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

Shan Chen[1], Samuel Jero[2], Matthew Jagielski[3], Alexandra Boldyreva[1], and Cristina Nita-Rotaru[3]

[1]Georgia Institute of Technology
shanchen@gatech.edu sasha@gatech.edu
[2]Purdue University
sjero@sjero.net
[3]Northeastern University
jagielski.m@husky.neu.edu c.nitarotaru@neu.edu

Abstract

Secure channel establishment protocols such as TLS are some of the most important cryptographic protocols, enabling the encryption of Internet traffic. Reducing the latency (the number of interactions between parties) in such protocols has become an important design goal to improve user experience. The most important protocols addressing this goal are the just-released TLS 1.3, which is likely to see deployment in the near future, and QUIC, a secure transport protocol from Google that is available 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 privacy 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 crypto cores of the protocols. Equipped with our new models we provide a detailed comparison of 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 provably 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 results will help protocol designers in their future protocol analyses and that our results will help practitioners better understand the advantages and limitations of novel secure channel establishment protocols.

1 Introduction

MOTIVATION. Nowadays, more than half of all Internet traffic is encrypted according to a 2017 EFF report [24], with Google reporting that 93% of its traffic is encrypted as of January 2019 [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 TLS. The session key establishment with TLS today involves 3 round-trip times (RTTs) of end-to-end communication, including the cost 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 [40], while a typical RTT on a connection from New York to London is 70ms [26].

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 [51] (the just-released successor to the current TLS 1.2 standard) and Google's QUIC [53].

With TLS 1.3, it is possible to reduce the number of RTTs (prior to sending encrypted data) during session resumption to 1. This reduction 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. The remaining 1-RTT during session resumption is due to the aforementioned TCP connection. However, one recent optimization for TCP, called TCP Fast Open (TFO) [50, 14] 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.

Like TLS 1.3, Google's QUIC uses weaker initial keys, under which data can be encrypted earlier, and a token saved from previous communication between the parties. But unlike TLS, QUIC operates over UDP 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.

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
TFO+TLS 1.3	2-RTT	**0-RTT**
UDP+QUIC	1-RTT	**0-RTT**
UDP+QUIC[TLS]	1-RTT	**0-RTT**

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 the new design for QUIC [27] (which we refer to as QUIC[TLS] [57] to indicate that it borrows the key exchange from TLS 1.3) over UDP. It is clear that the last three win in terms of the number of interactions. But how does their security compare?

At first glance, the question is easy to answer. Recent works have done formal security analyses of TLS 1.3 [34, 8, 18, 15, 35, 19, 39, 22, 17, 7, 16, 10] and Google's QUIC [21, 41]. Most works confirm that (the cryptographic cores of) both protocols are provably secure under reasonable computational assumptions. Moreover, as shown in [41, 22], 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, the 0-RTT keys do not provide forward secrecy and the 0-RTT data suffers from replay attacks. 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 [41] analyzing the IP spoofing (source validation) of QUIC, did not consider packet-level availability attacks. 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, to the best of our knowledge, 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 data transmission functionality is [41], we take their *Quick Connections (QC)* protocol definition and *Quick Authenticated and Confidential Channel Establishment (QACCE)* security model as our starting point.

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 set during each session. Then, we extend the *Quick Authenticated and Confidential Channel Establishment (QACCE)* security model [41] 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, and the latter is novel and is for packet-level security.

Like most security models, we consider a very powerful attacker who can initiate communications between honest parties, can intercept, inject, drop, or modify the exchanged packets, and can adaptively learn parties' stage keys or adaptively corrupt them to learn their long-term keys and secret states. The attacker can also have prior knowledge of the exchanged data. However, the attacker should not be able to prevent clients from establishing final session keys without noticing the attacker's involvement (Server Authentication) or using these keys to achieve a secure channel with data privacy and integrity (Channel Security). These standard security goals are captured by our first model.

For the second model that deals with packet-level availability attacks, we first follow QACCE [41] to consider IP-spoofing prevention (also known as address validation) and further extend it to additionally capture 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 authenticitiy 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 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 [41], 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. This makes such attacks more harmful and their treatment more important. Finally, we formalize the new goal of Reset Authentication to deal with attacks forging a reset packet to abruptly terminate an honest party's session.

3

Equipped with our new models, we study the security and availability functionalities provided by TFO+TLS 1.3, UDP+QUIC, and UDP+QUIC[TLS]. 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 achieves IP-spoofing prevention and we show that this extends to our stronger notion. As for TFO+TLS 1.3, its IP-spoofing prevention 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 block cipher is a pseudorandom function.

Regarding SC Header Integrity, we show that while UDP+QUIC is 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 [54, 31, 4, 13, 37, 36, 29, 49, 12, 25, 48, 43, 58, 30], such as TCP flow control manipulation, TCP acknowledgement injection, etc.

We next show that neither protocol 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. 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 [41], ServerReject Triggering is harder to detect and more harmful in this sense. We show that TFO+TLS 1.3, on the other hand, achieves KE Payload Integrity.

We further show that neither TFO+TLS 1.3 nor UDP+QUIC provide Reset Authentication, justifying the TCP Reset attack [58] relevant for TFO+TLS 1.3 and the PublicReset attack for UDP+QUIC. For completeness, we recall the results from [41, 22] showing that neither protocol provides forward secrecy for the keys encrypting 0-RTT data and that this data can be replayed.

We finally show that the new UDP+QUIC[TLS] protocol achieves the strongest security of 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, while 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 (based on the unpredictability of its reset tokens).

Our results are summarized in Table 2 in Section 5. Even though QUIC 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 novel 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, and QUIC in Section 2. Sections 3 and B 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 summarizes our findings about how security of the three layered protocols compare. 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 and TCP with the Fast Open optimization (TFO); and application, represented by TLS or QUIC.

2.1 TLS 1.3 over TFO

TCP Fast Open. TCP Fast Open (TFO) is an optimization to 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 Fig. 5 in Appendix A, 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 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 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 Fig. 1, 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 0, 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 Fig. 2), 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 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 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 [51] 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 Fig. 1, fourth 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
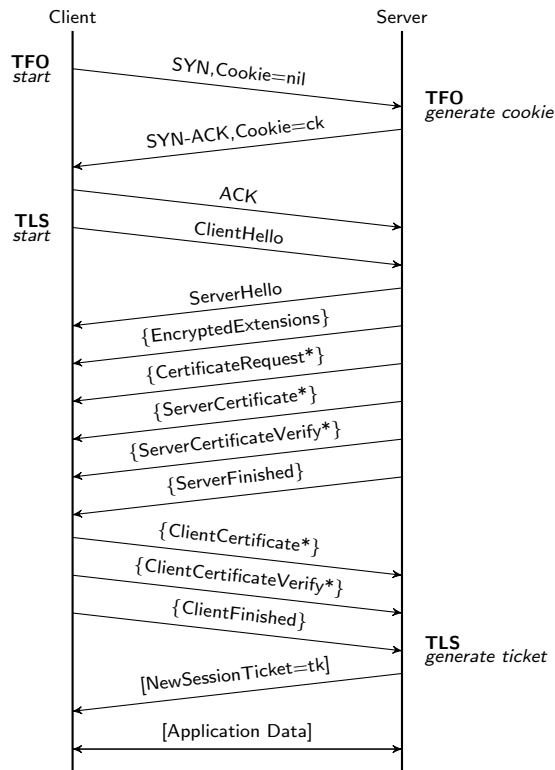
Figure 1: TFO+TLS 1.3 (EC)DHE 2-RTT full handshake. * indicates optional messages. {} and [] respectively indicate messages protected with initial and final keys.

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 mes-

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

sages. 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 TCP Fast Open 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 Fig. 1), 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 connection and include the TLS 1.3 `ClientHello` message in the SYN packet. The TLS `ClientHello` message can use the cached TLS session ticket to perform a 0-RTT resumption handshake. Thus, the TCP and TLS 1.3 connections are established at the same time, as shown in Fig. 2.

2.2 QUIC over UDP

UDP. UDP [46] is an extremely simple transport protocol providing unreliable datagram delivery, the 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 Fig. 6 in Appendix A) 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. Quick UDP Internet Connections (QUIC) is a transport protocol developed by Google and implemented by Chrome and Google servers since 2013 [53]. It now provides service for the majority of requests by Chrome to Google properties [56]. 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.
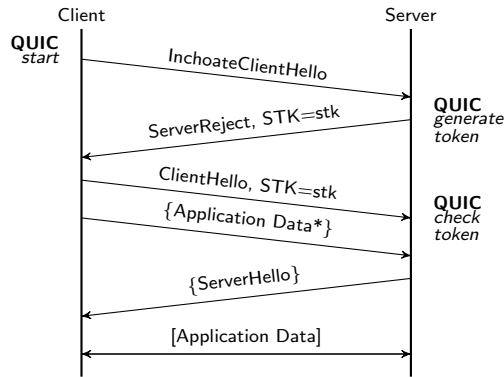
7

Figure 3: UDP+QUIC 1-RTT full handshake. * indicates optional messages. {} and [] repectively indicate messages protected with initial and final keys.


Figure 4: UDP+QUIC 0-RTT resumption handshake. * indicates optional messages. {} and [] repectively indicate messages protected with initial and final keys.

