Replay Attacks on Zero Round-Trip Time: The Case of the TLS 1.3 Handshake Candidates

Marc Fischlin   Felix Günther

Cryptoplexity, Technische Universität Darmstadt, Germany
`marc.fischlin@cryptoplexity.de`, `guenther@cs.tu-darmstadt.de`

February 2, 2017

Abstract. We investigate security of key exchange protocols supporting so-called zero round-trip time (0-RTT), enabling a client to establish a fresh provisional key without interaction, based only on cryptographic material obtained in previous connections. This key can then be already used to protect early application data, transmitted to the server before both parties interact further to switch to fully secure keys. Two recent prominent examples supporting such 0-RTT modes are Google's QUIC protocol and the latest drafts for the upcoming TLS version 1.3.

We are especially interested in the question how replay attacks, enabled through the lack of contribution from the server, affect security in the 0-RTT case. Whereas the first proposal of QUIC uses state on the server side to thwart such attacks, the latest version of QUIC and TLS 1.3 rather accept them as inevitable. We analyze what this means for the key secrecy of both the preshared-key-based 0-RTT handshake in draft-14 of TLS 1.3 as well as the Diffie–Hellman-based 0-RTT handshake in TLS 1.3 draft-12. As part of this we extend previous security models to capture such cases, also shedding light on the limitations and options for 0-RTT security under replay attacks.

Contents

1   Introduction

Key exchange protocols are among the most widely used cryptographic protocols today, incorporated, e.g., in the TLS, SSH, IPsec, and QUIC protocols. They serve the purpose of establishing a (potentially authenticated) secret key between two parties in a network. While efficiency has always been a relevant aspect for such protocols, optimization traditionally focused on the cryptographic operations, which for a long time dominated the overall cost (in time) for executions. With the technological progress in speed of computation, but also advances and, equally important, the deployment of elliptic-curve cryptography, researchers and practitioners managed to reduce the cost of (even asymmetric) cryptographic operations drastically over the last decades. As a result, the communication complexity has become a more and more dominant factor for the overall efficiency of key exchange protocols.

1.1   Zero Round-Trip Time

While steadily increasing bandwidth on the Internet renders the data complexity aspect of communication subordinate, speed of light prepares to set a definitive lower bound for the time a message needs to be sent back and forth between two parties (called *round-trip time*). Reducing the round complexity has hence become a major design criteria in the last years, with several low-latency designs for key exchange proposed by researchers [PZS⁺13, KW16, HJLS15, WTSB16] as well as by practitioners. Prominent practical examples are in particular Google's QUIC protocol [QUI] incorporated into the Chrome browser and the upcoming TLS version 1.3 [Res16e], the latter being based on the OPTLS key exchange protocol by Krawczyk and Wee [KW16]. Those designs set out to establish an initial key in zero round-trip time (0-RTT) that allows one party (usually the client) to send "early" data already along with the first key exchange message to a (previously visited) server.

Without the server being able to contribute, it is well understood that such an approach cannot achieve equally strong security guarantees for the initial key as classical key exchange protocols are able to provide with a full round-trip (and hence contributions from both parties). In particular, the initial key cannot provide (forward) secrecy in a setting where no state is shared between sessions and all but the ephemeral keying material is compromised after the key exchange run. The common strategy is, hence, that both parties switch to a stronger key (e.g., achieving forward secrecy) after the server contributed in a second step of the key exchange and protect any further communication under this key.

Diffie–Hellman-based 0-RTT.   One main concept to derive a 0-RTT key based on a Diffie–Hellman-style key exchange and to later upgrade to a stronger, forward-secret key, is shared by both recent prominent instances QUIC and TLS 1.3 (up to `draft-12` [Res16b]).[1] From a high-level perspective (i.e., omitting necessary mechanisms to protect, e.g., against replays or man-in-the-middle attacks which both protocols employ), this concept works as follows. Prior to the actual key exchange, the client is assumed to have talked to the server before and, in that communication, obtained a so-called *server configuration*. Cryptographically speaking, this configuration includes a *semi-static* Diffie–Hellman share g^s, for which the server stores the secret exponent s for a certain time. In QUIC, authentication of this server configuration is via an (offline) signed structure announced by the server, in TLS 1.3 it is signed (online) during a prior handshake.

Within its first message in subsequent executions, the client then sends an ephemeral Diffie–Hellman share g^x, derives the 0-RTT key K_1 as (a function of) $(g^s)^x = g^{xs}$, and is hence immediately able to, e.g., send encrypted data under the key. The server then computes the same key as $(g^x)^s$ (enabling decryption of the 0-RTT data) and responds with its own ephemeral share g^y for the stronger shared key. Both parties

---

[1] We refer here to the (EC)DHE 0-RTT variant in TLS 1.3 draft `draft-ietf-tls-tls13-12` and the original QUIC proposal `Rev 20130620`, see also our comment in Section 1.3 about the status of these documents.

derive the full key K_2 as (a function of) g^{xy}, which can then enjoy forward secrecy in the sense that it remains secure even if g^s or the parties long-term secrets are later compromised.

Preshared-key-based 0-RTT. Another concept for establishing a key in zero round-trip time is based on pre-shared keys (PSKs) and, from `draft-13` [Res16c] on, forms the basis of the only 0-RTT handshake mode specified for TLS 1.3 (i.e., the option for Diffie–Hellman-based 0-RTT was deferred in `draft-13`). Here, the 0-RTT key K_1 is derived from a previously established secret key (e.g., in TLS 1.3 a key established for session resumption in a regular handshake). The client can perform this computation without interaction with the server and hence is able to immediately send encrypted data under K_1. Later, both parties update a full key K_2 derived from the pre-shared secret and further exchanged material, e.g., fresh Diffie–Hellman shares to ensure forward secrecy.

1.2   The Problem with Replays and How It Is (Not) Solved in QUIC and TLS 1.3

The standard approach in key exchange protocols to prevent a man-in-the-middle attacker from replaying messages in order to make a party derive the same key twice is two include a nonce in both the client's and the server's messages and let the nonce contribute to the derived key. For a 0-RTT key exchange, which is essentially a one-pass (i.e., one-message) key exchange protocol [BWM99], messages (and hence keys) are—at first glance—inevitably replayable[2].

The QUIC protocol side-stepped the replay problem in its original cryptographic design [LC13] (called `Rev 20130620` here) by demanding the server to store all nonces seen in a so-called "strike register"— restricted in size by a server-specific "orbit" prefix and current time contained in the nonces—and rejecting any recurring nonce. As security analyses confirmed [FG14, LJBN15], this approach indeed allows to establish a secure 0-RTT key which is non-replayable in the sense that no adversary can make a party derive the same key twice. However, while this approach can succeed to prevent *replays on the key-exchange level* (in terms of preventing double-derivation of keys), it inevitably fails to prevent *(logical) replays of the actual data exchanged*, in particular when it comes to real-world settings where a server entity is implemented in a cluster of, potentially distributed, servers, as we explain next. Let us stress that this problem with replays is independent of whether the 0-RTT key exchange is based on Diffie–Hellman or on preshared keys.

As discovered by Daniel Kahn Gillmor in the discussion around the upcoming TLS version 1.3 [Res15b], any 0-RTT anti-replay mechanism deployed at the key exchange level becomes void when combined within an overall channel protocol that aims to provide reliable delivery of data messages (like, e.g., QUIC or TLS). The reason is that such a protocol will resend rejected 0-RTT data under the second (final) key derived automatically in order to ensure delivery. A generic attacker can hence, for any client sending 0-RTT key-exchange messages together with some encrypted data, make this data being delivered twice in the following attack, also illustrated in Figure 1 (see [Res15b] for a more detailed description of the attack).

The attacker first conveys the client's 0-RTT messages and encrypted data to the server (which processes it), but drops the server's key exchange response. It then forces the server to lose its state, e.g., through rebooting, and presents the same messages again to the server. The server, with knowledge about its reset, has to conservatively decline the 0-RTT part of the key exchange for security reasons, but will reply with its own key exchange contribution for the final key which the attacker now simply relays to the client. The client derives the final key and, to ensure reliable delivery, *sends the desired data again* under this key, which the server will hence decrypt and process a second time. This constitutes a replay of the contained application data and might, e.g., result in a web transaction being processed twice.

---

[2]We use the notion of replays interchangeably for both messages and the keys computed based on those replayed messages.
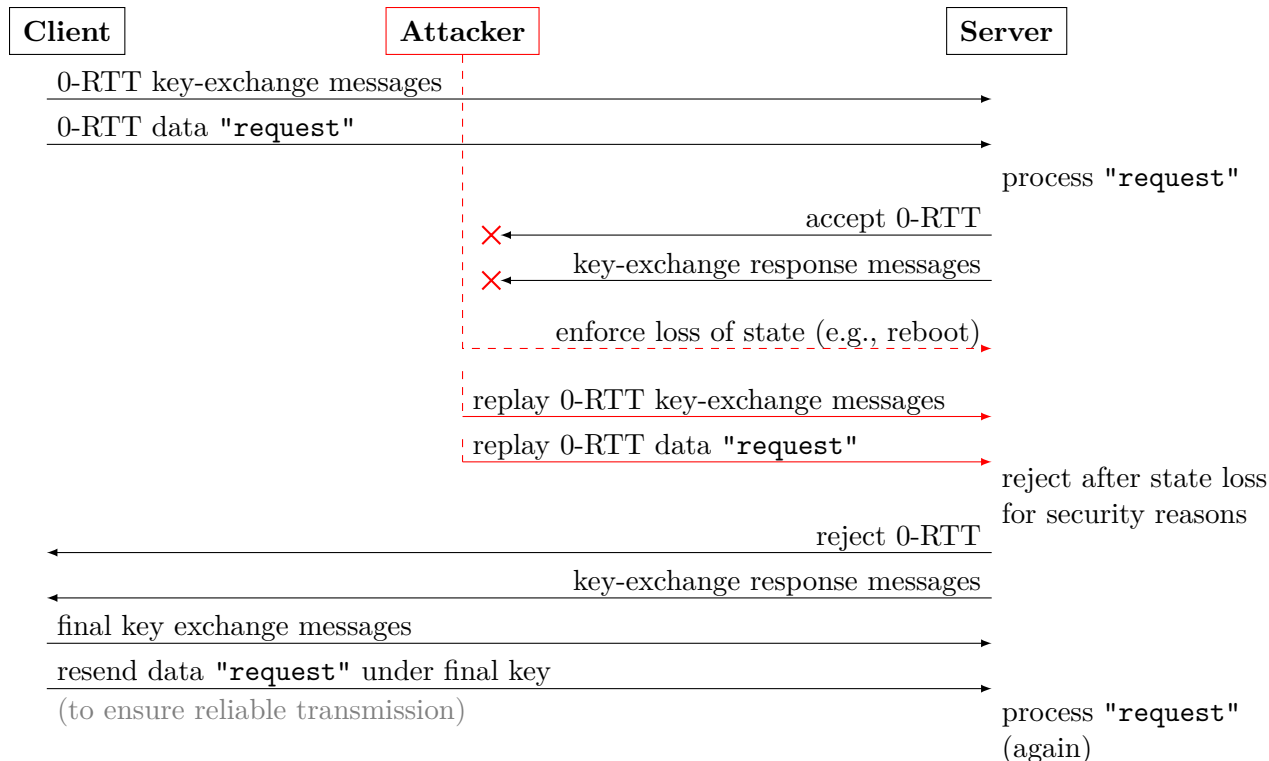
Figure 1: Generic replay attack discovered by Daniel Kahn Gillmor in the IETF TLS working group discussion around TLS 1.3 [Res15b]. The 0-RTT data `"request"` could, e.g., be an HTTP request `"POST /buy-something"`.

Note that the contrived requirement that the attacker is able to reboot the server (while the client keeps waiting for a response) vanishes in a real-world scenario with distributed server clusters, where the attacker instead simply forwards the 0-RTT messages to two servers and drops the first server's response. The described attack hence in particular affects the cryptographic design of QUIC, which (among others) specifically targets settings with distributed clusters. Holding up the originally envisioned 0-RTT full replay protection being impossible, Langley and Chang write in the specification of July 2015 [LC15] (`Rev 20150720`) that this design is "destined to die" and will be replaced by (an adapted version of) the TLS 1.3 handshake. We, however, argue here that QUIC's strategy in `Rev 20130620` still supports some kind of replay resistance, only at a different level. TLS 1.3, in contrast, forgoes any protection mechanisms and instead accepts replays as inevitable (on the channel level). Developers using TLS 1.3 are supposed to be provided with a different API call for sending 0-RTT data [Res16e, Appendix B.1], indicating its replayability, and responsible for taking replays into account for such data.

There is, then, a significant conceptual gap between replays (of key-exchange messages and keys) on the key-exchange level, and the replay of user data faced on the level of the overall secure channel protocol in the 0-RTT setting. While the former can effectively be prevented within the key exchange protocol, this does not necessarily prevent the latter which can be (and in practice is) induced by the network stack of the channel actively and automatically re-sending (presumably) rejected 0-RTT data under the main key. The latter type of logical, network-stack replays is hence fundamentally beyond of what key exchange protocols can protect against.

1.3 Our Contribution

In this work, we reconsider the derivation of 0-RTT keys in the domain of multi-stage key exchange protocols [FG14], designed to capture protocols establishing sequences of keys in an intertwined way within a single run. We particularly focus on the security of 0-RTT keys and the question of replays. Within this context, we analyze the preshared-key-based (PSK-based) 0-RTT handshake mode of draft `draft-ietf-tls-tls13-14` (short: `draft-14`) for the upcoming TLS version 1.3 as well as the recently abandoned Diffie–Hellman-based (DH-based) 0-RTT handshake mode in its last specified form in `draft-ietf-tls-tls13-12` (short: `draft-12`).[3] While doing so, we discuss the commonalities and differences between (the original version of) QUIC and the two TLS 1.3 modes in this respect. We stress that, while TLS 1.3 `draft-14` already is a relatively mature protocol specification, it remains a draft and still contains un- or underspecified parts. Our inquiry of the TLS 1.3 0-RTT handshakes hence should be conceived as an early (affirmative) discussion of their cryptographic strength, but cannot constitute a definitive analysis.

As mentioned, the Diffie–Hellman-based ((EC)DHE) 0-RTT handshake was removed from the TLS 1.3 draft specification in `draft-13`, leaving only a PSK-based 0-RTT mode (with or without additional Diffie–Hellman exchange) in the latest drafts. Still, the (EC)DHE 0-RTT variant is much closer to the QUIC and OPTLS proposals, and it may be used as a TLS extension [Res16f], especially since it provides some kind of forward secrecy [Kra16]. We hence also provide an analysis of the DH-based variant to enable a comparison of the security guarantees provided, but focus on the PSK-based 0-RTT handshake specified for draft `draft-14`.

In more detail, our contributions are fourfold.

Comparison of QUIC and TLS 1.3. We point out, in passing and explicitly in Section 8, how the designs of QUIC and TLS 1.3 differ in the way of handling 0-RTT, replay attacks, and data, and how this affects the security. This testifies that, although there may be an agreement on the general goals which should be achieved with 0-RTT, the technical details can vary significantly. One major difference has already been discussed above, carving out that both protocols treat replays differently. Another difference is that QUIC basically restarts the key exchange for an invalid (rejected) 0-RTT request, whereas TLS 1.3 instead only skips over to the regular handshake part. Both protocols also employ different approaches to derive the session keys: While QUIC uses the early key for transmitting both key-exchange messages and application data, TLS 1.3 uses a more versatile approach to create early keys for designated purposes and thus achieves stronger security guarantees and better modularity.

Multi-stage key exchange with replayable 0-RTT keys. As the original QUIC key exchange `Rev 20130620` [LC13] ensures non-replayability (on the key-exchange level), the analysis by Fischlin and Günther in the multi-stage setting [FG14] did not consider replays. The TLS 1.3 0-RTT handshake candidates (both `draft-12` DH-based and `draft-14` PSK-based), however, do not aim at preventing replays even on the key-exchange level, motivated by that any such measure would be defeated by the active, logical replay of data ensuring reliable delivery, as discussed above.

We hence extend (in Section 3) the previous multi-stage key exchange models used to analyze QUIC [FG14] and the full and preshared-key handshakes of TLS 1.3 [DFGS15a, DFGS16] in several aspects. First of all, we introduce the distinction between replayable and non-replayable stages (and, hence, keys) which, in the case of TLS 1.3, allows us to capture that the 0-RTT key exchange messages of a client session can be replayed to multiple server sessions which will all derive the same key. In order to capture the effects of exposures of the semi-static keys used in a DH-based handshake to non-interactively derive the 0-RTT keys, we allow them to be compromised and define how this affects both 0-RTT keys (which will

---

[3]Since our analysis, several follow-up draft versions of TLS 1.3 have been published, maintaining the PSK-based 0-RTT handshake mode.

be compromised) and non–0-RTT keys (which are required to remain secure). We additionally distinguish between keys used to protect application data only (called external keys) and keys which are also used within the key exchange, e.g., to encrypt key exchange messages (called internal keys). Such distinction was previously only made informally for the final key(s) derived [DFGS15b, DFGS16]. In the TLS 1.3 0-RTT handshakes however, also intermediate keys (namely, the 0-RTT early-data application key tk_{ead}) are only used externally, a setting for which our notion provides a cleaner separation.

Security analysis of the TLS 1.3 draft-14 PSK/PSK-(EC)DHE 0-RTT and draft-12 (EC)DHE 0-RTT handshakes. We then apply our model (in Sections 5 and 7) to analyze the PSK and PSK-(EC)DHE 0-RTT handshake modes specified in TLS 1.3 `draft-14`, which we describe in Section 4 first, as well as the (EC)DHE 0-RTT handshake mode of `draft-12`, described in Section 6. The other specified handshake modes of TLS 1.3, full (EC)DHE and pre-shared key, have already been analyzed by Dowling et al. [DFGS16] for the previous (relatively close) `draft-10` [Res15a]. Our analysis shows that all three 0-RTT handshakes are secure (multi-stage) key exchange protocols, establishing random-looking keys. In particular, the two 0-RTT keys derived to protect the early handshake messages and application data, tk_{ehs} resp. tk_{ead}, achieve the desired unilateral resp. mutual authentication, and are—as expected—replayable. Furthermore, we confirm that the second parts of the handshakes (essentially a full (EC)DHE resp. a regular PSK/PSK-(EC)DHE handshake), achieve security similar to that attested by Dowling et al. [DFGS16] for `draft-10`. Applying concepts established by Fischlin and Günther [FG14] and Dowling et al. [DFGS15a], we show that security holds for the different authentication options of TLS 1.3 running in parallel and that all keys derived are independent in the sense of that leaking one of them does not affect any other key.

Our security results hold under mostly standard cryptographic assumptions like the unforgeability of the signature resp. MAC scheme, collision resistance of the hash function, and pseudorandomness properties of the key derivation function. The handshakes' security further relies on the (plain) pseudorandom-function oracle Diffie–Hellman (PRF-ODH) assumption (introduced and used earlier for the analysis of several Diffie–Hellman–based modes of TLS 1.2 [JKSS12, KPW13]). Notably, for technical reasons that we detail in our proof, we furthermore need to employ a slightly strengthened, double-sided variant of the PRF-ODH assumption (which we define under the name of msPRF-ODH in Section 2) for the analysis of the (EC)DHE 0-RTT handshake (in `draft-12`).

Composition with external keys. The distinction between external(-only) and internal keys finally allows us to establish slightly more general, cleaner composition results (in Section 9). We recall that Fischlin and Günther [FG14] lifted the Bellare–Rogaway compositional result by Brzuska et al. [BFWW11] to the multi-stage setting (later extended by Dowling et al. [DFGS15a, DFGS16]). On a high level, their result specifies sufficient conditions for a multi-stage key exchange protocol such that the protocol generically composes the established session keys with any symmetric-key protocol. Our refined model determines more clearly which derived session keys can be possibly amenable to this generic composition result, establishing the key being *external* as a necessary condition. We also capture the influence of replays on generic composition, establishing *non-replayability* as a further condition.

1.4 Related Work

Our work builds upon and extends the multi-stage key exchange models by Fischlin and Günther [FG14] (used to analyze QUIC) and Dowling et al. [DFGS15a, DFGS15b, DFGS16] (used to analyze the full (EC)DHE and preshared-key handshakes of several prior TLS 1.3 drafts).

The practical requirement for reduced round-trip time, nowadays exemplified in the recent designs of QUIC and TLS 1.3, has already appeared in the works about MQV [BWM99] as well as HMQV [Kra05]

and its one-pass version [HK11]. The idea has later been discussed more formally under the notion of non-interactive key exchange (NIKE) [CKS09, FHKP13]. The difference to 0-RTT key exchange is that NIKEs describe protocols which establish a single session key which is not used within the key exchange. Of course, the key is replayable in the above sense, and security requirements usually disallow trivial attacks on such replayable keys. Another difference to 0-RTT protocols is that for NIKEs there is usually no notion of semi-static keys, i.e., since no further interaction takes place the parties cannot authenticate additional cryptographic keys with limited life span in the subsequent steps.

Krawczyk and Wee [KW16, KW15] recently introduced the OPTLS protocol, which forms the clean and elegant cryptographic core of the TLS 1.3 handshake modes, as well as analyzed its security in the Canetti–Krawczyk model [CK01], and also proposed its one-pass version [HK11] for the 0-RTT key. Focusing on the security of the two keys derived in OPTLS (corresponding to the early-data and application traffic keys tk_{ead} and tk_{app} in the TLS 1.3 0-RTT handshakes) separately, the coherence notion of key independence is beyond the scope of their analysis (though mentioned informally), as is the compositional security of these keys. While remarking that a PRF-ODH-like assumption could potentially be used for the proof, Krawczyk and Wee employ different assumptions for their analysis, namely the Gap-DH assumption in the random oracle model. Furthermore, OPTLS does not include a PSK-based 0-RTT mode and TLS 1.3 also extends the OPTLS protocol in order to include client authentication, resumption, encryption of handshake messages, and further exported keying material; aspects that we take into account in our analysis of the 0-RTT handshake modes.

A recent work by Hale et al. [HJLS15] introduces a simplified security model for low-latency key exchange with two session keys, one early key and one final key, as in case of QUIC. For their security model they introduce a notion called strong key independence which basically says that revealing the early key does not violate secrecy of the final key (or vice versa). This seems to be exactly in the same spirit as the notion of key independence introduced earlier in [FG14] and also used here. Furthermore, Hale et al. [HJLS15] point out that QUIC does not provide strong key independence. This has already been discussed in [FG14], inciting the authors in [FG14] to also propose a slight modification of QUIC, called QUICi, which achieves their notion of key independence via modifying the key derivation steps only. Hale et al. [HJLS15] then give a new generic construction secure in their model, including strong key independence, based on non-interactive key exchange. Finally, as discussed in [HJLS15], the back then available `draft-08` still had a not fully specified 0-RTT mode, leaving the implications of the results to the current draft of TLS 1.3 unclear. Remarkably, the model allows for replay attacks on the 0-RTT keys, as a consequence confining the admissible tests on keys, even though such replay attacks are excluded in QUIC via strike registers. Moreover, the latest drafts of TLS 1.3 generate more than two session keys, as considered in the model of [HJLS15].

Also recently, Cremers et al. [CHSvdM16] presented a tool-supported analysis of TLS 1.3 `draft-10` including 0-RTT mode, resumption, and delayed client authentication, discovering an attack on the interaction between the preshared-key handshake and the (expected) specification of client authentication.

Finally, our work is part of a substantial effort of the security research community in analyzing draft versions of TLS 1.3 prior to its standardization. We refer to Paterson and van der Merwe [PvdM16] for an overview over these analyses and a discussion of TLS 1.3's proactive standardization process.

2   Preliminaries

The key derivation in TLS 1.3 consists of a complex schedule of operations (cf. Figures 2 and 3) based on the HKDF key derivation function by Krawczyk [Kra10] which in turn is grounded on the HMAC scheme [BCK96, KBC97]. In particular for the analysis of the 0-RTT keys established in TLS 1.3 `draft-12` and `draft-14` we rely—besides the common assumptions about the collision resistance of hash functions,

unforgeability of signatures resp. MACs, or pseudorandomness properties of the key derivation function—on two more specific cryptographic assumptions about the security of (the extraction step of) HKDF, which we describe in the following.

The first assumption is induced by an intermediary HKDF extraction step in the TLS 1.3 `draft-14` key schedule where a value ES is derived from a pre-shared key PSK as $\text{ES} \leftarrow \text{HKDF.Extract}(0, \text{PSK}) = \text{HMAC}(0, \text{PSK})$. While this step can in general accommodate non-uniform pre-shared keys, those in our analysis are always uniformly random values $\text{PSK} \in \{0,1\}^\lambda$ and hence $\text{HMAC}(0, \text{PSK})$ is supposed to work as a uniform mapping from $\{0,1\}^\lambda$ to $\{0,1\}^\lambda$. This leads us to introducing the according assumption on HMAC (denoted $\mathsf{HMAC}(0,\$)\text{-}\$$) that $\mathsf{HMAC}(0, x)$ for an unknown, uniformly random value $x \xleftarrow{\$} \{0,1\}^\lambda$ is computationally indistinguishable from another uniformly random value $y \xleftarrow{\$} \{0,1\}^\lambda$. We formalize this assumption as follows.

Definition 2.1 ($\mathsf{HMAC}(0,\$)\text{-}\$$ assumption). *Let* HMAC *be the HMAC function as defined in [BCK96],* $\mathcal{A}$ *be a PPT algorithm, and* $x, y \xleftarrow{\$} \{0,1\}^\lambda$ *be two independent, uniformly random values.*

We define the advantage function

$$\mathsf{Adv}_{\mathsf{HMAC},\mathcal{A}}^{\mathsf{HMAC}(0,\$)\text{-}\$} := \left| \Pr\left[\mathcal{A}(1^\lambda, \mathsf{HMAC}(0, x)) = 1 \right] - \Pr\left[\mathcal{A}(1^\lambda, y) = 1 \right] \right|$$

and say that the $\mathsf{HMAC}(0,\$)\text{-}\$$ *assumption holds for* HMAC *if for any* $\mathcal{A}$ *the advantage function is negligible (as a function in* λ*).*

The second assumption, denoted $\mathsf{msPRF\text{-}ODH}$, is a double-sided variant of the pseudorandom-function oracle-Diffie–Hellman ($\mathsf{PRF\text{-}ODH}$) assumption [JKSS12], which itself is an adaptation of the oracle Diffie–Hellman assumption introduced by Abdalla et al. [ABR01] to the PRF setting. The $\mathsf{PRF\text{-}ODH}$ assumption has been used to analyze the security of the previous TLS version 1.2 DHE handshake by Jager et al. [JKSS12], the TLS 1.3 `draft-10` (EC)DHE and PSK handshakes by Dowling et al. [DFGS16] (in a single-query variant which we also employ in our analysis of the `draft-14` PSK-(EC)DHE 0-RTT handshake), and further TLS 1.2 handshake variants by Krawczyk et al. [KPW13] (in a multi-query variant forming the basis for our $\mathsf{msPRF\text{-}ODH}$ definition).[4] We make use of the $\mathsf{msPRF\text{-}ODH}$ assumption in the analysis of the `draft-12` (EC)DHE 0-RTT handshake in Section 7 where we also discuss its context and necessity in more detail.

