

Simple Security Definitions for and Constructions of 0-RTT Key Exchange

Britta Hale¹(✉), Tibor Jäger², Sebastian Lauer³, and Jörg Schwenk³

¹ NTNU, Norwegian University of Science and Technology, Trondheim, Norway
britta.hale@ntnu.no

² Paderborn University, Paderborn, Germany
tiber.jager@upb.de

³ Horst Görtz Institute, Ruhr-University Bochum, Bochum, Germany
{sebastian.lauer,joerg.schwenk}@rub.de

Abstract. Zero Round-Trip Time (0-RTT) key exchange protocols allow for the transmission of cryptographically protected payload data without requiring the prior exchange of messages of a cryptographic key exchange protocol. The 0-RTT KE concept was first realized by Google in the QUIC Crypto protocol, and a 0-RTT mode has been intensively discussed for inclusion in TLS 1.3.

In 0-RTT KE two keys are generated, typically using a Diffie-Hellman key exchange. The first key is a combination of an ephemeral client share and a long-lived server share. The second key is computed using an ephemeral server share and the same ephemeral client share.

In this paper, we propose simple security models, which catch the intuition behind known 0-RTT KE protocols; namely that the first (resp. second) key should remain indistinguishable from a random value, even if the second (resp. first) key is revealed. We call this property *strong key independence*. We also give the first constructions of 0-RTT KE which are provably secure in these models, based on the generic assumption that secure non-interactive key exchange (NIKE) exists (This work was partially supported by a STSM Grant from COST Action IC1306).

Keywords: Foundations · Low-latency key exchange · 0-RTT protocols · Authenticated key exchange · Non-interactive key exchange · QUIC · TLS 1.3.

1 Introduction

Efficiency, in terms of messages to be exchanged before a key is established, is a growing consideration for internet protocols today. Basically, the first generation of internet key exchange protocols did not care too much about efficiency, since secure connections were considered to be the exception rather than the rule: SSL (versions 2.0 and 3.0) and TLS (versions 1.0, 1.1, and 1.2) require 2 round-trip times (RTT) for key establishment before the first cryptographically-protected payload data can be sent. With the increased use of encryption,¹ efficiency is

¹ For example, initiatives like Let's Encrypt (<https://letsencrypt.org/>).

of escalating importance for protocols like TLS. Similarly, the older IPsec IKE version v1 needs between 3 RTT (aggressive mode + quick mode) and 4.5 RTT (main mode + quick mode). This was soon realized to be problematic, and in IKEv2 the number of RTTs was reduced to 2.

The QUIC Protocol. Fundamentally, the discussion on low-latency key exchange (aka. LLKE, zero-RTT or 0-RTT key exchange) was opened when Google proposed the QUIC protocol.² QUIC (cf. Fig. 1) achieves low-latency by caching a signed server configuration file on the client side, which contains a medium-lived Diffie-Hellman (DH) share $Y_0 = g^{y_0}$.³

When a client wishes to establish a connection with a server and possesses a valid configuration file of that server, it chooses a fresh ephemeral DH share $X = g^x$ and computes a temporal key k_1 from $g^{y_0 x}$. Using this key k_1 , the client can encrypt and authenticate data to be sent to the server, together with X . In response, the server sends a fresh DH share $Y = g^y$ and computes a session key k_2 from g^{xy} , which is used for all subsequent data exchanges.