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 64bit connection identifier referred to as `cid`, and a variable length packet number. 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 Fig. 3, 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 which is valid and then sends a `ClientHello` indicating the `scfg` its 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 secure 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 the client's `ClientHello` which contains a newly generated Diffie-Hellman share. Once the client receives this message, client and server derive and begin using the new forward secure encryption key.

For the client that has connected to a server before, it can instead initiate a resumption connection. This consists of only the last two steps of a full connection, sending the `ClientHello` and `ServerHello` messages as shown in Fig. 4.

2.3 QUIC with TLS 1.3 Key Exchange over UDP

A new version of QUIC [27], 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 [57]. We provide more details (e.g., about its new stateless reset feature) in Section 5.

3 Preliminaries

Notations. Let $\{0,1\}^*$ denote the set of all finite-length binary strings (including the empty string ε) and $\{0,1\}^n$ denote the set of n-bit binary strings. $[n]$ denotes the set of integers $\{1, 2, \ldots, n\}$. For a finite set $\mathcal{R}$, let $|\mathcal{R}|$ denote its size and $r \xleftarrow{\$} \mathcal{R}$ 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. $y \xleftarrow{\$} F(x)$) 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 assume the public keys used in our analysis are supported by a *public key infrastructure (PKI)* and do not consider certificates or certificate checks explicitly. In other words, we assume each public key is certified and bound to the corresponding party's identity.

PRF and AEAD. In Appendix B 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). Accordingly, there we provide the definitions for the corresponding advantages: $\mathbf{Adv}_F^{\mathsf{prf}}(A), \mathbf{Adv}_{\mathsf{sAEAD}}^{\mathsf{aead}\text{-}al}(A)$. We also refer to [52] for the syntax and security definitions of a nonce-based AEAD scheme.

4 msACCE Protocol and its Security

In this section, we define the syntax and two security models for *Multi-Stage Authenticated and Confidential Channel Establishment (msACCE)* protocols.

9

4.1 Protocol Syntax

Our msACCE protocol is an extension to the *Quick Connection (QC)* protocol proposed by Lychev *et al.* [41] and the *Multi-Stage Key Exchange (MSKE)* protocol proposed by Fischlin and Günther [21] (and further developed by [18, 19, 39, 22]). Even though the authors of [41] 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 (explicitly, unlike QC), session resets, and header-only packets exchanged in secure channel phases. The detailed protocol syntax is defined below.

A msACCE protocol is an interactive protocol between a client and a server. They establish keys in one or more stages and exchange messages encrypted and decrypted with these keys. Messages are exchanged via *packets*. A packet consists of source and destination IP addresses[1] $\mathrm{IP_s}, \mathrm{IP_d} \in \{0,1\}^{32} \cup \{0,1\}^{64}$, a header, and a payload. Each party P has a unique IP address IP_P.

The protocol is associated with the security parameter $\lambda \in \mathbb{N}_+$, a key generation algorithm Kg that takes as input 1^λ and outputs a public and secret key pair, a header space[2] (for transport and application layers) $\mathcal{H} \subseteq \{0,1\}^*$, a payload space $\mathcal{PD} \subseteq \{0,1\}^*$, header and payload spaces $\mathcal{H}_{\mathsf{rst}} \subseteq \mathcal{H}, \mathcal{PD}_{\mathsf{rst}} \subseteq \mathcal{PD}$ for reset packets (described later), a resumption state space $\mathcal{RS} \subseteq \{0,1\}^*$, a stateful AEAD scheme[3] $\mathsf{sAEAD} = (\mathsf{sG}, \mathsf{sE}, \mathsf{sD})$ (with a key space $\mathcal{K} = \{0,1\}^\lambda$, 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*[4] message spaces $\mathcal{M}_{\mathsf{KE}}, \mathcal{M}_{\mathsf{SC}}, \mathcal{M}_{\mathsf{pRST}} \subseteq \mathcal{M}$ with $\mathcal{M}_{\mathsf{KE}}, \mathcal{M}_{\mathsf{SC}}$ for messages encrypted during key exchange and secure channel phases respectively and $\mathcal{M}_{\mathsf{pRST}}$ for pre-reset messages (described later) encrypted in a secure channel phase, and a server configuration generation function $\mathsf{scfg\_gen}$ described below.

The protocol's execution is associated with the universal notion of time divided into discrete periods $\tau_1, \tau_2, \ldots$. During its execution, both parties can keep states that are initialized to the empty string ε. In the beginning of each time period, the protocol may periodically update each server's configuration state $\mathtt{scfg}$ with $\mathsf{scfg\_gen}$ (which takes as input 1^λ, a server secret key, and a time period, then outputs a server configuration state). Otherwise, $\mathsf{scfg\_gen}$ is undefined and without loss of generality the protocol is executed within a single time period.

A *reset* packet enables a sender, who lost its session state due to some error condition (e.g., server reboots, denial-of-service attacks, etc.), to abruptly terminate a session with the receiver. A *pre-reset* message (e.g., a reset token in QUIC[TLS]) is sent to the receiver in a secure channel phase[5] before the sender loses its state in order to authenticate the sender's reset packet. Each session has *at most one* pre-reset message for each party. 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 *single* 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[6] $\mathcal{M}^{\mathsf{snd}} = (M_1, \ldots, M_l), M_i \in$

---

[1] 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.

[2] Some protocol header fields (e.g., port numbers, checksums, etc.) can be excluded if they are not the focus of the security analysis.

[3] To fit TLS 1.3's encryption scheme, unlike QACCE we model QUIC's encryption scheme as a more general stateful AEAD scheme rather than a nonce-based one.

[4] Disjointness is a reasonable assumption as practical protocols (such as those in Table 1) enforce different leading bits for different types of messages.

[5] 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.

[6] For simplicity, we consider transportation of *atomic* messages rather than a data *stream* that can be modeled as a

$\mathcal{M}_{\mathsf{SC}}, l \in \mathbb{N}$ (where the total message length $|\mathcal{M}^{\mathsf{snd}}|$ is polynomial in λ and $\mathcal{M}^{\mathsf{snd}}$ can be empty) as well as the other party's IP address. In a full session, the server takes as input its associated public and secret key pair (generated by running $\mathsf{Kg}(1^\lambda)$) and the client takes the server's public key as input. In a resumption session, each party additionally takes as input its own resumption state $rs \in \mathcal{RS}$ (set in the associated full session). In either case, the client sends the first packet to start the session.

A D-stage msACCE protocol consists of $D \in \mathbb{N}_+$ 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 can be $\perp$, i.e., unused.[7] If this is the final stage in a full session, each party can send additional messages[8] in $\mathcal{M}_{\mathsf{KE}}$ 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. The client uses k_c^d 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) $\mathcal{M}_i^{\mathsf{rcv}} = (M_1', \ldots, M_{l_i'}')$, $l_i' \in \mathbb{N}$, $M_i' \in \mathcal{M}_{\mathsf{SC}}$.

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 $\mathcal{M}^{\mathsf{snd}}$, in which the messages are sent among D stages according to any partitioning $\mathcal{M}^{\mathsf{snd}} = \mathcal{M}_1^{\mathsf{snd}}, \ldots, \mathcal{M}_D^{\mathsf{snd}}$, is equal to the other party's total output message list $\mathcal{M}^{\mathsf{rcv}} = \mathcal{M}_1^{\mathsf{rcv}}, \ldots, \mathcal{M}_D^{\mathsf{rcv}}$, in which the message order is preserved. Each party terminates its session upon receiving the other party's reset packet.

REMARK. With our more general protocol syntax, the ACCE [28] and QC [41] 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 packet authentication. Our models do not consider the key exchange and secure channel phases independently, as was the case for some previous QUIC and TLS 1.3 security analyses [21, 18, 19, 39, 22], 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 [41, 19].

1) msACCE Standard Security Model:

In this msACCE standard (msACCE-std) security model, we consider the standard security goals such as server authentication[9] and channel security (which captures data privacy and integrity) for msACCE protocols. Our msACCE-std model is very similar to the standard security portion of the QACCE model [41], but extends it to capture more (rather than two) stages and use a more general stateful encryption scheme to fit both TLS 1.3 and QUIC.

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.

---

stream-based channel [23] and later extended to capture multiplexing [45].

[7] This captures the case where a 0-RTT key only consists of a client encryption key while the server encryption key does not exist.

[8] This captures the post-handshake key exchange messages that are used for session resumption, post-handshake authentication, key update, etc.

[9] Our msACCE-std model focuses on the most common server authentication, but can be extended to mutual authentication, e.g., as described in [34].

Our 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 $\pi_P^1, \pi_P^2, \ldots$, where π_P^i models P executing a protocol instance in session $i \in \mathbb{N}_+$.

Matching Conversations. As part of the security model, *matching conversations* are used to model entity authentication, session key confirmation, and handshake integrity. A client (resp. server) oracle has a matching conversation with a server (resp. client) oracle if and only if both session oracles observe the same[10] *session identifier sid* defined according to the protocol specifications and security goals. Note that a msACCE protocol may have two different session identifiers in full and resumption modes, but for simplicity we use the same notation *sid*. Compared to the general definition of matching conversations [6, 28], *sid* is often defined as a *subset* of the whole communication transcript. For instance, QUIC's *sid* in QACCE [41] is defined as the second-round key exchange messages, i.e., `ClientHello` and `ServerHello`, while the first-round messages are excluded to allow for valid but different source-address tokens or signatures. Similarly, TLS 1.2's *sid* in ACCE [34] is defined as the first three key exchange messages, while the rest are excluded to allow for valid but different encrypted `Finished` messages.