Definition 2.2 ($\mathsf{msPRF\text{-}ODH}$ assumption). *Let* $\mathbb{G} = \langle g \rangle$ *be a cyclic group of prime order* q *with generator* g, $\mathsf{PRF} \colon \mathbb{G} \times \{0,1\}^* \to \{0,1\}^\lambda$ *be a pseudorandom function with keys in* $\mathbb{G}$, *input strings from* $\{0,1\}^*$, *and output strings of length* λ, *let* $b \in \{0,1\}$ *be a bit, and* $\mathcal{A}$ *be a PPT algorithm.*

We define the following $\mathsf{msPRF\text{-}ODH}$ *security game* $G_{\mathsf{PRF},\mathbb{G},\mathcal{A}}^{\mathsf{msPRF\text{-}ODH},b}$:

Setup. *The challenger chooses* $v \xleftarrow{\$} \mathbb{Z}_q$ *at random and gives* g^v *to* $\mathcal{A}$.

Query 1. *In the next phase* $\mathcal{A}$ *can ask queries of the form* $(g^u, x) \in (\mathbb{G}, \{0,1\}^*)$ *which the challenger answers with the value* $y \leftarrow \mathsf{PRF}((g^u)^v, x)$.[5]

Challenge. *At some point* $\mathcal{A}$ *asks a challenge query* $\hat{x} \in \{0,1\}^*$ *on which the challenger chooses* $\hat{u} \xleftarrow{\$} \mathbb{Z}_q$ *at random, sets* $\hat{y}_0 \leftarrow \mathsf{PRF}(g^{\hat{u}v}, \hat{x})$ *and* $\hat{y}_1 \xleftarrow{\$} \{0,1\}^\lambda$, *and answers with* $(g^{\hat{u}}, \hat{y}_b)$.

---

[4] Jager et al. [JKSS12] were able use the weaker single-query variant of the $\mathsf{PRF\text{-}ODH}$ assumption as they simulated protocol steps for an ephemeral Diffie–Hellman key only. Since we employ the assumption for the (re-usable) semi-static key, a stronger multi-query variant is necessary, as it was likewise to prove security of the non-ephemeral DH ciphersuite in TLS 1.2 [KPW13]. Moreover, we also need to enable a single query involving the ephemeral Diffie–Hellman share, rendering the assumption double-sided.

[5] We require that the first element is in $\mathbb{G}$ and hence write it as g^u, although $\mathcal{A}$ does not necessarily know u.

Query 2. *Again, $\mathcal{A}$ can ask queries of the form $(g^u, x) \in (\mathbb{G}, \{0,1\}^*)$ which the challenger answers with the value $y \leftarrow \mathsf{PRF}((g^u)^v, x)$, except that $\mathcal{A}$ is not allowed to query the pair $(g^{\hat{u}}, \hat{x})$.*

Additionally, and this is where our version differs from the previous $\mathsf{PRF\text{-}ODH}$ assumption, adversary $\mathcal{A}$ can ask one distinct query of the form $(g^{\hat{v}}, x) \in (\mathbb{G}, \{0,1\}^*)$ for $g^{\hat{v}} \neq g^v$ which the challenger answers with the value $y \leftarrow \mathsf{PRF}((g^{\hat{v}})^{\hat{u}}, x)$.

Guess. *Eventually, $\mathcal{A}$ stops and outputs a bit b' which is also the game output, denoted by $G^{\mathsf{msPRF\text{-}ODH},b}_{\mathsf{PRF},\mathbb{G},\mathcal{A}}$.*

We define the advantage function

$$\mathsf{Adv}^{\mathsf{msPRF\text{-}ODH}}_{\mathsf{PRF},\mathbb{G},\mathcal{A}} := \left| \Pr\left[G^{\mathsf{msPRF\text{-}ODH},0}_{\mathsf{PRF},\mathbb{G},\mathcal{A}} = 1 \right] - \Pr\left[G^{\mathsf{msPRF\text{-}ODH},1}_{\mathsf{PRF},\mathbb{G},\mathcal{A}} = 1 \right] \right|$$

and, assuming a sequence of groups in dependency of the security parameter, we say that the $\mathsf{msPRF\text{-}ODH}$ assumption holds for PRF with keys from $(\mathbb{G}_\lambda)_\lambda$ if for any $\mathcal{A}$ the advantage function is negligible (as a function in λ).

3   Modeling Replayable 0-RTT in Multi-Stage Key Exchange

In this section we recap the *multi-stage key exchange* model introduced by Fischlin and Günther [FG14], and later extended by Dowling et al. [DFGS15a, DFGS16], and augment it for capturing replayable 0-RTT keys. To capture security of the `draft-12` (EC)DHE 0-RTT handshake (where authentication is provided through long-term signing keys) as well as the `draft-14` PSK-based 0-RTT handshakes, we treat both the public-key (MSKE) and the preshared-secret (MS-PSKE) variant of the model. We start by outlining the original multi-stage setting and the modifications we introduce. We then define the security model formally.

3.1   Outline of the Model for Multi-Stage Key Exchange

The multi-stage key exchange model follows the game-based paradigm of Bellare and Rogaway [BR94]. That is, the adversary controls the network over which the parties communicate, giving the adversary the power to read and alter protocol messages in transmission. For this the adversary can call a NewSession oracle (starting a new session of a specified honest party) and a Send oracle (which delivers some message to a specified party). To argue about the secrecy of keys the adversary may make (multiple) Test queries for some stage of the protocol to either receive the corresponding session key of that stage or to obtain an independent random key instead. In the multi-stage setting one must restrict the set of admissible Test queries to avoid trivial attacks, e.g., if a tested session key is used in a later stage of the execution.

To model leakage of long-term secrets and, through this, forward secrecy of keys (i.e., security after long-term secret compromise) we also grant the adversary access to a Corrupt oracle which returns the corresponding secret key. Moreover, we (independently) capture leakage of semi-static keys (used in Diffie–Hellman-based 0-RTT key derivation as, e.g., TLS 1.3 `draft-12` (EC)DHE 0-RTT) through a separate RevealSemiStaticKey oracle. In our model we do not consider leakage of internal state of the parties (such as randomness or master secrets) but note that one can in principle enhance the model further to capture such attacks. Still, we allow for leakage of session keys, modeling insecure usage of such keys in follow-up protocol steps (of the key exchange protocol itself or in subsequent communication protocols). As in the common Bellare–Rogaway model we must prohibit compromise of secrets in sessions partnered with tested sessions. Here, partners are identified via session identifiers.

Session identifiers in the multi-stage setting are more elaborate than in the single-stage Bellare–Rogaway case. Depending on the protocol, revealing a session key via a Reveal query may render the

subsequent session key insecure, e.g., if contributions to the session key of the following stage are sent authenticated under the key of the current stage (as in QUIC [LC15, FG14]). A protocol which can tolerate such leakage is called (session-)key independent, else it is called key dependent.

For TLS 1.3, Dowling et al. [DFGS15a] refined the original model [FG14] (beyond introducing the preshared-secret variant) further to cover the various authentication properties of the different handshakes and also of the different stages. For this they distinguish between unauthenticated, unilaterally authenticated, and mutually authenticated stages, and treat security of multiple sessions running in parallel with different authentication modes. They also implemented the common TLS property of post-specified peers [CK02] via wildcards '∗' for intended communication partners, which can be set once throughout the protocol run, e.g., after the party has verified the certificate of the partner.

Another change from [FG14] to [DFGS15a] concerns the notion of contributive identifiers for the case of sessions with unauthenticated partners. Since the adversary can potentially impersonate the unauthenticated partner, keys in such sessions cannot be secure in general. Still, such keys should be considered secure as long as the full contribution to the key clearly stems from an honest party (even though this honest party may never complete its execution to output a session identifier). Contributive identifiers allow to specify such partnered contributions.

Both works [FG14, DFGS15a] follow the approach of Brzuska et al. [BFWW11] to split the security requirements into one for Match security, capturing among others uniqueness of session identifiers, correspondence of session identifiers and keys, and linking contributive identifiers to session identifiers, and into the common requirement for key secrecy.

3.2 Adding 0-RTT to Multi-Stage Protocols

To capture 0-RTT in the multi-stage setting we augment the model in [DFGS15a, DFGS16] by the following points:

1. We introduce the notion of replayable stages.

2. We allow exposure of the semi-static keys used for establishing 0-RTT keys. Note that the exposure of pre-shared keys (used for PSK-based 0-RTT) is already captured through Corrupt queries.

3. We distinguish between external session keys (used for protecting the application layer only) and internal session keys (which can be used within the key exchange protocol). Note that since, by default, session keys are always output by key exchange protocols, internal keys may thus also be used further in applications, of course.

As mentioned in the introduction, replayable stages in a multi-stage key exchange protocol are basically those stages in which an adversary can force more than two sessions to share the same session identifiers and session keys by replaying previous interactions. Note that for 0-RTT, replayability is inevitable for stateless parties, whereas Google's QUIC protocol in version `Rev 20130620` thwarts such replay attacks via strike registers storing information about previous connections.

In our model we assume that the protocol specifies stages as replayable or non-replayable. The latter type of stage leaves the original security properties untouched. The replayable kind of stage allows for multiple collisions among session identifiers and keys, such that we need to relax the notion of Match security for such stages. Key secrecy should be not affected by replayability because the adversary may be able to foist the same key on multiple sessions, but the key itself should still look random.

The other modification refers to exposure of semi-static keys. Recall that such keys are used in TLS 1.3 (up to `draft-12`) to enable 0-RTT for the DH-based mode by having the client mix a fresh ephemeral key to such a semi-static key. The life span of such a semi-static key may range over multiple sessions, and we therefore leave the choice of which of these keys to use (and when to create it) to the adversary.

This is modeled by augmenting the NewSession query by a field for specifying the semi-static key, and by having a NewSemiStaticKey command to let a party create a fresh semi-static key. Leakage of semi-static keys is captured by adding a RevealSemiStaticKey query to the model through which the adversary learns the corresponding secret key. The previous works [FG14, DFGS15a] have not yet supported such leakage, even though [FG14] already introduced the equivalent idea of a temporary key for analyzing QUIC.

Note that we separate leakage of the semi-static key by RevealSemiStaticKey queries from revealing the long-term secret via Corrupt queries. This corresponds to the setting where a semi-static key may actually no longer be used by a server and its secret key been irrevocably erased. The adversary can, of course, always mount RevealSemiStaticKey commands before a Corrupt request, thus enhancing the adversary's capabilities and strengthening the security claims.

The explicit distinction between internal and external session keys implements a cleaner way to deal with early derived keys used exclusively in, e.g., the TLS record protocol. In contrast to QUIC, where only the final session key is not used in the key exchange messages, the 0-RTT step of TLS 1.3 allows both parties to immediately establish the early handshake and early application-data traffic key tk_{ehs} and tk_{ead}, with a clear separation of concerns. The former early handshake key is used to protect the handshake messages (internally), and the early application-data key is only used to protect external application data sent by the client before finishing the full handshake.

3.3 Preliminaries

In contrast to previous works formalizing multi-stage key exchange [FG14, DFGS15a, DFGS16], we explicitly separate some *protocol-specific* properties (as, e.g., various authentication flavors) from *session-specific* properties (as, e.g., the state of a running session). We represent protocol-specific properties as a vector $(M, AUTH, USE, REPLAY)$ that captures the following:

- $M \in \mathbb{N}$: the number of stages (i.e., the number of keys derived)[6]

- $AUTH \subseteq \{unauth, unilateral, mutual\}^M$: the set of supported authentication properties (for each stage). As in [DFGS15a] we call stages and keys *unauthenticated* if they provide no authentication for either communication partner, *unilaterally authenticated* if they authenticate only the responder (server) side, and *mutually authenticated* if they authenticate both communication partners.

- $USE \in \{internal, external\}^M$: the usage indicator for each stage, where USE_i indicates the usage of the stage-i key. Here, an internal key is used within the key exchange protocol (but possibly also externally), whereas an external key must not be used within the protocol, making the latter potentially amenable to generic composition (see Section 9).

- $REPLAY \in \{replayable, nonreplayable\}^M$: the replayability indicator for each stage, where $REPLAY_i$ indicates whether the i-th stage is replayable in the sense that a party can easily force identical communication and thus identical session identifiers and keys in this stage (e.g., re-sending the same data in 0-RTT stages). Note that the adversary, however, should still not able to distinguish such a replayed key from a random one. Note that, from a security viewpoint, the usage of replayable stages should ideally be small, whereas such stages usually come with an efficiency benefit.

To give a concrete example, the TLS 1.3 `draft-14` PSK 0-RTT handshake is a multi-stage key exchange protocol with the following properties:

---

[6]We fix a maximum stage M only for ease of notation. Note that M can be arbitrarily large in order to cover protocols where the number of stages is not bounded a-priori. Also note that, for technical convenience, stages and session keys may be "back to back," without further protocol interactions between parties.

- it has five stages ($M = 5$),

- it provides mutual authentication for all five keys derived ($\mathsf{AUTH} = \{(\mathsf{mutual}, \mathsf{mutual}, \mathsf{mutual}, \mathsf{mutual}, \mathsf{mutual})\}$),

- it uses keys of stages 1 and 3 internally within the key exchange protocol and the other keys only externally, exporting them to other protocols $\mathsf{USE} = (\mathsf{internal}, \mathsf{external}, \mathsf{internal}, \mathsf{external}, \mathsf{external})$),

- and the initial two stages are replayable ($\mathsf{REPLAY} = (\mathsf{replayable}, \mathsf{replayable}, \mathsf{nonreplayable}, \mathsf{nonreplayable}, \mathsf{nonreplayable})$).

To formalize sessions we denote by $\mathcal{U}$ the set of *identities* used to model the participants in the system, each identified by some $U \in \mathcal{U}$ Sessions of a protocol are uniquely identified (on the administrative level of the model) using a *label* $\mathsf{label} \in \mathsf{LABELS} = \mathcal{U} \times \mathcal{U} \times \mathbb{N}$, where (U, V, k) indicates the k-th local session of identity U (the session *owner*) with V as the intended communication *partner*.

In the public-key variant of the model (MSKE), each identity U is associated with a certified *long-term* public key pk_U and secret key sk_U. In the preshared-secret setting (MS-PSKE), a session instead holds a key index for the preshared secret pss (and its unique identifier psid) used. The challenger maintains vectors $\vec{\mathsf{pss}}_{U,V}$ and $\vec{\mathsf{psid}}_{U,V}$ of preshared secrets created on adversary demand, with the k-th entry indicating the k-th secret resp. corresponding identifier shared by parties U and V.

In addition to the long-term keys, parties may (in the public-key setting) also hold certified *semi-static* key pairs $(\mathsf{sspk}, \mathsf{sssk})$, each identified by a semi-static key identifier sskid.[7] Semi-static keys moreover are associated with some auxiliary data sskaux which may for example carry the data structure in which a party learns the semi-static key (in TLS 1.3, this is the `ServerConfiguration` message and other identifiers). Finally, a flag $\mathsf{st}_{\mathsf{ssk},\mathsf{sskid}} \in \{\mathsf{fresh}, \mathsf{revealed}\}$ indicates whether a semi-static key has been revealed to the adversary or not. This flag is convenient since in our model semi-static keys are linked to replayable stages, especially to 0-RTT stages (as is common practice), such that we consider the disclosure of such keys inevitably rendering these session key to be insecure.

For each session, a tuple with the following information is maintained as an entry in the *session list* List_S, where values in square brackets $[\,]$ indicate the default/initial value. Some variables have values for each stage $i \in \{1, \ldots, M\}$.

- $\mathsf{label} \in \mathsf{LABELS}$: the unique (administrative) session label

- $U \in \mathcal{U}$: the session owner

- $V \in (\mathcal{U} \cup \{*\})$: the intended communication partner, where the distinct wildcard symbol '$*$' stands for "unknown identity" and can be set to a specific identity in $\mathcal{U}$ *once* by the protocol

- $\mathsf{role} \in \{\mathsf{initiator}, \mathsf{responder}\}$: the session owner's role in this session

- $\mathsf{auth} \in \mathsf{AUTH}$: the intended authentication type (for each stage) from the set of supported authentication properties AUTH, where auth_i indicates the authentication level in stage $i > 0$

- $\mathsf{st}_{\mathsf{exec}} \in (\mathsf{RUNNING} \cup \mathsf{ACCEPTED} \cup \mathsf{REJECTED})$: the state of execution $[\mathsf{running}_0]$, where $\mathsf{RUNNING} = \{\mathsf{running}_i \mid i \in \mathbb{N}_0\}$, $\mathsf{ACCEPTED} = \{\mathsf{accepted}_i \mid i \in \mathbb{N}\}$, $\mathsf{REJECTED} = \{\mathsf{rejected}_i \mid i \in \mathbb{N}\}$

- $\mathsf{stage} \in \{0, \ldots, M\}$: the current stage $[0]$, where stage is incremented to i when $\mathsf{st}_{\mathsf{exec}}$ reaches $\mathsf{accepted}_i$ resp. $\mathsf{rejected}_i$

---

[7]While previous multi-stage key exchange models [FG14, DFGS15a] denoted those keys as "temporary keys," we adopt the more common notation of "semi-static keys" here which is also used in TLS 1.3.

- $\mathsf{sid} \in (\{0,1\}^* \cup \{\bot\})^{\mathsf{M}}$: sid_i $[\bot]$ indicates the session identifier in stage $i > 0$

- $\mathsf{cid} \in (\{0,1\}^* \cup \{\bot\})^{\mathsf{M}}$: cid_i $[\bot]$ indicates the contributive identifier in stage $i > 0$

- $\mathsf{K} \in (\{0,1\}^* \cup \{\bot\})^{\mathsf{M}}$: K_i $[\bot]$ indicates the established session key in stage $i > 0$

- $\mathsf{st_{key}} \in \{\mathsf{fresh}, \mathsf{revealed}\}^{\mathsf{M}}$: $\mathsf{st_{key},i}$ $[\mathsf{fresh}]$ indicates the state of the session key in stage $i > 0$

- $\mathsf{tested} \in \{\mathsf{true}, \mathsf{false}\}^{\mathsf{M}}$: test indicator tested_i $[\mathsf{false}]$, where true means that K_i has been tested

In the public-key (MSKE) variant, $\mathsf{List_S}$ furthermore contains the following entries:

- sskid_U: the key identifier for the semi-static key pair $(\mathsf{sspk}, \mathsf{sssk})$ used by the session owner ($\bot$ if no key is used)

- sskid_V: the semi-static key identifier for the communication partner ($\bot$ if no key is used)

In the preshared-secret (MS-PSKE) variant, $\mathsf{List_S}$ instead contains the following extra entries:

- $k \in \mathbb{N}$: the index of the preshared secret used in a protocol run with the communication partner

- $\mathsf{pss} \in (\{0,1\}^* \cup \{\bot\})$: the preshared secret to be used in the session

- $\mathsf{psid} \in (\{0,1\}^* \cup \{\bot\})$: the preshared secret identifier of the preshared secret to be used in the session

By convention adding a not fully specified tuple $(\mathsf{label}, U, V, \mathsf{role}, \mathsf{auth}, \mathsf{sskid}_U, \mathsf{sskid}_V)$ resp. $(\mathsf{label}, U, V, \mathsf{role}, \mathsf{auth}, k, \mathsf{pss}, \mathsf{psid})$ to $\mathsf{List_S}$ sets all other entries to their default value. We furthermore write, e.g., $\mathsf{label.sid}$ as shorthand for the element sid in the tuple with (unique) label label in $\mathsf{List_S}$.

3.4 Adversary Model

We consider a probabilistic polynomial-time adversary $\mathcal{A}$ which controls the communication between all parties, enabling interception, injection, and dropping of messages.

As in [FG14, DFGS15a] our model below treats forward and "regular" secrecy simultaneously. A key is called forward secret if it stays secure even if the long-term secrets involved in its derivation are later revealed. We say that a protocol provides *stage-k forward secrecy* if keys from the k-th stage on are forward secret. Note that forward secrecy refers (only) to disclosure of long-term secrets: if, in the public-key variant, semi-static keys are revealed we in any case expect keys to look random, except for the ones in replayable stages. Vice versa, we expect that the compromise of long-term secrets alone (i.e., without also exposing the semi-static key involved) does not affect keys in (forward-secret) replayable stages. Our experiment furthermore includes a flag lost (initialized to false) that captures admissibility of adversarial interactions and conditions where the adversary trivially loses (such as both revealing and testing the session key in partnered sessions).

The adversary interacts with the protocol via the following queries, which mostly (and sometimes in verbatim) operate as in [FG14, DFGS15a]:

- $\mathsf{NewSecret}(U, V, k, \mathsf{psid})$: This query is only available in the preshared-secret (MS-PSKE) variant. Generates a fresh preshared secret with identifier psid shared as k-th secret between parties U and V. If there already is a k-th entry in $\vec{\mathsf{pss}}_{U,V}$ or if psid is already registered for any other pss, return $\bot$. The latter ensures global uniqueness of the psid value. Otherwise, sample pss uniformly at random and store pss and psid as the k-th entry in $\vec{\mathsf{pss}}_{U,V}$ and $\vec{\mathsf{pss}}_{V,U}$ resp. in $\vec{\mathsf{psid}}_{V,U}$ and $\vec{\mathsf{psid}}_{V,U}$.

- NewSemiStaticKey$(U, \mathsf{sskaux}_{\mathsf{pre}})$: This query is only available in the public-key (MSKE) variant. Generates a new semi-static key pair $(\mathsf{sspk}, \mathsf{sssk})$ for identity U with associated (protocol-defined) auxiliary data sskaux (which might include some adversarially pre-specified parts $\mathsf{sskaux}_{\mathsf{pre}}$) and a (unique) new identifier sskid. Set $\mathsf{st}_{\mathsf{ssk,sskid}} \leftarrow \mathsf{fresh}$ and return the tuple $(\mathsf{sskid}, \mathsf{sspk}, \mathsf{sskaux})$.

- NewSession$(U, V, \mathsf{role}, \mathsf{auth}, \mathsf{sskid}_U, \mathsf{sskid}_V)$ or NewSession$(U, V, \mathsf{role}, \mathsf{auth}, k)$: The first query is only used in the public-key (MSKE) variant, the second query only in the preshared-secret (MS-PSKE) variant. Creates a new session with a (unique) new label label for participant identity U with role role, having V as intended partner (potentially unspecified, indicated by $V = *$) and aiming at authentication type $\mathsf{auth} \in \mathsf{AUTH}$.

  In the public-key variant, sskid_U and sskid_V indicate the semi-static key identifiers used by each side. Either semi-static key identifier can also be left unspecified ($\mathsf{sskid}_U = \bot$ resp. $\mathsf{sskid}_V = \bot$), indicating that the according party does not use such key. In the preshared secret variant, k indicates the key index of the shared pss between U and V. If no such pss has been registered, return $\bot$. Otherwise, set $\mathsf{label.pss}$ to pss and $\mathsf{label.psid}$ to the corresponding k-th entry of $\vec{\mathsf{psid}}_{U,V}$. Add $(\mathsf{label}, U, V, \mathsf{role}, \mathsf{auth}, \mathsf{sskid}_U, \mathsf{sskid}_V)$ resp. $(\mathsf{label}, U, V, \mathsf{role}, \mathsf{auth}, k, \mathsf{pss}, \mathsf{psid})$ to $\mathsf{List_S}$ and return label.

- Send(label, m): Sends a message m to the session with label label.

  If there is no tuple with label label in $\mathsf{List_S}$, return $\bot$. Otherwise, run the protocol on behalf of U on message m and return the response and the updated state of execution $\mathsf{label.st_{exec}}$. As a special case, if $\mathsf{label.role} = \mathsf{initiator}$ and $m = \mathsf{init}$, the protocol is initiated (without any input message).

  If, during the protocol execution, the state of execution changes to $\mathsf{accepted}_i$ for some i, the protocol execution is immediately suspended and $\mathsf{accepted}_i$ is returned as result to the adversary. The adversary can later trigger the resumption of the protocol execution by issuing a special Send$(\mathsf{label}, \mathsf{continue})$ query. For such a query, the protocol continues as specified, with the party creating the next protocol message and handing it over to the adversary together with the resulting state of execution $\mathsf{st_{exec}}$. We note that this is necessary to allow the adversary to test such a key, before it may be used immediately in the response and thus cannot be tested anymore for triviality reasons.

  If the state of execution changes to $\mathsf{label.st_{exec}} = \mathsf{accepted}_i$ for some i and there is a partnered session label' in $\mathsf{List_S}$ (i.e., $\mathsf{label.sid}_i = \mathsf{label}'.\mathsf{sid}_i$) with $\mathsf{label}'.\mathsf{st}_{\mathsf{key},i} = \mathsf{revealed}$, then, for key independence, $\mathsf{label.st}_{\mathsf{key},i}$ is set to $\mathsf{revealed}$ as well, whereas for key-dependent security, all $\mathsf{label.st}_{\mathsf{key},i'}$ for $i' \geq i$ are set to $\mathsf{revealed}$. The former corresponds to the case that session keys of partnered sessions should be considered revealed as well, the latter implements that for key dependency all subsequent keys are potentially available to the adversary, too.

  If the state of execution changes to $\mathsf{label.st_{exec}} = \mathsf{accepted}_i$ for some i and there is a partnered session label' in $\mathsf{List_S}$ (i.e., $\mathsf{label.sid}_i = \mathsf{label}'.\mathsf{sid}_i$) with $\mathsf{label}'.\mathsf{tested}_i = \mathsf{true}$, then set $\mathsf{label.tested}_i \leftarrow \mathsf{true}$ and (only if $\mathsf{USE}_i = \mathsf{internal}$) $\mathsf{label.K}_i \leftarrow \mathsf{label}'.\mathsf{K}_i$. This ensures that, if the partnered session has been tested before, subsequent Test queries for the session are answered accordingly and, in case it is used internally, this session's key K_i is set consistently[8].

  If the state of execution changes to $\mathsf{label.st_{exec}} = \mathsf{accepted}_i$ for some i and the intended communication partner $V \neq *$ is corrupted, then set $\mathsf{label.st}_{\mathsf{key},i} \leftarrow \mathsf{revealed}$.

