{UDP+QUIC}}^{\text{cs-1}}(\mathcal{B})$ . This is because  $\mathcal{B}$  can simulate  $\mathcal{A}$ ’s security experiment perfectly by answering the `Pack` queries using `Encrypt` with same inputs and forwarding the `Deliver` queries to `Decrypt`, and if  $\mathcal{A}$  outputs  $(P, i)$  and wins then  $\mathcal{B}$  will also win by getting the secret bit  $b_p^i$  from `Decrypt`.

**Reset Authentication.** UDP+QUIC does not achieve this security notion because, similar to TCP Reset, its reset packet `PublicReset` is not authenticated either. In the following availability attack, an efficient adversary  $\mathcal{A}$  can always get  $\text{Adv}_{\text{UDP+QUIC}}^{\text{rst-auth}}(\mathcal{A}) = 1$ :

*PublicReset Attack.*  $\mathcal{A}$  can first make two session oracles complete a handshake using `Connect`, `Send` queries, then use `Pack`, `Deliver` queries to let them exchange secure channel packets. By observing these packet headers,  $\mathcal{A}$  can easily forge a valid (plaintext) reset packet by setting its `PUBLIC_FLAG_RESET` bit to 1 and the remaining packet fields to reasonable values (which is easy because it simply contains the connection ID `cid`, the sequence number of the rejected packet, and a nonce to prevent replay). This attack will cause similar effects as described in the TCP Reset attack. Note that this vulnerability is fixed in QUIC[TLS] shown below.

### 5.3 QUIC[TLS] over UDP

#### 1) Protocol Description:

As mentioned in the Background, QUIC[TLS] replaces QUIC’s key exchange with the TLS 1.3 key exchange, i.e., the stage keys are set in the same way as TLS 1.3. Therefore, the UDP+QUIC[TLS] 2-RTT full handshake is a 2-stage msACCE protocol in the full mode and the UDP+QUIC[TLS] 0-RTT resumption handshake is a 3-stage msACCE protocol in the resumption mode.

The UDP+QUIC[TLS] header is similar to the UDP+QUIC header and we can likewise omit the entire UDP header for our analysis. The detailed description of the QUIC[TLS] header is omitted here and referred to [36]. A reset packet in UDP+QUIC[TLS] looks indistinguishable from a non-reset secure channel packet, but accepting it leads to a session reset. `scfg_gen` is undefined.

QUIC[TLS] enforces different frame types for encrypted key exchange, secure channel, and pre-reset messages. For simplicity, we define  $\mathcal{M}_{\text{KE}}, \mathcal{M}_{\text{SC}}, \mathcal{M}_{\text{PRST}}$  as consisting of bit strings differing in their first two bits. The stateful AEAD scheme  $\text{sAEAD}_{\text{QUIC[TLS]}}$  used by QUIC[TLS] to encrypt packets is essentially the same as  $\text{sAEAD}_{\text{QUIC}}$ , except that the nonce of the underlying nonce-based AEAD is computed as the exclusive OR of the secret  $iv$  (part of the key generated in  $\text{sAEAD}_{\text{QUIC[TLS]}}$ ) and the packet number.<sup>19</sup> As with  $\text{sAEAD}_{\text{QUIC}}$ , the level-1 AEAD security of  $\text{sAEAD}_{\text{QUIC[TLS]}}$  is reduced to the security of its underlying nonce-based AEAD, following a proof very similar to that of Theorem 3 in Appendix B. The associated data space  $\mathcal{AD}$  contains the entire QUIC[TLS] header.

QUIC[TLS] also provides address validation with a secure token generated by the server, similar to the case in Google’s QUIC. We discuss QUIC[TLS]’s stateless reset mechanism later in the security analysis of Reset Authentication and refer to [36, 69] for the detailed UDP+QUIC[TLS] handshake messages and key generations.

#### 2) Security Results:

<sup>19</sup>Note that QUIC[TLS] applies a header protection mechanism on the encrypted packet to further hide the header fields used to form a nonce, which provides stronger nonce-hiding security as proposed by [9]. For simplicity, we do not consider this new feature and assume header protection will not weaken the security of the encrypted packet, as our analysis focuses on the packet header’s integrity other than its confidentiality. We refer to [23] for a security analysis of QUIC[TLS]’s record layer that covers header protection.

UDP+QUIC[TLS]’s session identifier  $\text{sid}_{\text{QUIC}[\text{TLS}]}$  is defined as all key exchange messages (that could be encrypted) from `ClientHello` to `ServerFinished`, excluding the header fields. By construction, UDP+QUIC[TLS] inherits the `msACCE-std` security from TLS 1.3, but uses `sAEADQUIC[\text{TLS}]` for encryption. That is, it achieves level-1 Channel Security and implicitly achieves level-4 authentication in the same way as UDP+QUIC. The IP-Spoofing Prevention of UDP+QUIC[TLS] follows an address validation scheme very similar to UDP+QUIC. In particular, if the token is generated with an authenticated encryption scheme, then the IP-Spoofing Prevention security is reduced to the encryption scheme’s authenticity security, as with UDP+QUIC. However, such an address validation scheme suffers from the same kind of availability attack against KE Payload Integrity as `ServerReject Triggering` for UDP+QUIC, where the adversary replaces the address-validation token with a random value to downgrade a 0-RTT resumption connection. As noted in [69], an adversary can also modify the unauthenticated ACK frames in the Initial packets without being detected. Furthermore, UDP+QUIC[TLS] achieves SC Header Integrity in the same way as UDP+QUIC. We are only left to show its security of KE Header Integrity and Reset Authentication.

**KE Header Integrity.** UDP+QUIC[TLS] does not achieve these security notions because its first-round Initial packets (see [36]) are not authenticated. For instance, an efficient adversary  $\mathcal{A}$  can get  $\text{Adv}_{\text{UDP+QUIC}[\text{TLS}]}^{\text{int-keh}}(\mathcal{A}) = 1$  as follows.  $\mathcal{A}$  first queries `Connect`( $\pi_C^i, S$ ) to get  $\pi_C^i$ ’s Initial packet ( $\text{IP}_C, \text{IP}_S, H, pd$ ). Then, as described in [69],  $\mathcal{A}$  can decrypt this packet with its destination connection ID  $cid$  in  $H$ , change it to another value  $cid'$ , and re-encrypt the whole packet with this new  $cid'$ . The resulting packet ( $\text{IP}_C, \text{IP}_S, H', pd'$ ), where  $H \neq H'$ , is valid and will be accepted by a new server oracle  $\pi_S^j$  without being detected by the client. However, this is only a theoretical attack with no practical impact.

**Reset Authentication.** In UDP+QUIC[TLS], the stateless reset works as follows. One party generates a 128-bit reset token using its static key and a random 64-bit QUIC connection ID as input. Then this token (carried within the pre-reset message) is sent to the other party in a secure channel phase. Later, the same party that generated this token can perform a stateless reset by regenerating the token and sending it to the other party in clear (via a reset packet).