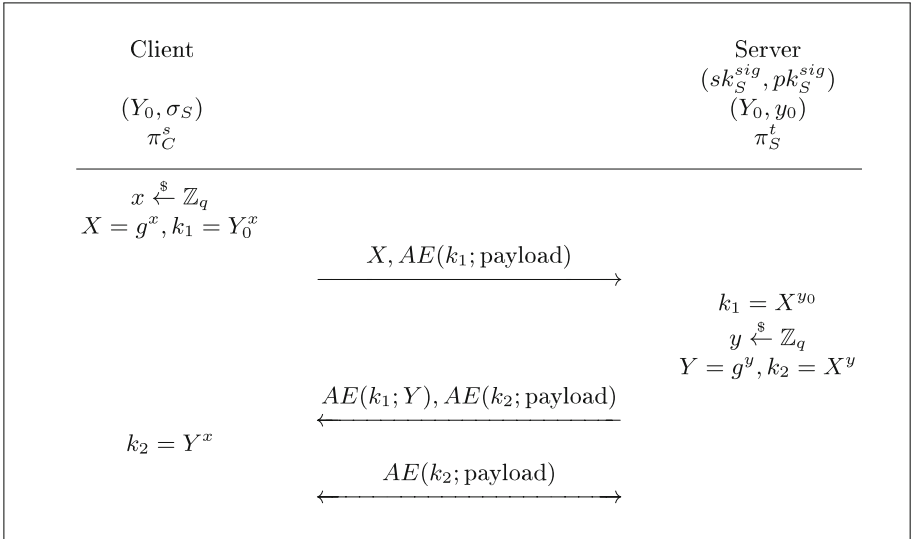


Fig. 1. Google’s QUIC protocol (simplified) with cached server key configuration file (Y_0, σ_S) . AE denotes a symmetric authenticated encryption algorithm (e.g., AES-GCM), (sk_S^{sig}, pk_S^{sig}) denotes the server’s long-term signing keys, and π_S^t (resp. π_C^s) denotes the oracle at server S executing the single t -th instance of the protocol (resp. for client).

² See <https://www.chromium.org/quic>.

³ If the client does not have a valid file, it has to be requested from the server, which increases the number of RTTs by 1, but may then be re-used for future sessions.

TLS 1.3. Early TLS 1.3 drafts, e.g. `draft-ietf-tls-tls13-08` [25], contained a 0-RTT key exchange mode where a QUIC-like `ServerConfiguration` message is cached by the client. The current version `draft-ietf-tls-tls13-18` [26] follows a different approach, where the initial key establishment between a client and a server is never 0-RTT. Instead, it defines a method to establish a new session based on the secret key of a previous session. Even though this is also called “0-RTT” in the current TLS 1.3 specification, it is rather a “0-RTT session resumption” protocol, but does not allow for 0-RTT key establishment. Most importantly, the major difference between the approach of the current TLS 1.3 draft in comparison to a “real” 0-RTT key exchange protocol is that the former requires storing of *secret* key information on the client between sessions. In contrast, a 0-RTT key establishment protocol does not require secret information to be stored between sessions.

Facebook’s Zero Protocol. Very recently, the social network Facebook announced that it is currently experimenting with a 0-RTT KE protocol called *Zero*.⁴ Zero is very similar to QUIC, except that it uses another nonce and encryption of the `ServerHello` message. It is noteworthy that the main difference between Zero and QUIC was introduced in order to prevent an attack discovered by Facebook, which has been reported to Google and meanwhile been fixed in QUIC, too. We believe that this is a good example that shows the demand of simple security definitions and provably-secure constructions for such protocols.

Security Goals. 0-RTT KE protocols like QUIC have ad-hoc designs that aim at achieving three goals: (1) 0-RTT encryption, where ciphertext data can already be sent together with the first handshake message; (2) perfect forward secrecy (PFS), where all ciphertexts exchanged after the second handshake message will remain secure even after the (static or semi-static) private keys of the server have been leaked, and (3) key independence, where “knowledge” about one of the two symmetric keys generated should not endanger the security of the other key.

Strong Key Independence. Intuitively, a 0-RTT KE protocol should achieve *strong key independence* between k_1 and k_2 ; if any one of the two keys is leaked at any time, the other key should still be indistinguishable from a random value. In all known security models, this intuition would be formalized as follows: if the adversary \mathcal{A} asks a `Reveal` query for k_1 , he is still allowed to ask a `Test` query for k_2 , and vice versa. If the two keys are computationally independent from each other (which also includes computations on the different protocol messages), then the adversary should have only a negligible advantage in answering the `Test` query correctly.