- Reveal(label, i): Reveals $\mathsf{label.K}_i$, the session key of stage i in the session with label label.

---

[8]Note that for internal keys this implicitly assumes the following property of the later-defined Match security: Whenever two partnered sessions both accept a key in some stage, these keys will be equal.

If there is no session with label label in $\mathsf{List_S}$, or $\mathsf{label.stage} < i$, or $\mathsf{label.tested}_i = \mathsf{true}$, then return $\bot$. Otherwise, set $\mathsf{label.st_{key}}_{,i}$ to revealed and provide the adversary with $\mathsf{label.K}_i$.

If there is a partnered session $\mathsf{label'}$ in $\mathsf{List_S}$ (i.e., $\mathsf{label.sid}_i = \mathsf{label'.sid}_i$) with $\mathsf{label'.stage} \geq i$, then $\mathsf{label'.st_{key}}_{,i}$ is set to revealed as well. This means the i-th session keys of all partnered sessions (if established) are considered revealed too.

As above, in the case of key-dependent security, since not yet established future keys depend on the revealed key, we cannot ensure their security anymore (neither in this session in question, nor in partnered sessions). Therefore, if $\mathsf{label.stage} = i$, set $\mathsf{label.st_{key}}_{,j} = \mathsf{revealed}$ for all $j > i$, as they depend on the revealed key. For the same reason, if a partnered session $\mathsf{label'}$ ($\mathsf{label.sid}_i = \mathsf{label'.sid}_i$) has $\mathsf{label'.stage} = i$, then set $\mathsf{label'.st_{key}}_{,j} = \mathsf{revealed}$ for all $j > i$. Note that if however $\mathsf{label'.stage} > i$, then keys $\mathsf{label'.K}_j$ for $j > i$ derived in the partnered session are not considered to be revealed by this query since they have been accepted previously, i.e., prior to K_i being revealed in this query.

- RevealSemiStaticKey(sskid): This query is only available in the public-key (MSKE) variant. If there exists a semi-static key pair $(\mathsf{sspk}, \mathsf{sssk})$ with identifier sskid, set $\mathsf{st_{sssk,sskid}} \leftarrow \mathsf{revealed}$, and output sssk.

  Furthermore, for each session label with $\mathsf{label.sskid}_U = \mathsf{sskid}$ or $\mathsf{label.sskid}_V = \mathsf{sskid}$ and all replayable stages $i \in \{1, \dots, \mathsf{M}\}$ with $\mathsf{REPLAY}_i = \mathsf{replayable}$, set $\mathsf{label.st_{key}}_{,i}$ to revealed. That is, any replayable stage's session key in a session that uses the revealed semi-static key is considered to be disclosed.

- Corrupt(U) or Corrupt(psid): The first query is only used in the public-key (MSKE) variant, the second query only in the preshared-secret (MS-PSKE) variant. Provide the adversary with the corresponding long-term secret, i.e., sk_U (MSKE) resp. pss corresponding to psid (MS-PSKE). No further queries are allowed to sessions owned by U (MSKE) resp. to any session label with $\mathsf{label.psid} = \mathsf{psid}$ (MS-PSKE).

  In the non-forward-secret case, for each session label owned by U (MSKE) resp. holding $\mathsf{label.psid} = \mathsf{psid}$ (MS-PSKE) and for all $i \in \{1, \dots, \mathsf{M}\}$, set $\mathsf{label.st_{key}}_{,i}$ to revealed. In this case, all (previous and future) session keys are considered to be disclosed.

  In the case of stage-j forward secrecy, $\mathsf{st_{key}}_{,i}$ of such sessions label is instead set to revealed only if $i < j$ or if $i > \mathsf{stage}$. This means that session keys before the j-th stage (where forward secrecy kicks in) as well as keys that have not yet been established are potentially disclosed.

  Independent of the forward secrecy aspect, in the case of key-dependent security, setting the relevant key states to revealed for some stage i is done by internally invoking $\mathsf{Reveal}(\mathsf{label}, i)$, ignoring the response and also the restriction that a call with $i > \mathsf{stage}$ would immediately return $\bot$. This ensures that follow-up revocations of keys that depend on the revoked keys are carried out correctly.

- Test(label, i): Tests the session key of stage i in the session with label label. In the security game this oracle is given a uniformly random test bit b_test as state which is fixed throughout the game.

  If there is no session with label label in $\mathsf{List_S}$ or if $\mathsf{label.st_{exec}} \neq \mathsf{accepted}_i$ or $\mathsf{label.tested}_i = \mathsf{true}$, return $\bot$. If there is a partnered session $\mathsf{label'}$ in $\mathsf{List_S}$ (i.e., $\mathsf{label.sid}_i = \mathsf{label'.sid}_i$) with $\mathsf{label'.st_{exec}} \neq \mathsf{accepted}_i$, set the 'lost' flag to $\mathsf{lost} \leftarrow \mathsf{true}$. This ensures that keys can only be tested once and if they have just been accepted but not used yet, including ensuring any partnered session that may have already established this key has not used it.

  If $\mathsf{label.auth}_i = \mathsf{unauth}$ or if $\mathsf{label.auth}_i = \mathsf{unilateral}$ and $\mathsf{label.role} = \mathsf{responder}$, but there is no session $\mathsf{label'}$ (for $\mathsf{label} \neq \mathsf{label'}$) in $\mathsf{List_S}$ with $\mathsf{label.cid}_i = \mathsf{label'.cid}_i$, then set $\mathsf{lost} \leftarrow \mathsf{true}$. This ensures that having an honest contributive partner is a prerequisite for testing responder sessions

in an unauthenticated or unilaterally authenticated stage and for testing an initiator session in an unauthenticated stage.[9]

Otherwise, set $\mathsf{label.tested}_i$ to true. If the test bit b_test is 0, sample a key $\mathsf{K} \xleftarrow{\$} \mathcal{D}$ at random from the session key distribution $\mathcal{D}$. If $b_\mathsf{test} = 1$, let $\mathsf{K} \leftarrow \mathsf{label.K}_i$ be the real session key. If $\mathsf{USE}_i = \mathsf{internal}$ (i.e., the tested i-th key is indicated as being used internally), set $\mathsf{label.K}_i \leftarrow \mathsf{K}$, i.e., we substitute an *internally* used session key by the random and independent test key K which is also used for consistent future deployments *within* the key exchange protocol. In contrast, *externally used* session keys are not replaced by random ones, the adversary only receives the real (in case $b_\mathsf{test} = 1$) or random (in case $b_\mathsf{test} = 0$) key. This distinction between internal and external keys for Test queries emphasizes that external keys are not supposed to be used within the key exchange (and hence there is no need to register the tested random key in the protocol's session key field) while internal keys will be used (and hence the tested random key must be deployed in the remaining protocol steps for consistency).

Moreover, if there exists a partnered session label' which has also just accepted the i-th key (i.e., $\mathsf{label.sid}_i = \mathsf{label}'.\mathsf{sid}_i$ and $\mathsf{label.st_{exec}} = \mathsf{label}'.\mathsf{st_{exec}} = \mathsf{accepted}_i$), then also set $\mathsf{label}'.\mathsf{tested}_i \leftarrow \mathsf{true}$ and (only if $\mathsf{USE}_i = \mathsf{internal}$) $\mathsf{label}'.\mathsf{K}_i \leftarrow \mathsf{label.K}_i$ to ensure consistency (of later tests and (internal) key usage) in the special case that both label and label' are in state $\mathsf{accepted}_i$ and, hence, either of them can be tested first.

Return K.

Secret compromise paradigm. Our model captures the leakage of long-term secret keys and output session keys as well as the leakage of semi-static keys (in the public-key variant), following the paradigm of Bellare and Rogaway [BR94] to capture the compromise of long(er)-lived secret inputs and key outputs of a key exchange protocol. We note that other classical key exchange models by Canetti and Krawczyk [CK01] resp. LaMacchia et al. [LLM07] further capture the leakage of internal session state or ephemeral secret inputs, which we do not, and against most of which TLS 1.3 also does not aim to protect.

Revisiting and adopting the discussion on the TLS 1.3 full handshake by Dowling et al. [DFGS15a], this means that in the context of the TLS 1.3 0-RTT handshakes we consider the leakage of:

- *Long-term keys* (i.e., the signing key of the server or client and their preshared (resumption) secrets). This is allowed since usage of keys over a long time period induces a substantial risk of compromise, spawning the notion of forward secrecy. Leakage of long-term keys is modeled by the $\mathsf{Corrupt}$ query.

- *Session keys* (i.e., tk_{ehs}, tk_{ead}, tk_{hs}, tk_{app}, RMS, and EMS).
  This is allowed since session keys are the actual output of a key exchange, used in a follow-up protocol (e.g., for encryption, resumption, or key export in TLS 1.3) or internally in a subsequent key exchange step (treated in our key independence notion). Leakage of session keys is modeled by the Reveal query.

- *Semi-static keys* (i.e., s in g^s for the DH-based 0-RTT mode).
  This is allowed since these keys are meant to be used (by servers) in several connections with multiple clients and over a potentially significant time span. Furthermore, their leakage affect the security of derived replayable keys. Leakage of semi-static keys is modeled by the $\mathsf{RevealSemiStaticKey}$ query.

We do not permit the leakage of:

---

[9]Note that $\mathsf{List_S}$ entries are only created for honest sessions, i.e., sessions generated by $\mathsf{NewSession}$ queries.

- *Ephemeral secrets* (e.g., the randomness in the signature algorithm or the (ephemeral) Diffie–Hellman exponents x and y).

  This is disallowed as TLS 1.3 does not aim (neither achieves) being secure against compromises of this kind.[10]

- *Internal values / session state* (e.g., the master secret MS or internal MAC keys FS_C, FS_S).

  This is disallowed as, again, TLS 1.3 does not provide protection against such leakage.

3.5 Security of Multi-Stage Key Exchange Protocols

The security properties for multi-stage key exchange protocols are almost identical with those given for the TLS 1.3 full and resumption handshake analysis by Dowling et al. [DFGS15a]; split into two games following Fischlin and Günther [FG14] and Brzuska et al. [BFWW11, Brz13]. On the one hand, Match security ensures that the session identifiers sid effectively match the partnered sessions. On the other hand, Multi-Stage security ensures Bellare–Rogaway-like key secrecy.

For the analysis of the TLS 1.3 0-RTT handshakes, most changes in the model are already reflected in the specification of the adversarial queries. Beyond introducing the new protocol-specific properties and the RevealSemiStaticKey query in the definition, we only extend the Match security definition in order to allow for multiple sessions being partnered in replayable stages.

3.5.1 Match Security

The notion of Match security—which we adapt from [DFGS15a, DFGS16] to capture replayable stages—ensures soundness of the session identifiers sid, i.e., that they properly identify the partnered sessions in the sense that

1. sessions with the same session identifier for some stage hold the same key at that stage,

2. sessions with the same session identifier for some stage agree on that stage's authentication level,

3. sessions with the same session identifier for some stage share the same contributive identifier at that stage,

4. sessions are partnered with the intended (authenticated) participant, and for mutual authentication share the same key index,

5. session identifiers do not match across different stages, and

6. at most two sessions have the same session identifier at any non-replayable stage.

The Match security game $G_{KE,\mathcal{A}}^{Match}$ thus is defined as follows.

Definition 3.1 (Match security). *Let* KE *be a multi-stage key exchange protocol with properties* (M, AUTH, USE, REPLAY) *and a PPT adversary* $\mathcal{A}$ *interacting with* KE *via the queries defined in Section 3.4 in the following game* $G_{KE,\mathcal{A}}^{Match}$*:*

---

[10]To be precise, the OPTLS key exchange protocol underlying the TLS 1.3 handshake actually maintains secrecy of the application traffic key tk_{app} even under exposure of the *server's* ephemeral Diffie–Hellman exponent y, given that the client's exponent x as well as the server's semi-static key exponent s remain secret, as analyzed by Krawczyk and Wee [KW15, KW16]. We omitted capturing one-sided ephemeral secret leakage in our model in order to not further increase its complexity, but conjecture a similar result can be obtained (for tk_{app}, RMS, and EMS) in the multi-stage setting for the TLS 1.3 (EC)DHE 0-RTT handshake in `draft-12`.

Setup. *In the public-key variant (MSKE), the challenger generates long-term public/private-key pairs for each participant $U \in \mathcal{U}$.*

Query. *The adversary $\mathcal{A}$ receives the generated public keys (MSKE) and has access to the queries* NewSecret, NewSemiStaticKey, NewSession, Send, Reveal, RevealSemiStaticKey, *and* Corrupt.

Stop. *At some point, the adversary stops with no output.*

We say that $\mathcal{A}$ wins the game, denoted by $G^{\mathsf{Match}}_{\mathsf{KE},\mathcal{A}} = 1$, if at least one of the following conditions hold:

1. *There exist two distinct labels* label, label$'$ *such that* label.sid$_i$ = label$'$.sid$_i \neq \bot$ *for some stage $i \in \{1, \ldots, \mathsf{M}\}$,* label.st$_{\mathsf{exec}} \neq$ rejected$_i$, *and* label$'$.st$_{\mathsf{exec}} \neq$ rejected$_i$, *but* label.K$_i \neq$ label$'$.K$_i$. *(Different session keys in some stage of partnered sessions.)*

2. *There exist two distinct labels* label, label$'$ *such that* label.sid$_i$ = label$'$.sid$_i \neq \bot$ *for some stage $i \in \{1, \ldots, \mathsf{M}\}$, but* label.auth$_i \neq$ label$'$.auth$_i$. *(Different authentication types in some stage of partnered sessions.)*

3. *There exist two distinct labels* label, label$'$ *such that* label.sid$_i$ = label$'$.sid$_i \neq \bot$ *for some stage $i \in \{1, \ldots, \mathsf{M}\}$, but* label.cid$_i \neq$ label$'$.cid$_i$ *or* label.cid$_i$ = label$'$.cid$_i = \bot$. *(Different or unset contributive identifiers in some stage of partnered sessions.)*

4. *There exist two distinct labels* label, label$'$ *such that* label.sid$_i$ = label$'$.sid$_i \neq \bot$ *for some stage $i \in \{1, \ldots, \mathsf{M}\}$,* label.auth$_i$ = label$'$.auth$_i \in \{\mathsf{unilateral}, \mathsf{mutual}\}$, label.role = initiator, *and* label$'$.role = responder, *but* label.$V \neq$ label$'$.U *or (only if* label.auth$_i$ = mutual*)* label.$U \neq$ label$'$.V *or (only for MS-PSKE and if* label.auth$_i$ = mutual*)* label.$k \neq$ label$'$.k. *(Different intended authenticated partner or (only MS-PSKE) different key indices in mutual authentication.)*

5. *There exist two (not necessarily distinct) labels* label, label$'$ *such that* label.sid$_i$ = label$'$.sid$_j \neq \bot$ *for some stages $i, j \in \{1, \ldots, \mathsf{M}\}$ with $i \neq j$. (Different stages share the same session identifier.)*

6. *There exist three distinct labels* label, label$'$, label$''$ *such that* label.sid$_i$ = label$'$.sid$_i$ = label$''$.sid$_i \neq \bot$ *for some stage $i \in \{1, \ldots, \mathsf{M}\}$ with* REPLAY$_i$ = nonreplayable. *(More than two sessions share the same session identifier in a non-replayable stage.)*

We say KE *is* Match-*secure if for all PPT adversaries $\mathcal{A}$ the following advantage function is negligible in the security parameter:*

$$\mathsf{Adv}^{\mathsf{Match}}_{\mathsf{KE},\mathcal{A}} := \Pr\left[G^{\mathsf{Match}}_{\mathsf{KE},\mathcal{A}} = 1\right].$$

3.5.2 Multi-Stage Security

The second notion, Multi-Stage security, captures Bellare–Rogaway-like key secrecy in the multi-stage setting as follows.

Definition 3.2 (Multi-Stage security). *Let* KE *be a multi-stage key exchange protocol with key distribution $\mathcal{D}$ and properties $(\mathsf{M}, \mathsf{AUTH}, \mathsf{USE}, \mathsf{REPLAY})$, and $\mathcal{A}$ a PPT adversary interacting with* KE *via the queries defined in Section 3.4 within the following game $G^{\mathsf{Multi\text{-}Stage},\mathcal{D}}_{\mathsf{KE},\mathcal{A}}$:*

Setup. *The challenger chooses the test bit $b_{\mathsf{test}} \xleftarrow{\$} \{0, 1\}$ at random and sets* lost $\leftarrow$ false. *In the public-key variant (MSKE), it furthermore generates long-term public/private-key pairs for each participant $U \in \mathcal{U}$.*

Query. *The adversary $\mathcal{A}$ receives the generated public keys (MSKE) and has access to the queries* NewSecret, NewSemiStaticKey, NewSession, Send, Reveal, RevealSemiStaticKey, Corrupt, *and* Test*. Note that such queries may set* lost *to* true*.*

Guess. *At some point, $\mathcal{A}$ stops and outputs a guess b.*

Finalize. *The challenger sets the 'lost' flag to* lost $\leftarrow$ true *if there exist two (not necessarily distinct) labels* label, label$'$ *and some stage* $i \in \{1, \ldots, \mathsf{M}\}$ *such that* label.sid$_i$ = label$'$.sid$_i$, label.st$_{\mathsf{key},i}$ = revealed, *and* label$'$.tested$_i$ = true*. (Adversary has tested and revealed the key in a single session or in two partnered sessions.)*

We say that $\mathcal{A}$ wins the game, denoted by $G_{\mathsf{KE},\mathcal{A}}^{\mathsf{Multi\text{-}Stage},\mathcal{D}} = 1$, if $b = b_{\mathsf{test}}$ and lost = false*. Note that the winning conditions are independent of key dependency, forward secrecy, and authentication properties of* KE*, as those are directly integrated in the affected (*Reveal *and* Corrupt*) queries and the finalization step of the game; for example,* Corrupt *is defined differently for non-forward-secrecy versus stage-j forward secrecy.*

We say KE *is* Multi-Stage*-secure in a key-dependent resp. key-independent and non-forward-secret resp. stage-j-forward-secret manner with concurrent authentication types* AUTH*, key usage* USE*, and replayability property* REPLAY *if* KE *is* Match*-secure and for all PPT adversaries $\mathcal{A}$ the following advantage function is negligible in the security parameter:*

$$\mathsf{Adv}_{\mathsf{KE},\mathcal{A}}^{\mathsf{Multi\text{-}Stage},\mathcal{D}} := \Pr\left[G_{\mathsf{KE},\mathcal{A}}^{\mathsf{Multi\text{-}Stage},\mathcal{D}} = 1 \right] - \frac{1}{2}.$$

4 The TLS 1.3 draft-14 PSK and PSK-(EC)DHE 0-RTT Handshake Protocols

Starting from `draft-13`, TLS 1.3 only specifies PSK-based 0-RTT handshake modes, abandoning the (EC)DHE-based variant predominant in earlier drafts. We hence first analyze the preshared-key–based variants, PSK(-only) 0-RTT and PSK-(EC)DHE 0-RTT, as specified in TLS 1.3 `draft-14` [Res16d]. In the PSK 0-RTT mode, keys are solely derived from a beforehand established pre-shared key (usually the resumption master secret RMS derived in a full TLS 1.3 handshake). In the PSK-(EC)DHE 0-RTT mode, (elliptic curve) Diffie–Hellman shares additionally enter the key derivation. Our analysis of the purely Diffie–Hellman-based (EC)DHE 0-RTT mode (of `draft-12`) is stated later in Sections 6 and 7.

The TLS 1.3 0-RTT handshake protocols can be conceptually subdivided into four phases:

Key exchange. In the key exchange phase, parties negotiate the ciphersuites and key-exchange parameters to be used and establish shared key material as well as traffic keys to encrypt the remaining handshake.

0-RTT. In the 0-RTT (data) phase, which is interleaved with the key exchange phase, the client can send application data already in its first flight. For this purpose, traffic keys for encrypting the early handshake and application data are established.[11]

Server parameters. In the server parameters phase, further handshake parameters (as, e.g., whether client authentication is demanded) are fixed by the server.

Authentication. In the authentication phase, both the server and client can (based on the aspired authentication) authenticate, verify that they share the same view of the handshake, and derive (authenticated) application traffic keys.

Client

$$\text{PSK} \leftarrow \text{HKDF.Expand}(\text{RMS}, \text{label}_1)$$
$$\text{rctxt} \leftarrow \text{HKDF.Expand}(\text{RMS}, \text{label}_2)$$

ClientHello: $r_c \xleftarrow{\$} \{0,1\}^{256}$
+ ClientEarlyData: ticket_age
+ ClientPreSharedKey: psk_id$_1, \dots$
$[+$ ClientKeyShare: $X \leftarrow g^x]^{\dagger}$ →

$$\text{ES} \leftarrow \text{HKDF.Extract}(0, \text{PSK})$$
$$H_1 \leftarrow \text{H}(\text{CH})||\text{H}(\text{rctxt})$$
$$\text{ETS} \leftarrow \text{HKDF.Expand}(\text{ES}, \text{label}_3||H_1)$$
$$tk_{ehs} \leftarrow \text{HKDF.Expand}(\text{ETS}, \text{label}_4) \qquad \text{stage 1}$$

(ClientFinished$_0$):
$$\text{FS}_{0\text{-RTT}} \leftarrow \text{HKDF.Expand}(\text{ETS}, \text{label}_5)$$
CF$_0 \leftarrow$ HMAC(FS$_{0\text{-RTT}}, H_1$) →

check CF$_0 =$ HMAC(FS$_{0\text{-RTT}}, H_1$)
$$tk_{ead} \leftarrow \text{HKDF.Expand}(\text{ETS}, \text{label}_6) \qquad \text{stage 2}$$

record layer (application data), using AEAD with key tk_{ead} ⇢

Server

ServerHello: $r_s \xleftarrow{\$} \{0,1\}^{256}$
+ ServerEarlyData
+ ServerPreSharedKey: psk_id
$[+$ ServerKeyShare: $Y \leftarrow g^y]^{\dagger}$ ←

$$H_2 \leftarrow \text{H}(\text{CH}||\text{SH})||\text{H}(\text{rctxt})$$
$[\text{DHE} \leftarrow Y^x]^{\dagger} \qquad [\text{DHE} \leftarrow 0]^{\diamond} \qquad [\text{DHE} \leftarrow X^y]^{\dagger}$
$$\text{HS} \leftarrow \text{HKDF.Extract}(\text{ES}, \text{DHE})$$
$$\text{HTS} \leftarrow \text{HKDF.Expand}(\text{HS}, \text{label}_7||H_2)$$
$$tk_{hs} \leftarrow \text{HKDF.Expand}(\text{HTS}, \text{label}_8) \qquad \text{stage 3}$$

{EncryptedExtensions}
$$H_3 \leftarrow \text{H}(\text{CH}||\text{SH}||\text{EE})||\text{H}(\text{rctxt})$$
$$\text{FS}_S \leftarrow \text{HKDF.Expand}(\text{HTS}, \text{label}_9)$$
{ServerFinished}:
SF $\leftarrow$ HMAC(FS$_S, H_3$) ←

check SF $=$ HMAC(FS$_S, H_3$)
$$H_4 \leftarrow \text{H}(\text{CH}||\text{SH}||\text{EE}||\text{SF})||\text{H}(\text{rctxt})$$
$$\text{FS}_C \leftarrow \text{HKDF.Expand}(\text{HTS}, \text{label}_5)$$
{ClientFinished}:
CF $\leftarrow$ HMAC(FS$_C, H_4$) →

check CF $=$ HMAC(FS$_C, H_4$)
$$H_5 \leftarrow \text{H}(\text{CH}||\text{SH}||\text{EE}||\text{SF}||\text{CF})||\text{H}(\text{rctxt})$$
$$\text{MS} \leftarrow \text{HKDF.Extract}(\text{HS}, 0)$$
$$\text{TS} \leftarrow \text{HKDF.Expand}(\text{MS}, \text{label}_{10}||H_4)$$
$$tk_{app} \leftarrow \text{HKDF.Expand}(\text{TS}, \text{label}_{11}) \qquad \text{stage 4}$$
$$\text{EMS} \leftarrow \text{HKDF.Expand}(\text{MS}, \text{label}_{12}||H_5) \qquad \text{stage 5}$$

record layer (application data), using AEAD with key tk_{app} ⇠⇢

Protocol flow legend

MSG: Y	TLS 1.3 message MSG containing Y
(MSG)	message MSG AEAD-encrypted with tk_{ehs}
{MSG}	message MSG AEAD-encrypted with tk_{hs}
+ MSG	message sent as extension within previous message
$[\dots]^{\dagger}$	message/computation only in PSK-(EC)DHE 0-RTT
$[\dots]^{\diamond}$	message/computation only in PSK 0-RTT