Partners. We say a client oracle and a server oracle are each other's *partner* if they observe the same first-stage session identifier sid_1 (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 partners if sid_1 consists of only message(s) sent from the client oracle, which can be replayed to the same[11] server to establish multiple (identical) first-stage keys. Therefore, a session oracle's partner may not be its final unique communication partner. Instead, the real partner is the session oracle with which the oracle has a matching conversation.

Security Experiments. In the beginning of the experiments, run $\mathsf{Kg}(1^\lambda)$ for all servers to generate the public and secret key pairs and initialize the global states of all parties and the local states of all session oracles. In the beginning of each time period, run $\mathsf{scfg\_gen}$ (if defined) for each server to update its configuration state scfg. We assume that both the server oracles and the adversary A are aware of the current time period. 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 channel security experiment is associated with an authentication level $al \in [4]$. Each oracle π_P^i at stage d is associated with a random bit $b_P^{i,d} \xleftarrow{\$} \{0,1\}$. Each oracle π_P^i has a global state $\widetilde{m}$ (initialized to $\bot$) that stores its pre-reset message. The adversary A is given all public keys and the IP addresses associated with all parties and then interacts with the session oracles via the following queries:

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

This query asks π_C^i to output the first packet that it would send to S in a full session according to the protocol if π_C^i was not *used* (i.e., as input of previous $\mathsf{Connect}, \mathsf{Resume}, \mathsf{Send}$ queries). This output packet is returned to A instead of S. After this query, we say S is the *target server* of π_C^i.

This query allows the adversary to ask a specified client oracle to start a full session with a specified server.

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

This query asks π_C^i to output the first packet that it, taking $\pi_C^{i'}$'s resumption state as input, would send to S in a resumption session according to the protocol and returns this packet to A, if π_C^i was not

---

[10]As discussed in [28], two session oracles having matching conversations with each other may not observe the same transcript due to the gap between one oracle sending a message and the other receiving it. We can use *symmetric* session identifiers to define matching conversations because our msACCE-std model focuses only on server authentication and we require session identifiers to exclude, if any, a client oracle's last key exchange message(s) sent immediately before it sets its session key.

[11]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 will be rejected.

used and $\pi_C^{i'}$ has set its resumption state in a previous full session with its target server S. Otherwise, it returns $\perp$.

This query allows the adversary to ask a specified client oracle to start a resumption session with a specified server oracle to resume a specified full session between the two parties, if the associated previous client oracle has set its resumption state.

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

This query sends pkt to π_P^i and returns its response if π_P^i is in a key exchange phase, otherwise, returns $\perp$.

This query allows the adversary to send any packet to a specified session oracle and get its response in a key exchange phase.

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

This query returns π_P^i's (perhaps unset) stage-d key k^d. After this query, we say k^d was *revealed*.

This query allows the adversary to learn any stage key of a specified session oracle.

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

This query returns S's secret key and all its current states including its `scfg` and resumption states (for all full sessions involving S) in the current time period. After this query, we say S was *corrupted*.

This query allows the adversary to learn the long-term secret along with all current states of a specified server.

- Encrypt(π_P^i, d, ad, m_0, m_1), for $P \in \mathcal{P}, i \in [N], d \in [D], ad \in \mathcal{AD}, m_0, m_1 \in \mathcal{M}_{\mathsf{SC}} \cup \mathcal{M}_{\mathsf{pRST}} \cup \{\mathtt{rst}\}$.

This query proceeds as follows:

1: if m_0, m_1 are of different types (i.e., $\mathcal{M}_{\mathsf{SC}}$ or $\mathcal{M}_{\mathsf{pRST}}$ or $\{\mathtt{rst}\}$) or $|m_0| \neq |m_1|$ or π_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 $m_0 = m_1 = \mathtt{rst}$, return $\widetilde{m}$

3: if $m_0, m_1 \in \mathcal{M}_{\mathsf{pRST}}$, return $\perp$ if $\widetilde{m} \neq \perp$ or set $\widetilde{m} \leftarrow m_{b_P^{i,d}}$ otherwise

4: (upon setting each encryption stage key, initialize $st_e \in \mathcal{ST}, u \leftarrow 0, sent \leftarrow \varepsilon$)

5: $u \leftarrow u + 1, (sent.ct_u, st_e') \xleftarrow{\$} \mathsf{sE}(k_p^d, ad, m_{b_P^{i,d}}, st_e)$

6: $(sent.ad_u, st_e) \leftarrow (ad, st_e')$

7: return $sent.ct_u$

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 $b_P^{i,d}$, a random bit associated with the specified session oracle at the specified stage. This query also stores the pre-reset message and returns it when the input is $\mathtt{rst}$. (Recall that each oracle has at most one pre-reset message, so this query stores only the first pre-reset message and rejects others.)

- Decrypt(π_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 proceeds as follows:

1: if π_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 $st_d \in \mathcal{ST}, v \leftarrow 0, rcvd \leftarrow \varepsilon, \mathtt{outofsync} \leftarrow 0$)

3: $v \leftarrow v + 1, rcvd.ct_v \leftarrow ct, (m, st_d') \leftarrow \mathsf{sD}(k_p^d, ad, ct, st_d)$

4: $(rcvd.ad_v, st_d) \leftarrow (ad, st_d')$

5: if $m \notin \mathcal{M}_{\mathsf{SC}} \cup \mathcal{M}_{\mathsf{pRST}}$, set $m \leftarrow \perp$

6: if $(al = 4) \wedge \mathsf{cond}_4$ or $(al \leq 3) \wedge (m \neq \perp) \wedge \mathsf{cond}_{al}$,
   set $\mathtt{outofsync} \leftarrow 1$

7: if $\mathtt{outofsync} = 1$, return $b_P^{i,d}$, otherwise, return $\perp$

This query allows the adversary to specify any associated data and any ciphertext to be decrypted by the *partner(s)* of the specified session oracle at the specified stage, then get the secret bit $b_P^{i,d}$ if and only if this query is "out-of-sync", otherwise, it still gets $\perp$ (to avoid trivial wins). The "out-of-sync" condition (see line 6) captures different authentication levels. For conciseness, we list only the authentication conditions for level 1 and 4 (refer to [9] for level 2 and 3) as follows:

$\mathsf{cond}_1 = (\nexists w : (ct = sent.ct_w) \wedge (ad = sent.ad_w))$

$$\mathsf{cond}_4 \ = (u < v) \lor (ct \neq sent.ct_v) \lor (ad \neq sent.ad_v)$$

Note that cond_1 corresponds to the lowest authentication level (e.g., for the stateful AEAD scheme in QUIC) that only guarantees no forgeries, while cond_4 corresponds to the highest authentication level (e.g., for the stateful AEAD scheme in TLS 1.3) that prevents forgeries, replays, reordering, and dropping.

Advantage Measures. An adversary A against a msACCE protocol Π in msACCE-std has the following advantage measures.

• *Server Authentication.* We define $\mathbf{Adv}_{\Pi}^{\mathsf{s\text{-}auth}}(A)$ as the probability that there exist a client oracle π_C^i and its target server S such that the following holds:

1. π_C^i has set its session key;

2. S was not corrupted before π_C^i set its session key;

3. No interim keys of π_C^i or its partner(s) were revealed;[12]

4. There is no unique server oracle π_S^j with which π_C^i has a matching conversation.

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

1. If $P = S \in \mathcal{S}$, π_S^i has a matching conversation with a client oracle π_C^j; if $P = C \in \mathcal{C}$, denote S as π_C^i's target server;

2. S was not corrupted before π_P^i set its last stage key; If *forward secrecy* is not required for the d-th stage keys, S was not corrupted in the same time period associated with π_P^i;

3. No stage keys of π_P^i or its partner(s) were revealed;

4. If two different pre-reset messages were queried in the d-th stage, later no $\mathsf{Encrypt}(\pi_P^i, \cdot, \cdot, \mathtt{rst}, \mathtt{rst})$ queries were made.

The above captures the attacks in which the adversary compromises the privacy or integrity of secure channel messages without revealing stage keys or revealing the hidden pre-reset message or corrupting the server before the client set its last stage key (which may not be the session key). If the stage key at the target stage is not supposed to provide forward secrecy, the adversary is further restricted not to corrupt the server during the same associated time period of the target session.

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 (a.k.a. source authentication) 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

---

[12]More precisely, we allow revealing other stage keys (if any), except the first stage key, of a π_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 π_C^i. Similar condition relaxation also holds for our other security notions.

phases. Note that, unlike the availability attacks shown in [41], 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 they are more harmful in this sense. Finally, our model captures malicious *undetectable* session resets in a secure channel phase with Reset Authentication.

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, matching conversations, and partners.