The Reset Authentication security of UDP+QUIC[TLS] can be reduced to its level-1 Channel Security and the PRF security of the reset token generation function  $F$  (which can be instantiated with an HMAC [7] for instance), as shown in the following theorem with the proof in Appendix G:

**Theorem 2** *For any efficient adversary  $\mathcal{A}$  making at most  $q_D$  Deliver queries, there exist efficient adversaries  $\mathcal{B}$  and  $\mathcal{C}$  such that:*

$$\text{Adv}_{\text{UDP+QUIC}[\text{TLS}]}^{\text{rst-auth}}(\mathcal{A}) \leq |\mathcal{P}| \text{Adv}_F^{\text{prf}}(\mathcal{B}) + |\mathcal{P}| \text{Adv}_{\text{UDP+QUIC}[\text{TLS}]}^{\text{cs-1}}(\mathcal{C}) + \frac{|\mathcal{P}|q_D}{2^{|\text{rtk}|}} + \frac{|\mathcal{P}|N^2}{2^{|\text{cid}|}},$$

where  $|\mathcal{P}|$  is the number of parties,  $N$  is the maximum number of sessions for each party,  $|\text{rtk}|$  is the bit length of a reset token, and  $|\text{cid}|$  is the bit length of a connection ID.

REMARK. Recall that  $|\text{rtk}| = 128$  and  $|\text{cid}| = 64$ . Similar to the remark following Theorem 1, regarding the last probability term, any attacker has to observe roughly  $2^{32}$  online sessions in order to find a connection ID collision. To forge (basically replay) a valid reset packet, the attacker has to make roughly  $2^{32}$  online `Pack`( $\cdot, \cdot, \text{rst}$ ) queries to the target party to keep record of the reset packets for those observed sessions since each session is typically alive for a short time, which is an online attack that is easy to detect and suppress.<sup>20</sup> Note that the attacker does not get to choose which session to be maliciously reset since a collision can occur at any time.

<sup>20</sup>If the target party can open  $2^{32}$  concurrent sessions, then the attacker only needs to make a single online query. However, we assume in practice this is unlikely since there are at most 65535 TCP sockets per IP address.

## 6 Conclusion

Our work is the first to provide a thorough, formal, and fine-grained security comparison of the most efficient secure channel establishment protocols on the market today. By including packet-level attacks in our analysis, our results shed light on how the reliability, flow control, and congestion control of TFO+TLS 1.3, UDP+QUIC, and UDP+QUIC[TLS] compare besides their basic security, in adversarial settings.

We show that availability functionalities provided by transport-layer protocols like TCP can be compromised due to lack of packet-level authentication, which may undermine the performance of their supporting application-layer protocols. To protect against availability attacks, secure channel establishment protocols should better implement and authenticate their own transport functionalities like QUIC does. Besides, the key exchange packet integrity should also be scrutinized to avoid serious undetectable availability attacks.

We hope that our models will help protocol designers in their future protocol analyses and that our results will help practitioners better understand the advantages and limitations of secure channel establishment protocols.

## Acknowledgments

We thank the anonymous reviewers for their useful comments. This paper is based upon work supported by the National Science Foundation under Grant No. 1422794.

## References

- [1] HTTPS encryption on the web - Google transparency report. <https://transparencyreport.google.com/https/overview>. Accessed: 2020-10-22.
- [2] Josh Aas. Let's Encrypt: Looking forward to 2019. <https://letsencrypt.org/2018/12/31/looking-forward-to-2019.html>, 2018.
- [3] Michel Abdalla, Mihir Bellare, and Phillip Rogaway. The oracle Diffie-Hellman assumptions and an analysis of DHIES. In *Cryptographers Track at the RSA Conference*, pages 143–158. Springer, 2001.
- [4] Raz Abramov and Amir Herzberg. TCP Ack storm DoS attacks. In *IFIP International Information Security Conference*, pages 29–40, 2011.
- [5] Nimrod Aviram, Kai Gellert, and Tibor Jager. Session resumption protocols and efficient forward security for TLS 1.3 0-RTT. In *EUROCRYPT 2019*, pages 117–150. Springer, 2019.
- [6] Manuel Barbosa and Pooya Farshim. Security analysis of standard authentication and key agreement protocols utilising timestamps. In *International Conference on Cryptology in Africa*, pages 235–253. Springer, 2009.
- [7] Mihir Bellare, Ran Canetti, and Hugo Krawczyk. Keying hash functions for message authentication. In *CRYPTO 1996*, pages 1–15. Springer, 1996.
- [8] Mihir Bellare, Tadayoshi Kohno, and Chanathip Namprempre. Breaking and provably repairing the SSH authenticated encryption scheme: A case study of the encode-then-encrypt-and-mac paradigm. *ACM Transactions on Information and System Security (TISSEC)*, 7(2):206–241, 2004.
- [9] Mihir Bellare, Ruth Ng, and Björn Tackmann. Nonces are noticed: Aead revisited. In *Annual International Cryptology Conference*, pages 235–265. Springer, 2019.
- [10] Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In *CRYPTO 1993*, pages 232–249. Springer, 1993.
- [11] Karthikeyan Bhargavan, Bruno Blanchet, and Nadim Kobeissi. Verified models and reference implementations for the TLS 1.3 standard candidate. In *2017 IEEE Symposium on Security and Privacy*, pages 483–502. IEEE, 2017.