Key schedule legend

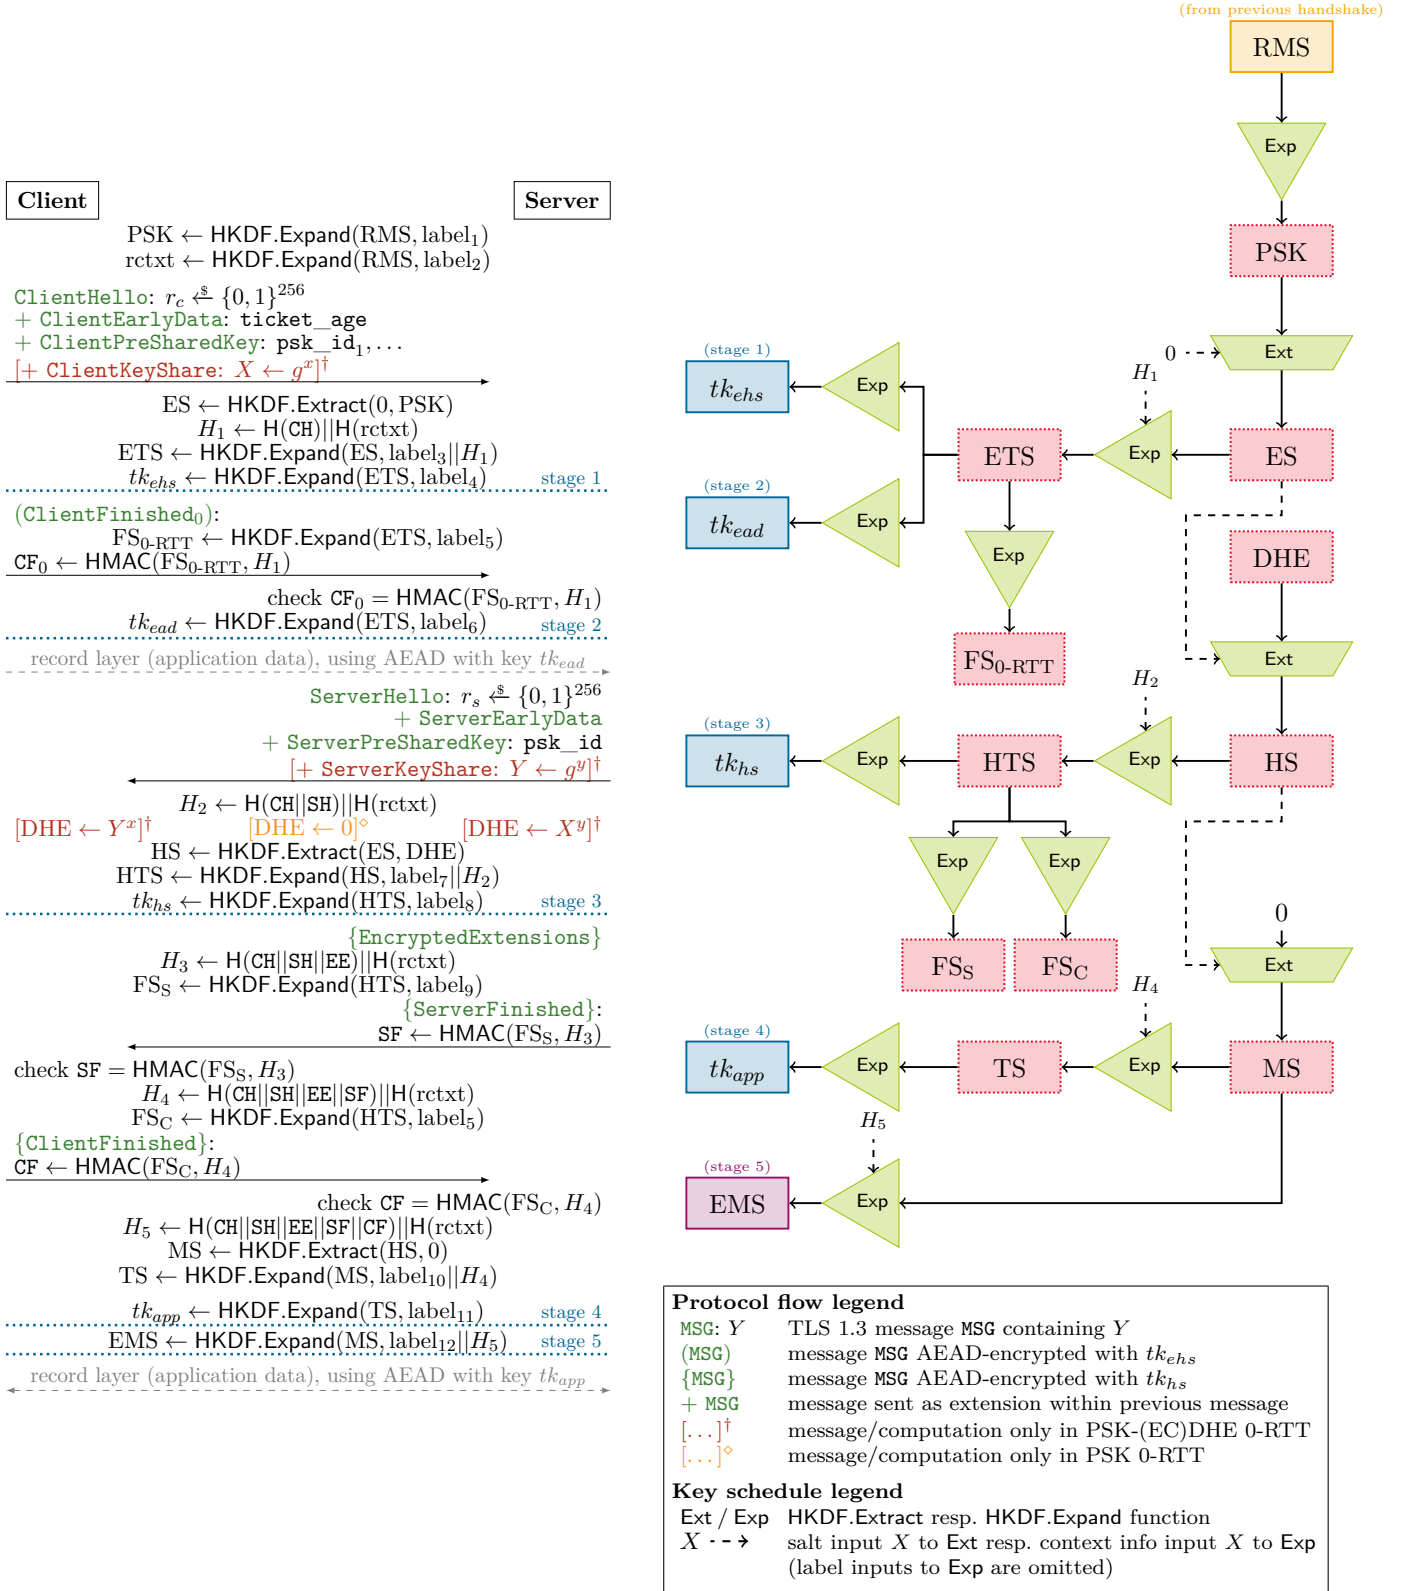|	
Ext / Exp	HKDF.Extract resp. HKDF.Expand function
$X \dashrightarrow$	salt input X to Ext resp. context info input X to Exp (label inputs to Exp are omitted)

Figure 2: The TLS 1.3 `draft-14` PSK and PSK-(EC)DHE 0-RTT handshake protocols (left) and key schedule (right).

We illustrate the protocol flow (with the cryptographically relevant computations) and the key schedule for both the PSK(-only) and PSK-(EC)DHE 0-RTT handshakes in Figure 2. The handshake messages are the following:

- `ClientHello` (CH)/`ServerHello` (SH) contains random nonces r_c / r_s of the respective parties, as well as negotiation parameters (supported versions and ciphersuites). Several extensions can be sent along with the `Hello` messages; the PSK and PSK-(EC)DHE 0-RTT handshakes require the following two resp. three extensions to be included.

- `ClientEarlyData` (CEAD)/`ServerEarlyData` (SEAD) are extensions sent to announce a 0-RTT hand-shake. The client includes the (masked) age `ticket_age` of the (ticket issuing the) used resumption secret. The server signals accepting the 0-RTT exchange with an empty `ServerEarlyData` extension.

  TLS 1.3 `draft-14` [Res16d, Section 4.2.6.2] recommends that servers should use the `ticket_age` value to check that client messages are not replayed. Depending on how well clocks are synchronized, this can prevent delayed replays, but not immediate replays. We do not rely on this check in our analysis but conservatively treat the 0-RTT key exchange messages as arbitrarily replayable.

- `ClientPreSharedKey` (CPSK)/`ServerPreSharedKey` (SPSK) are extensions in which the client an-nounces one (or multiple) pre-shared key identifier(s) (`psk_id`), of which the server selects one to be used as the pre-shared secret (`pss`) in the handshake. Focusing on 0-RTT handshakes only, we only consider the case where client and server agree on the first announced `psk_id`, the one used to derive 0-RTT keys.

- `ClientKeyShare` (CKS)/`ServerKeyShare` (SKS) are extensions sent only for the PSK-(EC)DHE 0-RTT handshake. They contain the ephemeral Diffie–Hellman shares $X = g^x$ resp. $Y = g^y$ for several (in case of the client) or one group (the server decided for).

Based on these messages both sides can already derive the 0-RTT keys. First, from the shared pre-shared secret (the resumption master secret established in a previous handshake) $\mathsf{pss} = \mathrm{RMS}$ a pre-shared key PSK and a resumption context value rctxt are derived using the Expand component of HKDF [Kra10]. In the key derivation, PSK serves as the starting secret while rctxt binds the derived keys to the previous handshake that established RMS. Then, the early secret ES is computed from PSK using HKDF.Extract. Via an intermediate (expanded) early traffic secret ETS both the 0-RTT handshake and application traffic keys tk_{ehs} and tk_{ead} are finally expanded.

For the two HKDF functions we use the following common notation: Function HKDF.Extract(XTS, SKM) obtains as input a (not necessarily secret and potentially fixed) extractor salt XTS and some source key material SKM, and outputs a pseudorandom key PRK. Function HKDF.Expand(PRK, $CTXinfo$) ob-tains as input a pseudorandom key PRK (here: the output of the Extract step) and some (potentially empty) context information $CTXinfo$, and outputs some key material KM.[12] Both functions are based on HMAC [BCK96].

The client then completes its first flight by sending a 0-RTT `Finished` message, sent encrypted un-der tk_{ehs}:

- `ClientFinished`$_0$ (CF$_0$) consists of an HMAC (message authentication code) value which is com-puted using the 0-RTT finished secret FS$_{0\text{-RTT}}$ on the (hashed) 0-RTT messages and the resumption context rctxt.

---

[11]For comparison, omitting the 0-RTT phase essentially yields the TLS 1.3 full resp. PSK-based handshake.

[12]We assume the third, output-length parameter L in the Expand function to be fixed to $L = \lambda$ for our security parameter λ and hence always omit it.

Following $\texttt{ClientFinished}_0$, the client can use tk_{ead} to encrypt and send 0-RTT application data.[13]

After receiving the client's first flight, the server sends its $\texttt{ServerHello}$ message along with the indicated extensions. At this point (resp. after receiving $\texttt{ServerHello}$ for the client) both sides extract from ES the handshake secret HS (incorporating the joint Diffie–Hellman share DHE $= g^{xy}$ in the PSK-(EC)DHE 0-RTT handshake). Again via first expanding an intermediate handshake traffic secret HTS, the handshake traffic key tk_{hs} is derived.

Server and client then complete the handshake by sending the following messages encrypted under tk_{hs}:

- $\texttt{EncryptedExtensions}$ (EE), sent by the server, allows to specify further extensions.

- $\texttt{ClientFinished}$ (CF)/$\texttt{ServerFinished}$ (SF) contain an HMAC value over the handshake hash, keyed with the client resp. server finished secret $\text{FS}_\text{C}/\text{FS}_\text{S}$ which are both expanded from the handshake traffic secret HTS.

At the end of the handshake, the master secret MS is extracted from HS and used to expand the application traffic key tk_{app} (via an intermediate traffic secret TS), used to protect the (non–0-RTT) application data sent, as well as the exporter master secret EMS which can be used to derive further key material outside of TLS.

On client authentication, 0.5-RTT data, and post-handshake messages. When analyzing the (PSK-based) 0-RTT handshake candidates for TLS 1.3, we focus on the main components of the handshake and hence do not capture the following more advanced options specified in $\texttt{draft-14}$.

First, the server can optionally ask the *client to authenticate* (beyond the shared secret key) by sending a public-key certificate and signing the transcript (i.e., by signature-based authentication as employed in the (EC)DHE-based handshakes of TLS 1.3).[14] We omit this option in our analysis but note that our multi-stage key exchange model can in principle be augmented to capture combined authentication under multiple long-term secrets.

Second, instead of deriving the application traffic key tk_{app} at the end of the handshake (as depicted in Figure 2), the server might already do so after sending the $\texttt{ServerFinished}$ message in order to send so-called *0.5-RTT data* directly following his flight, i.e., without waiting for the $\texttt{ClientFinished}$ response. We omit analyzing this variant of the handshake but expect that results for it with potentially weaker authentication guarantees for tk_{app} can be obtained in our model.

Third, TLS 1.3 introduces *post-handshake messages* that can be sent (potentially long) after the initial handshake was completed in order to update the used traffic key, authenticate the client, or issue tickets for session resumption. Here, we focus on the main handshake and do not consider post-handshake messages.

5 Security of the TLS 1.3 draft-14 PSK and PSK-(EC)DHE 0-RTT Handshakes

Our security analysis of the TLS 1.3 $\texttt{draft-14}$ PSK and PSK-(EC)DHE 0-RTT handshakes ($\texttt{draft-14-PSK-0RTT}$ resp. $\texttt{draft-14-PSK-DHE-0RTT}$) is carried out in the preshared-secret variant (MS-PSKE) of the multi-stage key exchange model. We begin with stating the protocol-specific properties $(\mathsf{M}, \mathsf{AUTH}, \mathsf{USE}, \mathsf{REPLAY})$ mostly shared by both handshakes:

---

[13]The server may decide to not derive any 0-RTT keys (and not accept any 0-RTT data). In that case it would, in our model, simply set the first two session identifiers sid_1, sid_2 and keys K_1, K_2 to $\bot$ and continue with deriving the third key.

[14]There is also discussion to further include signature-based server authentication in the PSK-based 0-RTT handshakes [Res16a].

- $\mathsf{M} = 5$: the PSK-based 0-RTT handshakes have five stages (deriving, in that order, keys tk_{ehs}, tk_{ead}, tk_{hs}, tk_{app}, and EMS).

- the authentication properties AUTH differ between the PSK(-only) and the PSK-(EC)DHE 0-RTT handshakes:

    - for PSK 0-RTT, $\mathsf{AUTH} = \{(\mathsf{mutual}, \mathsf{mutual}, \mathsf{mutual}, \mathsf{mutual}, \mathsf{mutual})\}$: all keys established are mutually authenticated (wrt. the established pre-shared secret).
    - for PSK-(EC)DHE 0-RTT, $\mathsf{AUTH} = \{(\mathsf{mutual}, \mathsf{mutual}, \mathsf{unauth}, \mathsf{mutual}, \mathsf{mutual})\}$: the handshake traffic key tk_{hs} is unauthenticated, all other keys are mutually authenticated (wrt. the established pre-shared secret).[15]

- $\mathsf{USE} = (\mathsf{internal}, \mathsf{external}, \mathsf{internal}, \mathsf{external}, \mathsf{external})$: the (0-RTT and main) handshake traffic keys tk_{ehs} and tk_{hs} are used to protect messages within the handshake while the application traffic keys tk_{ead} and tk_{app} as well as the exporter master secret EMS are only used externally.

- $\mathsf{REPLAY} = (\mathsf{replayable}, \mathsf{replayable}, \mathsf{nonreplayable}, \mathsf{nonreplayable}, \mathsf{nonreplayable})$: the 0-RTT stages 1 and 2 are replayable, the other stages are not.

Both TLS 1.3 `draft-14` PSK-based 0-RTT handshakes enjoy key independence for all keys. Expectedly, the PSK(-only) 0-RTT handshake provides no forward secrecy. The PSK-(EC)DHE 0-RTT handshake instead ensures forward secrecy for the non–0-RTT keys (i.e., from stage 3 on), but not for the 0-RTT keys.

Session matching is defined via the following session identifiers, consisting of the unencrypted messages exchanged up to each stage:

$\mathsf{sid}_1 = (\texttt{ClientHello})$,

$\mathsf{sid}_2 = (\mathsf{sid}_1, \text{``EAD''})$,

$\mathsf{sid}_3 = (\texttt{ClientHello}, \texttt{ServerHello})$,

$\mathsf{sid}_4 = (\texttt{ClientHello}, \texttt{ServerHello}, \texttt{EncryptedExtensions}, \texttt{ServerFinished})$, and

$\mathsf{sid}_5 = (\texttt{ClientHello}, \texttt{ServerHello}, \texttt{EncryptedExtensions}, \texttt{ServerFinished}, \texttt{ClientFinished})$.

Here, `Hello` messages also comprise the sent `EarlyData`, `KeyShare`, and `PreSharedKey` extensions. We remark that, as for the analysis of the TLS 1.3 full and resumption handshake [DFGS15a], we too define the session identifiers over the *unencrypted* messages. This diverges from the common practice to set the session identifier as the concatenation of the (here encrypted) protocol transmissions, but is necessary to achieve key independence in the multi-stage security for such protocols.

For the contributive identifiers, we need to ensure that a server session can in any case be tested when receiving an honest client contribution (even if that client never receives the `ServerHello` response), analogously to the full handshake analysis in [DFGS15a]. Hence, for stage 3, on sending resp. receiving the `ClientHello` message, client resp. server initially sets $\mathsf{cid}_3 = (\texttt{ClientHello})$ and subsequently, on receiving (resp. sending) the `ServerHello` message, extend it to $\mathsf{cid}_3 = (\texttt{ClientHello}, \texttt{ServerHello})$. All other contributive identifiers are set to $\mathsf{cid}_i = \mathsf{sid}_i$ when the respective stage's session identifier is set.

We are now ready to state our security results for the PSK and PSK-DHE 0-RTT handshakes of TLS 1.3 `draft-14`. Naturally, the proof aspects concerning the non–0-RTT parts of the handshakes are structurally close to the proofs for the `draft-10` PSK-based handshakes by Dowling et al. [DFGS16], but need to take the modified key schedule into account.

---

[15]Although including the pre-shared secret in the derivation of tk_{hs}, as the involved Diffie–Hellman shares are only authenticated after its derivation, tk_{hs} cannot enjoy both forward secrecy and mutual authentication. We remark that alternatively to considering tk_{hs} being unauthenticated but forward-secret (a security property close to the notion of "weak (perfect) forward secrecy" [Kra05]), one might instead also consider tk_{hs} to be non–forward-secret but mutually authenticated.

5.1 PSK(-only) 0-RTT Handshake

Theorem 5.1 (Match security of `draft-14-PSK-0RTT`). *The* `draft-14` *PSK 0-RTT handshake is* Match-*secure: for any efficient adversary $\mathcal{A}$ we have*

$$\mathsf{Adv}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}}^{\mathsf{Match}} \leq n_s^2 \cdot 2^{-|nonce|},$$

where n_s is the maximum number of sessions and $|nonce| = 256$ is the bit-length of the nonces.

Proof. We need to prove six properties for Match security.

1. *Sessions with the same session identifier for some stage hold the same key at that stage.*
   Containing the `ClientHello` message, all session identifiers fix the used pre-shared secret identifier `psk_id` and hence the used secret pss = RMS, determining the values PSK and rctxt. Each stage's session identifier moreover contains all messages included in the (handshake) hashes used in the key derivation of that stage's key, which is hence uniquely determined by the session identifier.

2. *Sessions with the same session identifier for some stage agree on that stage's authentication level.*
   This trivially holds as the PSK 0-RTT handshake fixes mutual authentication for all stages.

3. *Sessions with the same session identifier for some stage share the same contributive identifier.*
   For stages $i \in \{1, 2, 4, 5, 6\}$ this follows immediately from $\mathsf{sid}_i = \mathsf{cid}_i$. For stage 3, the contributive identifier is set to its final value $\mathsf{cid}_3 = \mathsf{sid}_3$ when both sides set the session identifier.

4. *Sessions are partnered with the intended (authenticated) participant and (for mutual authentication) share the same key index.*
   As honest sessions only used their own pre-shared secret identifier `psk_id`, this value included (via `ClientHello`) in all session identifiers ensures agreement of both the intended partner and key index.

5. *Session identifiers do not match across different stages.*
   This holds trivially since session identifiers sid_1 and sid_3–sid_5 contain distinct (non-optional) messages and sid_2 includes a separating identifier.

6. *At most two sessions have the same session identifier at any non-replayable stage.*
   We only need to consider the non-replayable stages 3–5 here.[16] The according session identifiers contain (through the `Hello` messages) randomly chosen nonces n_c and n_s (of bit-length $|nonce| = 256$) from each side, one of which a third session would need to pick by coincidence. This probability can be upper-bounded by $n_s^2 \cdot 2^{-|nonce|}$ for n_s being the maximum number of sessions. □

Theorem 5.2 (Multi-Stage security of `draft-14-PSK-0RTT`). *The* `draft-14` *PSK 0-RTT handshake is* Multi-Stage-*secure in a key-independent and non–forward-secret manner with properties* (M, AUTH, USE, REPLAY) *given above. Formally, for any efficient adversary $\mathcal{A}$ against the* Multi-Stage *security there exist efficient algorithms $\mathcal{B}_1, \ldots, \mathcal{B}_9$ such that*

$$
\begin{aligned}
\mathsf{Adv}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}}^{\mathsf{Multi\text{-}Stage},\mathcal{D}} \leq 5 n_s \cdot \Big(&\mathsf{Adv}_{\mathsf{H},\mathcal{B}_1}^{\mathsf{COLL}} + n_p \cdot \Big(\mathsf{Adv}_{\mathsf{HKDF.Expand},\mathcal{B}_2}^{\mathsf{PRF\text{-}sec}} + \mathsf{Adv}_{\mathsf{HMAC},\mathcal{B}_3}^{\mathsf{HMAC}(0,\$)\text{-}\$} + \mathsf{Adv}_{\mathsf{HMAC},\mathcal{B}_4}^{\mathsf{PRF\text{-}sec}} \\
&+ \mathsf{Adv}_{\mathsf{HKDF.Expand},\mathcal{B}_5}^{\mathsf{PRF\text{-}sec}} + \mathsf{Adv}_{\mathsf{HMAC},\mathcal{B}_6}^{\mathsf{PRF\text{-}sec}} + \mathsf{Adv}_{\mathsf{HKDF.Expand},\mathcal{B}_7}^{\mathsf{PRF\text{-}sec}} \\
&+ \mathsf{Adv}_{\mathsf{HKDF.Expand},\mathcal{B}_8}^{\mathsf{PRF\text{-}sec}} + \mathsf{Adv}_{\mathsf{HKDF.Expand},\mathcal{B}_9}^{\mathsf{PRF\text{-}sec}} \Big) \Big),
\end{aligned}
$$

where n_s is the maximum number of sessions and n_p is the maximum number of pre-shared secrets.

---

[16]Observe that an adversary can indeed replay the client's first messages to multiple server sessions, resulting in the same session identifier and derived keys.

Proof. We first restrict the adversary $\mathcal{A}$ to a single Test query. By a hybrid argument (following the detailed one for the full handshake proof by Dowling et al. [DFGS15b, Appendix A]) this reduces $\mathcal{A}$'s advantage by a factor at most $1/5n_s$ for the five stages in the at most n_s sessions. It also allows to speak about *the* session label tested at stage i, which we now know in advance.

Our proof then proceeds via the following sequence of games.

Game 0. We begin with G_0 identical to the Multi-Stage game restricted to a single test query:

$$\mathsf{Adv}^{G_0}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} = \mathsf{Adv}^{\text{1-Multi-Stage}}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}}.$$

Game 1. In a first step, we exclude hash collisions by aborting whenever during the execution of honest sessions the same hash value under hash function H is computed for two distinct inputs. By having an algorithm $\mathcal{B}_1$ act as the challenger in Game 0 and output, when they occur, these two inputs as a collision for H, we can bound the probability of aborting by $\mathcal{B}_1$'s advantage in breaking the collision resistance of H:

$$\mathsf{Adv}^{G_0}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} \leq \mathsf{Adv}^{G_1}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} + \mathsf{Adv}^{\mathsf{COLL}}_{\mathsf{H},\mathcal{B}_1}.$$

Game 2. Next, we guess the (index $\texttt{psk\_id}$ of the) pre-shared secret pss employed in the tested session (i.e., the resumption master secret RMS used for this PSK handshake). Aborting on an incorrect guess, this reduces the advantage of $\mathcal{A}$ by a factor of at most the number of pre-shared secrets n_p:

$$\mathsf{Adv}^{G_1}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} \leq n_p \cdot \mathsf{Adv}^{G_2}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}}.$$

We can now, one at a time, replace the outputs of HKDF.Expand and HKDF.Extract evaluations using RMS and derived keys by random values, leading to a sequence of according advantage bounds for their PRF security or randomness bounds of the underling HMAC function.

Game 3. We begin by replacing any HKDF.Expand application using $\mathsf{pss} = \mathsf{RMS}$ by evaluations of a (lazy-sampled) random function, which in particular leads to PSK and rctxt being replaced by random values $\widetilde{\mathsf{PSK}}, \widetilde{\mathsf{rctxt}} \xleftarrow{\$} \{0,1\}^\lambda$ in the tested (and any partnered) session.

The introduced difference in the advantage of $\mathcal{A}$ can be bounded by an adversary $\mathcal{B}_2$ against the PRF security of HKDF.Expand as follows. Algorithm $\mathcal{B}_2$ simulates Game 2 faithfully, but uses its PRF oracle for evaluating HKDF.Expand under RMS. Note that all keys derived in the PSK 0-RTT handshake are non–forward-secret and hence any (successful) adversary $\mathcal{A}$ cannot issue a Corrupt query on $\mathsf{pss} = \mathsf{RMS}$ used in the tested session. The employed pre-shared key is hence an unknown and uniformly random value to $\mathcal{A}$ and hence $\mathcal{B}_2$ perfectly simulates Game 2 in case its oracle computes HKDF.Expand and Game 3 in case its oracle is a random function.

This leads to following bound:

$$\mathsf{Adv}^{G_2}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} \leq \mathsf{Adv}^{G_3}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_2}.$$

Game 4. Next, we replace values ES computed as $\mathsf{HKDF.Extract}(0, \widetilde{\mathsf{PSK}})$ by a random value $\widetilde{\mathsf{ES}} \xleftarrow{\$} \{0,1\}^\lambda$, in particular in the tested session and partnered sessions.

Recall that $\mathsf{HKDF.Extract}(XTS, SKM)$ is defined as $\mathsf{HMAC}(XTS, SKM)$ [Kra10]. Assuming that for any polynomial-time algorithm $\mathcal{B}_3$ it is computationally hard to distinguish $\mathsf{HMAC}(0, SKM)$ from $X \xleftarrow{\$} \{0,1\}^\lambda$ for uniformly random chosen values $SKM \in \{0,1\}^\lambda$, we can bound this step by $\mathcal{B}_3$'s distinguishing advantage which we denote by $\mathsf{Adv}^{\mathsf{HMAC}(0,\$)\text{-}\$}_{\mathsf{HMAC},\mathcal{B}_3}$. For this, we let $\mathcal{B}_3$ simulate Game 3 as the challenger,

but using its challenge as $\text{ES} = \text{HKDF.Extract}(0, \widetilde{\text{PSK}})$. In case this challenge is $\text{HMAC}(0, SKM)$ for $SKM \in \{0,1\}^\lambda$, this simulates Game 3, if the challenge is a uniformly random value $X \xleftarrow{\$} \{0,1\}^\lambda$, this simulates Game 4.

We can hence bound this step as

$$\text{Adv}^{G_3}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} \leq \text{Adv}^{G_4}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} + \text{Adv}^{\text{HMAC}(0,\$)\text{-}\$}_{\text{HMAC},\mathcal{B}_3}.$$

Game 5. We next replace evaluations of HKDF.Expand keyed with $\widetilde{\text{ES}}$ as well as HKDF.Extract using $\widetilde{\text{ES}}$ as salt by random functions. This in particular replaces, in the tested and partnered sessions, the early traffic and handshake secret in these sessions by random values $\widetilde{\text{ETS}}, \widetilde{\text{HS}} \xleftarrow{\$} \{0,1\}^\lambda$.