Security Experiments. Consider the same experiment setups as in msACCE-std, except that no random bit $b_P^{i,d}$ is needed. The adversary A is given all the public parameters and interacts with the session oracles via the same Connect, Resume, Send, Reveal, Corrupt queries as in the msACCE-std model[13], 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 always returns $\perp$. If $\text{cmp} = 1$, π_C^i and π_S^j establish a *complete* full session privately without showing their communication to the adversary. If $\text{cmp} = 0$, π_C^i and π_S^j establish a *partial* full session privately such that the last packet sent from π_C^i right before π_S^j sets its first stage key is blocked.

This query allows the adversary to establish a complete or partial full session between any client and server oracles without observing their communication. By taking an additional flag cmp as input, this query extends the QACCE Connprivate query [41] to model IP-spoofing attacks happening in both *full* and resumption sessions.

• Pack(π_P^i, ad, m), for $P \in \mathcal{P}, i \in [N], ad \in \mathcal{AD}, m \in \mathcal{M}_{\mathsf{SC}} \cup \{\texttt{prst}, \texttt{rst}\}$.

This query returns $\perp$ if π_P^i is not in a secure channel phase. If $m \in \mathcal{M}_{\mathsf{SC}}$, it asks π_P^i to output the packet that it would send to its partner(s) for the specified associated data ad and message m according to the protocol, then returns this packet. If $m = \texttt{prst}$, π_P^i generates its pre-reset message (if any, hidden from the adversary), encrypts it with the specified associated data ad, and outputs the resulting packet, then this packet is returned. (Recall that each oracle has at most one pre-reset message, so $m = \texttt{prst}$ is allowed to be queried at most once per session oracle.) If $m = \texttt{rst}$, this query asks π_P^i to output its reset packet (if any) and returns it.

This query allows the adversary to specify any associated data and any message in a secure channel phase, then get the packet output by the specified session oracle. The adversary can also specify a session oracle to get the packet resulting from encrypting the session oracle's pre-reset message (which the adversary does not know) or get its reset packet.

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

This query delivers pkt to π_P^i and returns its response if π_P^i is in a secure channel phase, otherwise, returns $\perp$.

This query allows the adversary to deliver any packet to a specified session oracle and get its response in a secure channel phase.

Advantage Measures. An adversary A against a msACCE protocol Π in msACCE-pauth has the following associated advantage measures.

• *IP-Spoofing Prevention.* We define $\mathbf{Adv}_\Pi^{\mathsf{ipsp}}(A)$ as the probability that there exist a client oracle π_C^i and a server oracle π_S^j such that the following holds:

1. π_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 π_S^j set its first stage key;

3. The only allowed queries concerning both C and S in the time period associated with π_S^j are:

   - Connprivate$(\pi_C^x, \pi_S^y, \cdot)$ for any $x, y \in [N]$, and

---

[13]Note that Encrypt and Decrypt queries are not needed because msACCE-pauth does not consider data privacy explicitly.

15

- $\mathsf{Send}(\pi_S^y, (\mathrm{IP}_C, \mathrm{IP}_S, \cdot, \cdot))$ for any $y \in [N]$, where $(\mathrm{IP}_C, \mathrm{IP}_S, \cdot, \cdot)$ is the last packet received by π_S^y right before it sets 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.

• *KE Header Integrity.* We define $\mathbf{Adv}_\Pi^{\mathsf{int\text{-}keh}}(A)$ as the probability that there exist a client oracle π_C^i and a server oracle π_S^j such that the following holds:

1. π_C^i has set its session key and has a matching conversation with π_S^j;

2. S was not corrupted before π_C^i set its session key;

3. No interim keys of π_C^i or its partner(s) were revealed;

4. In a key exchange phase before π_C^i set its session key, π_C^i (resp. π_S^j) accepted a packet with a new header that was not output by π_S^j (resp. π_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 $\mathbf{Adv}_\Pi^{\mathsf{int\text{-}kep}}(A)$ as the probability that there exist a client oracle π_C^i and a server oracle π_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 π_C^i set its session key, π_C^i (resp. π_S^j) accepted a packet with a new payload that was not output by π_S^j (resp. π_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.

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

4. In the secure channel phase of the d-th stage, π_P^i accepted a non-reset packet with a new header that was not output by its partner(s) (via Pack queries), or π_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 $\mathbf{Adv}_\Pi^{\mathsf{rst\text{-}auth}}(A)$ as the probability that A outputs (P, i, d) such that the same (1)∼(3) conditions as in the Channel Security notion hold and the following holds:

4. In the secure channel of the d-th stage, π_P^i accepted a packet output by a $\mathsf{Pack}(\cdot, \cdot, \mathtt{prst})$ query to its partner $\pi_{P'}^j$. Later (in the d-th or a later stage), π_P^i accepted a reset packet but A made no $\mathsf{Pack}(\pi_{P'}^j, \cdot, \mathtt{rst})$ queries.

The above captures the attacks in which the adversary forges a valid reset packet without breaking any Channel Security conditions. Note that such attacks are *undetectable* by the accepting party, as opposed to a network attacker that simply drops packets.

We say a msACCE protocol Π achieves a security notion in our msACCE security models if the associated advantage is negligible (in λ) for any *probabilistic-polynomial-time (PPT) A*.

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. As mentioned in the Introduction, by [22] results, no protocol 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

	TLS 1.3 +TFO	QUIC +UDP	QUIC[TLS] +UDP
0-RTT Key Forward Secrecy [22]	✗	✗	✗
0-RTT Data Anti-Replay [22]	✗	✗	✗
Server Authentication	✓	✓	✓
Channel Security	✓	✓	✓
IP-Spoofing Prevention	✓	✓	✓
KE Header Integrity	✗	✗	✗
KE Payload Integrity	✓	✗	✗
SC Header Integrity	✗	✓	✓
Reset Authentication	✗	✗	✓

5.1   TLS 1.3 over TFO

1) Protocol: Referring to the msACCE protocol syntax, a TFO+TLS 1.3 2-RTT full handshake (see Fig. 1) is a 2-stage msACCE protocol in the full mode and a 0-RTT resumption handshake (see Fig. 2) 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 [14, 51], the TFO+TLS 1.3 header contains the TCP header (see Fig. 5 in Appendix A). 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 fields: a 32-bit sequence number `sqn`, a 32-bit acknowledgment number `ack`, a 4-bit data offset `off`, a 6-bit reserved field `resvd`, a 6-bit control bits field `ctrl`, a 16-bit window `window`, a 16-bit urgent pointer `urgp`, a variable-length ($\leq$ 320-bit) padded options `opt`. For encrypted packets, $\mathcal{H}$ additionally contains the 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. Note that `scfg_gen` is undefined.

TLS 1.3 enforces different content types for encrypted key exchange and secure channel messages. For simplicity, we define $\mathcal{M}_{\mathsf{KE}}$ and $\mathcal{M}_{\mathsf{SC}}$ as consisting of bit strings differing in their first bits. $\mathcal{M}_{\mathsf{pRST}} = \varnothing$. In Appendix C, we define TLS 1.3's stateful AEAD scheme $\mathsf{sAEAD}_{\text{TLS}} = (\mathsf{sG}, \mathsf{sE}, \mathsf{sD})$ based on the underlying nonce-based AEAD scheme $\mathsf{AEAD} = (\mathsf{G}, \mathsf{E}, \mathsf{D})$ (instantiated with AES-GCM [42] or others as documented in [51]).

We refer to Appendix E for the remaining details of TFO and refer to [22, 10] for the detailed descriptions of TLS 1.3 handshake messages and key generations in earlier TLS 1.3 drafts as well as [51] for the latest updates.

2) Security: TFO+TLS 1.3's session identifier sid_{TLS} is defined as all key exchange messages from `ClientHello` to `ServerFinished`, excluding TCP headers and IP addresses. The 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 [20, 22, 39] 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 D, we show one can adapt their security results to prove TLS 1.3's Server Authentication and level-4 Channel Security in our msACCE-std model, by additionally relying on the level-4 AEAD security of $\mathsf{sAEAD}_{\text{TLS}}$ (which can be reduced to the nonce-based AEAD security of the underlying AEAD as shown in [17]).

The 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, an AES-128 block cipher, as a PRF $F : \{0,1\}^n \times \{0,1\}^\lambda \to \{0,1\}^n$, we have the following theorem with the proof in Appendix E:

Theorem 1 *For any PPT adversary A making at most q* Send *queries, there exists a PPT adversary B such that:*

$$\mathbf{Adv}_{TFO+TLS\ 1.3}^{\mathsf{ipsp}}(A) \leq |\mathcal{S}|\mathbf{Adv}_F^{\mathsf{prf}}(B) + \frac{q}{\min\{2^{|\mathsf{sqn}|}, 2^n\}} \ .$$

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 a PPT adversary A can always get $\mathbf{Adv}_{\text{TFO+TLS 1.3}}^{\mathsf{int\text{-}keh}}(A) = 1$:

TFO Cookie Removal. A can first make $\pi_C^{i'}$ complete a full handshake with $\pi_S^{j'}$ (via Connect, Send queries), then query $\mathsf{Resume}(\pi_C^i, S, i')$ ($i' < i$) to get the output packet $(\text{IP}_C, \text{IP}_S, H, pd)$, which is a SYN packet with a TFO cookie. 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 π_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 [14]. 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 sid_{TLS} consists of the payloads of all key exchange packets exchanged between the communicating parties before the client set its session key. That is, for any client oracle that has a matching conversation with any server oracle, by definition they observe the same sid_{TLS} and hence no key exchange packet payload can be modified, i.e., $\textbf{Adv}^{\text{int-kep}}_{\text{TFO+TLS 1.3}}(A) = 0$ for any PPT adversary A.