- [12] Karthikeyan Bhargavan, Cédric Fournet, Markulf Kohlweiss, Alfredo Pironti, Pierre-Yves Strub, and Santiago Zanella-Béguelin. Proving the TLS handshake secure (as it is). In *CRYPTO 2014*, pages 235–255. Springer, 2014.
- [13] Colin Boyd and Britta Hale. Secure channels and termination: The last word on TLS. In *International Conference on Cryptology and Information Security in Latin America*, pages 44–65. Springer, 2017.
- [14] Colin Boyd, Britta Hale, Stig Frode Mjølsnes, and Douglas Stebila. From stateless to stateful: Generic authentication and authenticated encryption constructions with application to TLS. In *Cryptographers Track at the RSA Conference*, pages 55–71. Springer, 2016.
- [15] Jacqueline Brendel, Marc Fischlin, and Felix Günther. Breakdown resilience of key exchange protocols: NewHope, TLS 1.3, and Hybrids. In *European Symposium on Research in Computer Security*, pages 521–541. Springer, 2019.
- [16] Jacqueline Brendel, Marc Fischlin, Felix Günther, and Christian Janson. PRF-ODH: Relations, instantiations, and impossibility results. In *CRYPTO 2017*, pages 651–681. Springer, 2017.
- [17] Yue Cao, Zhiyun Qian, Zhongjie Wang, Tuan Dao, Srikanth V Krishnamurthy, and Lisa M Marvel. Off-path TCP exploits: Global rate limit considered dangerous. In *USENIX Security Symposium*, pages 209–225, 2016.
- [18] Y. Cheng, J. Chu, S. Radhakrishnan, and A. Jain. TCP Fast Open. RFC 7413, December 2014.
- [19] C. Cremers, M. Horvat, S. Scott, and T. v. Merwe. Automated analysis and verification of TLS 1.3: 0-RTT, resumption and delayed authentication. In *2016 IEEE Symposium on Security and Privacy*, pages 470–485, 2016.
- [20] Cas Cremers, Marko Horvat, Jonathan Hoyland, Sam Scott, and Thyla van der Merwe. A comprehensive symbolic analysis of TLS 1.3. In *2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1773–1788. ACM, 2017.
- [21] Britt Cyr, Jeremy Dorfman, Ryan Hamilton, Jana Iyengar, Fedor Kouranov, Charles Krasic, Jo Kulik, Adam Langley, Jim Roskind, Robbie Shade, et al. QUIC wire layout specification. [https://docs.google.com/document/d/1WJvyZf1A02pq77y0Lbp9NsGjC1CHetAXV8I0fQe-B\\_U/edit](https://docs.google.com/document/d/1WJvyZf1A02pq77y0Lbp9NsGjC1CHetAXV8I0fQe-B_U/edit), 2016.
- [22] Antoine Delignat-Lavaud, Cédric Fournet, Markulf Kohlweiss, Jonathan Protzenko, Aseem Rastogi, Nikhil Swamy, Santiago Zanella Béguelin, Karthikeyan Bhargavan, Jianyang Pan, and Jean Karim Zinzindohoue. Implementing and proving the TLS 1.3 record layer. In *2017 IEEE Symposium on Security and Privacy*, pages 463–482. IEEE Computer Society, 2017.
- [23] Antoine Delignat-Lavaud, Cédric Fournet, Bryan Parno, Jonathan Protzenko, Tahina Ramananandro, Jay Bosamiya, Joseph Lallemand, Itsaka Rakotonirina, and Yi Zhou. A security model and fully verified implementation for the ietf quic record layer.
- [24] David Derler, Tibor Jager, Daniel Slamanig, and Christoph Striecks. Bloom filter encryption and applications to efficient forward-secret 0-RTT key exchange. In *EUROCRYPT 2018*, pages 425–455. Springer, 2018.
- [25] Benjamin Dowling, Marc Fischlin, Felix Günther, and Douglas Stebila. A cryptographic analysis of the TLS 1.3 handshake protocol candidates. In *2015 ACM SIGSAC Conference on Computer and Communications Security*, pages 1197–1210. ACM, 2015.
- [26] Benjamin Dowling, Marc Fischlin, Felix Günther, and Douglas Stebila. A cryptographic analysis of the TLS 1.3 draft-10 full and pre-shared key handshake protocol. Cryptology ePrint Archive, Report 2016/081, 2016. <https://eprint.iacr.org/2016/081>.
- [27] Benjamin Dowling, Marc Fischlin, Felix Günther, and Douglas Stebila. A cryptographic analysis of the TLS 1.3 handshake protocol. *Journal of Cryptology*, 2020.
- [28] M. Fischlin and F. Günther. Replay attacks on zero round-trip time: The case of the TLS 1.3 handshake candidates. In *2017 IEEE European Symposium on Security and Privacy*, pages 60–75. IEEE, 2017.
- [29] M. Fischlin, F. Günther, G. Azzurra Marson, and K. G Paterson. Data is a stream: Security of stream-based channels. In *CRYPTO 2015*, pages 545–564. Springer, 2015.
- [30] Marc Fischlin and Felix Günther. Multi-stage key exchange and the case of Google’s QUIC protocol. In *2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 1193–1204. ACM, 2014.

- [31] Gennie Gebhart. Tipping the scales on HTTPS: 2017 in review. <https://www.eff.org/deeplinks/2017/12/tipping-scales-https>, 2017.
- [32] Yossi Gilad and Amir Herzberg. Off-path attacking the web. In *WOOT 2012*, pages 41–52, 2012.
- [33] Fernando Gont. Security assessment of the Transmission Control Protocol. Technical Report CPNI Technical Note 3/2009, Centre for the Protection of National Infrastructure, 2009.
- [34] Felix Günther, Britta Hale, Tibor Jager, and Sebastian Lauer. 0-RTT key exchange with full forward secrecy. In *EUROCRYPT 2017*, pages 519–548. Springer, 2017.
- [35] Felix Günther and Sogol Mazaheri. A formal treatment of multi-key channels. In *CRYPTO 2017*, pages 587–618. Springer, 2017.
- [36] J. Iyengar and M. Thomson. QUIC: A UDP-based multiplexed and secure transport. <https://quicwg.org/base-drafts/draft-ietf-quic-transport.html>. Accessed: 2020-10-22.
- [37] Tibor Jager, Florian Kohlar, Sven Schäge, and Jörg Schwenk. On the security of TLS-DHE in the standard model. In *CRYPTO 2012*, pages 273–293. Springer, 2012.
- [38] S. Jero, H. Lee, and C. Nita-Rotaru. Leveraging state information for automated attack discovery in transport protocol implementations. In *IEEE/IFIP International Conference on Dependable Systems and Networks*, pages 1–12, 2015.
- [39] Samuel Jero, Endadul Hoque, David Choffnes, Alan Mislove, and C. Nita-Rotaru. Automated attack discovery in TCP congestion control using a model-guided approach. In *Network and Distributed Systems Security Symposium (NDSS)*, 2018.
- [40] Laurent Joncheray. A simple active attack against TCP. In *USENIX Security Symposium*, 1995.
- [41] Tadayoshi Kohno, Adriana Palacio, and John Black. Building secure cryptographic transforms, or how to encrypt and mac. 2003. <https://eprint.iacr.org/2003/177>.
- [42] Hugo Krawczyk, Kenneth G Paterson, and Hoeteck Wee. On the security of the TLS protocol: A systematic analysis. In *CRYPTO 2013*, pages 429–448. Springer, 2013.
- [43] Hugo Krawczyk and Hoeteck Wee. The OPTLS protocol and TLS 1.3. In *2016 IEEE European Symposium on Security and Privacy*, pages 81–96. IEEE, 2016.
- [44] V. A. Kumar, P. S. Jayalekshmy, G. K. Patra, and R. P. Thangavelu. On remote exploitation of TCP sender for low-rate flooding denial-of-service attack. *IEEE Communications Letters*, 13(1):46–48, 2009.
- [45] Aleksandar Kuzmanovic and Edward Knightly. Low-rate TCP-targeted denial of service attacks and counter strategies. *IEEE/ACM Transactions on Networking*, 14(4):683–696, 2006.
- [46] Adam Langley and Wan-Teh Chang. QUIC crypto. [https://docs.google.com/document/d/1g5nIXAIkN\\_Y-7XJW5K45Ib1Hd\\_L2f5LTaDUDwvZ5L6g/edit](https://docs.google.com/document/d/1g5nIXAIkN_Y-7XJW5K45Ib1Hd_L2f5LTaDUDwvZ5L6g/edit), 2016.
- [47] Xinyu Li, Jing Xu, Zhenfeng Zhang, Dengguo Feng, and Honggang Hu. Multiple handshakes security of TLS 1.3 candidates. In *2016 IEEE Symposium on Security and Privacy*, pages 486–505. IEEE, 2016.
- [48] Greg Linden. Make data useful. <https://sites.google.com/site/glinden/Home/StanfordDataMining.2006-11-29.ppt>, 2006.
- [49] R. Lychev, S. Jero, A. Boldyreva, and C. Nita-Rotaru. How secure and quick is QUIC? provable security and performance analyses. In *2015 IEEE Symposium on Security and Privacy*, pages 214–231, 2015.
- [50] David A McGrew and John Viega. The security and performance of the Galois/Counter Mode (GCM) of operation. In *International Conference on Cryptology in India*, pages 343–355. Springer, 2004.
- [51] Alfred Menezes and Berkant Ustaoglu. Comparing the pre-and post-specified peer models for key agreement. In *Australasian Conference on Information Security and Privacy*, pages 53–68. Springer, 2008.
- [52] Robert Morris. A weakness in the 4.2 BSD Unix TCP/IP software. Technical report, AT&T Bell Laboratories, 1985.
- [53] Kenneth G Paterson, Thomas Ristenpart, and Thomas Shrimpton. Tag size does matter: Attacks and proofs for the TLS record protocol. In *EUROCRYPT 2011*, pages 372–389. Springer, 2011.
- [54] Christopher Patton and Thomas Shrimpton. Partially specified channels: The TLS 1.3 record layer without elision. In *2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 1415–1428. ACM, 2018.