Observe that in both replaced Extract and Expand evaluations, $\widetilde{\text{ES}}$ is (by definition of HKDF) used to key the HMAC function, applied to H_1 (when expanding ETS) resp. to a fixed value 0 (when extracting HS), i.e., distinct inputs. We can hence bound the difference in the advantage of $\mathcal{A}$ introduced by this step by the advantage of an adversary $\mathcal{B}_4$ against the PRF security of HMAC, having $\mathcal{B}_4$ use its PRF oracle to compute the replaced Expand and Extract evaluations as detailed for Game 3. The resulting bound is thus:

$$\text{Adv}^{G_4}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} \leq \text{Adv}^{G_5}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} + \text{Adv}^{\text{PRF-sec}}_{\text{HMAC},\mathcal{B}_4}.$$

Note that from now on, $\widetilde{\text{ETS}}$ is independent of any value computed in sessions that are not partnered (in stage 1 and 2) with the tested session: as such sessions do not hold the same $\text{sid}_1/\text{sid}_2$ and hence not the same $\texttt{ClientHello}$ and as, by Game 1, there are no hash collisions, no non-partnered session will compute the same hash value H_1 which hence serves as a unique label in the tested and partnered sessions.

Game 6. As the next step, we replace HKDF.Expand evaluations keyed with $\widetilde{\text{ETS}}$ (in the tested and partnered session) by a lazy-sample random function, in particular replacing in those sessions the early handshake and data traffic keys as well as the client's 0-RTT finished secret by random values $\widetilde{tk_{ehs}}, \widetilde{tk_{ead}}, \widetilde{\text{FS}}_{\text{0-RTT}} \xleftarrow{\$} \{0,1\}^\lambda$.

Similar to Game 3 the advantage difference induced for $\mathcal{A}$ by this step can be bound by the advantage of an adversary $\mathcal{B}_5$ against the PRF security of HKDF.Expand:

$$\text{Adv}^{G_5}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} \leq \text{Adv}^{G_6}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} + \text{Adv}^{\text{PRF-sec}}_{\text{HKDF.Expand},\mathcal{B}_5}.$$

Game 7. We again in parallel replace both HKDF.Expand and HKDF.Extract evaluations, this time keyed resp. salted with $\widetilde{\text{HS}}$ by random function, replacing the handshake traffic secret and master secret with random values $\widetilde{\text{HTS}}, \widetilde{\text{MS}} \xleftarrow{\$} \{0,1\}^\lambda$, in particular in the tested and partnered sessions.

As for Game 5, both evaluations are HMAC invocations keyed with $\widetilde{\text{HS}}$ and applied to distinct values (H_2 resp. 0). We can hence likewise bound the advantage difference introduced by the PRF security of HMAC as

$$\text{Adv}^{G_6}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} \leq \text{Adv}^{G_7}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} + \text{Adv}^{\text{PRF-sec}}_{\text{HMAC},\mathcal{B}_6}.$$

Again, as sid_3 uniquely determines the message inputs to hash value H_2 entering the derivation of HTS and, by Game 1 there are no hash collisions, $\widetilde{\text{HTS}}$ is independent of values computed in sessions not partnered with the tested session in stage 3.

Game 8. We now replace evaluations of HKDF.Expand using $\widetilde{\text{HTS}}$ (in the tested and partnered sessions) by a random function, leading to random values $\widetilde{tk_{hs}}, \widetilde{\text{FS}}_{\text{S}}, \widetilde{\text{FS}}_{\text{C}} \xleftarrow{\$} \{0,1\}^\lambda$ in those sessions. This step is again bounded by the PRF security of HKDF.Expand:

$$\text{Adv}^{G_7}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} \leq \text{Adv}^{G_8}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} + \text{Adv}^{\text{PRF-sec}}_{\text{HKDF.Expand},\mathcal{B}_7}.$$

Game 9. Next, we replace HKDF.Expand evaluations keyed with $\widetilde{MS}$ by a random function, in particular leading to uniformly random values $\widetilde{TS}, \widetilde{EMS} \overset{\$}{\leftarrow} \{0,1\}^\lambda$ in the tested and partnered sessions. These are moreover independent of any other values computed in sessions not partnered in stages 4 and 5, due to Game 1 and sid_4 and sid_5 fixing the inputs to H_4 and H_5. Again,

$$\mathsf{Adv}^{G_8}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} \leq \mathsf{Adv}^{G_9}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_8}.$$

Game 10. Finally, we replace the HKDF.Expand evaluations using $\widetilde{TS}$ (in the tested and partnered sessions) by a random function, resulting in a random application traffic key $\widetilde{tk_{app}} \overset{\$}{\leftarrow} \{0,1\}^\lambda$, again bounded by PRF security:

$$\mathsf{Adv}^{G_9}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} \leq \mathsf{Adv}^{G_{10}}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_9}.$$

In Game 10, all keys derived in the tested session ($\widetilde{tk_{ehs}}$, $\widetilde{tk_{ead}}$, $\widetilde{tk_{hs}}$, $\widetilde{tk_{app}}$, and $\widetilde{EMS}$) are now chosen uniformly at random, making the Test query independent of the test bit b_{test}.

Observe furthermore that replaying a `ClientHello` message to multiple server sessions leads to all these sessions being partnered to the originating client session (and hence prevents Reveal queries). In contrast, sessions that are not partnered with the tested session (even if using the same pre-shared secret) by definition hold different session identifiers and hence use different handshake hashes (due to Game 1) as label inputs to HKDF.Expand in the key derivation. The resulting keys therefore are independent random values themselves, uncorrelated with the keys established in the tested session.

Therefore, $\mathcal{A}$ learns no information on the test bit b_{test} and hence

$$\mathsf{Adv}^{G_{10}}_{\texttt{draft-14-PSK-0RTT},\mathcal{A}} \leq 0. \qquad \square$$

5.2 PSK-(EC)DHE 0-RTT Handshake

Theorem 5.3 (Match security of `draft-14-PSK-DHE-0RTT`)**.** *The* `draft-14` *PSK-(EC)DHE 0-RTT handshake is* Match*-secure: for any efficient adversary $\mathcal{A}$ we have*

$$\mathsf{Adv}^{\mathsf{Match}}_{\texttt{draft-14-PSK-DHE-0RTT},\mathcal{A}} \leq n_s^2 \cdot 1/q \cdot 2^{-|nonce|},$$

where n_s is the maximum number of sessions, q is the group order, and $|nonce| = 256$ is the bit-length of the nonces.

Proof. For conditions 2–5 of Match security, the arguments are as for the Match security of the PSK(-only) 0-RTT handshake (cf. Theorem 5.1). We hence focus on conditions 1 and 6.

1. *Sessions with the same session identifier for some stage hold the same key at that stage.*
   Beyond the arguments from Theorem 5.1, the `ClientHello` and `ServerHello` messages contained in all session identifiers from stage 3 on also fix the chosen Diffie–Hellman shares g^x and g^y. This ensures agreement on HS and hence also on the keys for stages 3–5 derived (also) from these Diffie–Hellman shares.

6. *At most two sessions have the same session identifier at any non-replayable stage.*
   Again we can focus on the non-replayable stages 3–5 here. For the same argument as in Theorem 5.1, three (honest) sessions sharing the same session identifier requires (at least) two sessions pick the same nonce and, for PSK-(EC)DHE, also the same Diffie–Hellman share. We can upper-bound this probability by $n_s^2 \cdot 1/q \cdot 2^{-|nonce|}$, where n_s is the maximum number of sessions. $\qquad \square$

Theorem 5.4 (Multi-Stage security of `draft-14-PSK-DHE-0RTT`). *The* `draft-14` *PSK-(EC)DHE 0-RTT handshake is* Multi-Stage*-secure in a key-independent and stage-3-forward-secret manner with properties* (M, AUTH, USE, REPLAY) *given above. Formally, for any efficient adversary $\mathcal{A}$ against the* Multi-Stage *security there exist efficient algorithms $\mathcal{B}_1, \ldots, \mathcal{B}_{16}$ such that*

$$
\begin{aligned}
\mathsf{Adv}^{\mathsf{Multi\text{-}Stage},\mathcal{D}}_{\texttt{draft-14-PSK-DHE-0RTT},\mathcal{A}} \leq 5n_s \cdot \Big(& \mathsf{Adv}^{\mathsf{COLL}}_{\mathsf{H},\mathcal{B}_1} \\
& + n_p \cdot \Big(\mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_2} + \mathsf{Adv}^{\mathsf{HMAC}(0,\$)\text{-}\$}_{\mathsf{HMAC},\mathcal{B}_3} \\
& \quad + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HMAC},\mathcal{B}_4} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_5} \Big) \\
& + n_s \cdot n_p \cdot \Big(\mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_6} + \mathsf{Adv}^{\mathsf{HMAC}(0,\$)\text{-}\$}_{\mathsf{HMAC},\mathcal{B}_7} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HMAC},\mathcal{B}_8} \\
& \quad + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HMAC},\mathcal{B}_9} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_{10}} + \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{HMAC},\mathcal{B}_{11}} \Big) \\
& + n_s \cdot n_p \cdot \Big(\mathsf{Adv}^{\mathsf{PRF\text{-}ODH}}_{\mathsf{HKDF.Extract},\mathbb{G},\mathcal{B}_{12}} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HMAC},\mathcal{B}_{13}} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_{14}} \\
& \quad + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_{15}} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_{16}} \Big) \Big),
\end{aligned}
$$

where n_s is the maximum number of sessions and n_p is the maximum number of pre-shared secrets.

Proof. Again we first restrict the adversary $\mathcal{A}$ to a single Test query (for which we now know label and stage in advance), inducing a security loss of at most $5n_s$ by a hybrid argument.

Game 1. We next exclude hash collisions by aborting the game whenever the challenger would (in honest sessions) compute the same hash value for distinct inputs. As in Game 1 for Theorem 5.2 the caused difference in the advantage can be bounded by that of an adversary $\mathcal{B}_1$ against the collision resistance of the hash function:

$$
\mathsf{Adv}^{\mathsf{1\text{-}Multi\text{-}Stage}}_{\texttt{draft-14-PSK-DHE-0RTT},\mathcal{A}} \leq \mathsf{Adv}^{\mathsf{1\text{-}Multi\text{-}Stage,no\text{-}coll}}_{\texttt{draft-14-PSK-DHE-0RTT},\mathcal{A}} + \mathsf{Adv}^{\mathsf{COLL}}_{\mathsf{H},\mathcal{B}_1}.
$$

Case separation. Our proof then treats the following three (disjoint) cases separately:

  A. the adversary tests a stage-1 or stage-2 key,

  B. the adversary tests a stage-i key for $i \in \{3, 4, 5\}$ in a session without honest contributive partner in the third stage (i.e., there does not exist a session label′ with label.cid$_3$ = label′.cid$_3$ when the Test query is issued on label), and

  C. the adversary tests a stage-i key for $i \in \{3, 4, 5\}$ in a session with honest contributive partner in the third stage.

Case A. Test in Stage 1–2

In the first proof case we are concerned with a Test query on a 0-RTT key (in stage 1 or 2). As both stages are non–forward-secret, we know that is this case no Corrupt query can have been issued for the pre-shared secret pss employed in the tested session (neither before nor after the Test query), as the adversary would otherwise lose. This allows us to apply the same proof strategy as for the Multi-Stage security of the `draft-14` PSK(-only) 0-RTT handshake in the proof Theorem 5.2. Via the very same sequence of games G_2–G_6 (starting after excluding hash collisions) we reach a game where both 0-RTT keys tk_{ehs}

and tk_{ead} are replaced by independent and uniformly random values (leaving the adversary $\mathcal{A}$ no change to win). The introduced differences in advantage of $\mathcal{A}$ are bound as for Theorem 5.2 by

$$\mathsf{Adv}^{\text{1-Multi-Stage,no-coll,test 1–2}}_{\texttt{draft-14-PSK-DHE-0RTT},\mathcal{A}} \leq n_p \cdot \left(\mathsf{Adv}^{\text{PRF-sec}}_{\text{HKDF.Expand},\mathcal{B}_2} + \mathsf{Adv}^{\text{HMAC}(0,\$)\text{-}\$}_{\text{HMAC},\mathcal{B}_3} + \mathsf{Adv}^{\text{PRF-sec}}_{\text{HMAC},\mathcal{B}_4} + \mathsf{Adv}^{\text{PRF-sec}}_{\text{HKDF.Expand},\mathcal{B}_5} \right),$$

where $\mathcal{B}_2, \ldots, \mathcal{B}_5$ are the reduction algorithms given in the proof of Theorem 5.2.

Case B. Test in Stage 3–5 without Contributive Stage-3 Partner

We now consider the case that the Test query is issued on stage 3–5 in a (client or server) session without honest contributive partner in stage 3. Since stage 3 is unauthenticated, testing this stage actually leads to immediately losing the Multi-Stage game, hence we can focus on stages 4 and 5. As we will see, given the HMAC values in the exchanged `Finished` messages are unforgeable, we can ultimately exclude that such Test queries are issued via the following sequence of games.

Game B.0. We begin with the Multi-Stage game restricted to a single Test query in stage 3–5 without contributive stage-3 partner, where collisions are excluded:

$$\mathsf{Adv}^{G_{B.0}}_{\texttt{draft-14-PSK-DHE-0RTT},\mathcal{A}} = \mathsf{Adv}^{\text{1-Multi-Stage,no-coll,stage 3–5,no } \text{cid}_3\text{-partner}}_{\texttt{draft-14-PSK-DHE-0RTT},\mathcal{A}}.$$

Game B.1. We now introduce an abortion of the game as soon as a session accepts in stage 4 without honest contributive partner in stage 3. Denoting this event as $\mathsf{abort}^{G_{B.1},\mathcal{A}}_{acc}$ we can bound the induced advantage difference of $\mathcal{A}$ as

$$\left| \mathsf{Adv}^{G_{B.0}}_{\texttt{draft-14-PSK-DHE-0RTT},\mathcal{A}} - \mathsf{Adv}^{G_{B.1}}_{\texttt{draft-14-PSK-DHE-0RTT},\mathcal{A}} \right| \leq \Pr[\mathsf{abort}^{G_{B.1},\mathcal{A}}_{acc}].$$

Observe that we can immediately bound $\mathsf{Adv}^{G_{B.1}}_{\texttt{draft-14-PSK-DHE-0RTT},\mathcal{A}} \leq 0$, as the game aborts before $\mathcal{A}$ has the chance to issue a Test query (recall that Test queries may only be issued to stage 4 or 5 of session without contributive stage-3 partner). We hence continue by bounding $\Pr[\mathsf{abort}^{G_{B.1},\mathcal{A}}_{acc}]$ in the remaining game sequence.

Game B.2. Next, we guess the first session that accepts in stage 4 without honest contributive stage-3 partner (i.e., the session causing $\mathsf{abort}^{G_{B.1},\mathcal{A}}_{acc}$) and abort if we guessed incorrectly. Observe that Game B.2 equals Game B.1 for a correct guess and we can hence bound

$$\Pr[\mathsf{abort}^{G_{B.1},\mathcal{A}}_{acc}] \leq n_s \cdot \Pr[\mathsf{abort}^{G_{B.2},\mathcal{A}}_{acc}],$$

where n_s is the maximum number of sessions.

Moreover, no Corrupt query can have been issued to the guessed session (or any other session using the same pre-shared secret pss): On the one hand, sessions stop execution on Corrupt and thus no such query could have been issued before the guessed session accepted in stage 4 (as otherwise it would not have accepted). On the other hand, the game aborts when the guessed session accepts in stage 4, so there is no chance for $\mathcal{A}$ to issue a Corrupt query afterwards. The pre-shared secret pss employed in the guessed session hence remains an unknown, random value for $\mathcal{A}$ which we can leverage in the following games.

Game B.3. We can now first guess the pre-shared secret $\mathsf{pss} = \mathsf{RMS}$ employed in the guessed session. Aborting on an incorrect guess introduces a factor of at most the number of pre-shared secrets n_p:

$$\Pr[\mathsf{abort}^{G_{B.2},\mathcal{A}}_{acc}] \leq n_p \cdot \Pr[\mathsf{abort}^{G_{B.3},\mathcal{A}}_{acc}].$$

Game B.4. Applying to the guessed session the steps introduced in the Games 3, 4, 5, 7, and 8 from the proof of draft-14 PSK(-only) Multi-Stage security (Theorem 5.2), we (in particular) replace the values PSK, ES, HS, HTS, and finally $\mathrm{FS_S}$ and $\mathrm{FS_C}$ by uniformly random values sampled from $\{0,1\}^\lambda$. Denoting the resulting game by Game B.4, the introduced advantage difference is bounded as

$$\Pr[\mathsf{abort}_{acc}^{G_{B.3},\mathcal{A}}] \leq \Pr[\mathsf{abort}_{acc}^{G_{B.4},\mathcal{A}}]$$

$$+ \mathsf{Adv}_{\mathsf{HKDF.Expand},\mathcal{B}_6}^{\mathsf{PRF\text{-}sec}} + \mathsf{Adv}_{\mathsf{HMAC},\mathcal{B}_7}^{\mathsf{HMAC}(0,\$)\text{-}\$} + \mathsf{Adv}_{\mathsf{HMAC},\mathcal{B}_8}^{\mathsf{PRF\text{-}sec}} + \mathsf{Adv}_{\mathsf{HMAC},\mathcal{B}_9}^{\mathsf{PRF\text{-}sec}} + \mathsf{Adv}_{\mathsf{HKDF.Expand},\mathcal{B}_{10}}^{\mathsf{PRF\text{-}sec}},$$

where $\mathcal{B}_6, \ldots, \mathcal{B}_{10}$ are the algorithms $\mathcal{B}_2, \mathcal{B}_3, \mathcal{B}_4, \mathcal{B}_6$, and $\mathcal{B}_7$ given for Games 3, 4, 5, 7, and 8 in the proof of Theorem 5.2.

At this point, both server and client finished secrets used in the guessed session are replaced by uniformly random random values $\widetilde{\mathrm{FS_S}}$ resp. $\widetilde{\mathrm{FS_C}}$. As the final step, we show how this allows to turn any adversary triggering the $\mathsf{abort}_{acc}^{G_{B.4},\mathcal{A}}$ event in Game B.4 into an (EUF-CMA) MAC forger $\mathcal{B}_{11}$ for HMAC. To this extent, we let $\mathcal{B}_{11}$ act as the challenger in Game B.4, but instead of computing HMAC values using keys $\widetilde{\mathrm{FS_S}}$ or $\widetilde{\mathrm{FS_C}}$ on its own, relay them to MAC oracles of two EUF-CMA security instances for HMAC. As $\widetilde{\mathrm{FS_S}}$ and $\widetilde{\mathrm{FS_C}}$ are uniformly random values, this provides a sound simulation of Game B.4 for $\mathcal{A}$.

Recall that $\mathsf{abort}_{acc}^{G_{B.4},\mathcal{A}}$ is triggered as soon as the first (guessed) session accepts in stage 4 without contributive partner in stage 3. This in particular means that there is no session holding the same $\mathsf{cid}_3 = (\texttt{ClientHello}, \texttt{ServerHello})$ value. However, in order for the guessed session to accept in stage 4, it must have received (within $\texttt{ServerFinished}$ or $\texttt{ClientFinished}$) an HMAC value on a hash covering also the messages contained in cid_3. As no session holds the same cid_3, and as there are no hash collisions (by Game 1), no honest session will have issued this HMAC value. It hence was not queried to the MAC oracle in the EUF-CMA security game (used for client resp. server messages) which allows $\mathcal{B}_{11}$ to output it as a valid forgery in the respective game. Algorithm $\mathcal{B}_{11}$ hence being successful whenever $\mathcal{A}$ triggers $\mathsf{abort}_{acc}^{G_{B.4},\mathcal{A}}$, this finally bounds

$$\Pr[\mathsf{abort}_{acc}^{G_{B.4},\mathcal{A}}] \leq \mathsf{Adv}_{\mathsf{HMAC},\mathcal{B}_{11}}^{\mathsf{EUF\text{-}CMA}}.$$

In summary, this provides the following advantage bounds for this proof case:

$$\mathsf{Adv}_{\texttt{draft-14-PSK-DHE-0RTT},\mathcal{A}}^{\text{1-Multi-Stage,no-coll,stage 3–5,no cid}_3\text{-partner}} \leq n_s \cdot n_p \cdot \Big(\mathsf{Adv}_{\mathsf{HKDF.Expand},\mathcal{B}_6}^{\mathsf{PRF\text{-}sec}} + \mathsf{Adv}_{\mathsf{HMAC},\mathcal{B}_7}^{\mathsf{HMAC}(0,\$)\text{-}\$} + \mathsf{Adv}_{\mathsf{HMAC},\mathcal{B}_8}^{\mathsf{PRF\text{-}sec}}$$

$$+ \mathsf{Adv}_{\mathsf{HMAC},\mathcal{B}_9}^{\mathsf{PRF\text{-}sec}} + \mathsf{Adv}_{\mathsf{HKDF.Expand},\mathcal{B}_{10}}^{\mathsf{PRF\text{-}sec}} + \mathsf{Adv}_{\mathsf{HMAC},\mathcal{B}_{11}}^{\mathsf{EUF\text{-}CMA}} \Big).$$

Case C. Test in Stage 3–5 with Contributive Stage-3 Partner

In the last proof case we treat test sessions accepting in stages 3–5 that have an honest contributive partner in stage 3. In contrast to Case B this in particular allows the (unauthenticated) stage-3 key to be tested. Our proof strategy is geared towards leveraging the availability of honest Diffie–Hellman shares g^x and g^y (through honest cid_3-partnering) as source of randomness (unknown to $\mathcal{A}$) which ensures (forward) secrecy of the keys derived from it in stages 3–5, even if the involved pre-shared secret pss is corrupted.

Game C.0. Our initial game is the Multi-Stage game restricted to a single Test query in stage 3–5 with contributive stage-3 partner, where collisions are excluded:

$$\mathsf{Adv}_{\texttt{draft-14-PSK-DHE-0RTT},\mathcal{A}}^{G_{C.0}} = \mathsf{Adv}_{\texttt{draft-14-PSK-DHE-0RTT},\mathcal{A}}^{\text{1-Multi-Stage,no-coll,stage 3–5,cid}_3\text{-partner}}.$$

Game C.1. We first guess the (index of the) session contributively partnered in stage 3 with the tested session and abort on an incorrect guess, inducing a factor of n_s (the numbers of sessions):

$$\mathsf{Adv}_{\texttt{draft-14-PSK-DHE-0RTT},\mathcal{A}}^{G_{C.0}} \leq n_s \cdot \mathsf{Adv}_{\texttt{draft-14-PSK-DHE-0RTT},\mathcal{A}}^{G_{C.1}}.$$

Game C.2. Next, we guess the pre-shared secret $\mathsf{pss} = \mathrm{RMS}$ used in the tested session (and abort on an incorrect guess), reducing the advantage of $\mathcal{A}$ by a factor of n_p (the number of pre-shared secrets):

$$\mathsf{Adv}^{G_{C.1}}_{\texttt{draft-14-PSK-DHE-ORTT},\mathcal{A}} \leq n_p \cdot \mathsf{Adv}^{G_{C.2}}_{\texttt{draft-14-PSK-DHE-ORTT},\mathcal{A}}.$$

Game C.3. Knowing the (honest) session that contributes in particular the Diffie–Hellman share to stage 3 of the tested session and the used pre-shared secret in advance now allows us to encode a Diffie–Hellman challenge in the shares g^x and g^y used at the tested session.

If the tested session is a client session, we know that both it and the partnered session guessed in Game C.1 hold the same shares g^x, g^y. If the tested session however is a server session, we are not ensured that the contributive (client) partner from Game C.1 will receive the test session's share g^y unmodified.[17] It might instead receive an adversarially-controlled value $g^{y'}$ for which we then, for a correct simulation, need to be able to compute $g^{xy'}$ (while knowing neither x nor y'). To this extent, we model the security of the HKDF.Extract function deriving HS using ES as salt and $\mathrm{DHE} = g^{xy}$ as source key material using the PRF-ODH assumption [JKSS12] (in its single-query variant). This allows us to replace HS in the tested session with a random value while still being able to compute the (potentially different) handshake secret HS' derived from $g^{xy'}$ (for an arbitrary $g^{y'} \neq g^y$).

More precisely, in Game C.3 we replace the handshake secret HS derived from ES and $\mathrm{DHE} = g^{xy}$ in the tested and (potentially) its partnered session by a uniformly random value $\widetilde{\mathrm{HS}} \xleftarrow{\$} \{0,1\}^\lambda$. The introduced advantage difference for $\mathcal{A}$ can be bound by the advantage of an algorithm $\mathcal{B}_{12}$ against the PRF-ODH security of HKDF.Extract (using ES as salt and $\mathrm{DHE} = g^{xy}$ as source key material). To this extent, $\mathcal{B}_{12}$ outputs ES (precomputed from the test session's pre-shared secret guessed in Game C.2) as the PRF challenge label. It then employs the obtained Diffie–Hellman shares as values g^x and g^y in the `ClientKeyShare` resp. `ServerKeyShare` message of the tested session and the contributive stage-3 partner session (guessed in Game C.1). It furthermore uses the obtained PRF challenge value as the handshake secret HS in the tested session and, potentially, the partnered session (obtaining the same Diffie–Hellman shares). In case the guessed contributive partner session from Game C.1 is a client session and obtains, within `ServerKeyShare`, a value $g^{y'} \neq g^y$, $\mathcal{B}_{12}$ leverages its (single) PRF-ODH query to compute HS from $g^{xy'}$ (without knowing x or y').

Recall that $\mathcal{B}_{12}$ is free to replace values g^x and g^y in the tested and contributive-partnered session at will as $\mathcal{A}$ can only passively observe them. The simulation $\mathcal{B}_{12}$ provides to $\mathcal{A}$ hence equals Game C.2 in case the PRF challenge value $\mathcal{B}_{12}$ obtains is computed as $\mathsf{HKDF.Extract}(\mathrm{ES}, g^{xy})$ while, if the challenge is a uniformly random value it equals Game C.3. Therefore,

$$\mathsf{Adv}^{G_{C.2}}_{\texttt{draft-14-PSK-DHE-ORTT},\mathcal{A}} \leq \mathsf{Adv}^{G_{C.3}}_{\texttt{draft-14-PSK-DHE-ORTT},\mathcal{A}} + \mathsf{Adv}^{\mathsf{PRF\text{-}ODH}}_{\mathsf{HKDF.Extract},\mathbb{G},\mathcal{B}_{12}}.$$

Observe that in this reduction we do not rely on (the secrecy of) the early secret ES at all. This in particular allows ES to be known to $\mathcal{A}$ through issuing a Corrupt query on the pre-shared secret $\mathsf{pss} = \mathrm{RMS}$ involved in the tested session at any time, thereby ensuring forward secrecy (from stage 3 on).