SC Header Integrity. TFO+TLS 1.3 does not achieve this security notion again because of the unauthenticated TCP headers. A PPT adversary A can get $\textbf{Adv}^{\text{int-h}}_{\text{TFO+TLS 1.3}}(A) = 1$ 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 [54, 31, 4, 13, 37, 36, 29, 49, 12, 25, 48, 43, 58, 30]. Some of the practical attacks are described in Appendix F.

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 a PPT adversary A always gets $\textbf{Adv}^{\text{rst-auth}}_{\text{TFO+TLS 1.3}}(A) = 1$:

TCP Reset Attack. 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, 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 [58] 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: Referring to the msACCE protocol syntax, an UDP+QUIC 1-RTT full handshake (see Fig. 3) is a 2-stage msACCE protocol in the full mode and a 0-RTT resumption handshake (see Fig. 4) 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 [53, 46, 38], the UDP+QUIC header contains the UDP header (see Fig. 6 in Appendix A) 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 only affects the length of the QUIC header and payload. We thus can completely omit the UDP header and define the header space $\mathcal{H}$ as containing the following fields: an 8-bit public flag `flag`, a 64-bit connection ID `cid`, a variable-length ($\leq$ 48-bit) sequence number `sqn`, 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. $\mathcal{M}_{\text{pRST}} = \varnothing$. In Appendix C, we define QUIC's stateful AEAD scheme $\text{sAEAD}_{\text{QUIC}} = (\text{sG}, \text{sE}, \text{sD})$ based on the underlying nonce-based AEAD scheme $\text{AEAD} = (\text{G}, \text{E}, \text{D})$ (instantiated with AES-GCM [42]).

We refer to [41] for the detailed descriptions of `scfg_gen` and QUIC handshake messages and key generations.

2) Security: UDP+QUIC's session identifier sid_{QUIC} 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 G. Note that UDP+QUIC only achieves level-1 Channel Security, but, as discussed in [41], QUIC implicitly prevents packet reordering by authenticating `sqn` in the packet header. It also prevents replays and dropping 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 [41], 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 [41] 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, a PPT adversary A can get $\mathbf{Adv}^{\text{int-keh}}_{\text{UDP+QUIC}}(A) = 1$ as follows. A can first query $\mathsf{Connect}(\pi^i_C, S)$ to get the output packet $(\text{IP}_C, \text{IP}_S, H, pd)$, then modify the `flag` and `sqn` fields of H to get a new header $H' \neq H$ that only changes `sqn`'s length but not its value. The resulting packet will be accepted by a new server oracle π^j_S. 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 a PPT adversary A can get $\mathbf{Adv}^{\text{int-kep}}_{\text{UDP+QUIC}}(A) \approx 1$:

ServerReject Triggering. A can first let $\pi^{i'}_C$ complete a full handshake with $\pi^{j'}_S$ with $\mathsf{Connect}, \mathsf{Send}$ queries, then query $\mathsf{Resume}(\pi^i_C, S, i')$ $(i' < i)$ to get the output `ClientHello` packet. 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 π^j_S 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 0-RTT data. Note that this attack is hard to detect because π^i_C may think its original `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 PPT adversary A there exists a PPT adversary B such that $\mathbf{Adv}^{\text{int-h}}_{\text{UDP+QUIC}}(A) \leq \mathbf{Adv}^{\text{cs-1}}_{\text{UDP+QUIC}}(B)$.

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, a PPT adversary A can always get $\mathbf{Adv}^{\text{rst-auth}}_{\text{UDP+QUIC}}(A) = 1$:

PublicReset Attack. A can first make two session oracles complete a handshake using $\mathsf{Connect}, \mathsf{Send}$ queries, then use $\mathsf{Pack}, \mathsf{Deliver}$ queries to let them exchange secure channel packets. By observing these packet headers, A can easily forge a valid (plaintext) reset packet by setting its PUB-LIC_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: As mentioned in the Background, QUIC[TLS] replaces QUIC's key exchange with the TLS 1.3 key exchange. So, as with TLS 1.3, a UDP+QUIC[TLS] 2-RTT full handshake is a 2-stage msACCE protocol in the full mode and a 0-RTT resumption handshake is a 3-stage msACCE protocol in the resumption mode. The stage keys are set in the same way as in TLS 1.3.

The header fields (as specified in [27]) are similar to those in UDP+QUIC. Reset packets are defined as those whose first two header bits are 01. scfg_gen is undefined. UDP+QUIC[TLS] also enforces different frame types for encrypted key exchange, secure channel, and pre-reset messages. For simplicity, we define $\mathcal{M}_{\mathsf{KE}}, \mathcal{M}_{\mathsf{SC}}, \mathcal{M}_{\mathsf{pRST}}$ as consisting of bit strings differing in their first two bits. UDP+QUIC[TLS]'s stateful encryption scheme is the same as $\mathsf{sAEAD}_{\mathrm{QUIC}}$ based on the underlying nonce-based AEAD scheme $\mathsf{AEAD} = (\mathsf{G}, \mathsf{E}, \mathsf{D})$ (instantiated with AES-GCM [42] or others as documented in [51]).

QUIC[TLS] still provides source 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 [27, 57] for the detailed UDP+QUIC[TLS] handshake messages and key generations.

2) Security: UDP+QUIC[TLS]'s session identifier $sid_{\mathrm{QUIC[TLS]}}$ is defined as sid_{TLS}. By construction, UDP+QUIC[TLS] inherits the msACCE-std security from TLS 1.3 (but using QUIC's underlying encryption scheme). That is, it achieves level-1 Channel Security and implicitly achieves level-4 authentication as discussed before. UDP+QUIC[TLS] has a similar source-validation token scheme as QUIC. If the token is generated with an authenticated encryption scheme, the IP-Spoofing Prevention security of UDP+QUIC[TLS] can be reduced to the encryption scheme's authenticity security. However, such a source-validation scheme suffers from an availability attack against KE Payload Integrity similar to ServerReject Triggering for UDP+QUIC, where the adversary replaces the source-validation token with a random value to downgrade a 0-RTT resumption connection. As noted in [57], 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 [27]) are not fully authenticated. For instance, a PPT adversary A can get $\mathbf{Adv}_{\mathrm{UDP+QUIC[TLS]}}^{\mathsf{int\text{-}keh}}(A) = 1$ as follows. A first queries $\mathsf{Connect}(\pi_C^i, S)$ to get π_C^i's Initial packet $(\mathrm{IP}_C, \mathrm{IP}_S, H, pd)$. Then, as described in [57], A can decrypt this packet with its Destination Connection ID $\mathtt{dcid}$ in H, change it to another value $\mathtt{dcid}'$, and re-encrypt the whole packet with this new $\mathtt{dcid}'$. The resulting packet $(\mathrm{IP}_C, \mathrm{IP}_S, H', pd')$, where $H \neq H'$, is valid and will be accepted by a new server oracle π_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 $\mathtt{cid}$ 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 $F : \{0,1\}^{|\mathtt{cid}|} \times \{0,1\}^\lambda \rightarrow \{0,1\}^n$ as shown in the theorem below with the proof in Appendix H:

Theorem 2 *For any PPT adversary A delivering at most q forged reset packets (via $\mathsf{Deliver}$ queries), there exist PPT adversaries B and C such that:*

$$\mathbf{Adv}_{\mathrm{UDP+QUIC[TLS]}}^{\mathsf{rst\text{-}auth}}(A) \leq |\mathcal{P}|\mathbf{Adv}_F^{\mathsf{prf}}(B) + \mathbf{Adv}_{\mathrm{UDP+QUIC[TLS]}}^{\mathsf{cs\text{-}1}}(C) + \frac{|\mathcal{P}|N^2}{2^{|\mathtt{cid}|}} + \frac{q}{2^n} \ .$$

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 found that availability functionalities provided by transport-layer protocols like TCP can be easily compromised without packet-level authentication, which may undermine the performance of their supporting application-layer protocols. To protect against availability attacks, new 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 model will help protocol designers in their future protocol analyses and that our results will help practitioners better understand the advantages and limitations of novel secure channel establishment protocols.

Acknowledgments

References

[1] HTTPS encryption on the web - Google transparency report, 2018.

[2] Let's encrypt: Looking forward to 2019, December 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] 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.

[6] Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In *Annual International Cryptology Conference*, pages 232–249. Springer, 1993.

[7] Karthikeyan Bhargavan, Bruno Blanchet, and Nadim Kobeissi. Verified models and reference implementations for the tls 1.3 standard candidate. In *Security and Privacy (SP)*, pages 483–502. IEEE, 2017.

[8] Karthikeyan Bhargavan, Cédric Fournet, Markulf Kohlweiss, Alfredo Pironti, Pierre-Yves Strub, and Santiago Zanella-Béguelin. Proving the TLS handshake secure (as it is). In *Proceedings of CRYPTO*, 2014.