- [55] Jon Postel. User Datagram Protocol. RFC 768 (Standard), 1980.
- [56] Jon Postel. Transmission Control Protocol. RFC 793 (Standard), 1981.
- [57] Zhiyun Qian and Z. Morley Mao. Off-path TCP sequence number inference attack - how firewall middleboxes reduce security. In *2012 IEEE Symposium on Security and Privacy*, pages 347–361, 2012.
- [58] Zhiyun Qian, Z. Morley Mao, and Yinglian Xie. Collaborative TCP sequence number inference attack: how to crack sequence number under a second. In *2012 ACM SIGSAC Conference on Computer and Communications Security*, pages 593–604. ACM, 2012.
- [59] S. Radhakrishnan, Y. Cheng, J. Chu, A. Jain, and B. Raghavan. TCP Fast Open. In *COncference on emerging Networking EXperiments and Technologies*, page 21. ACM, 2011.
- [60] E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446, August 2018.
- [61] Eric Rescorla and Tim Dierks. The Transport Layer Security (TLS) Protocol Version 1.2. RFC 5246, August 2008.
- [62] Phillip Rogaway. Authenticated-encryption with associated-data. In *2002 ACM SIGSAC Conference on Computer and Communications Security*, pages 98–107. ACM, 2002.
- [63] Jim Roskind. QUIC: Design document and specification rationale. [https://docs.google.com/document/d/1RNHkx\\_VvKWYwg6Lr8SZ-saqsQx7rFV-ev2jRFUoVD34/edit](https://docs.google.com/document/d/1RNHkx_VvKWYwg6Lr8SZ-saqsQx7rFV-ev2jRFUoVD34/edit), 2013.
- [64] Stefan Savage, Neal Cardwell, David Wetherall, and Tom Anderson. TCP congestion control with a misbehaving receiver. *ACM SIGCOMM Computer Communication Review*, 29(5), 1999.
- [65] Jörg Schwenk. Modelling time for authenticated key exchange protocols. In *European Symposium on Research in Computer Security*, pages 277–294. Springer, 2014.
- [66] Victor Shoup. Sequences of games: a tool for taming complexity in security proofs. *IACR Cryptol. ePrint Arch.*, 2004:332, 2004.
- [67] Ahren Studer and Adrian Perrig. The Coremelt attack. In *European Symposium on Research in Computer Security*, pages 37–52, 2009.
- [68] Ian Swett. QUIC deployment experience @Google. <https://www.ietf.org/proceedings/96/slides/slides-96-quic-3.pdf>, 2016.
- [69] M. Thomson and S. Turner. Using Transport Layer Security (TLS) to secure QUIC. <https://quicwg.org/base-drafts/draft-ietf-quic-tls.html>. Accessed: 2020-10-22.
- [70] Verizon Enterprise Solutions. Monthly IP latency data — Verizon Enterprise Solutions. <http://www.verizonenterprise.com/about/network/latency/>. Accessed: 2020-10-22.
- [71] Paul Watson. Slipping in the window: TCP reset attacks. Technical report, CanSecWest, 2004.

## A Preliminary Definitions

### A.1 Pseudorandom Functions

For a function family  $F : \{0, 1\}^\lambda \times \{0, 1\}^n \rightarrow \{0, 1\}^m$ , consider the following security experiment associated with an adversary  $\mathcal{A}$ . In the beginning, sample a bit  $b \xleftarrow{\$} \{0, 1\}$ . If  $b = 0$ ,  $\mathcal{A}$  is given oracle access, i.e., can make queries, to  $F_k(\cdot) = F(k, \cdot)$  where  $k \xleftarrow{\$} \{0, 1\}^\lambda$ . If  $b = 1$ ,  $\mathcal{A}$  is given oracle access to  $f(\cdot)$  that maps elements from  $\{0, 1\}^n$  to  $\{0, 1\}^m$  uniformly at random. In the end,  $\mathcal{A}$  outputs a bit  $b'$  as a guess of  $b$ . The advantage of  $\mathcal{A}$  is defined as  $\mathbf{Adv}_F^{\text{prf}}(\mathcal{A}) = |\Pr[b' = 1|b = 0] - \Pr[b' = 1|b = 1]|$ , which measures  $\mathcal{A}$ 's ability to distinguish  $F_k$  (with random  $k$ ) from a random function  $f$ .

$F$  is a *pseudorandom function (PRF)* if:

- For any  $k \in \{0, 1\}^\lambda$  and  $x \in \{0, 1\}^n$ , there exists a polynomial-time (in  $\lambda$ ) algorithm to compute  $F(k, x)$ ; and
- For any efficient adversary  $\mathcal{A}$ ,  $\mathbf{Adv}_F^{\text{prf}}(\mathcal{A})$  is sufficiently small (e.g., roughly  $2^{-\lambda}$ ).

## A.2 Stateful Authenticated Encryption with Associated Data

We follow [41, 14] in extending the stateful authenticated encryption notion of Bellare *et al.* [8] to capture a hierarchy of stateful AEAD security notions based on different authentication levels. The following definitions are the same as [14], except that we exclude the length-hiding property proposed by Paterson *et al.* [53] for conciseness.