We complete this proof case by applying to the tested session (and its potential partner) the steps described for the Games 7, 8, 9, and 10 from the proof of `draft-14` PSK(-only) Multi-Stage security (Theorem 5.2). Thereby, we in particular replace the session keys for stages 3–5, tk_{hs}, tk_{app}, and EMS, by independent and uniformly random values sampled from $\{0,1\}^\lambda$. This leaves the adversary $\mathcal{A}$ with no chance to determine b_{test} better than through guessing and hence bounds its advantage in Game C.3 as

$$\mathsf{Adv}^{G_{C.3}}_{\texttt{draft-14-PSK-DHE-ORTT},\mathcal{A}} \leq \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HMAC},\mathcal{B}_{13}} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_{14}} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_{15}} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_{16}}.$$

---

[17]Observe that the server's MAC value in `ServerFinished` is only processed after deriving the stage-3 key.

where $\mathcal{B}_{13}, \ldots, \mathcal{B}_{16}$ are the algorithms $\mathcal{B}_6$, $\mathcal{B}_7$, $\mathcal{B}_8$, and $\mathcal{B}_9$ given for Games 7, 8, 9, and 10 in the proof of Theorem 5.2.

To conclude, the advantage bounds for this proof case sum up to:

$$\mathsf{Adv}^{\text{1-Multi-Stage,no-coll,stage 3–5,cid}_3\text{-partner}}_{\text{draft-14-PSK-DHE-0RTT},\mathcal{A}} \leq n_s \cdot n_p \cdot \Big(\mathsf{Adv}^{\text{PRF-ODH}}_{\text{HKDF.Extract},\mathbb{G},\mathcal{B}_{12}} + \mathsf{Adv}^{\text{PRF-sec}}_{\text{HMAC},\mathcal{B}_{13}} + \mathsf{Adv}^{\text{PRF-sec}}_{\text{HKDF.Expand},\mathcal{B}_{14}}$$

$$+ \mathsf{Adv}^{\text{PRF-sec}}_{\text{HKDF.Expand},\mathcal{B}_{15}} + \mathsf{Adv}^{\text{PRF-sec}}_{\text{HKDF.Expand},\mathcal{B}_{16}} \Big). \qquad \square$$

6 The TLS 1.3 draft-12 (EC)DHE 0-RTT Handshake Protocol

The latest TLS 1.3 drafts do not specify a Diffie–Hellman-based ((EC)DHE) 0-RTT handshake anymore, the last draft doing so is `draft-12` [Res16b]. We nevertheless provide a security analysis of this 0-RTT mode (as specified in `draft-12`) for two reasons: For one, it is much closer to the QUIC and OPTLS protocols and our analysis hence enables a comparison with those designs. For another, it provides slightly stronger forward secrecy properties [Kra16] as reflected in our analysis and may (for that or other reasons) be re-established as a TLS 1.3 extension [Res16f].

On a high level, the (EC)DHE 0-RTT handshake goes through the same four phases as the PSK-based 0-RTT modes: key exchange, 0-RTT, server parameters, and authentication (cf. Section 5). A notable difference though is that, in `draft-12`, the client may perform signature-based authentication in the 0-RTT step. We provide the protocol flow (and cryptographic computations) as well as the key schedule of the (EC)DHE 0-RTT handshake in Figure 3 and explain the handshake messages in the following.

- `ClientHello` (CH)/`ServerHello` (SH) transmit—as for the PSK 0-RTT handshake—the randomly chosen nonces r_c / r_s of the respective parties, as well as negotiation parameters (supported versions and ciphersuites). Both messages furthermore can include various extension fields; (EC)DHE 0-RTT in particular requires the following two extensions.

- `ClientEarlyData` (CEAD)/`ServerEarlyData` (SEAD) are extensions sent to announce a 0-RTT handshake. For `draft-12` (EC)DHE 0-RTT, the client includes an identifier `config_id` for a previously obtained server configuration (including a semi-static public key $S = g^s$ for which the server holds the secret exponent s) along with a matching ciphersuite used for deriving (and encrypting under) the 0-RTT keys.[18] The server signals accepting the 0-RTT exchange with an empty `ServerEarlyData` extension.

- `ClientKeyShare` (CKS)/`ServerKeyShare` (SKS) are extension fields that contain the ephemeral Diffie–Hellman shares $X = g^x$ resp. $Y = g^y$ for several (in case of the client) or one group (the server decided for).

After sending its `ClientHello` and extensions, the client can already derive one of the two main secret inputs for key derivation, namely the *static secret* SS, as the Diffie–Hellman shared value g^{xs}.[19] The initial 0-RTT keys are then derived using HKDF in the extract-then-expand paradigm [Kra10]: first, an extracted value xSS is computed from which, in a second step, the 0-RTT handshake and 0-RTT application data traffic keys tk_{ehs} and tk_{ead} are expanded. Here, we use the same notation for the two HKDF functions as introduced in Section 4.

---

[18]TLS 1.3 does generally not provide protection against replay of 0-RTT data between multiple connections, but allows inclusion of an optional context value within the `CEAD` message to implement unique per-connection configuration identifiers delivered out-of-band as an anti-replay measurement outside of TLS (cf. [Res16b, Section 6.3.2.5.2]). We do not consider this special mechanism in our analysis.

[19]Note that the TLS 1.3 key schedule changed significantly from `draft-12` to `draft-13` and hence in particular differs from that in our analysis of the `draft-14` PSK-based 0-RTT handshakes in Sections 4 and 5.
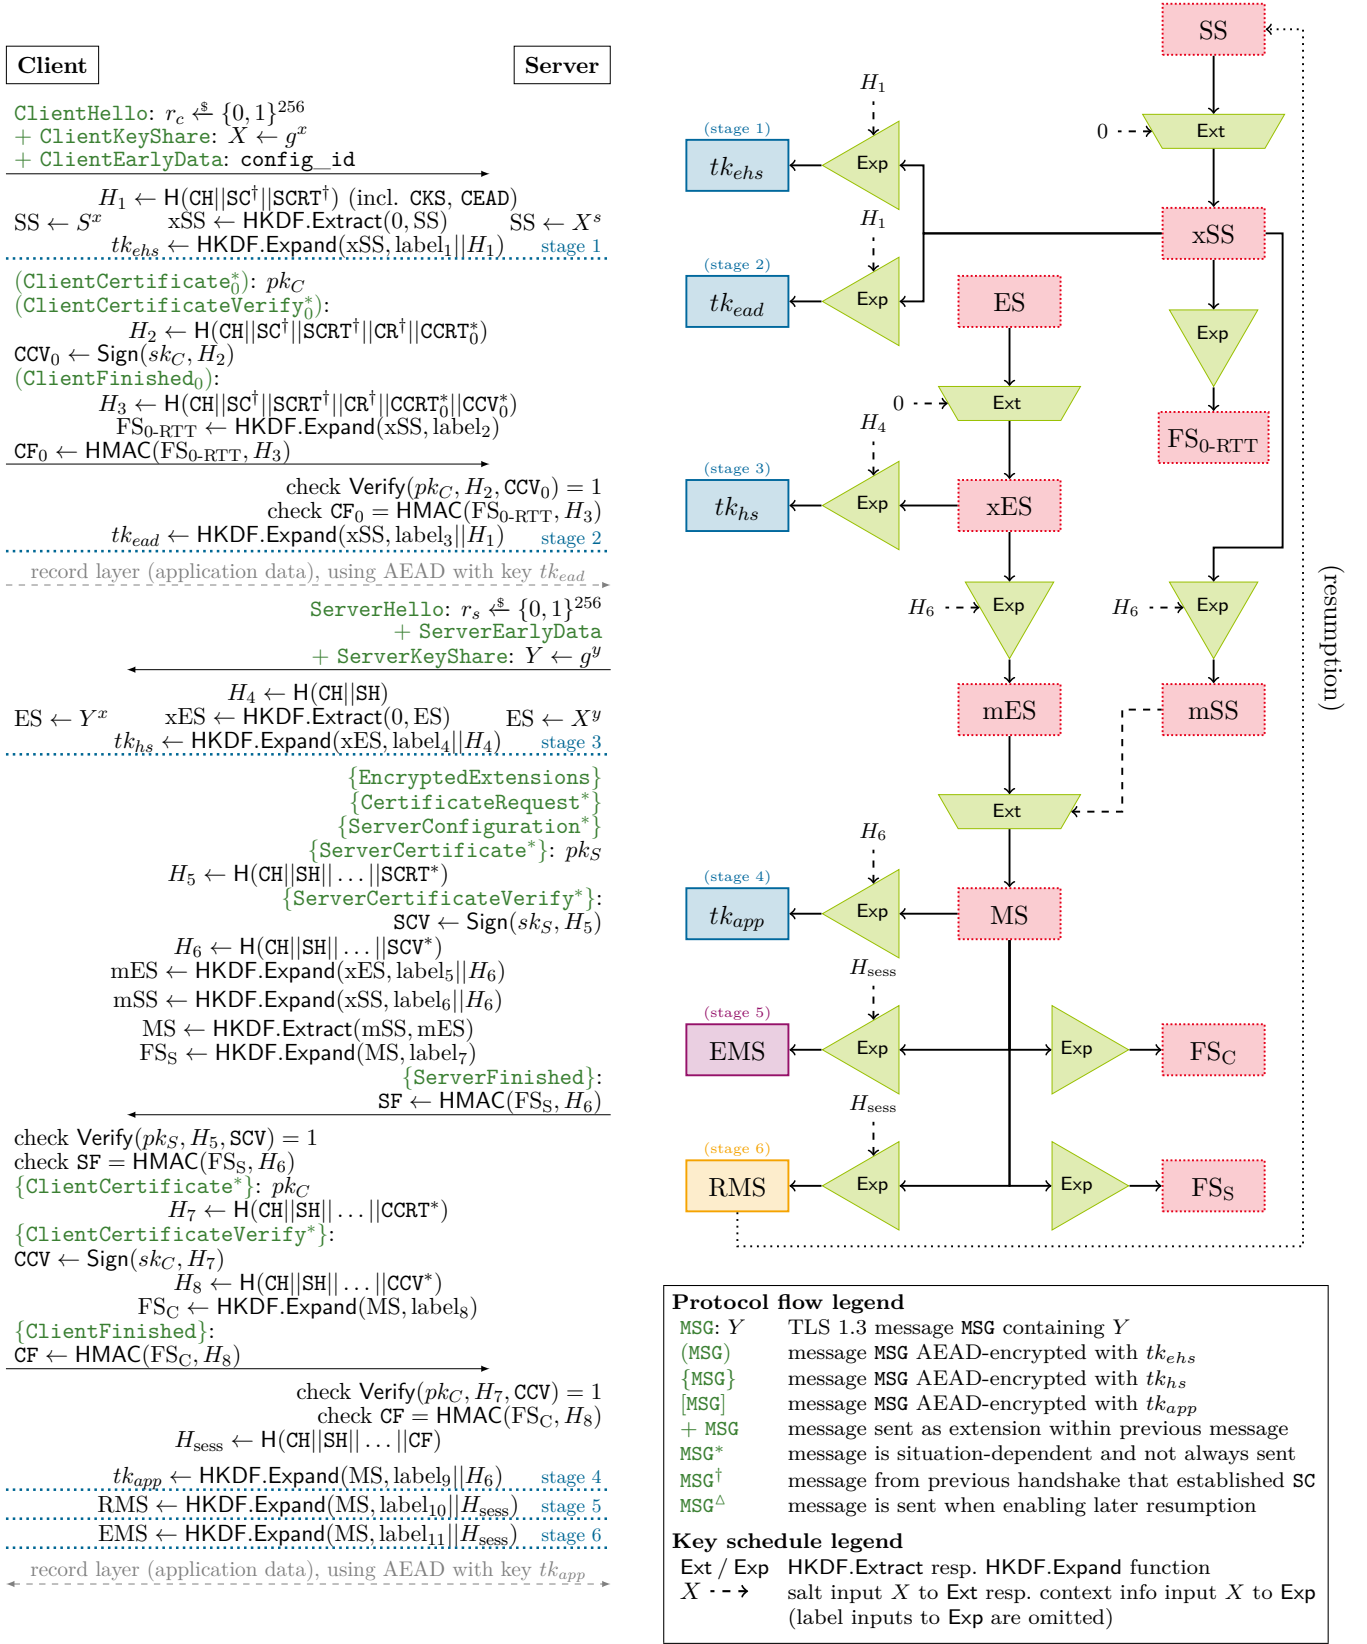
Client   **Server**

ClientHello: $r_c \xleftarrow{\$} \{0,1\}^{256}$
+ ClientKeyShare: $X \leftarrow g^x$
+ ClientEarlyData: config_id

$\quad H_1 \leftarrow \mathsf{H}(CH||SC^\dagger||SCRT^\dagger)$ (incl. CKS, CEAD)
$SS \leftarrow S^x \qquad xSS \leftarrow \mathsf{HKDF.Extract}(0, SS) \qquad SS \leftarrow X^s$
$\qquad tk_{ehs} \leftarrow \mathsf{HKDF.Expand}(xSS, label_1||H_1) \qquad$ stage 1

(ClientCertificate$_0^*$): pk_C
(ClientCertificateVerify$_0^*$):
$\quad H_2 \leftarrow \mathsf{H}(CH||SC^\dagger||SCRT^\dagger||CR^\dagger||CCRT_0^*)$
$CCV_0 \leftarrow \mathsf{Sign}(sk_C, H_2)$
(ClientFinished$_0$):
$\quad H_3 \leftarrow \mathsf{H}(CH||SC^\dagger||SCRT^\dagger||CR^\dagger||CCRT_0^*||CCV_0^*)$
$\quad FS_{0\text{-}RTT} \leftarrow \mathsf{HKDF.Expand}(xSS, label_2)$
$CF_0 \leftarrow \mathsf{HMAC}(FS_{0\text{-}RTT}, H_3)$

$\qquad\qquad$ check $\mathsf{Verify}(pk_C, H_2, CCV_0) = 1$
$\qquad\qquad$ check $CF_0 = \mathsf{HMAC}(FS_{0\text{-}RTT}, H_3)$
$\quad tk_{ead} \leftarrow \mathsf{HKDF.Expand}(xSS, label_3||H_1) \qquad$ stage 2

$\cdots\cdots$ record layer (application data), using AEAD with key tk_{ead} $\cdots\cdots$

$\qquad$ ServerHello: $r_s \xleftarrow{\$} \{0,1\}^{256}$
$\qquad\qquad$ + ServerEarlyData
$\qquad\qquad$ + ServerKeyShare: $Y \leftarrow g^y$

$\qquad H_4 \leftarrow \mathsf{H}(CH||SH)$
$ES \leftarrow Y^x \qquad xES \leftarrow \mathsf{HKDF.Extract}(0, ES) \qquad ES \leftarrow X^y$
$\qquad tk_{hs} \leftarrow \mathsf{HKDF.Expand}(xES, label_4||H_4) \qquad$ stage 3

$\qquad\qquad\qquad$ {EncryptedExtensions}
$\qquad\qquad\qquad$ {CertificateRequest*}
$\qquad\qquad\qquad$ {ServerConfiguration*}
$\qquad\qquad\qquad$ {ServerCertificate*}: pk_S
$H_5 \leftarrow \mathsf{H}(CH||SH||\ldots||SCRT^*)$
$\qquad\qquad\qquad$ {ServerCertificateVerify*}:
$\qquad\qquad\qquad\qquad SCV \leftarrow \mathsf{Sign}(sk_S, H_5)$
$H_6 \leftarrow \mathsf{H}(CH||SH||\ldots||SCV^*)$
$mES \leftarrow \mathsf{HKDF.Expand}(xES, label_5||H_6)$
$mSS \leftarrow \mathsf{HKDF.Expand}(xSS, label_6||H_6)$
$MS \leftarrow \mathsf{HKDF.Extract}(mSS, mES)$
$FS_S \leftarrow \mathsf{HKDF.Expand}(MS, label_7)$
$\qquad\qquad\qquad$ {ServerFinished}:
$\qquad\qquad\qquad\qquad SF \leftarrow \mathsf{HMAC}(FS_S, H_6)$

check $\mathsf{Verify}(pk_S, H_5, SCV) = 1$
check $SF = \mathsf{HMAC}(FS_S, H_6)$
{ClientCertificate*}: pk_C
$\qquad H_7 \leftarrow \mathsf{H}(CH||SH||\ldots||CCRT^*)$
{ClientCertificateVerify*}:
$CCV \leftarrow \mathsf{Sign}(sk_C, H_7)$
$\qquad H_8 \leftarrow \mathsf{H}(CH||SH||\ldots||CCV^*)$
$\qquad FS_C \leftarrow \mathsf{HKDF.Expand}(MS, label_8)$
{ClientFinished}:
$CF \leftarrow \mathsf{HMAC}(FS_C, H_8)$

$\qquad\qquad$ check $\mathsf{Verify}(pk_C, H_7, CCV) = 1$
$\qquad\qquad$ check $CF = \mathsf{HMAC}(FS_C, H_8)$
$\qquad H_{sess} \leftarrow \mathsf{H}(CH||SH||\ldots||CF)$

$\quad tk_{app} \leftarrow \mathsf{HKDF.Expand}(MS, label_9||H_6) \qquad$ stage 4
$\quad RMS \leftarrow \mathsf{HKDF.Expand}(MS, label_{10}||H_{sess}) \qquad$ stage 5
$\quad EMS \leftarrow \mathsf{HKDF.Expand}(MS, label_{11}||H_{sess}) \qquad$ stage 6

$\cdots\cdots$ record layer (application data), using AEAD with key tk_{app} $\cdots\cdots$

Protocol flow legend

MSG: Y	TLS 1.3 message MSG containing Y
(MSG)	message MSG AEAD-encrypted with tk_{ehs}
{MSG}	message MSG AEAD-encrypted with tk_{hs}
[MSG]	message MSG AEAD-encrypted with tk_{app}
+ MSG	message sent as extension within previous message
MSG*	message is situation-dependent and not always sent
MSG†	message from previous handshake that established SC
MSG$^\triangle$	message is sent when enabling later resumption

Key schedule legend

Ext / Exp	HKDF.Extract resp. HKDF.Expand function
$X \dashrightarrow$	salt input X to Ext resp. context info input X to Exp (label inputs to Exp are omitted)

Figure 3: The TLS 1.3 `draft-12` (EC)DHE with 0-RTT handshake protocol (left) and key schedule (right).

34

The following three messages of the 0-RTT phase are then sent encrypted using tk_{ehs}, which is unilaterally (server-side) authenticated. They are conceptually similar to the client's authentication messages sent later (as in the full handshake), but marked with a subscript 0 for distinction.

- ClientCertificate$_0$ (CCRT$_0$) contains the client's public-key certificate and is optionally sent if the client wishes to authenticate its 0-RTT data.

- ClientCertificateVerify$_0$ (CCV$_0$) is sent only if ClientCertificate$_0$ is sent and contains a digital signature over the *handshake hash* (the hash of all sent and received handshake messages at this point, i.e., here only the client's messages CH (incl. extensions) and CCRT$_0$), as well as the server configuration used by the client, and further messages from the previous handshake in which the client received this server configuration.

- ClientFinished$_0$ (CF$_0$) contains an HMAC value, a message authentication code computed over (a hash of) the same messages as for CCV$_0$ and the CCV$_0$ message itself; keyed with the 0-RTT finished secret FS$_{0\text{-RTT}}$ derived from xSS.

At this point, the client can send early (0-RTT) application data, encrypted under tk_{ead} (which is server- and optionally also client-authenticated).

After receiving the client's first flight, the server can derive the same 0-RTT keys using its stored configuration for the semi-static key g^s and process the client's 0-RTT handshake and application data. The server can also decide to reject 0-RTT keys and data (setting sid_1, sid_2, K_1, and K_2 to $\perp$ in our model); it must in particular do so if the server configuration identifier config_id sent by the client is invalid, unknown, or expired. In either case, the server continues the handshake by sending out its ServerHello message and extensions.

Both parties can then compute the second secret input, the *ephemeral secret* ES, as the shared Diffie–Hellman value g^{xy}, and an extracted value xES from which the (unauthenticated) handshake traffic key tk_{hs} is expanded. The key tk_{hs} is then used to encrypt the remaining handshake:

- EncryptedExtensions (EE) allows to specify further extensions.

- CertificateRequest (CR) is sent by the server if it demands client authentication.

- ServerConfiguration (SC) is optionally sent to provide the client with a new server configuration (cryptographically a semi-static Diffie–Hellman share $g^{s'}$) which the client can later use for another (EC)DHE 0-RTT handshake.

- ServerCertificate (SCRT)/ClientCertificate (CCRT) contain the client's resp. server's public-key certificate; each is sent if the according party is supposed to authenticate.

- ServerCertificateVerify (SCV)/ClientCertificateVerify (CCV) are digital signatures over the current handshake hash at the point when they are sent.

After the ServerCertificateVerify message is sent resp. received, both parties derive the master secret MS, extracted from the intermediate expanded versions mES and mSS of the (extracted) ephemeral and static secrets xES and xSS).

- ClientFinished (CF)/ServerFinished (SF) contain an HMAC value over the handshake hash using the client resp. server finished secret FS$_C$/FS$_S$ as the key; both finished secrets are derived from the master secret MS.

At this point, three final keys are derived from the master secret through HKDF expansion steps: the application traffic key tk_{app} for protecting the (non–0-RTT) application data sent[20], the resumption master secret RMS for later preshared-key–based session resumption, and the exporter master secret EMS from which further key material for usage outside of TLS can be derived.

As for the PSK-based 0-RTT handshakes (cf. Section 6), our analysis of the draft-12 (EC)DHE 0-RTT handshake focuses on its core components and we hence do not treat 0.5-RTT data and post-handshake messages.

7 Security of the TLS 1.3 draft-12 (EC)DHE 0-RTT Handshake

We conduct our security analysis of the TLS 1.3 draft-12 (EC)DHE 0-RTT handshake (draft-12-(EC)DHE-0RTT) in the public-key variant (MSKE) of the multi-stage key exchange model. First, we need to state the (formalized) protocol-specific properties (such as the available authentication modes, its replayability properties, etc.) as well as how session matching is defined via the session and contributive identifiers for each stage. The properties are captured as follows:

- $M = 6$: the 0-RTT handshake has six stages (deriving, in that order, keys tk_{ehs}, tk_{ead}, tk_{hs}, tk_{app}, RMS, and EMS).

- $\mathsf{AUTH} = \{(\mathsf{unilateral}, \mathsf{auth}_{ead}, \mathsf{unauth}, \mathsf{auth}_{fin}, \mathsf{auth}_{fin}, \mathsf{auth}_{fin}) \mid \mathsf{auth}_{ead} \in \{\mathsf{unilateral}, \mathsf{mutual}\}, \mathsf{auth}_{fin} \in \{\mathsf{unauth}, \mathsf{unilateral}, \mathsf{mutual}\}\}$: the first-stage key is always unilaterally authenticated, the second-stage (early-data) key can bei unilaterally or mutually authenticated, the traffic handshake key is always unauthenticated, and the final three keys (tk_{app}, RMS, and EMS) share the same authentication property which can be no authentication, unilateral authentication, or mutual authentication.

- $\mathsf{USE} = (\mathsf{internal}, \mathsf{external}, \mathsf{internal}, \mathsf{external}, \mathsf{external}, \mathsf{external})$: the early-data and regular handshake traffic keys tk_{ehs} and tk_{hs} are used within the handshake, whereas the early-data and main application traffic keys tk_{ead} and tk_{app} as well as the resumption and exporter master secret RMS and EMS are used only externally.

- $\mathsf{REPLAY} = (\mathsf{replayable}, \mathsf{replayable}, \mathsf{nonreplayable}, \mathsf{nonreplayable}, \mathsf{nonreplayable}, \mathsf{nonreplayable})$: the early-data stages 1 and 2 are replayable, the other stages are not.

As we will see, the TLS 1.3 draft-12 0-RTT handshake furthermore enjoys key independence and forward secrecy (wrt. compromise of long-term secrets) for all keys. For the forward secrecy of the early-data (0-RTT) keys tk_{ehs} and tk_{ead}, recall that our model treats compromises of long-term and semi-static secrets independently through the Corrupt resp. RevealSemiStaticKey query. While those keys remain (forward) secret after a long-term key compromise, they are replayable and hence become insecure when the involved semi-static key is revealed.[21]

For session matching, we define the session identifiers for the first four stages to be the unencrypted messages sent and received that enter the handshake hash for the respective key's derivation, including (for

---

[20]More precisely, TLS 1.3 derives an intermediate *traffic secret* from which the actual application traffic key tk_{app} is expanded. This is done in order to allow for later key updates, where first an updated traffic secret is computed from which then the new traffic key is derived. We do not capture this key update mechanism and hence omit the intermediate traffic secret derivation for simplicity.

[21]For this reason, our result in particular does not contradict the statement in draft-12 that "[t]his [0-RTT] data is not forward secret, because it is encrypted solely with the server's semi-static (EC)DH share" [Res16b, Section 6.2.2]. The draft merely requires a forward-secret key to be resilient against compromises of both long-term *and* semi-static keys.

the early-data stages) the `ServerConfiguration` and accompanying messages the client received earlier:

$$\mathsf{sid}_1 = (\texttt{ClientHello}, \texttt{ServerConfiguration}^\dagger, \texttt{ServerCertificate}^\dagger, \texttt{CertificateRequest}^\dagger),$$

$$\mathsf{sid}_2 = (\texttt{ClientHello}, \texttt{ServerConfiguration}^\dagger, \texttt{ServerCertificate}^\dagger, \texttt{CertificateRequest}^\dagger,$$
$$\texttt{ClientCertificate}_0^*, \texttt{ClientCertificateVerify}_0^*, \texttt{ClientFinished}_0),$$

$$\mathsf{sid}_3 = (\texttt{ClientHello}, \texttt{ServerHello}), \quad \text{and}$$

$$\mathsf{sid}_4 = (\texttt{ClientHello}, \texttt{ServerHello}, \texttt{EncryptedExtensions}, \texttt{CertificateRequest}^*,$$
$$\texttt{ServerConfiguration}^*, \texttt{ServerCertificate}^*, \texttt{ServerCertificateVerify}^*, \texttt{ServerFinished},$$
$$\texttt{ClientCertificate}^*, \texttt{ClientCertificateVerify}^*, \texttt{ClientFinished}).$$

Here, `Hello` messages always contain the according `KeyShare` and `EarlyData` extensions, components marked with † are those enabling the 0-RTT exchange received by the client in an earlier handshake, and starred (*) components are not present in all authentication modes. For the two final stages, we define the session identifiers to contain a distinguishing label beyond the full stage-4 identifier, namely $\mathsf{sid}_5 = (\mathsf{sid}_4, \text{"RMS"})$ and $\mathsf{sid}_6 = (\mathsf{sid}_4, \text{"EMS"})$.