[9] 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.

[10] Jacqueline Brendel, Marc Fischlin, and Felix Günther. Breakdown resilience of key exchange protocols and the cases of newhope and tls 1.3. Cryptology ePrint Archive, Report 2017/1252, 2017.

[11] Jacqueline Brendel, Marc Fischlin, Felix Günther, and Christian Janson. Prf-odh: Relations, instantiations, and impossibility results. In *Annual International Cryptology Conference*, pages 651–681. Springer, 2017.

[12] 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*, 2016.

[13] Centre for the Protection of National Infrastructure. Security assessment of the transmission control protocol. Technical Report CPNI Technical Note 3/2009, Centre for the Protection of National Infrastructure, 2009.

[14] Y. Cheng, J. Chu, S. Radhakrishnan, and A. Jain. TCP Fast Open. RFC 7413 (Experimental), December 2014.

[15] 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 (SP)*, volume 00, pages 470–485, 2016.

[16] Cas Cremers, Marko Horvat, Jonathan Hoyland, Sam Scott, and Thyla van der Merwe. A comprehensive symbolic analysis of tls 1.3. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1773–1788. ACM, 2017.

[17] 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, SP 2017*, pages 463–482. IEEE Computer Society, 2017.

[18] Benjamin Dowling, Marc Fischlin, Felix Günther, and Douglas Stebila. A cryptographic analysis of the tls 1.3 handshake protocol candidates. In *ACM SIGSAC Conference on Computer and Communications Security*, CCS '15, pages 1197–1210, New York, NY, USA, 2015. ACM.

[19] 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.

[20] Benjamin James Dowling. *Provable security of internet protocols*. PhD thesis, Queensland University of Technology, 2017.

[21] Marc Fischlin and Felix Günther. Multi-stage key exchange and the case of google's quic protocol. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 1193–1204. ACM, 2014.

[22] Marc Fischlin and Felix Günther. Replay attacks on zero round-trip time: The case of the tls 1.3 handshake candidates. In *Security and Privacy (EuroS&P), 2017 IEEE European Symposium on*, pages 60–75. IEEE, 2017.

[23] Marc Fischlin, Felix Günther, Giorgia Azzurra Marson, and Kenneth G Paterson. Data is a stream: Security of stream-based channels. In *Annual Cryptology Conference*, pages 545–564. Springer, 2015.

[24] Gennie Gebhart. Tipping the scales on https: 2017 in review, December 2017.

[25] Yossi Gilad and Amir Herzberg. Off-path attacking the web. In *WOOT*, pages 41–52, 2012.

[26] IP Latency Statistics — Verizon Enterprise Solutions. Verizon Enterprise Solutions, 2018.

[27] J. Iyengar and M. Thomson. Quic: A udp-based multiplexed and secure transport, January 2019.

[28] Tibor Jager, Florian Kohlar, Sven Schäge, and Jörg Schwenk. On the security of tls-dhe in the standard model. In *Advances in Cryptology–CRYPTO 2012*, pages 273–293. Springer, 2012.

[29] 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*, 2015.

[30] 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.

[31] Laurent Joncheray. A simple active attack against TCP. In *USENIX Security Symposium*, 1995.

[32] Tadayoshi Kohno, Adriana Palacio, and John Black. Building secure cryptographic transforms, or how to encrypt and mac. 2003.

[33] Hugo Krawczyk. Cryptographic extraction and key derivation: The hkdf scheme. In *Annual Cryptology Conference*, pages 631–648. Springer, 2010.

[34] Hugo Krawczyk, Kenneth G Paterson, and Hoeteck Wee. On the security of the tls protocol: A systematic analysis. In *Advances in Cryptology–CRYPTO 2013*, pages 429–448. Springer, 2013.

[35] Hugo Krawczyk and Hoeteck Wee. The optls protocol and tls 1.3. In *Security and Privacy (EuroS&P), 2016 IEEE European Symposium on*, pages 81–96. IEEE, 2016.

[36] 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.

[37] 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.

[38] A Langley and W Chang. Quic crypto, 2016.

[39] Xinyu Li, Jing Xu, Zhenfeng Zhang, Dengguo Feng, and Honggang Hu. Multiple handshakes security of tls 1.3 candidates. In *Security and Privacy (SP), 2016 IEEE Symposium on*, pages 486–505. IEEE, 2016.

[40] Greg Linden. Make data useful, 2006.

[41] Robert Lychev, Samuel Jero, Alexandra Boldyreva, and Cristina Nita-Rotaru. How secure and quick is quic? provable security and performance analyses. In *Security and Privacy (SP), 2015 IEEE Symposium on*, pages 214–231. IEEE, 2015.

[42] 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.

[43] Robert Morris. A weakness in the 4.2 BSD unix TCP/IP software. Technical report, AT&T Bell Leboratories, 1985.

[44] Kenneth G Paterson, Thomas Ristenpart, and Thomas Shrimpton. Tag size does matter: Attacks and proofs for the tls record protocol. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 372–389. Springer, 2011.

[45] Christopher Patton and Thomas Shrimpton. Partially specified channels: The TLS 1.3 record layer without elision. In *ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2018.

[46] Jon Postel. User datagram protocol. RFC 768 (Standard), 1980.

[47] Jon Postel. Transmission control protocol. RFC 793 (Standard), 1981.

[48] Zhiyun Qian and Z. Morley Mao. Off-path TCP sequence number inference attack - how firewall middleboxes reduce security. In *IEEE Symposium on Security and Privacy*, pages 347–361, 2012.

[49] Zhiyun Qian, Z. Morley Mao, and Yinglian Xie. Collaborative TCP sequence number inference attack: how to crack sequence number under a second. In *ACM Conference on Computer and Communications Security*, 2012.

[50] Sivasankar Radhakrishnan, Yuchung Cheng, Jerry Chu, Arvind Jain, and Barath Raghavan. Tcp fast open. In *Proceedings of the Seventh COnference on emerging Networking EXperiments and Technologies*, page 21. ACM, 2011.

[51] Eric Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446, August 2018.

[52] Phillip Rogaway. Authenticated-encryption with associated-data. In *Proceedings of the 9th ACM conference on Computer and communications security*, pages 98–107. ACM, 2002.

[53] Jim Roskind. Quic(quick udp internet connections): Multiplexed stream transport over udp. *Technical report, Google*, 2013.

[54] Stefan Savage, Neal Cardwell, David Wetherall, and Tom Anderson. TCP congestion control with a misbehaving receiver. *ACM SIGCOMM Computer Communication Review*, 29(5), 1999.

[55] Ahren Studer and Adrian Perrig. The coremelt attack. In *European Symposium on Research in Computer Security*, pages 37–52, 2009.

[56] Ian Swett. QUIC deployment experience @Google. `https://www.ietf.org/proceedings/96/slides/slides-96-quic-3.pdf`, 2016.

[57] M. Thomson and S. Turner. Using transport layer security (tls) to secure quic, January 2019.

[58] Paul Watson. Slipping in the window: TCP reset attacks. Technical report, CanSecWest, 2004.

A   Transport Protocol Headers

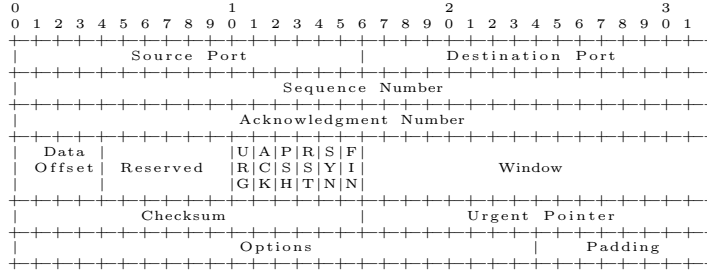Fig. 5 presents the TCP header while Fig. 6 presents the UDP header.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|          Source Port          |       Destination Port        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Sequence Number                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Acknowledgment Number                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Data  |           |U|A|P|R|S|F|                               |
| Offset| Reserved  |R|C|S|S|Y|I|            Window             |
|       |           |G|K|H|T|N|N|                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Checksum            |         Urgent Pointer        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Options                    |    Padding     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 5: TCP header. [47]

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|          Source Port          |       Destination Port        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|            Length             |           Checksum            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```
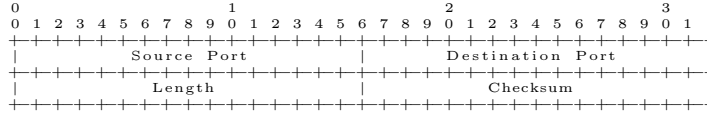
Figure 6: UDP header. [46]

B Preliminary Definitions

B.1 Pseudorandom Functions

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

F is a *pseudorandom function (PRF)* if the following holds:

- For any $k \in \{0,1\}^\lambda$ and $x \in \{0,1\}^n$, there exists a polynomial-time (in λ) algorithm to compute $F(k, x)$;

- For any PPT adversary A, $\mathbf{Adv}_F^{\mathsf{prf}}(A)$ is negligible in λ.