**Syntax.** A stateful AEAD scheme  $\text{sAEAD}$  is a 4-tuple  $(\text{sG}, \text{sl}, \text{sE}, \text{sD})$  associated with a message space  $\mathcal{M} \subseteq \{0, 1\}^*$ , an associated data space  $\mathcal{AD} \subseteq \{0, 1\}^*$ , and a state space  $\mathcal{ST} \subseteq \{0, 1\}^*$ .  $\text{sG}$  is a probabilistic algorithm that samples a random key  $k$  from a finite non-empty key space  $\mathcal{K}$ .  $\text{sl}$  is an algorithm that initializes the encryption and decryption states  $st_e, st_d$ .  $\text{sE}$  is a probabilistic encryption algorithm that takes as input  $k \in \mathcal{K}, ad \in \mathcal{AD}, m \in \mathcal{M}$  and  $st_e$  and outputs a ciphertext  $ct \in \{0, 1\}^*$  with an updated  $st_e$ .  $\text{sD}$  is a deterministic decryption algorithm that takes as input  $k \in \mathcal{K}, ad \in \mathcal{AD}, ct \in \{0, 1\}^*$  and  $st_d$  and outputs  $m \in \mathcal{M} \cup \{\perp\}$  with an updated  $st_d$ . The *correctness* requires that, for any  $k \in \mathcal{K}$  sampled by  $\text{sG}$ , any  $st_e = st_e^0, st_d = st_d^0$  initialized by  $\text{sl}$ , and any sequence of encryptions  $\{(ct_{i+1}, st_e^{i+1}) \stackrel{\$}{\leftarrow} \text{sE}(k, ad_i, m_i, st_e^i)\}_{i \geq 0}$ , the sequence of decryptions  $\{(m'_{i+1}, st_d^{i+1}) \leftarrow \text{sD}(k, ad, E(k, ad_i, ct_i, st_d^i))\}_{i \geq 0}$  satisfies  $m_i = m'_i, i \geq 0$ .

**Security.** Consider the following experiment with an authentication level  $al \in [4]$ . In the beginning, run  $\text{sG}$  to generate a key  $k$  and run  $\text{sl}$  to initialize  $st_e, st_d$ . Sample  $b \stackrel{\$}{\leftarrow} \{0, 1\}$  and set  $(u, v, \text{outofsync}) \leftarrow (0, 0, 0)$ . Then, the adversary  $\mathcal{A}$  is given access to the following oracles:

Enc $(ad, m_0, m_1)$ :

- 1:  $u \leftarrow u + 1, (sent.ct_u, st_e') \stackrel{\$}{\leftarrow} \text{sE}(k, ad, m_b, st_e)$
- 2:  $(sent.ad_u, st_e) \leftarrow (ad, st_e')$ , return  $sent.ct_u$

Dec $(ad, ct)$ :

- 1: if  $b = 0$ , return  $\perp$
- 2:  $v \leftarrow v + 1, (m, st'_d) \leftarrow \text{sD}(k, ad, ct, st_d)$
- 3:  $(rcvd.ad_v, st_d) \leftarrow (ad, st'_d)$
- 4: if  $(al = 4) \wedge \text{cond}_4$  or  $(al \leq 3) \wedge (m \neq \perp) \wedge \text{cond}_{al}$ ,<sup>21</sup>  
set  $\text{outofsync} \leftarrow 1$
- 5: if  $\text{outofsync} = 1$ , return  $m$ , otherwise, return  $\perp$

In the end,  $\mathcal{A}$  outputs a bit  $b'$ . The stateful AEAD scheme  $\text{sAEAD}$  is *secure* with authentication level  $al$  if and only if  $\text{Adv}_{\text{sAEAD}}^{\text{aead-}al}(\mathcal{A}) = |2 \Pr[b = b'] - 1|$  is sufficiently small (e.g., roughly  $2^{-\log |\mathcal{K}|}$ ) for any efficient adversary  $\mathcal{A}$ .

## B QUIC and TLS 1.3's Stateful AEAD Schemes and Their Security

### B.1 QUIC's Stateful AEAD Scheme and its Security

First, we show QUIC's stateful encryption scheme  $\text{sAEAD}_{\text{QUIC}}$  constructed from a nonce-based AEAD scheme  $\text{AEAD} = (\text{G}, \text{E}, \text{D})$  as follows.

sG $(\cdot)$ :

$k_e \stackrel{\$}{\leftarrow} \text{G}(), iv \stackrel{\$}{\leftarrow} \{0, 1\}^{32}$   
return  $(k_e, iv)$

sl $(\cdot)$ :

return  $(st_e, st_d) \leftarrow (\emptyset, \perp)$

sE $(k, ad, m, st_e)$ :

$(k_e, iv) \leftarrow k, (cid, pn) \leftarrow ad$   
if  $pn \in st_e$ : return  $(\perp, st_e)$   
 $c \leftarrow \text{E}(k_e, iv || pn, ad, m)$   
 $st_e \leftarrow st_e \cup \{pn\}$   
return  $(c, st_e)$

sD $(k, ad, ct, st_d)$ :

$(k_e, iv) \leftarrow k, (cid, pn) \leftarrow ad$   
 $m \leftarrow \text{D}(k_e, iv || pn, ad, ct)$   
return  $(m, \perp)$

<sup>21</sup>Authentication conditions  $\text{cond}_{al}$  are defined in the same way as in the msACCE-std Decrypt query.

Note that  $\text{sAEAD}_{\text{QUIC}}$  uses its encryption state to keep track of used nonces to avoid repeating and its decryption state is not used. To reduce its level-1 AEAD security to the underlying AEAD's nonce-based AEAD security, we first recall that the nonce-based AEAD security is defined as two separate parts, privacy and authenticity. For privacy, the adversary guesses the secret bit of a left-or-right encryption oracle but cannot make queries with a repeated nonce. The associated advantage is denoted by  $\text{Adv}_{\text{AEAD}}^{\text{ind-cpa}}(\mathcal{A})$ . For authenticity, the adversary tries to forge a valid ciphertext (together with a nonce and an associated data), given an encryption oracle (without the secret bit). The associated advantage is denoted by  $\text{Adv}_{\text{AEAD}}^{\text{int-ctxt}}(\mathcal{A})$ . Now, we are ready to prove the following theorem.