Ultimately this leads to the following research questions: *Do existing examples of 0-RTT KE protocols have strong key independence? Can we describe a generic way to construct 0-RTT KE protocols that provably achieve strong key independence?*

⁴ See <https://code.facebook.com/posts/608854979307125/building-zero-protocol-for-fast-secure-mobile-connections/>.

QUIC Does Not Provide Strong Key Independence. If an attacker \mathcal{A} is allowed to learn k_1 by a *Reveal*-query, then he is able to decrypt $AE(k_1; Y)$ and re-encrypt its own value $Y^* := g^{y^*}$ as $AE(k_1; Y^*)$. Furthermore, he can then compute the same $k_2 = X^{y^*}$ as the client oracle, and can thus distinguish between the “real” key and a “random” key chosen by the *Test* query. See [11] for more details on key dependency in QUIC.

Note that this theoretical attack does not imply that QUIC is insecure. It only shows that the authenticity of the server’s Diffie-Hellman share, which is sent in QUIC to establish k_2 , depends strongly on the security of key k_1 . Therefore QUIC does not provide strong key independence in the sense sketched above.

Previous Work on 0-RTT Key Exchange. The concept of 0-RTT key exchange was not developed in academia, but in industry – motivated by concrete practical demands of distributed applications. All previous works on 0-RTT KE [11, 23] conducted *a-posteriori* security analyses of the QUIC protocol, with tailored models. There are no foundational constructions as yet, and the relation to other cryptographic protocols and primitives is not yet well-understood.

At ACM CCS 2014, Fischlin and Günther [11] provided a formal definition of *multi-stage* key exchange protocols and used it to analyze the security of QUIC. Lychev *et al.* [23] gave an alternate analysis of QUIC, which considers both efficiency and security. They describe a security model which is bespoke to QUIC, adopting the complex, monolithic security model of [17] to the protocol’s requirements. Zhao [31] considers identity-concealed 0-RTT protocols, where user privacy is protected by hiding identities of users in a setting with mutual cryptographic authentication of both communicating parties. Günther *et al.* [14] extended the “puncturable encryption”-approach of Green and Miers [13] to show that even 0-RTT KE with full forward secrecy is possible, by evolving the secret key after each decryption. However, their construction is currently mainly of conceptual interest, as it is not yet efficient enough to be deployed at large scale in practice.

Security Model. In this paper, we use a variant of the Canetti-Krawczyk [7] security model. This family of security models is especially suited to protocols with only two message exchanges, with *one-round* key exchange protocols constituting the most important subclass. Popular examples of such protocols are MQV [22], HMQV [18], SMQV [27], KEA [21, 24], and NAXOS [20]. A comparison of different variants of the Canetti-Krawczyk model can be found in [9, 29].

The Importance of Simplicity of Security Models. Security models for key exchange protocols have to consider *active* adversaries that may modify, replay, inject, drop, etc., any message transmitted between communicating parties. They also need to capture *parallel* executions of multiple protocol sessions, potential reveals of earlier session keys, and adaptive corruptions of long-term secrets of parties. This makes even standard security models for key exchange *extremely* complex (in comparison to most other standard cryptographic primitives, like digital signatures or public-key encryption, for example).

Naturally, the novel primitive of 0-RTT KE requires formal security definitions. There are different ways to create such a model. One approach is to focus on *generality* of the model. Fischlin and Günther [11] followed this path, by defining *multi-stage* key exchange protocols, a generalization of 0-RTT KE. This approach has the advantage that it lays the foundation for the study of a very general class of interesting and novel primitives. However, its drawback is that this generality inherently also brings a huge *complexity* to the model. Clearly, the more complex the security model, the more difficult it becomes to devise new, simple, efficient, and provably-secure constructions. Moreover, proofs in complex models tend to be error-prone and less intuitive, because central technical ideas may be concealed in formal details that are required to handle the generality of the model.