B.2 Stateful Authenticated Encryption with Associated Data

We follow [32, 9] in extending the stateful authenticated encryption notion of Bellare *et al.* [5] to capture a hierarchy of stateful AEAD security notions based on different authentication levels. The following definitions are the same as [9], except that we exclude the length-hiding property proposed by Paterson *et al.* [44] for conciseness.

Syntax. A stateful AEAD scheme sAEAD is a three-tuple $(\mathsf{sG}, \mathsf{sE}, \mathsf{sD})$ associated with a key space $\mathcal{K} = \{0,1\}^\lambda$, a message space $\mathcal{M} \subseteq \{0,1\}^*$, an associated data space $\mathcal{AD} \subseteq \{0,1\}^*$, and a state space $\mathcal{ST} \subseteq \{0,1\}^*$. sG is a probabilistic algorithm that samples a random key from $\mathcal{K}$ and initializes the encryption and decryption states $st_e, st_d \in \mathcal{ST}$. 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. 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 $m \in \mathcal{M} \cup \{\bot\}$ with an updated st_d. The *correctness* requires that, for any $k \in \mathcal{K}, st_e = st_e^0, st_d = st_d^0$ sampled or initialized by sG and any sequence of encryptions $\{(ct_{i+1}, st_e^{i+1}) \xleftarrow{\$}$

$\mathsf{sE}(k, ad_i, m_i, st_e^i)\}_{i \geq 0}$, the sequence of decryptions $\{(m'_{i+1}, st_d^{i+1}) \leftarrow \mathsf{sD}(k, ad, \mathsf{E}(k, ad_i, ct_i, st_d^i))\}_{i \geq 0}$ satisfies $m_i = m'_i, i \geq 0$.

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

$\underline{\mathsf{Enc}(ad, m_0, m_1)}$:

1: $u \leftarrow u + 1, (sent.ct_u, st'_e) \xleftarrow{\$} \mathsf{sE}(k, ad, m_b, st_e)$
2: $(sent.ad_u, st_e) \leftarrow (ad, st'_e)$, return $sent.ct_u$

$\underline{\mathsf{Dec}(ad, ct)}$:

1: if $b = 0$, return $\perp$
2: $v \leftarrow v + 1, (m, st'_d) \leftarrow \mathsf{sD}(k, ad, ct, st_d)$
3: $(rcvd.ad_v, st_d) \leftarrow (ad, st'_d)$
4: if $(al = 4) \wedge \mathsf{cond}_4$ or $(al \leq 3) \wedge (m \neq \perp) \wedge \mathsf{cond}_{al}$,[14]
    set $\mathtt{outofsync} \leftarrow 1$
5: if $\mathtt{outofsync} = 1$, return m, otherwise, return $\perp$

In the end, A outputs a bit b'. The stateful AEAD scheme sAEAD is *secure* with authentication level al if and only if $\mathbf{Adv}_{\mathsf{sAEAD}}^{\mathsf{aead}\text{-}al}(A) = |2\Pr[b = b'] - 1|$ is negligible in λ for any PPT adversary A.

C   QUIC and TLS 1.3's Stateful AEAD Schemes and Their Security

C.1   QUIC's Stateful AEAD Scheme and its Security

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

$\underline{\mathsf{sG}()}$:

$\quad k_e \xleftarrow{\$} \mathsf{G}(), k_m \xleftarrow{\$} \{0, 1\}^{32}$
$\quad (st_e, st_d) \leftarrow (\varnothing, \perp)$
$\quad$ return (k_e, k_m)

$\underline{\mathsf{sD}(k, ad, ct, st_d)}$:

$\quad (k_e, k_m) \leftarrow k$
$\quad (\mathtt{cid}, \mathtt{sqn}) \leftarrow ad$
$\quad m \leftarrow \mathsf{D}(k_e, k_m \| \mathtt{sqn}, ad, ct)$
$\quad$ return $(m, \perp)$

$\underline{\mathsf{sE}(k, ad, m, st_e)}$:

$\quad (k_e, k_m) \leftarrow k$
$\quad (\mathtt{cid}, \mathtt{sqn}) \leftarrow ad$
$\quad$ if $\mathtt{sqn} \in st_e$,
$\quad\quad$ return $(\perp, st_e)$
$\quad c \leftarrow \mathsf{E}(k_e, k_m \| \mathtt{sqn}, ad, m)$
$\quad st_e \leftarrow st_e \cup \{\mathtt{sqn}\}$
$\quad$ return (c, st_e)

Note that $\mathsf{sAEAD}_{\mathrm{QUIC}}$ uses the encryption state to keep track of used nonces to avoid repeating and the decryption state is unused.

To reduce $\mathsf{sAEAD}_{\mathrm{QUIC}}$'s 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 $\mathbf{Adv}_{\mathsf{AEAD}}^{\mathsf{ind}\text{-}\mathsf{cpa}}(A)$. 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 $\mathbf{Adv}_{\mathsf{AEAD}}^{\mathsf{int}\text{-}\mathsf{ctxt}}(A)$. Now, we are ready to prove the following theorem.

Theorem 3 *For any PPT adversary A, there exist PPT adversaries B and C such that:*

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

---

[14]Authentication conditions cond_{al} are defined in the same way as in the msACCE-std $\mathsf{Decrypt}$ query.

Proof: Consider two games G_0 and G_1. G_0 is the real experiment for A and G_1 is the same as G_0 except that it will always return $\perp$ for Dec queries. Denote $\Pr_i$ as the advantage of A in G_i. $|\Pr_0 - \Pr_1|$ is bounded by the probability that A forges a new valid ciphertext given $b = 1$, which by definition is bounded by $\mathbf{Adv}_{\mathsf{AEAD}}^{\mathsf{int\text{-}ctxt}}(B)$ for some PPT adversary B. Then, note that according to the $\mathsf{sAEAD}_{\mathrm{QUIC}}$ construction nonces in AEAD encryption queries never repeat and G_1 can be simulated by an PPT adversary C against the nonce-based AEAD privacy security, which implies $\Pr_1 \leq \mathbf{Adv}_{\mathsf{AEAD}}^{\mathsf{ind\text{-}cpa}}(C)$. Therefore, we have $\mathbf{Adv}_{\mathsf{sAEAD}_{\mathrm{QUIC}}}^{\mathsf{aead\text{-}1}}(A) \leq \mathbf{Adv}_{\mathsf{AEAD}}^{\mathsf{int\text{-}ctxt}}(B) + \mathbf{Adv}_{\mathsf{AEAD}}^{\mathsf{ind\text{-}cpa}}(C)$. $\qquad\square$

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

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

$\underline{\mathsf{sG}():}$
  $k_e \overset{\$}{\leftarrow} \mathsf{G}(), k_m \overset{\$}{\leftarrow} \{0,1\}^n$
  $(st_e, st_d) \leftarrow (0, 0)$
  return (k_e, k_m)

$\underline{\mathsf{sE}(k, ad, m, st_e):}$
  $(k_e, k_m) \leftarrow k$
  $c \leftarrow \mathsf{E}(k_e, k_m \oplus st_e, ad, m)$
  $st_e \leftarrow st_e + 1$
  return (c, st_e)

$\underline{\mathsf{sD}(k, ad, ct, st_d):}$
  if $st_d = \perp$, return $(\perp, \perp)$
  $(k_e, k_m) \leftarrow k$
  $m \leftarrow \mathsf{D}(k_e, k_m \oplus st_d, ad, ct)$
  if $m = \perp$,
    $st_d \leftarrow \perp$
  otherwise,
    $st_d \leftarrow st_d + 1$
  return (m, st_d)

Note that in the above TLS's stateful encryption scheme, nonce repeating is prevented by the increasing counter kept by the encryption state st_e. Following a very similar argument as in the above proof of Theorem 3, one can show that the level-4 AEAD security of $\mathsf{sAEAD}_{\mathrm{TLS}}$ is also reduced to the nonce-based AEAD security of AEAD. This result has been proved by previous work (Theorem 3 in [17]), 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.

D 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 show a proof sketch below.

Previous works [20] and [22] respectively proved that the TLS 1.3 draft-16 (EC)DHE full handshake and draft-14 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 function, and pseudorandom function oracle Diffie-Hellman (PRF-ODH) assumption [28, 34, 11]. Their MSKE security, which captures only the key exchange phases, ensures the Bellare-Rogaway-style key secrecy [6] (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). These results derived the overall TLS 1.3 security using a *compositional* approach, i.e., composing a secure key exchange protocol (e.g., the TLS 1.3 handshake protocol) in the MSKE model with an arbitrary secure symmetric key protocol (e.g., the TLS 1.3 record protocol). However, as stated in [22], this generic composition result only works for key-independent, forward-secret, external, and non-replayable stage keys. In particular, it does not apply to the final session keys in full handshakes or the interim handshake keys because they are used internally in the key exchange phases. Besides, it does not apply to the 0-RTT keys, which are replayable and non-forward-secret. 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 a few TLS 1.3 updates and model differences as follows.

First, we show that the security results in [20, 22] for old TLS 1.3 drafts can be extended to the standard TLS 1.3 [51], i.e., the standard TLS 1.3 (EC)DHE full handshake and PSK-(EC)DHE 0-RTT handshake are secure in the MSKE model.