As for the PSK-based handshakes (cf. Section 5), we define the contributive identifiers such that a server session can be tested when receiving an honest client contribution. That is, for stage 3 (deriving the handshake traffic key), we let the client (resp. server) on sending (resp. receiving) the `ClientHello` message (including the `ClientKeyShare` and `ClientEarlyData` extension) initially set $\mathsf{cid}_3 = (\texttt{ClientHello})$ and subsequently, on receiving (resp. sending) the `ServerHello` message (incl. `SKS`, `SEAD`), extend it to $\mathsf{cid}_3 = (\texttt{ClientHello}, \texttt{ServerHello})$. The other contributive identifiers are set to $\mathsf{cid}_i = \mathsf{sid}_i$, for stages $i \in \{1,2\}$ when sending resp. receiving the 0-RTT `ClientFinished` message and for stages $i \in \{4,5,6\}$ by each party on sending its respective `Finished` message.

Finally, we capture as semi-static public resp. private keys the values g^s and s incorporated in the derivation of the static secret SS and 0-RTT key derivation; issued by servers and learned by clients via some `ServerConfiguration` (`SC`) message in an earlier (EC)DHE full or 0-RTT handshake. This message is sent together with a `ServerCertificate` (`SCRT`) and, optionally, `CertificateRequest` (`CR`) message (the latter enabling 0-RTT client authentication) and is signed within the `ServerCertificateVerify` message. In our model, we let the adversary control the generation of semi-static keys through the NewSemiStaticKey query, deciding whether a `CR` message is sent or not by setting the optional input $\mathsf{sskaux}_{\mathsf{pre}} = \texttt{CR}$. The NewSemiStaticKey query then samples an exponent value s at random to generate a new semi-static key, and outputs the auxiliary information $\mathsf{sskaux} = (\texttt{SC}, \texttt{SCRT})$ resp. $\mathsf{sskaux} = (\texttt{SC}, \texttt{SCRT}, \texttt{CR})$ along with $\mathsf{sspk} = g^s$ and a key identifier sskid. When instantiating a new session (through NewSession), the adversary controls which semi-static key a client session uses for the 0-RTT exchange via the sskid_V identifier and determines whether a server issues a new semi-static key in a `ServerConfiguration` message (and which key this is) via the server sessions' sskid_U identifier.[22] Finally, the RevealSemiStaticKey query allows the adversary to learn the secret exponent s for semi-static keys of its choice, at the price of not being allowed to test stage-1 and stage-2 (i.e., early-data) keys for which this semi-static key was used.

We are now able to provide our security results for the TLS 1.3 `draft-12` (EC)DHE 0-RTT handshake.

Theorem 7.1 (Match security of `draft-12-(EC)DHE-0RTT`). *The* `draft-12` *(EC)DHE 0-RTT handshake is* Match-*secure: for any efficient adversary $\mathcal{A}$ we have*

$$\mathsf{Adv}^{\mathsf{Match}}_{\texttt{draft-12-(EC)DHE-0RTT}, \mathcal{A}} \leq n_s^2 \cdot 1/q \cdot 2^{-|nonce|},$$

where n_s is the maximum number of sessions, q is the group order for the Diffie–Hellman problem, and $|nonce| = 256$ is the bit-length of the nonces.

---

[22]Note that in TLS 1.3 only servers hold semi-static keys. In particular, NewSession queries in our model will hence have the client's semi-static key identifier (sskid) parameter set to $\bot$.

Proof. We need to show that the six conditions for Match security hold:

1. *Sessions with the same session identifier for some stage hold the same key at that stage.*
   The session identifiers for stages 1–2 and stages 3–6 contain the client's resp. the client's and the server's Hello messages and hence fix the Diffie–Hellman shares g^x and g^s (through the unique semi-static key/con-figuration identifier $\mathsf{sskid} = \mathtt{config\_id}$) resp. also g^y. Therefore, coinciding session identifiers imply agreement on the static secret SS (in stages 1–2) and, for stage 3–4, also the ephemeral secret ES (i.e., the input keying material). As each session identifier in particular furthermore contains all messages exchanged that enter the handshake hash within the key derivations, the session keys are uniquely determined by the session identifier in each stage.

2. *Sessions with the same session identifier for some stage agree on that stage's authentication level.*
   For stages 1 and 3, the authentication level is fixed to unilateral resp. unauth. For the other stages, unilateral authentication is indicated (exactly) by the exchange of SCRT and SCV messages, whereas mutual authentication requires messages CCRT and CCV (as well as, for stages 4–6, also the server's messages). Hence, agreement on the session identifier (including theses messages) implies agreement on that stage's authentication.

3. *Sessions with the same session identifier for some stage share the same contributive identifier.*
   For stages $i \in \{1, 2, 4, 5, 6\}$ this follows immediately from $\mathsf{sid}_i = \mathsf{cid}_i$. For stage 3, note that when both sides set the session identifier, the contributive identifier is also set to its final value $\mathsf{cid}_3 = \mathsf{sid}_3$.

4. *Sessions are partnered with the intended (authenticated) participant.*
   As sessions of honest parties will not attest a different identity than their own in the SCRT and CCRT messages nor accept such a message for an identity different from the intended partner, agreement on these messages (which are included in the respective session identifiers for unilaterally and mutually authenticated stages) in particular implies agreement on each partner's identity.

5. *Session identifiers do not match across different stages.*
   This holds trivially since session identifiers sid_1–sid_4 contain distinct (non-optional) messages and sid_5 and sid_6 include a separating identifier.

6. *At most two sessions have the same session identifier at any non-replayable stage.*
   As stages 1 and 2 are replayable, we need to consider this condition only for the stages 3–6.[23] Observe that the session identifiers for those stages contain the client's and server's Hello message and, hence, for each side a randomly chosen nonce n_c (n_s) as well as a randomly chosen group element g^x (g^y). Therefore, in order for a third (client resp. server) session to agree on the same session identifier it needs to, at least, pick the same nonce and group element as the client resp. server, which can be upper bounded by the probability of any two parties colliding on the same nonce and group element. This probability can, again, be bounded from above by $n_s^2 \cdot 1/q \cdot 2^{-|nonce|}$, where n_s is the number of all (client or server) sessions, q is the group order, and $|nonce| = 256$ the bit-length of the nonces. □

Theorem 7.2 (Multi-Stage security of $\mathtt{draft\text{-}12\text{-}(EC)DHE\text{-}0RTT}$)**.** *The* $\mathtt{draft\text{-}12}$ *(EC)DHE 0-RTT handshake is* Multi-Stage*-secure in a key-independent and stage-1-forward-secret manner with properties* (M, AUTH, USE, REPLAY)*. Formally, for any efficient adversary $\mathcal{A}$ against the* Multi-Stage *security there exist*

---

[23]Note that, indeed, a server has no means to check whether the client's first-flight messages (being the only content of the first two stages' session identifiers) have been transmitted to another server before. This allows an adversary to replay those messages and, hence, make one client session be partnered with multiple server sessions, all deriving the same session key. As we will see in a moment, this still does not allow the adversary to break those keys' secrecy.

efficient algorithms $\mathcal{B}_1, \ldots, \mathcal{B}_{14}$ *such that*

$$
\begin{aligned}
\mathsf{Adv}^{\mathsf{Multi\text{-}Stage},\mathcal{D}}_{\mathtt{draft\text{-}12\text{-}(EC)DHE\text{-}0RTT},\mathcal{A}} \leq 6n_s \cdot \Big(& \mathsf{Adv}^{\mathsf{COLL}}_{\mathsf{H},\mathcal{B}_1} + n_u \cdot \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig},\mathcal{B}_2} \\
& + \mathsf{Adv}^{\mathsf{COLL}}_{\mathsf{H},\mathcal{B}_3} + n_s \cdot n_{ss} \cdot \Big(\mathsf{Adv}^{\mathsf{msPRF\text{-}ODH}}_{\mathsf{HKDF.Extract},\mathbb{G},\mathcal{B}_4} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_5} \Big) \\
& + \mathsf{Adv}^{\mathsf{COLL}}_{\mathsf{H},\mathcal{B}_6} + n_u \cdot \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig},\mathcal{B}_7} + \mathsf{Adv}^{\mathsf{COLL}}_{\mathsf{H},\mathcal{B}_8} + n_u \cdot \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig},\mathcal{B}_9} \\
& + \mathsf{Adv}^{\mathsf{COLL}}_{\mathsf{H},\mathcal{B}_{10}} + n_s \cdot \Big(\mathsf{Adv}^{\mathsf{PRF\text{-}ODH}}_{\mathsf{HKDF.Extract},\mathbb{G},\mathcal{B}_{11}} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_{12}} \\
& \qquad\qquad + \mathsf{Adv}^{\mathsf{st\text{-}Extract}}_{\mathsf{HKDF.Extract},\mathcal{B}_{13}} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_{14}} \Big) \Big),
\end{aligned}
$$

where n_u is the maximum number of users, n_s is the maximum number of sessions, and n_{ss} is the maximum number of semi-static keys.

Proof. First of all we consider that the adversary $\mathcal{A}$ makes a single Test query only. This reduces its advantage, based on a hybrid argument detailed for the full handshake proof [DFGS15b, Appendix A], by a factor at most $1/6n_s$ for the six stages in each of the n_s sessions. We can now speak about *the* session label tested at stage i, and that we know the index of the session and the stage in advance.

In our analysis, we then separately treat the (disjoint) cases that the adversary tests an early-data (stage-1 or stage-2) key or a regular key (stages 3–6). For tests on stages 1 and 2, we further distinguish between the following two (again disjoint) sub-cases that

A. the adversary tests a server session without honest (contributive) partner in the first stage (i.e., label.role = responder for the test session label and there exists no label′ ≠ label with label.cid$_1$ = label′.cid$_1$) and

B. the adversary tests a server session with honest (contributive) partner in the first stage or a client session (i.e., label.role = initiator or label.role = responder and there exists a label′ ≠ label with label.cid$_1$ = label′.cid$_1$).

For tests on stages 3–6, we split our analysis along the same (sub)cases C,D, and E used in the analysis of the full handshake mode [DFGS15a, DFGS15b, DFGS16] (for which, as we detail in the proof, the analysis closely follows the one for the draft-10 full handshake so that we moreover obtain the same security bounds):

C. the adversary tests a client session without honest contributive partner in the third stage (i.e., label.role = initiator for the test session label and there exists no label′ ≠ label with label.cid$_3$ = label′.cid$_3$),

D. the adversary tests a server session without honest contributive partner in the third stage (i.e., label.role = responder and there exists no label′ ≠ label with label.cid$_3$ = label′.cid$_3$), and

E. the tested session has an honest contributive partner in stage 3 (i.e., there exists label′ with label.cid$_3$ = label′.cid$_3$).

This allows us to split the adversary's advantage along these five cases:

$$\mathsf{Adv}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}}^{\textsf{Multi-Stage},\mathcal{D}} \leq 6n_s \cdot \Big(\mathsf{Adv}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}}^{\textsf{1-Multi-Stage, test 1–2, server without partner}}$$
$$+ \mathsf{Adv}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}}^{\textsf{1-Multi-Stage, test 1–2, server with partner/client}}$$
$$+ \mathsf{Adv}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}}^{\textsf{1-Multi-Stage, test 3–6, client without partner}}$$
$$+ \mathsf{Adv}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}}^{\textsf{1-Multi-Stage, test 3–6, server without partner}}$$
$$+ \mathsf{Adv}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}}^{\textsf{1-Multi-Stage, test 3–6, test with partner}} \Big).$$

For the proof of each case, we will proceed in a sequence of games: We start from the original Multi-Stage game with a single Test query (denoted 1-Multi-Stage), restricted to the case in question, and modify this game in each step, showing that the difference in the adversary's advantage between the two games can be bounded by complexity-theoretic assumptions. Finally, in the last game the advantage of $\mathcal{A}$ will be at most 0 and the advantage for the considered case hence bound by combination of the intermediate bounds.

Case A. Stage 1–2: Test Server without Stage-1 Partner

For the case that the adversary tests a server (responder) session in stage 1 or 2 without contributive partner in the first stage, we recall that $\mathsf{cid}_i = \mathsf{sid}_i$ for $i \in \{1,2\}$ and hence, as all messages in sid_1 are also contained in sid_2, the tested session also cannot have a partner in the second stage. In order to not lose immediately, the adversary can test responder session stages without contributive partner only if they are mutually authenticated. Since the stage-1 keys are unilaterally authenticated we can focus on the second stage to be tested and assume $\mathsf{label.auth}_2 = \mathsf{mutual}$.

Game A.0. We begin with the initial game $G_{A.0}$ which equals the Multi-Stage game with one Test query which must be issued on a stage-1 or stage-2 key of a server session without honest contributive identifier. Therefore,
$$\mathsf{Adv}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}}^{G_{A.0}} = \mathsf{Adv}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}}^{\textsf{1-Multi-Stage,test 1–2, server without partner}}.$$

Game A.1. Next, we exclude hash collisions in the execution by letting the challenger abort in case two honest sessions compute the same hash value for two distinct inputs in any evaluation of the hash function H. We can bound the probability of the game being aborted for that reason (i.e., the only way for $\mathcal{A}$ to distinguish the change) by advantage probability of an adversary $\mathcal{B}_1$ in breaking the collision resistance of H. For this, $\mathcal{B}_1$ can simply take the role of the challenger and, in case of the abort, output the colliding hash inputs. This way $\mathcal{B}_1$ wins if the game is aborted and hence

$$\mathsf{Adv}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}}^{G_{A.0}} \leq \mathsf{Adv}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}}^{G_{A.1}} + \mathsf{Adv}_{\mathsf{H},\mathcal{B}_1}^{\mathsf{COLL}}.$$

Game A.2. As the final change, we let the challenger abort the game if the tested server session receives a $\texttt{ClientCertificateVerify}_0$ message which contains a valid signature (under some public key pk_U) that no honest session ever computed.

We can bound the probability of such an event by the advantage of an adversary $\mathcal{B}_2$ against the unforgeability (in the sense of EUF-CMA) of the signature scheme Sig. In the reduction, $\mathcal{B}_2$ first needs to guess the identity $U \in \mathcal{U}$ under whose public key the obtained signature will verify, replacing that user's public key with the challenge public key in the EUF-CMA game and generating signatures using the signing

oracle. All other keys are generated by $\mathcal{B}_2$ during the game setup. When the tested session receives a valid $\mathtt{CCV}_0$ message causing an abort, $\mathcal{B}_2$ outputs the contained signature as forgery.

As the tested session has no honest partner, no honest client signed the contained handshake hash before. Otherwise, this client would agree on all messages up to $\mathtt{CCV}_0$ and also on $\mathtt{CF}_0$ (as these former messages uniquely determine the latter), and hence also on the session identifier. Moreover, no party with a different session identifier signed the same handshake hash, as this would constitute a hash collision which we excluded in Game A.1. Given that $\mathcal{B}_2$ correctly guessed the used public key pk_U (among the keys of the at most n_u users), its output constitutes a valid forgery and hence

$$\mathsf{Adv}^{G_{A.1}}_{\mathtt{draft\text{-}12\text{-}(EC)DHE\text{-}0RTT},\mathcal{A}} \leq \mathsf{Adv}^{G_{A.2}}_{\mathtt{draft\text{-}12\text{-}(EC)DHE\text{-}0RTT},\mathcal{A}} + n_u \cdot \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig},\mathcal{B}_2}.$$

At this point, we ensured that the tested session obtains a $\mathtt{ClientCertificateVerify}_0$ message issued by an honest client. This client hence agrees on all messages up to $\mathtt{CCV}_0$, which also determine $\mathtt{CF}_0$, and hence will hold the same session identifier sid_2. Therefore, the tested session cannot be without (contributive) partner and hence $\mathcal{A}$ cannot issue a Test query anymore at this point, leaving the test bit unknown to $\mathcal{A}$ and thus

$$\mathsf{Adv}^{G_{A.2}}_{\mathtt{draft\text{-}12\text{-}(EC)DHE\text{-}0RTT},\mathcal{A}} \leq 0.$$

Case B. Stage 1–2: Test Server with Stage-1 Partner or Client

In the second case for tests on early-data keys, we know that a tested server session always has a partnered session and that, for client sessions, the server is always authenticated (through the server configuration of some previous communication). A Test query can in this case be issued in any of the two stages.

Game B.0. We start with the unmodified initial game $G_{B.0}$:

$$\mathsf{Adv}^{G_{B.0}}_{\mathtt{draft\text{-}12\text{-}(EC)DHE\text{-}0RTT},\mathcal{A}} = \mathsf{Adv}^{\mathsf{1\text{-}Multi\text{-}Stage, \ test \ 1\text{-}2, \ server \ with \ partner/client}}_{\mathtt{draft\text{-}12\text{-}(EC)DHE\text{-}0RTT},\mathcal{A}}.$$

Game B.1. Our first modification then again excludes hash collisions by letting the challenger abort whenever two honest sessions on different inputs compute the same hash value under H. Like in Game A.1, the probability of this happening can be bounded by the advantage of an adversary $\mathcal{B}_3$ against H's collision resistance:[24]

$$\mathsf{Adv}^{G_{B.0}}_{\mathtt{draft\text{-}12\text{-}(EC)DHE\text{-}0RTT},\mathcal{A}} \leq \mathsf{Adv}^{G_{B.1}}_{\mathtt{draft\text{-}12\text{-}(EC)DHE\text{-}0RTT},\mathcal{A}} + \mathsf{Adv}^{\mathsf{COLL}}_{\mathsf{H},\mathcal{B}_3}.$$

Game B.2. Second, we guess the (index of the) client session involved in the test (i.e., the client session itself if the client-side is tested, or the client session partnered with the server session if the server side is tested) and abort if this guess is incorrect. Note that in both cases, this client session is an honest session simulated by the challenger. This can reduce the adversary's advantage by a factor of at most the number of sessions n_s:

$$\mathsf{Adv}^{G_{B.1}}_{\mathtt{draft\text{-}12\text{-}(EC)DHE\text{-}0RTT},\mathcal{A}} \leq n_s \cdot \mathsf{Adv}^{G_{B.2}}_{\mathtt{draft\text{-}12\text{-}(EC)DHE\text{-}0RTT},\mathcal{A}}.$$

Game B.3. Next, we additionally guess the (index of the) configuration identifier $\mathtt{config\_id}$ (i.e., in terms of our model, the semi-static key identifier sskid resp. the $\mathsf{NewSemiStaticKey}$ query through which

---

[24]In principle we could merge this game hop with the one in the first case to establish collision freeness of nonces once and for all, yielding a slightly better bound, but we prefer to present all cases in a closed form instead.

it is issued) the involved client session will use within its `ClientHello` message. Again, this reduces the advantage of $\mathcal{A}$ by a factor of at most the number of semi-static keys n_{ss}:

$$\mathsf{Adv}^{G_{B.2}}_{\texttt{draft-12-(EC)DHE-ORTT},\mathcal{A}} \leq n_{ss} \cdot \mathsf{Adv}^{G_{B.3}}_{\texttt{draft-12-(EC)DHE-ORTT},\mathcal{A}}.$$

Game B.4. At this point, we know in advance which semi-static key the involved client session will employ, enabling us to encode a Diffie–Hellman challenge in the static secret SS derived from the client's ephemeral share g^x and the server's semi-static share g^s (selected through the server configuration identifier).

As the server's semi-static key g^s is potentially used in more than one session, we need to be able to compute (keys from) further static secrets SS′ from the same g^s but a different $g^{x'}$, even when knowing neither s nor x'. Moreover, when the client session continues, it might obtain, within the `ServerKeyShare` message, an ephemeral Diffie–Hellman share $g^{y'}$ different from the value g^y chosen by the honest server session partnered in the early-data stages.[25] We hence, once, need to compute (keys from) the ephemeral secret ES (resp. xES) in the client session, even without knowing x or y'. To this extent, we model the HKDF.Extract function as a pseudorandom function keyed with elements from group $\mathbb{G}$ and employ the msPRF-ODH assumption (cf. Definition 2.2). The latter allows us to replace a HKDF/PRF value under SS with a random value while providing oracle access to evaluate the PRF under further keys SS′ = $g^{x's}$ (for $g^{x'} \neq g^x$) as well as, once, under a key ES′ = $g^{xy'}$ (for an arbitrary $g^{y'} \neq g^s$, without knowing x).

More in detail, in Game B.4 we replace the extracted static secret xSS by a uniformly random string $\widetilde{xSS} \xleftarrow{\$} \{0,1\}^\lambda$ in both the tested and its potentially partnered session(s), as well as in any session that sends or receives within the `ClientHello` message the same ephemeral key g^x and (identifier for the) semi-static key g^s (and ignore the SS value in those sessions). We can bound the difference in advantage of adversary $\mathcal{A}$ through this modification by the advantage of an algorithm $\mathcal{B}_4$ in winning the msPRF-ODH game as follows.

Initially, $\mathcal{B}_4$ receives a group element g^v in the msPRF-ODH game and immediately issues the challenge query $\hat{x} = 0$ for which it obtains a response $(g^{\hat{u}}, \hat{y})$ where $\hat{y}$ is either the value $\mathsf{PRF}(g^{\hat{u}v}, 0)$ or a uniformly random string. It then acts as the challenger in the Multi-Stage game for $\mathcal{A}$, choosing a test bit $b_{\mathsf{test}} \xleftarrow{\$} \{0,1\}$ at random and simulating it according to the description except for the following changes.

- When $\mathcal{A}$ issues the NewSemiStaticKey query guessed in Game B.3, algorithm $\mathcal{B}_4$ uses g^v for the returned semi-static public key sspk (implicitly setting $g^s = g^v$ for the tested session).

- For the tested session and its potential partnered session(s), $\mathcal{B}_4$ uses $g^{\hat{u}}$ as the ephemeral Diffie–Hellman share of the guessed tested client session resp. partnered client session of the tested server session, implicitly setting $g^x = g^{\hat{u}}$ for the tested session. Furthermore, we use xSS = $\hat{y}$ as the extracted semi-static secret in both the tested and potential partnered session(s) as well as in server sessions obtaining the same ephemeral g^x and previous `ServerConfiguration` for g^s.

- For any server session using the guessed semi-static key that obtains a client ephemeral Diffie–Hellman share $g^{x'} \neq g^x$, algorithm $\mathcal{B}_4$ does not compute SS explicitly but uses xSS ← $\mathsf{PRF}((g^{x'})^v, 0)$ directly as the response of a msPRF-ODH query $(g^{x'}, 0)$.

- For the client session tested or partnered with the tested server session in stage 1, algorithm $\mathcal{B}_4$ does not explicitly compute the ephemeral secret ES. Instead, it directly uses the response of a distinct msPRF-ODH query $(g^{y'}, 0)$ as xES ← $\mathsf{PRF}((g^{y'})^x, 0)$, where $g^{y'}$ is the server's key share obtained by the client within the `ServerKeyShare` message.

---

[25]Observe that, while the server might sign its share, this signature is only checked *after* the (always unauthenticated) handshake traffic key is computed and, hence, might contain an adversarially controlled server share $g^{y'}$.

In the special case that $g^{y'} = g^s$ (resulting in SS = ES for the client session), algorithm $\mathcal{B}_4$ however does not issue a query, but simply (re-)uses the challenge $\mathrm{xES} = \hat{y}$.

Finally, $\mathcal{B}_4$ outputs 1 if $\mathcal{A}$ wins in the Multi-Stage game and otherwise 0.

In case $\hat{y} = \mathsf{PRF}(g^{\hat{u}v}, 0)$ this approach equals Game B.3 while if $\hat{y}$ is a uniformly random value, it equals Game B.4. For this, let us see why $\mathcal{B}_4$ provides correct simulations for $\mathcal{A}$ in both cases.

First of all, recall that implicitly $g^s = g^v$ and $g^x = g^{\hat{u}}$ in the tested and partnered sessions, hence $\mathrm{xSS} \leftarrow \mathsf{PRF}(g^{\hat{u}v}, 0) = \mathsf{PRF}(g^{xs}, 0)$ is chosen as in the real Game B.3 in case $b = 0$ in the msPRF-ODH game (recall that HKDF.Extract is our PRF function). In case $b = 1$, the value $\mathrm{xSS} \xleftarrow{\$} \{0,1\}^\lambda$ is a randomly chosen element as specified for Game B.4.

Moreover, $\mathcal{B}_4$ is free to replace these two values at will, as $\mathcal{A}$ is only able to passively observe them and does not get to learn the discrete logarithms: For one, g^s needs to be unrevealed (i.e., $\mathsf{st_{ssk,sskid}} = \mathsf{fresh}$) in order for Test queries on 0-RTT keys derived from it to be allowed. This holds as a RevealSemiStaticKey query on a tested replayable stage would set the lost flag to true. At the same time, Corrupt queries do not reveal semi-static keys. Then, if the client side is tested, this side necessarily is honest and hence picks g^x, while, if the server side is tested, it must have an honest partnered client session, which means $\mathcal{A}$ cannot have modified the honestly picked ephemeral g^x on the way. Finally, both g^x and g^s are chosen independently of the other ephemeral and semi-static values (which $\mathcal{B}_4$, hence, can still select on its own), which in particular implies that $\mathcal{B}_4$ can detect a differing behavior of $\mathcal{A}$ in case, coincidentally, the same Diffie–Hellman shares are picked independently in another session.

Using the queries provided in the msPRF-ODH security game, $\mathcal{B}_4$ is moreover able to correctly compute keys from both further static secrets SS′ for server sessions using the same semi-static key g^s (without knowing s), as well as (once) the ephemeral secret ES′ in the involved client session (without knowing x). Hence, it can in particular correctly answer Reveal queries to any of these sessions.