Another approach is to devise a model which is tailored to the analysis of *one specific protocol*. For example, the complex, monolithic ACCE security model was developed in [17] to provide an *a posteriori* security analysis of TLS.⁵ A similar approach was followed by Lychev *et al.* [23], who adopted this model for an *a posteriori* analysis of QUIC, by defining the so-called *Q-ACCE* model. The notable drawback of this approach is that such tailor-made models tend to capture only the properties achieved by existing protocols, but not necessarily all properties that we would expect from a “good” 0-RTT KE protocol. In general, such tailor-made models do not, therefore, form a useful foundation for the creation of new protocols.

In this paper, we follow a different approach. We propose novel “bare-bone” security models for 0-RTT KE, which aim at capturing *all* (strong key independence and forward secrecy), but also *only* the properties intuitively expected from “good” 0-RTT KE protocols. We propose two different models. One considers the practically-relevant case of *server-only* authentication (where the client may or may not authenticate later over the established communication channel, similar in spirit to the server-only-authenticated ACCE model of [19]). The other considers traditional *mutual* cryptographic authentication of a client and server.

The reduced generality of our definitions – in comparison to the very general multi-stage security model of [11] – is intended. A model which captures *only*, but also *all* the properties expected from a “good” 0-RTT KE protocol allows us to devise relatively simple, foundational, and generic constructions of 0-RTT KE protocols with as-clean-as-possible security analyses.

Importance of Foundational Generic Constructions. Following [3], we use non-interactive key exchange (NIKE) [8, 12] in combination with digital signatures as a main building block.⁶ This yields the first examples of 0-RTT KE protocols with strong key independence, as well as the first constructions of 0-RTT KE from generic complexity assumptions. There are many advantages of such generic constructions:

⁵ A more modular approach was later proposed in [4].

⁶ Recall that digital signatures are implied by one-way functions, which in turn are implied by NIKE. Thus, essentially we only assume the existence of NIKE as a building block.

1. Generic constructions provide a better understanding of the structure of protocols. Since the primitives we use have abstract security properties, we can see precisely which abstract security requirements are needed to implement 0-RTT KE protocols.
2. They clarify the relations and implications between different types of cryptographic primitives.
3. They can be generically instantiated with building blocks based on different complexity assumptions. For example, if “post-quantum” security is needed, one can directly obtain a concrete protocol by using only post-quantum secure building blocks in the generic construction.

Usually generic constructions tend to involve more computational overhead than ad-hoc constructions. However, we note that our 0-RTT KE protocols can be instantiated relatively efficiently, given the efficient NIKE schemes of [12], for example.

Contributions. Contributions in this paper can be summarized as follows:

- *Simple security models.* We provide simple security models, which capture all properties that we expect from a “good” 0-RTT KE protocol, but only these properties. We consider both the “practical” setting with server-only authentication and the classical setting with mutual authentication.
- *First generic constructions.* We give intuitive, relatively simple, and efficient constructions of 0-RTT KE protocols in both settings.
- *First Non-DH instantiation.* Both QUIC and TLS 1.3 are based on DH key exchange. Our generic construction yields the first 0-RTT KE protocol which is not based on Diffie-Hellman (e.g., by instantiating the generic construction with the factoring-based NIKE scheme of Freire *et al.* [12]).
- *First 0-RTT KE with strong key independence.* Our 0-RTT KE protocols are the first to achieve strong key independence in the sense described above.
- *Well-established, general assumptions.* The construction is based on general assumptions, namely the existence of secure NIKE and digital signature schemes. For all building blocks we require only standard security properties.
- *Security in the Standard Model.* The security analysis is completely in the standard model, i.e. it is performed without resorting to the Random Oracle heuristic [1] and without relying on non-standard complexity assumptions.
- *Efficient instantiability.* Despite the fact that our constructions are generic, the resulting protocols can be instantiated relatively efficiently.