1) The multi-stage key generation procedures are updated in the just-released TLS 1.3. Recall that TLS 1.3 performs key derivation in the extract-then-expand paradigm [33] using the HMAC-based Extract-and-Expand Key Derivation Function (HKDF), which consists of two functions HKDF.Extract and HKDF.Expand. In particular, it first extracts an internal secret (e.g., early secret, handshake secret, and master secret), then expands it (twice) to derive the corresponding stage key. The latest standard TLS 1.3 performs an expand-then-extract procedure instead of a single extract procedure for the extraction of the handshake secret and master secret. However, these two additional expand steps do not affect the MSKE security because they only add a constant (single) query to HKDF.Expand, leading to a larger constant for its PRF advantage. Besides, compared to [20], such extra expand steps help TLS 1.3's MSKE security no longer rely on the PRF security of the underlying HMAC primitive of both HKDF functions.

2) The message flows of the PSK-(EC)DHE 0-RTT resumption handshake are updated in the just-released TLS 1.3. The 0-RTT Finished message is replaced by a pre-shared key (PSK) binder. They are both HMAC values generated with very similar procedures and have the same purpose, i.e., to authenticate the ClientHello message and to bind the current resumption session with the assoicated full session. Such a replacement does not affect the TLS 1.3's MSKE security. Besides, a new EndOfEarlyData message is added as an indicator to end 0-RTT data transmission. This is an empty handshake message independent of key generation so does not affect the security either.

Then, based on the above extended TLS 1.3 MSKE security, we can apply the security results in [39] to get the Multi-Level&Stage security of the combination of the TLS 1.3 full handshake and 0-RTT resumption handshake. Referring to their notions [39], our msACCE-std model focuses only on two modes, i.e., the (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.

Finally, we show that the above TLS 1.3 security result in the Multi-Level&Stage model [39] can be augmented to prove TLS 1.3's Server Authentication and level-4 Channel Security.

1) The above security result 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 make reducing KE Payload Integrity to Server Authentication easy. (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 their proof of the TLS 1.3 Multi-Level&Stage server authentication to replace handshake keys with independent and random values, then use $\mathsf{sAEAD}_\mathrm{TLS}$'s AEAD oracles to simulate encrypted key exchange messages in sid_TLS and the decryption of them. In this way, Server Authentication can be reduced to the TLS 1.3 Multi-Level&Stage server authentication and the AEAD security.

2) To prove level-4 Channel Security, we follow their proof of the TLS 1.3 Multi-Level&Stage security to replace all stage keys with independent and random values and then use the AEAD oracles to simulate encrypted key exchange messages and Encrypt, Decrypt queries. In this way, level-4 Channel Security can be reduced to the TLS 1.3 Multi-Level&Stage security and the level-4 AEAD security of $\mathsf{sAEAD}_\mathrm{TLS}$. Note that the AEAD oracles are also used to simulate post-handshake messages like NewSessionTicket. This bypasses the composition issue [19] faced by the MSKE model (and its extensions), in which the application keys in full handshakes cannot be composed with secure symmetric key protocols because these keys are used internally in the key exchange phase to encrypt NewSessionTicket messages.
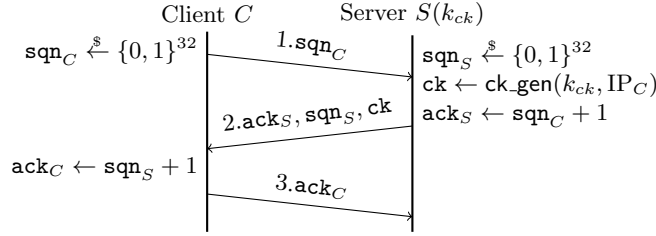
E   Proof of Theorem 1


Figure 7: TFO initial connection.

Figure 8: TFO 0-RTT resumption connection. * indicates optional messages.

The TFO protocol specifications are shown in Fig. 7 and Fig. 8, where the server samples $k_{ck} \xleftarrow{\$} \{0,1\}^\lambda$ in the beginning and then generates cookies with ck_gen:

ck_gen(k_{ck}, IP_C):
    return $F_{k_{ck}}(\mathrm{IP}_C \| \mathrm{IP}_S \| 0 \cdots 0)$

Proof: Consider a sequence of games $G_0, \ldots, G_{|\mathcal{S}|}$. G_0 is the real experiment for A and $G_{|\mathcal{S}|}$ uses random functions instead of the PRF F for all servers. The hybrid game G_i uses random functions for the first i servers and PRF for the last $|\mathcal{S}| - i$ servers. Denote Pr_i as the winning probability of A in G_i. By the PRF definition, for any $i \in [|\mathcal{S}|]$ there exists a PPT adversary B_i such that $|\mathrm{Pr}_{i-1} - \mathrm{Pr}_i| \le \mathbf{Adv}_F^{\mathsf{prf}}(B_i)$. Therefore, there exists a PPT adversary B such that $|\mathrm{Pr}_0 - \mathrm{Pr}_{|\mathcal{S}|}| \le |\mathcal{S}|\mathbf{Adv}_F^{\mathsf{prf}}(B)$.

Now we only need to bound $\mathrm{Pr}_{|\mathcal{S}|}$ by considering two cases. 1) A wins by sending a valid ACK packet. In this case, A must have generated a valid $\mathtt{ack}_C$ by correctly guessing the target server's TCP sequence number $\mathtt{sqn}_S$. The winning probability of each guess is exactly $1/2^{|\mathtt{sqn}|}$. 2) A wins by sending a valid SYN packet in resumption sessions. In this case, A must have forged a valid TFO cookie ck. The winning probability of each forgery is exactly $1/2^n$ because the TFO cookie generation functions are independent and truly random. By applying a union bound on the q queries, we have $\mathrm{Pr}_{|\mathcal{S}|} \le q/\min\{2^{|\mathtt{sqn}|}, 2^n\}$. □

F   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.

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. [30] 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 [54]. 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 [55].

G   UDP+QUIC's msACCE-std Security

It has been proven in [41] that QUIC is QACCE-secure in the random oracle model based on the unforgeability of the signature scheme, the computational Diffie-Hellman (DH) assumption [3], and the nonce-based AEAD security. Note that msACCE-std with $\mathsf{sAEAD}_{\mathrm{QUIC}}$ is semantically equivalent to QACCE with nonce-based AEAD and get_iv (defined in [41]), so their QACCE security results can be trivially adapted to show that UDP+QUIC achieves Server Authentication and level-1 Channel Security in our msACCE-std model. Note that msACCE-std security relies on the level-1 AEAD security of $\mathsf{sAEAD}_{\mathrm{QUIC}}$ instead of the nonce-based AEAD security of the underlying AEAD, but the former can be reduced to the latter as shown in Appendix C.

H   Proof of Theorem 2

Proof: Consider a sequence of games G_0, G_1, G_2, G_3 and let Pr_i denote the winning probability of A in game G_i.

G_0 is the real experiment for A, so $\mathrm{Pr}_0 = \mathbf{Adv}^{\mathsf{rst\text{-}auth}}_{\mathrm{UDP+QUIC[TLS]}}(A)$.

G_1 is the same as G_0 except that the connection IDs never repeat for each party. Since the probability of cid collision for each party is at most $N^2/2^{|\mathtt{cid}|}$, we have $|\mathrm{Pr}_0 - \mathrm{Pr}_1| \leq |\mathcal{P}|N^2/2^{|\mathtt{cid}|}$.

G_2 is the same as G_1 except that G_2 uses independent random functions f instead of the PRF F for reset token generation. Similar to the hybrid argument in the proof of Theorem 1, there exists a PPT adversary B such that $|\mathrm{Pr}_1 - \mathrm{Pr}_2| \leq |\mathcal{P}|\mathbf{Adv}^{\mathsf{prf}}_F(B)$.

Now consider a Channel Security adversary C that simulates G_2 for A. C can simulate Pack and Deliver queries with Encrypt and Decrypt queries. For a $\mathsf{Pack}(\cdot, \cdot, \mathtt{prst})$ query, C generates two random reset tokens $x, y \xleftarrow{\$} \{0,1\}^n$ and uses them to construct two pre-reset messages. Then, C queries Encrypt with these pre-reset messages as challenge messages, uses the output ciphertext to form a pre-reset packet, and sends it to A. This simulates G_2 perfectly except when a forged reset packet happens to match the wrong reset token (generated from x or y) in which case C does not know whether or not to accept the reset packet.

G_3 is the same as G_2 except that a reset packet generated by A never matches the wrong reset token. Since there are at most q reset packets forged by A, we have $|\Pr_2 - \Pr_3| \leq q/2^n$.

Now a Channel Security adversary C can perfectly simulates G_3. When A outputs (P, i, d), C returns a random bit if π_P^i rejected the forged reset packet (i.e., it does not match either of the reset tokens x or y generated by C for π_P^i's partner); returns 0 if it matches x; returns 1 if it matches y. Then, because mismatch does not happen, C's winning probability is at least $\Pr_3 + (1 - \Pr_3)/2$, i.e., C's advantage is at least $\Pr_3$. Therefore, $\Pr_3 \leq \mathbf{Adv}^{\mathsf{cs\text{-}1}}_{\mathrm{UDP+QUIC[TLS]}}(C)$.

Our proof is concluded with a union bound on $\Pr_i$s. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$