Therefore, the advantage (or winning probability) difference of $\mathcal{A}$ between Game B.3 and Game B.4 is, through $\mathcal{B}_4$'s output, transformed into a difference of outputting 1 in the two cases of the msPRF-ODH game and, hence, we can bound the former difference as

$$\mathsf{Adv}^{G_{B.3}}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}} \leq \mathsf{Adv}^{G_{B.4}}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}} + \mathsf{Adv}^{\mathsf{msPRF\text{-}ODH}}_{\mathsf{HKDF.Extract},\mathbb{G},\mathcal{B}_4}.$$

Game B.5. In the last step, we replace the HKDF.Expand evaluations keyed with $\widetilde{\mathrm{xSS}}$, in particular in the tested and matching sessions, by a (lazy-sampled) random function. This in particular results in the early-data handshake and application traffic keys tk_{ehs} and tk_{ead}, the 0-RTT finished secret $\mathrm{FS_{0\text{-}RTT}}$, and the expanded static secret mSS in the tested session being replaced by random values $\widetilde{tk_{ehs}}, \widetilde{tk_{ead}}, \widetilde{\mathrm{FS}_{0\text{-}RTT}}$, $\widetilde{\mathrm{mSS}} \xleftarrow{\$} \{0,1\}^\lambda$.

We can turn any adversary $\mathcal{A}$ distinguishing this change (with non-negligible probability) into an adversary $\mathcal{B}_5$ against the PRF security of HKDF.Expand. Again, $\mathcal{B}_5$ acts as the Multi-Stage challenger for $\mathcal{A}$, this time using its PRF oracle for any HKDF.Expand evaluation under key $\widetilde{\mathrm{xSS}}$, while performing evaluations under different keys on its own. In case the PRF oracle computes the real function, this simulation equals Game B.4; if the oracle computes a random function, it equals Game B.5. Moreover, the simulation is sound as $\widetilde{\mathrm{xSS}}$ is an independent random value (due to the change in Game B.4) and hence chosen like the key in the PRF security game.

Thus, $\mathcal{B}_5$'s distinguishing advantage in the PRF game bounds the difference of $\mathcal{A}$ in the two games:

$$\mathsf{Adv}^{G_{B.4}}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}} \leq \mathsf{Adv}^{G_{B.5}}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_5}.$$

With the change in Game B.5, both potentially tested keys $\widetilde{tk_{ehs}}$ and $\widetilde{tk_{ead}}$ are now chosen independently at random. Hence, a Test query on them does not reveal any information about the test bit b_{test} to the

43

adversary. Moreover, even if $\mathcal{A}$ replays the first messages of the involved client session to further server sessions or injects the client's ephemeral share g^x in a differently crafted `ClientHello` message, this does not allow it to distinguish the real from the random key. In the former case, all sessions receiving the same client 0-RTT messages will be partnered with the test session (and hence those keys cannot be revealed). In the latter case, the keys tk_{ehs} and tk_{ead} in those (non-partnered) sessions will be derived from the same key $\widetilde{\text{xSS}}$ but with a distinct handshake hash (due to Game B.1) and, hence, are independent random values themselves. Therefore, b_{test} remains unknown to $\mathcal{A}$ and thus

$$\mathsf{Adv}^{G_{B.5}}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}} \leq 0. \qquad \square$$

Case C. Stage 3–6: Test Client without Stage-3 Partner

This case can be proven as for the `draft-10` full (EC)DHE handshake [DFGS16], keeping in mind that (as for the 0-RTT messages), the `ServerCertificateVerify` message uniquely determines the `ServerFinished` message and hence still attest agreement on the session identifiers for `draft-12`, now containing the finished messages. Hence, the security bound established by Dowling et al. [DFGS16] (for corresponding adversaries $\mathcal{B}_6$ and $\mathcal{B}_7$) still applies:

$$\mathsf{Adv}^{\text{1-Multi-Stage, test 3–6, client without partner}}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}} \leq \mathsf{Adv}^{\mathsf{COLL}}_{\mathsf{H},\mathcal{B}_6} + n_u \cdot \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig},\mathcal{B}_7}.$$

Case D. Stage 3–6: Test Server without Stage-3 Partner

As for the previous case, the bound established in [DFGS16] (for corresponding adversaries $\mathcal{B}_8$ and $\mathcal{B}_9$) applies:

$$\mathsf{Adv}^{\text{1-Multi-Stage, test 3–6, server without partner}}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}} \leq \mathsf{Adv}^{\mathsf{COLL}}_{\mathsf{H},\mathcal{B}_8} + n_u \cdot \mathsf{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{Sig},\mathcal{B}_9}.$$

Case E. Stage 3–6: Test with Stage-3 Partner

For the last case, the following aspects need to be considered when adapting the full handshake proof [DFGS16] to the (full handshake part of the) 0-RTT handshake.

First, the ephemeral and static secrets $\text{ES} = g^{xy}$ and $\text{SS} = g^{xs}$ are now derived differently (instead of having $\text{ES} = \text{SS}$ as in the full handshake). The PRF-ODH challenge encoded in g^x and g^y hence only allows us to replace xES by a uniformly random element $\widetilde{\text{xES}} \stackrel{\$}{\leftarrow} \{0,1\}^\lambda$, while the derivation SS remains unmodified (in particular, $\mathcal{A}$ is allowed to reveal the semi-static key g^s at any time via a `RevealSemiStaticKey`).

Second, when replacing with random the master secret derived as $\text{MS} \leftarrow \mathsf{HKDF.Extract}(\text{mSS}, \widetilde{\text{mES}})$, mSS is derived from $\text{SS} = g^{xs}$ and hence may be known to the adversary through a `RevealSemiStaticKey` query on g^s. We thus cannot rely on the PRF security of Extract as in the full handshake proof. Instead, following the analysis of the core cryptographic protocol OPTLS by Krawczyk and Wee [KW16], we model $\mathsf{HKDF.Extract}$ as a strong extractor with $\widetilde{\text{mES}}$ as entropy source and mSS as (public) seed. The step of replacing MS with an independent random value $\widetilde{\text{MS}} \stackrel{\$}{\leftarrow} \{0,1\}^\lambda$ can accordingly be bounded by the corresponding distinguishing advantage $\mathsf{Adv}^{\text{st-Extract}}_{\mathsf{HKDF.Extract},\cdot}$.

Considering these changes, the advantage bounds induced by the proof steps beyond that for the strong extractor remain identical to those established in [DFGS16]:

$$\mathsf{Adv}^{\text{1-Multi-Stage, test 3–6, test with partner}}_{\texttt{draft-12-(EC)DHE-0RTT},\mathcal{A}} \leq \mathsf{Adv}^{\mathsf{COLL}}_{\mathsf{H},\mathcal{B}_{10}} + n_s \cdot \Big(\mathsf{Adv}^{\mathsf{PRF\text{-}ODH}}_{\mathsf{HKDF.Extract},\mathbb{G},\mathcal{B}_{11}} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_{12}}$$
$$+ \mathsf{Adv}^{\text{st-Extract}}_{\mathsf{HKDF.Extract},\mathcal{B}_{13}} + \mathsf{Adv}^{\mathsf{PRF\text{-}sec}}_{\mathsf{HKDF.Expand},\mathcal{B}_{14}} \Big).$$

8 Comparing the QUIC and TLS 1.3 0-RTT Handshakes

We emphasize two aspects here in which the TLS 1.3 design is superior to QUIC and strengthens the achievable (multi-stage) security both in terms of key independence and compositionality: For one thing, it derives separate keys for the different purposes (in particular, tk_{ehs} and tk_{ead} as well as tk_{hs} and tk_{app} for the encryption of (0-RTT resp. regular) handshake messages and data), enabling a cleaner key separation. For another thing, it establishes authenticity of the server's Diffie–Hellman share g^y through an explicit MAC (PSK-(EC)DHE 0-RTT) resp. signature ((EC)DHE 0-RTT) instead of through an authenticated encryption (under the 0-RTT key) in the data channel, rendering the security of one session key not relying on the secrecy of another.

Conversely, QUIC in its original version `Rev 20130620` achieves replay protection for the derived 0-RTT key on the key exchange level whereas TLS 1.3 does not (and hence, technically, TLS 1.3 satisfies only a weaker notion of security in that respect). As discussed in the beginning, this protection however may become void in the overall setting of secure channels when clients actively replay rejected 0-RTT data over the main channel.

Finally, the (abandoned) Diffie–Hellman-based and the (remaining) PSK-based 0-RTT handshakes in TLS 1.3 (as specified for `draft-12` resp. `draft-14`) differ in the forward-secrecy guarantees they provide for 0-RTT keys, as already pointed out by Krawczyk on the TLS mailing list [Kra16]. While in `draft-12` (EC)DHE 0-RTT those keys are forward secret (wrt. long-term (signing) key compromise) and succumb only to exposures of the semi-static key involved, no forward secrecy is provided in the PSK and PSK-(EC)DHE 0-RTT mode of `draft-14`. It is important to note, though, that preshared resumption secrets used in the PSK-based 0-RTT modes (treated as long-term secrets in our model) are usually much shorter-lived than public-key long-term signing keys, mitigating the effects of a compromise. Still, preshared keys have to be stored safely by both the server and the client—a challenging task in practice, especially on the client's side. Diffie–Hellman-based 0-RTT hence poses weaker requirements in that respect as the client here only has to store the public part of a semi-static key.

9 Composition

Key exchange protocols would be of limited use if applied in isolation; in general the derived keys are meant to be deployed in a follow-up (or overall) protocol. The most common application is of course the encryption (and authentication) of data sent between the two involved parties within a (cryptographic) channel protocol, with the TLS record protocol being a prime example. The TLS 1.3 handshakes (with or without 0-RTT) additionally derive further keys for different purposes, namely the resumption master secret RMS (in non-PSK modes) enabling follow-up abbreviated (preshared-key) handshakes and the exporter master secret EMS which can be used to derive additional key material. For both, the key usage in the cryptographic channel as well as the usage for other purposes, it is desirable to modularize the analysis, treating key exchange and the composed protocol(s) independently and devising automatically the security of the combined execution.

The approach inspires studying the generic compositional guarantees the TLS 1.3 handshake or, in general, a (multi-stage) key exchange protocol can provide. For classical (non–multi-stage) key exchange protocols in the Bellare–Rogaway model [BR94] this has been argued formally by Brzuska et al. [BFWW11], and lifted to the multi-stage setting by Fischlin and Günther [FG14]. Intuitively, their composition theorem attests that keys derived in a secure (multi-stage) key exchange protocol KE (satisfying certain additional conditions) can be securely used within *any* symmetric-key protocol Π. This result in particular subsumes usage in an arbitrary channel protocol. But, in case of TLS 1.3, it can also be used to argue security of using the resumption master secret established in a full handshake as pre-shared key for a later abbreviated

handshake, as done by Dowling et al. [DFGS15a, DFGS16]. Here, security of the composed protocol $\mathsf{KE}_i; \Pi$ is intuitively defined as the symmetric-key protocol Π being secure when using the stage-i keys established in the key exchange protocol KE (see [FG14, DFGS15a, DFGS16] for a formal definition).

We augment the multi-stage composition result by Fischlin and Günther [FG14] and extended to— in particular—the preshared-secret setting and multiple concurrent authentication modes by Dowling et al. [DFGS15a, DFGS16], in order to also capture replayability of keys—which are not generically composable—based on our extended multi-stage key exchange model. The distinction between internal and external (usage of) keys furthermore is eminently useful for defining composition, since it elegantly replaces the necessary informal restriction to final keys in prior theorems.

As its original proof [FG14, DFGS15a] applies with marginal changes, we only state the composition theorem and briefly discuss the necessary modifications to the proof.

Theorem 9.1 (Multi-stage composition). *Let* KE *be a* Multi-Stage-*secure key exchange protocol (in the public-key or preshared-secret setting) providing key independence and stage-j forward secrecy with properties* $(\mathsf{M}, \mathsf{AUTH}, \mathsf{USE}, \mathsf{REPLAY})$ *and key distribution* $\mathcal{D}$, *and that allows for efficient multi-stage session matching*[26]. *Let* Π *be a symmetric-key protocol that is secure w.r.t. some game* G_Π *and has a key generation algorithm that outputs keys with distribution* $\mathcal{D}$. *Then the composition* $\mathsf{KE}_i; \Pi$ *for any external and non-replayable stage $i \geq j$ (i.e.,* $\mathsf{REPLAY}_i = \mathsf{nonreplayable}$ *and* $\mathsf{USE}_i = \mathsf{external}$*) is secure w.r.t. the composed security game* $G_{\mathsf{KE}_i; \Pi}$. *Formally, for any efficient adversary* $\mathcal{A}$ *against* $G_{\mathsf{KE}_i; \Pi}$ *there exist efficient algorithms* $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ *such that*

$$\mathsf{Adv}^{G_{\mathsf{KE}_i; \Pi}}_{\mathsf{KE}_i; \Pi, \mathcal{A}} \leq \mathsf{Adv}^{\mathsf{Match}}_{\mathsf{KE}, \mathcal{B}_1} + n_s \cdot \mathsf{Adv}^{\mathsf{Multi\text{-}Stage}, \mathcal{D}}_{\mathsf{KE}, \mathcal{B}_2} + \mathsf{Adv}^{G_\Pi}_{\Pi, \mathcal{B}_3},$$

where n_s is the maximum number of sessions in the key exchange game.

Compared to the previous result [DFGS16], our theorem statement incorporates two changes which are reflected in the proof as follows.

First, we guarantee composition for any external key instead of only for final keys. Since for external keys the tested random key in the Test query of our model (for $b_{\mathsf{test}} = 0$) does not replace the actual key in subsequent protocol steps, after the hybrid proof step all stage-i session keys are random and independent of the key exchange (subgame), hence allowing for an independent treatment of the protocol subgame. (This requirement was satisfied in [FG14, DFGS15a, DFGS16] through demanding that the composed key is final and therefore by definition not used in the key exchange.) If a key is instead internal (i.e., is used already within the key exchange), we cannot hope for any *generic* compositional security for that key. To give an example, say that the key is used in the key exchange to produce a MAC value, then composing the key with the according MAC protocol would—in the general sense—be insecure as an adversary can simply present the MAC value seen in the key exchange as a valid "fresh" forgery in the MAC-security game.

Second, we require that a key must be non-replayable (which was inherently demanded for all keys in the previous multi-stage models). Most notably, security of symmetric-key protocols in general expectedly breaks down when the same secret key is used in two separate instances of a protocol (due to a replay). Moreover, in the public-key setting, non-replayability is a necessary condition for a composed key (similar to requiring the key to be forward secret, i.e., resilient to Corrupt queries) in order to not become compromised through a RevealSemiStaticKey query issued after the key was derived and, hence, potentially already used within the symmetric-key protocol. Recall that, in our model, we treat compromises of long-term and semi-static keys independently through the Corrupt resp. RevealSemiStaticKey query.

---

[26]Multi-stage session matching is a technical notion which essentially requires that it is publicly decidable from the transcript whether two sessions are partnered or not when given all keys derived so far (see [DFGS15b] for a formal definition).

Acknowledgments

References

[ABR01]   Michel Abdalla, Mihir Bellare, and Phillip Rogaway. The oracle Diffie-Hellman assumptions and an analysis of DHIES. In David Naccache, editor, *Topics in Cryptology – CT-RSA 2001*, volume 2020 of *Lecture Notes in Computer Science*, pages 143–158, San Francisco, CA, USA, April 8–12, 2001. Springer, Heidelberg, Germany. (Cited on page 9.)

[BCK96]   Mihir Bellare, Ran Canetti, and Hugo Krawczyk. Keying hash functions for message authentication. In Neal Koblitz, editor, *Advances in Cryptology – CRYPTO'96*, volume 1109 of *Lecture Notes in Computer Science*, pages 1–15, Santa Barbara, CA, USA, August 18–22, 1996. Springer, Heidelberg, Germany. (Cited on pages 8, 9, and 22.)

[BFWW11]   Christina Brzuska, Marc Fischlin, Bogdan Warinschi, and Stephen C. Williams. Composability of Bellare-Rogaway key exchange protocols. In Yan Chen, George Danezis, and Vitaly Shmatikov, editors, *ACM CCS 11: 18th Conference on Computer and Communications Security*, pages 51–62, Chicago, Illinois, USA, October 17–21, 2011. ACM Press. (Cited on pages 7, 11, 18, and 45.)

[BR94]   Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In Douglas R. Stinson, editor, *Advances in Cryptology – CRYPTO'93*, volume 773 of *Lecture Notes in Computer Science*, pages 232–249, Santa Barbara, CA, USA, August 22–26, 1994. Springer, Heidelberg, Germany. (Cited on pages 10, 17, and 45.)

[Brz13]   Christina Brzuska. *On the Foundations of Key Exchange*. PhD thesis, Technische Universität Darmstadt, Darmstadt, Germany, 2013. http://tuprints.ulb.tu-darmstadt.de/3414/. (Cited on page 18.)

[BWM99]   Simon Blake-Wilson and Alfred Menezes. Authenticated Diffie-Hellman key agreement protocols (invited talk). In Stafford E. Tavares and Henk Meijer, editors, *SAC 1998: 5th Annual International Workshop on Selected Areas in Cryptography*, volume 1556 of *Lecture Notes in Computer Science*, pages 339–361, Kingston, Ontario, Canada, August 17–18, 1999. Springer, Heidelberg, Germany. (Cited on pages 4 and 7.)

[CHSvdM16]   Cas Cremers, Marko Horvat, Sam Scott, and Thyla van der Merwe. Automated verification of TLS 1.3: 0-RTT, resumption and delayed authentication. In *2016 IEEE Symposium on Security and Privacy*, 2016. (Cited on page 8.)

[CK01]   Ran Canetti and Hugo Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In Birgit Pfitzmann, editor, *Advances in Cryptology – EUROCRYPT 2001*, volume 2045 of *Lecture Notes in Computer Science*, pages 453–474, Innsbruck, Austria, May 6–10, 2001. Springer, Heidelberg, Germany. (Cited on pages 8 and 17.)

[CK02]   Ran Canetti and Hugo Krawczyk. Security analysis of IKE's signature-based key-exchange protocol. In Moti Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of

Lecture Notes in Computer Science, pages 143–161, Santa Barbara, CA, USA, August 18–22, 2002. Springer, Heidelberg, Germany. `http://eprint.iacr.org/2002/120/`. (Cited on page 11.)

[CKS09]   David Cash, Eike Kiltz, and Victor Shoup. The twin Diffie-Hellman problem and applications. *Journal of Cryptology*, 22(4):470–504, October 2009. (Cited on page 8.)

[DFGS15a] Benjamin Dowling, Marc Fischlin, Felix Günther, and Douglas Stebila. A cryptographic analysis of the TLS 1.3 handshake protocol candidates. In Indrajit Ray, Ninghui Li, and Christopher Kruegel:, editors, *ACM CCS 15: 22nd Conference on Computer and Communications Security*, pages 1197–1210, Denver, CO, USA, October 12–16, 2015. ACM Press. (Cited on pages 6, 7, 10, 11, 12, 13, 14, 17, 18, 24, 39, and 46.)

[DFGS15b] Benjamin Dowling, Marc Fischlin, Felix Günther, and Douglas Stebila. A cryptographic analysis of the TLS 1.3 handshake protocol candidates. Cryptology ePrint Archive, Report 2015/914, 2015. `http://eprint.iacr.org/2015/914`. (Cited on pages 7, 26, 39, and 46.)

[DFGS16]   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. `http://eprint.iacr.org/2016/081`. (Cited on pages 6, 7, 9, 10, 11, 12, 18, 24, 39, 44, and 46.)

[FG14]   Marc Fischlin and Felix Günther. Multi-stage key exchange and the case of Google's QUIC protocol. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *ACM CCS 14: 21st Conference on Computer and Communications Security*, pages 1193–1204, Scottsdale, AZ, USA, November 3–7, 2014. ACM Press. (Cited on pages 4, 6, 7, 8, 10, 11, 12, 13, 14, 18, 45, and 46.)

[FHKP13]   Eduarda S. V. Freire, Dennis Hofheinz, Eike Kiltz, and Kenneth G. Paterson. Non-interactive key exchange. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013: 16th International Conference on Theory and Practice of Public Key Cryptography*, volume 7778 of *Lecture Notes in Computer Science*, pages 254–271, Nara, Japan, February 26 – March 1, 2013. Springer, Heidelberg, Germany. (Cited on page 8.)

[HJLS15]   Britta Hale, Tibor Jager, Sebastian Lauer, and Jörg Schwenk. Speeding: On low-latency key exchange. Cryptology ePrint Archive, Report 2015/1214, 2015. `http://eprint.iacr.org/2015/1214`. (Cited on pages 3 and 8.)

[HK11]   Shai Halevi and Hugo Krawczyk. One-pass HMQV and asymmetric key-wrapping. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *PKC 2011: 14th International Conference on Theory and Practice of Public Key Cryptography*, volume 6571 of *Lecture Notes in Computer Science*, pages 317–334, Taormina, Italy, March 6–9, 2011. Springer, Heidelberg, Germany. (Cited on page 8.)

[JKSS12]   Tibor Jager, Florian Kohlar, Sven Schäge, and Jörg Schwenk. On the security of TLS-DHE in the standard model. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology – CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 273–293, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Heidelberg, Germany. (Cited on pages 7, 9, and 32.)

[KBC97]   H. Krawczyk, M. Bellare, and R. Canetti. HMAC: Keyed-Hashing for Message Authentication. RFC 2104 (Informational), February 1997. Updated by RFC 6151. (Cited on page 8.)

[KPW13] Hugo Krawczyk, Kenneth G. Paterson, and Hoeteck Wee. On the security of the TLS protocol: A systematic analysis. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 429–448, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany. (Cited on pages 7 and 9.)

[Kra05] Hugo Krawczyk. HMQV: A high-performance secure Diffie-Hellman protocol. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 546–566, Santa Barbara, CA, USA, August 14–18, 2005. Springer, Heidelberg, Germany. (Cited on pages 7 and 24.)

[Kra10] Hugo Krawczyk. Cryptographic extraction and key derivation: The HKDF scheme. In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 631–648, Santa Barbara, CA, USA, August 15–19, 2010. Springer, Heidelberg, Germany. (Cited on pages 8, 22, 26, and 33.)

[Kra16] Hugo Krawczyk. [TLS] Call for consensus: Removing DHE-based 0-RTT. https://mailarchive.ietf.org/arch/msg/tls/xmnvrKEQkEbD-u8HTeQkyitmclY, March 2016. posting in above thread. (Cited on pages 6, 33, and 45.)

[KW15] Hugo Krawczyk and Hoeteck Wee. The OPTLS protocol and TLS 1.3. Cryptology ePrint Archive, Report 2015/978, 2015. http://eprint.iacr.org/2015/978. (Cited on pages 8 and 18.)

[KW16] Hugo Krawczyk and Hoeteck Wee. The OPTLS protocol and TLS 1.3. In *2016 IEEE European Symposium on Security and Privacy*, pages 81–96. IEEE, March 2016. (Cited on pages 3, 8, 18, and 44.)

[LC13] Adam Langley and Wan-Teh Chang. QUIC Crypto, June 2013. Revision 20130620. (Cited on pages 4 and 6.)

[LC15] Adam Langley and Wan-Teh Chang. QUIC Crypto. https://docs.google.com/document/d/1g5nIXAIkN_Y-7XJW5K45IblHd_L2f5LTaDUDwvZ5L6g/, July 2015. Revision 20150720. (Cited on pages 5 and 11.)

[LJBN15] Robert Lychev, Samuel Jero, Alexandra Boldyreva, and Cristina Nita-Rotaru. How secure and quick is QUIC? Provable security and performance analyses. In *2015 IEEE Symposium on Security and Privacy*, pages 214–231, San Jose, CA, USA, May 17–21, 2015. IEEE Computer Society Press. (Cited on page 4.)

[LLM07] Brian A. LaMacchia, Kristin Lauter, and Anton Mityagin. Stronger security of authenticated key exchange. In Willy Susilo, Joseph K. Liu, and Yi Mu, editors, *ProvSec 2007: 1st International Conference on Provable Security*, volume 4784 of *Lecture Notes in Computer Science*, pages 1–16, Wollongong, Australia, November 1–2, 2007. Springer, Heidelberg, Germany. (Cited on page 17.)

[PvdM16] Kenneth G. Paterson and Thyla van der Merwe. Reactive and proactive standardisation of TLS. In Lidong Chen, David McGrew, and Chris Mitchell, editors, *SSR 2016*, volume 10074 of *Lecture Notes in Computer Science*, pages 160–186. Springer, December 2016. (Cited on page 8.)

[PZS+13]   W. Michael Petullo, Xu Zhang, Jon A. Solworth, Daniel J. Bernstein, and Tanja Lange. MinimaLT: minimal-latency networking through better security. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 13: 20th Conference on Computer and Communications Security*, pages 425–438, Berlin, Germany, November 4–8, 2013. ACM Press. (Cited on page 3.)

[QUI]   QUIC, a multiplexed stream transport over UDP. https://www.chromium.org/quic. (Cited on page 3.)

[Res15a]   E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3 – draft-ietf-tls-tls13-10. https://tools.ietf.org/html/draft-ietf-tls-tls13-10, October 2015. (Cited on page 7.)

[Res15b]   Eric Rescorla. 0-RTT and Anti-Replay (IETF TLS working group mailing list). https://www.ietf.org/mail-archive/web/tls/current/msg15594.html, March 2015. (Cited on pages 4 and 5.)

[Res16a]   Eric Rescorla. Should it be possible to do 0-RTT with the server signing (pull request #443). https://github.com/tlswg/tls13-spec/issues/443, April 2016. (Cited on page 23.)

[Res16b]   Eric Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3 – draft-ietf-tls-tls13-12. https://tools.ietf.org/html/draft-ietf-tls-tls13-12, March 2016. (Cited on pages 3, 33, and 36.)

[Res16c]   Eric Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3 – draft-ietf-tls-tls13-13. https://tools.ietf.org/html/draft-ietf-tls-tls13-13, May 2016. (Cited on page 4.)

[Res16d]   Eric Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3 – draft-ietf-tls-tls13-14. https://tools.ietf.org/html/draft-ietf-tls-tls13-14, 2016. (Cited on pages 20 and 22.)

[Res16e]   Eric Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3 – draft-ietf-tls-tls13-18. https://tools.ietf.org/html/draft-ietf-tls-tls13-18, October 2016. (Cited on pages 3 and 5.)

[Res16f]   Eric Rescorla. TLS 1.3 — draft-ietf-tls-tls13-12 (presentation at ietf 95 meeting). https://www.ietf.org/proceedings/95/slides/slides-95-tls-2.pdf, April 2016. (Cited on pages 6 and 33.)

[WTSB16]   David J. Wu, Ankur Taly, Asim Shankar, and Dan Boneh. Privacy, discovery, and authentication for the internet of things. In Ioannis G. Askoxylakis, Sotiris Ioannidis, Sokratis K. Katsikas, and Catherine A. Meadows, editors, *ESORICS 2016: 21st European Symposium on Research in Computer Security, Part II*, volume 9879 of *Lecture Notes in Computer Science*, pages 301–319, Heraklion, Greece, September 26–30, 2016. Springer, Heidelberg, Germany. (Cited on page 3.)