Full Version of this Paper. Due to space limitations, we have to defer several results to the full version of this paper [15]. This includes the full proof of Theorem 1, the Definition and Security Model for a 0-RTT protocol under mutual authentication (0-RTT-M), a construction of a 0-RTT-M protocol along with its security model and its security proof.

2 Preliminaries

For our construction in Sect. 4, we need signature schemes and non-interactive key exchange (NIKE) protocols. Here we summarize the definitions of these two primitives and their security from the literature.

2.1 Digital Signatures

A digital signature scheme consists of three polynomial-time algorithm $\text{SIG} = (\text{SIG.Gen}, \text{SIG.Sign}, \text{SIG.Vfy})$. The key generation algorithm $(sk, pk) \stackrel{\$}{\leftarrow} \text{SIG.Gen}(1^\lambda)$ generates a public verification key pk and a secret signing key sk on input of security parameter λ . Signing algorithm $\sigma \stackrel{\$}{\leftarrow} \text{SIG.Sign}(sk, m)$ generates a signature for message m . Verification algorithm $\text{SIG.Vfy}(pk, \sigma, m)$ returns 1 if σ is a valid signature for m under key pk , and 0 otherwise.

Consider the following security experiment played between a challenger \mathcal{C} and an adversary \mathcal{A} .

1. The challenger generates a public/secret key pair $(sk, pk) \stackrel{\$}{\leftarrow} \text{SIG.Gen}(1^\lambda)$, the adversary receives pk as input.
2. The adversary may query arbitrary messages m_i to the challenger. The challenger replies to each query with a signature $\sigma_i = \text{SIG.Sign}(sk, m_i)$. Here i is an index, ranging between $1 \leq i \leq q$ for some $q \in \mathbb{N}$. Queries can be made adaptively.
3. Eventually, the adversary outputs a message/signature pair (m, σ) .

Definition 1. We define the advantage on an adversary \mathcal{A} in this game as

$$\text{Adv}_{\text{SIG}, \mathcal{A}}^{\text{sEUF-CMA}}(\lambda) := \Pr \left[(m, \sigma) \stackrel{\$}{\leftarrow} \mathcal{A}^{\mathcal{C}(\lambda)}(pk) : \begin{array}{l} \text{SIG.Vfy}(pk, \sigma, m) = 1, \\ (m, \sigma) \neq (m_i, \sigma_i) \forall i \end{array} \right].$$

SIG is strongly secure against existential forgeries under adaptive chosen-message attacks (*sEUF-CMA*), if $\text{Adv}_{\text{SIG}, \mathcal{A}}^{\text{sEUF-CMA}}(\lambda)$ is a negligible function in λ for all probabilistic polynomial-time adversaries \mathcal{A} .

Remark 1. Signatures with sEUF-CMA security can be constructed generically from any EUF-CMA-secure signature scheme and chameleon hash functions [6, 28].

2.2 Secure Non-interactive Key Exchange

Definition 2. A non-interactive key exchange (NIKE) scheme consists of two deterministic algorithms $(\text{NIKE.Gen}, \text{NIKE.Key})$.

$\text{NIKE.Gen}(1^\lambda, r)$ takes a security parameter λ and randomness $r \in \{0, 1\}^\lambda$. It outputs a key pair (pk, sk) . We write $(pk, sk) \stackrel{\$}{\leftarrow} \text{NIKE.Gen}(1^\lambda)$ to denote that $\text{NIKE.Gen}(1^\lambda, r)$ is executed with uniformly random $r \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$.

$\text{NIKE.Key}(sk_i, pk_j)$ is a deterministic algorithm which takes as input a secret key sk_i and a public key pk_j , and outputs a key $k_{i,j}$.