**Theorem 3** *For any efficient adversary  $\mathcal{A}$ , there exist efficient adversaries  $\mathcal{B}$  and  $\mathcal{C}$  such that:*

$$\text{Adv}_{\text{sAEAD}_{\text{QUIC}}}^{\text{aead-1}}(\mathcal{A}) \leq \text{Adv}_{\text{AEAD}}^{\text{int-ctxt}}(\mathcal{B}) + \text{Adv}_{\text{AEAD}}^{\text{ind-cpa}}(\mathcal{C}).$$

**Proof:** Consider two games (i.e., experiments)  $G_0$  and  $G_1$  and let  $\text{Pr}_i$  denote the winning probability of  $\mathcal{A}$  in  $G_i$ .  $G_0$  is the level-1 AEAD experiment and  $G_1$  is the same as  $G_0$  except that it always returns  $\perp$  for Dec queries. Following the game-hopping technique as illustrated in [66],  $|\text{Pr}_0 - \text{Pr}_1|$  is bounded by the probability that  $\mathcal{A}$  forges a new valid ciphertext given  $b = 1$ , which is by definition bounded by  $\text{Adv}_{\text{AEAD}}^{\text{int-ctxt}}(\mathcal{B})$  for some efficient adversary  $\mathcal{B}$ . Then, note that according to the  $\text{sAEAD}_{\text{QUIC}}$  construction nonces in AEAD encryption queries never repeat and  $G_1$  can be simulated by an efficient adversary  $\mathcal{C}$  against the nonce-based AEAD privacy security, which implies  $\text{Pr}_1 \leq \text{Adv}_{\text{AEAD}}^{\text{ind-cpa}}(\mathcal{C})$ .

Therefore, by a union bound, we have  $\text{Adv}_{\text{sAEAD}_{\text{QUIC}}}^{\text{aead-1}}(\mathcal{A}) \leq \text{Adv}_{\text{AEAD}}^{\text{int-ctxt}}(\mathcal{B}) + \text{Adv}_{\text{AEAD}}^{\text{ind-cpa}}(\mathcal{C})$ .  $\square$

REMARK. It is not hard to see that level-1 AEAD security is the best  $\text{sAEAD}_{\text{QUIC}}$  can achieve. Consider an adversary  $\mathcal{A}$  against the level- $al$  ( $al > 1$ ) AEAD security of  $\text{sAEAD}_{\text{QUIC}}$ .  $\mathcal{A}$  can easily set  $\text{outofsync} \leftarrow 1$  by querying Dec twice with the same ciphertext output by a previous Enc query, since  $\text{sAEAD}_{\text{QUIC}}$  does not prevent replays.

## B.2 TLS 1.3's Stateful AEAD Scheme and its Security

Next, we show TLS 1.3's stateful encryption scheme  $\text{sAEAD}_{\text{TLS}}$  constructed from a nonce-based AEAD scheme  $\text{AEAD} = (G, E, D)$  as follows (where  $n$  denotes the nonce length):

$\begin{array}{l} \text{sG}(): \\ k_e \xleftarrow{\$} G(), iv \xleftarrow{\$} \{0, 1\}^n \\ \text{return } (k_e, iv) \end{array}$	$\begin{array}{l} \text{sE}(k, ad, m, st_e): \\ (k_e, iv) \leftarrow k \\ c \leftarrow E(k_e, iv \oplus st_e, ad, m) \\ st_e \leftarrow st_e + 1 \\ \text{return } (c, st_e) \end{array}$	$\begin{array}{l} \text{sD}(k, ad, ct, st_d): \\ \text{if } st_d = \perp: \text{return } (\perp, \perp) \\ (k_e, iv) \leftarrow k \\ m \leftarrow D(k_e, iv \oplus st_d, ad, ct) \\ \text{if } m = \perp: st_d \leftarrow \perp \\ \text{otherwise: } st_d \leftarrow st_d + 1 \\ \text{return } (m, st_d) \end{array}$
$\begin{array}{l} \text{sl}(): \\ \text{return } (st_e, st_d) \leftarrow (0^{64}, 0^{64}) \end{array}$		

Note that in the above TLS 1.3's stateful encryption scheme, nonce repeating is prevented by the 64-bit increasing counter kept by the encryption state  $st_e$ . Following an argument very similar to the above proof of Theorem 3, one can show that the level-4 AEAD security of  $\text{sAEAD}_{\text{TLS}}$  is also reduced to the nonce-based AEAD security of  $\text{AEAD}$ . This result has been proved by previous work (Theorem 3 in [22]), but their stateful AEAD security definition is slightly different from ours. For instance, in their game the adversary needs to distinguish ciphertexts from random, while in our game the adversary distinguishes ciphertexts of two messages.

## C TFO+TLS 1.3's msACCE-std Security

Due to the high similarity among the abundant TLS 1.3 proofs in the MSKE model (and its extensions) and a security proof in our msACCE-std model, we only provide a proof sketch below.

A recent work [27] proved that the TLS 1.3 (EC)DHE 2-RTT full handshake and PSK-(EC)DHE 0-RTT resumption handshake are secure in the MSKE model based on the collision resistance of the hash function, unforgeability of the signature and MAC schemes, PRF security of the key derivation functions, and pseudorandom function oracle Diffie-Hellman (PRF-ODH) assumption [37, 42, 16]. Their MSKE security, which captures only the key exchange phases, ensures the Bellare-Rogaway-style key secrecy [10] (i.e., the stage keys are indistinguishable from random ones) with various authentication properties (for which our msACCE-std model focuses on the unilateral server authentication). In order to derive the overall TLS 1.3 security, the stage keys established by the handshake should be *composable*, i.e., safely used in any symmetric key protocol (e.g., the TLS 1.3 record protocol). However, as stated in [27], this generic composition result only works for stage keys that are external and non-replayable. In particular, it does not apply to the final session keys of the TLS 1.3 full handshakes or the interim handshake keys, which are used internally within the handshakes; it does not apply to the 0-RTT keys either, which are replayable. In order to adjust their security results to prove TLS 1.3’s Server Authentication and level-4 Channel Security in our model, we need to address the model differences as follows.

First, based on the above TLS 1.3 MSKE security, one can adapt the security results in [47] to derive the Multi-Level&Stage security of the combination of the TLS 1.3 full handshake and 0-RTT resumption handshake. Referring to their notions, our msACCE-std model focuses only on two modes (i.e., the TLS 1.3 (EC)DHE full handshake and PSK-(EC)DHE 0-RTT resumption handshake) and two levels (i.e., one level of full handshakes followed by one level of 0-RTT resumption handshakes). Note that [47] essentially treated a session ticket (carried within `NewSessionTicket`) as an opaque PSK identifier that leaks no information about the PSK used for session resumption. However, TLS 1.3 session tickets can also be implemented as a self-encrypted and self-authenticated value; in this case, its Multi-Level&Stage security further relies on the authenticated encryption security of the underlying ticket generation scheme. Furthermore, [47] did not differentiate between PSK and the resumption master secret (RMS) used to derive it, where RMS was established in a full handshake; but this is easy to fix by resorting to the PRF security of the key derivation function.

Then, we outline how the above TLS 1.3’s Multi-Level&Stage security can be augmented to prove TFO+TLS 1.3’s Server Authentication and level-4 Channel Security in our msACCE-std model.