We say that a NIKE scheme is correct, if for all $(pk_i, sk_i) \xleftarrow{\$} \text{NIKE.Gen}(1^\lambda)$ and $(pk_j, sk_j) \xleftarrow{\$} \text{NIKE.Gen}(1^\lambda)$ holds that $\text{NIKE.Key}(sk_i, pk_j) = \text{NIKE.Key}(sk_j, pk_i)$.

A NIKE scheme is used by d parties P_1, \dots, P_d as follows. Each party P_i generates a key pair $(pk_i, sk_i) \leftarrow \text{NIKE.Gen}(1^\lambda)$ and publishes pk_i . In order to compute the key shared by P_i and P_j , party P_i computes $k_{i,j} = \text{NIKE.Key}(sk_i, pk_j)$. Similarly, party P_j computes $k_{j,i} = \text{NIKE.Key}(sk_j, pk_i)$. Correctness of the NIKE scheme guarantees that $k_{i,j} = k_{j,i}$.

CKS-Light Security. The *CKS-light* security model for NIKE protocols is relatively simplistic and compact. We choose this model because other (more complex) NIKE security models like *CKS*, *CKS-heavy*, and *m-CKS-heavy* are polynomial-time equivalent to *CKS-light*. See [12] for more details.

Security of a NIKE protocol NIKE is defined by a game NIKE played between an adversary \mathcal{A} and a challenger. The challenger takes a security parameter λ and a random bit b as input and answers all queries of \mathcal{A} until she outputs a bit b' . The challenger answers the following queries for \mathcal{A} :

- **RegisterHonest**(i). \mathcal{A} supplies an index i . The challenger runs $\text{NIKE.Gen}(1^\lambda)$ to generate a key pair (pk_i, sk_i) and records the tuple $(\text{honest}, pk_i, sk_i)$ for later and returns pk_i to \mathcal{A} . This query may be asked *at most twice* by \mathcal{A} .
- **RegisterCorrupt**(pk_i). With this query \mathcal{A} supplies a public key pk_i . The challenger records the tuple $(\text{Corrupt}, pk_i)$ for later.
- **GetCorruptKey**(i, j). \mathcal{A} supplies two indices i and j where pk_i was registered as corrupt and pk_j as honest. The challenger runs $k \leftarrow \text{NIKE.Key}(sk_j, pk_i)$ and returns k to \mathcal{A} .
- **Test**(i, j). The adversary supplies two indices i and j that were registered honestly. Now the challenger uses bit b : if $b = 0$, then the challenger runs $k_{i,j} \leftarrow \text{NIKE.Key}(pk_i, sk_j)$ and returns the key $k_{i,j}$. If $b = 1$, then the challenger samples a random element from the key space, records it for later, and returns the key to \mathcal{A} .

The game NIKE outputs 1, denoted by $\text{NIKE}_{\text{NIKE}}^{\mathcal{A}}(\lambda) = 1$, if $b = b'$ and 0 otherwise. We say \mathcal{A} wins the game if $\text{NIKE}_{\text{NIKE}}^{\mathcal{A}}(\lambda) = 1$.

Definition 3. For any adversary \mathcal{A} playing the above NIKE game against a NIKE scheme NIKE , we define the advantage of winning the game NIKE as

$$\text{Adv}_{\text{NIKE}, \mathcal{A}}^{\text{CKS-light}}(\lambda) = \left| 2 \cdot \Pr \left[\text{NIKE}_{\text{NIKE}}^{\mathcal{A}}(\lambda) = 1 \right] - 1 \right|.$$

Let λ be a security parameter, NIKE be a NIKE protocol and \mathcal{A} an adversary. We say NIKE is a *CKS-light-secure NIKE protocol*, if for all probabilistic polynomial-time adversaries \mathcal{A} , the function $\text{Adv}_{\text{NIKE}, \mathcal{A}}^{\text{CKS-light}}(\lambda)$ is a negligible function in λ .