1) The above Multi-Level&Stage security guarantees server authentication, i.e., a client oracle that has set its final session key must share the same session identifier with a unique partner server oracle. However, their session identifier is defined as *unencrypted* key exchange messages in order to capture key independence (i.e., revealing independent stage keys in the same session does not break the unrevealed stage key’s secrecy). We instead use a “real” encrypted session identifier to simplify our model and facilitate the reduction from KE Payload Integrity to Server Authentication. (Note that an unencrypted session identifier may correspond to many valid encrypted session identifiers but KE Payload Integrity requires no modification in the encrypted payload). To prove Server Authentication, we need to follow the proof of TLS 1.3’s Multi-Level&Stage server authentication to replace stage keys with independent random values, then use `sAEADTLS`’s AEAD oracles to simulate encrypted key exchange messages in `sidTLS` and the decryption of them. In this way, TFO+TLS 1.3’s Server Authentication can be reduced to its Multi-Level&Stage security and the AEAD security of `sAEADTLS`.

2) To prove TFO+TLS 1.3’s level-4 Channel Security, we follow the proof of TLS 1.3’s Multi-Level&Stage security to replace all stage keys with independent random values and then use `sAEADTLS`’s AEAD oracles to simulate encrypted key exchange messages and `Encrypt`, `Decrypt` queries. The simulation is perfect except that time periods are not captured by the MSKE-type models. However, by the definition of Channel Security, advancement of time periods affects only non-forward-secret 0-RTT keys. In particular, `Encrypt` queries related to 0-RTT keys can be allowed even if the server is corrupted, but the corruption must occur after the server gets reconfigured (via a `NextTP` query); this is fine because the server reconfiguration of TFO+TLS 1.3 erases or refreshes the static state (resumption states or the session ticket encryption key) used to recover RMS that derives the 0-RTT key. Therefore, TFO+TLS 1.3’s level-4 Channel Security can be reduced to TLS 1.3’s Multi-Level&Stage

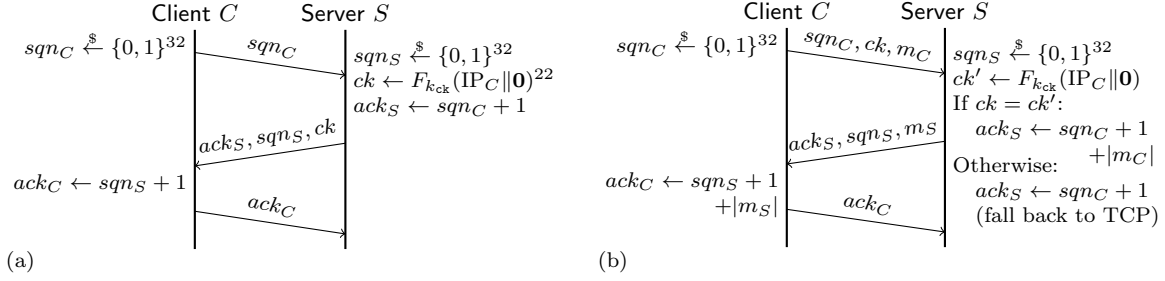


Figure 6: TFO initial connection (a) and 0-RTT resumption connection (b).

security and the level-4 AEAD security of  $\text{sAEAD}_{\text{TLS}}$ . Note that the AEAD oracles are also used to simulate (encrypted) post-handshake messages like `NewSessionTicket`. This bypasses the composition issue [26] faced by the MSKE model (and its extensions), where the application keys (which we call final session keys) in full handshakes cannot be composed with secure symmetric key protocols because those keys are used internally in the handshake to encrypt `NewSessionTicket` messages.

## D Proof of Theorem 1

The TFO initial connection and 0-RTT resumption connection are shown in Figure 6.

**Proof:** Consider a sequence of games (i.e., experiments) and let  $\Pr_i, i \geq 0$  denote the winning probability of  $\mathcal{A}$  in **Game**  $i$ .

**Game 0:** This is the original IP-Spoofing Prevention experiment, so  $\Pr_0 = \text{Adv}_{\text{TFO+TLS } 1.3}^{\text{ipsp}}(\mathcal{A})$ .

**Game 1:** This game is identical to Game 0 except the challenger first guesses the target server  $S$  (and the associated time period) that accepts a spurious connection request and aborts if the guess is wrong. Since the probability of a correct guess is at least  $1/(|\mathcal{S}| + q_T)$ , we have  $\Pr_0 \leq (|\mathcal{S}| + q_T) \Pr_1$ .

**Game 2:** This game replaces the PRF  $F$  used by  $S$  for TFO cookie generation with a truly random function  $f$ . By the PRF definition, it is straightforward to construct an efficient adversary  $\mathcal{B}$  such that  $|\Pr_1 - \Pr_2| \leq \text{Adv}_F^{\text{prf}}(\mathcal{B})$ .

Now, in Game 2, the TFO cookies generated by  $S$  are independent from each other for each client. We can bound  $\Pr_2$  by considering two cases. 1)  $\mathcal{A}$  wins by sending a valid ACK packet in a full session. In this case,  $\mathcal{A}$  must have generated a valid  $ack_C$  by correctly guessing the target server's TCP sequence number  $sqn_S$ . The winning probability of each guess is exactly  $1/2^{|\text{sqn}|}$ . 2)  $\mathcal{A}$  wins by sending a valid SYN packet in a resumption session. In this case,  $\mathcal{A}$  must have forged a valid TFO cookie  $ck \in \{0, 1\}^{64}$ . The winning probability of each forgery is exactly  $1/2^{|\text{ck}|}$  because the TFO cookie generation function is now a truly random function. By applying a union bound on the  $q_S$  queries and noting that  $|\text{sqn}| = 32 < 64 = |\text{ck}|$ , we have  $\Pr_2 \leq q_S / \min\{2^{|\text{sqn}|}, 2^{|\text{ck}|}\} = q_S / 2^{|\text{sqn}|}$ .  $\square$

## E TCP Attacks

*TCP Flow Control Manipulation.* An adversary with access to the communication channel can impact TCP's flow control mechanism to decrease the sending rate or stall the connection by modifying TCP's `window` header field. This field controls the amount of received data the sender of this packet is prepared to buffer. By reducing this quantity, the throughput of the connection can be reduced and if it is set to zero the connection will completely stall.

<sup>22</sup>The padding is not applied for a IPv6 address.

One example of this attack would be to modify the window field to zero in a TCP packet containing a TLS-encrypted HTTP request. Since TCP headers are not authenticated, this modification will not be detected. As a result, when the server receives this request and attempts to send the response, it will believe that the client cannot currently accept any data and will delay sending the response. After some timeout, TCP will probe the client with a single packet of data to determine whether the window is still zero. If the adversary also modifies the responses to these probes, the connection will remain stalled indefinitely; otherwise, the connection will eventually recover after a lengthy delay.

*TCP Acknowledgment Injection.* An adversary who can observe a target connection and forge packets can inject new acknowledgment packets into the TCP connection. Acknowledgment packets have no data making them undetectable by either TLS or the application. However, they are used by congestion control to determine the allowed sending rate of a connection.

Injecting duplicate or very slowly increasing acknowledgments can be used to slow a target connection down drastically. [39] demonstrated a 12x reduction in throughput using this approach with the attacker required to expend only 40Kbps. This, of course, represents a significant performance degradation for a TFO+TLS 1.3 connection.

Injecting acknowledgments can also be used to dramatically increase the sending rate of a connection, turning it into a firehose that an attacker can point at their desired target. This is done by sending acknowledgments for data that has not been received yet, an attack known as Optimistic Ack [64]. This attack renders TCP insensitive to congestion and can completely starve competing flows. It could be used with great effect to cause denial of service against a server or the Internet infrastructure as a whole [67].

## F UDP+QUIC’s msACCE-std Security

It has been proven in [49] that QUIC is QACCE-secure in the random oracle model based on the unforgeability of the signature scheme, computational Diffie-Hellman assumption [3], and nonce-based AEAD security. Note that msACCE-std with sAEAD<sub>QUIC</sub> is semantically equivalent to QACCE with nonce-based AEAD and `get_iv` (defined in [49]), except that our model formalizes the advancement of time periods that was assumed by QACCE. More precisely, our model captures server-side local times (without time synchronization) while QACCE assumed a global one, but this difference does not affect the security results because QACCE essentially does not rely on time synchronization either. In other words, the same proofs in [49] can also show that QUIC is secure in the modified QACCE model that instead assumes local server times. As a result, their proofs can be easily adapted to show that UDP+QUIC achieves Server Authentication and level-1 Channel Security in our msACCE-std model. Note that one can also prove UDP+QUIC’s msACCE-std security by relying on the level-1 AEAD security of sAEAD<sub>QUIC</sub> instead of the underlying nonce-based AEAD security, where the former can be reduced to the latter as shown in Appendix B.

## G Proof of Theorem 2

**Proof:** Consider a sequence of games (i.e., experiments) and let  $\Pr_i, i \geq 0$  denote the winning probability of  $\mathcal{A}$  in **Game**  $i$ .

**Game 0:** This is the original Reset Authentication experiment, so  $\Pr_0 = \mathbf{Adv}_{\text{UDP+QUIC[TLS]}}^{\text{rst-auth}}(\mathcal{A})$ .

**Game 1:** This game is the same as Game 0 except it aborts if the connection IDs repeat for a party. Since the probability of such a collision for each party is at most  $N^2/2^{|\text{cid}|}$ , by a union bound we have  $|\Pr_0 - \Pr_1| \leq |\mathcal{P}|N^2/2^{|\text{cid}|}$ .

**Game 2:** The challenger proceeds as before except it first guesses the target party  $P$  (recall that  $\mathcal{A}$  outputs  $(P, i)$  in the end) and aborts if the guess is wrong. Since the probability of a correct guess is at least  $1/|\mathcal{P}|$ , we have  $\Pr_1 \leq |\mathcal{P}|\Pr_2$ .



**Game 3:** This game replaces the PRF  $F$  used by  $P$  for reset token generation with a truly random function  $f$ . By the PRF definition, it is straightforward to construct an efficient adversary  $\mathcal{B}$  such that  $|\Pr_2 - \Pr_3| \leq \mathbf{Adv}_F^{\text{prf}}(\mathcal{B})$ .

**Game 4:** This game replaces the pre-reset packet (if any) output by the target session oracle  $\pi_P^i$  (i.e., output by  $\text{Pack}(\pi_P^i, \cdot, \text{prst})$ ) with encryption of an arbitrarily fixed pre-reset message  $m_{\text{prst}}^* \in \mathcal{M}_{\text{prst}}$ . As illustrated below, there exists an efficient adversary  $\mathcal{C}$  against the level-1 Channel Security of UDP+QUIC[TLS] such that  $|\Pr_3 - \Pr_4| \leq \mathbf{Adv}_{\text{UDP+QUIC[TLS]}}^{\text{cs-1}}(\mathcal{C})$ .

$\mathcal{C}$  samples the static reset token generation keys for all parties and simulate  $\text{Pack}$  and  $\text{Deliver}$  queries with its  $\text{Encrypt}$  and  $\text{Decrypt}$  queries. In particular, for a  $\text{Pack}(\cdot, \cdot, \text{prst})$  query made to a session oracle of some party other than the target party  $P$ ,  $\mathcal{C}$  generates the reset token with  $F_k(\text{cid}) = \text{rtk}$  (where  $k$  is the token generation key of the specified party and  $\text{cid}$  is the connection ID used by the specified session oracle) and forms the pre-reset message  $m_{\text{prst}}$  that carries that token  $F_k(\text{cid})$ , then queries  $\text{Encrypt}$  with  $(m_{\text{prst}}, m_{\text{prst}})$  as the challenge message pair and uses the output ciphertext to form a pre-reset packet sent to  $\mathcal{A}$ . For a  $\text{Pack}(\cdot, \cdot, \text{prst})$  query made to a session oracle of the target party  $P$  but not the target session  $i$ , the token is instead generated with  $f(\text{cid}) \xleftarrow{\$} \{0, 1\}^{|\text{rtk}|}$  via lazy sampling and the rest is the same. For a  $\text{Pack}(\pi_P^i, \cdot, \text{prst})$  query (i.e., to the target session oracle),  $\mathcal{C}$  generates a random reset token  $f(\text{cid})$  (if  $f(\text{cid})$  is not yet set) and forms a pre-reset message  $m_{\text{prst}}$  that carries that token  $f(\text{cid})$ , then queries  $\text{Encrypt}$  with  $(m_{\text{prst}}, m_{\text{prst}}^*)$  as the challenge message pair and uses the output ciphertext to form a pre-reset packet sent to  $\mathcal{A}$ . The  $\text{Pack}(\cdot, \cdot, \text{rst})$  queries can be simulated in a similar way, except that no  $\text{Encrypt}$  query is needed and the reset token is used to form a reset packet sent to  $\mathcal{A}$ . When receiving a reset packet via a  $\text{Deliver}$  query,  $\mathcal{C}$  checks if the reset token carried in this packet matches the previously received reset token carried in the pre-reset packet and then accepts it if and only if the check passes. It is not hard to see that  $\mathcal{C}$  can perfectly simulate Game 3 in the left world ( $b_P^i = 0$ ) or Game 4 in the right world ( $b_P^i = 1$ ).  $\mathcal{C}$  outputs 1 if and only if  $\mathcal{A}$  wins.

Now, in Game 4, the random pre-reset message of the target  $\pi_P^i$  is independent from  $\mathcal{A}$ 's view and each guess is correct with probability  $1/2^{|\text{rtk}|}$ . By a union bound, we have  $\Pr_4 \leq q_D/2^{|\text{rtk}|}$ .  $\square$