3 0-RTT Key Exchange Protocols: Syntax and Security with Server-Only Authentication

In the model presented in this section, we give formal definitions for 0-RTT KE with strong key independence and main-key forward secrecy. We start with the case of server-only authentication, as it is the more important case in practice (in particular, server-only authentication will be the main operating mode of both QUIC and TLS 1.3).

3.1 Syntax and Correctness

Definition 4. *A 0-RTT key exchange scheme with server-only authentication consists of deterministic algorithms $(\text{Gen}^{\text{server}}, \text{KE}_{\text{init}}^{\text{client}}, \text{KE}_{\text{refresh}}^{\text{client}}, \text{KE}_{\text{refresh}}^{\text{server}})$.*

- $\text{Gen}^{\text{server}}(1^\lambda, r) \rightarrow (pk, sk)$: A key generation algorithm that takes as input a security parameter λ and randomness $r \in \{0, 1\}^\lambda$ and outputs a key pair (pk, sk) . We write $(pk, sk) \stackrel{\$}{\leftarrow} \text{Gen}^{\text{server}}(1^\lambda)$ to denote that a pair (pk, sk) is the output of $\text{Gen}^{\text{server}}$ when executed with uniformly random $r \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$.
- $\text{KE}_{\text{init}}^{\text{client}}(pk_j, r_i) \rightarrow (k_{\text{tmp}}^{i,j}, m_i)$: An algorithm that takes as input a public key pk_j and randomness $r_i \in \{0, 1\}^\lambda$, and outputs a temporary key $k_{\text{tmp}}^{i,j}$ and a message m_i .
- $\text{KE}_{\text{refresh}}^{\text{server}}(sk_j, r_j, m_i) \rightarrow (k_{\text{main}}^{j,i}, k_{\text{tmp}}^{j,i}, m_j)$: An algorithm that takes as input a secret key sk_j , randomness r_j and a message m_i , and outputs a key $k_{\text{main}}^{j,i}$, a temporary key $k_{\text{tmp}}^{j,i}$ and a message m_j .
- $\text{KE}_{\text{refresh}}^{\text{client}}(pk_j, r_i, m_j) \rightarrow k_{\text{main}}^{i,j}$: An algorithm that takes as input a public key pk_j , randomness r_i , and message m_j , and outputs a key $k_{\text{main}}^{i,j}$.

We say that a 0-RTT key exchange scheme is correct, if for all $(pk_j, sk_j) \stackrel{\$}{\leftarrow} \text{Gen}^{\text{server}}(1^\lambda)$ and for all $r_i, r_j \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$ holds that

$$\Pr[k_{\text{tmp}}^{i,j} \neq k_{\text{tmp}}^{j,i} \text{ or } k_{\text{main}}^{i,j} \neq k_{\text{main}}^{j,i}] \leq \text{negl}(\lambda),$$

where $(k_{\text{tmp}}^{j,i}, m_i) \leftarrow \text{KE}_{\text{init}}^{\text{client}}(pk_j, r_i)$, $(k_{\text{main}}^{i,j}, k_{\text{tmp}}^{i,j}, m_j) \leftarrow \text{KE}_{\text{refresh}}^{\text{server}}(sk_j, r_j, m_i)$, and $k_{\text{main}}^{j,i} \leftarrow \text{KE}_{\text{refresh}}^{\text{client}}(pk_j, r_i, m_j)$.

A 0-RTT KE scheme is used by a set parties which are either clients C or servers S (cf. Fig. 2). Each server S_p has a generated key pair $(sk_p, pk_p) \stackrel{\$}{\leftarrow} \text{Gen}^{\text{server}}(1^\lambda, j)$ with published pk_p . The protocol is executed as follows:

1. The client oracle C_i chooses $r_i \in \{0, 1\}^\lambda$ and selects the public key of the intended partner S_j (which must be a server, otherwise this value is undefined). Then it computes $(k_{\text{tmp}}^{i,j}, m_i) \leftarrow \text{KE}_{\text{init}}^{\text{client}}(pk_j, r_i)$, and sends m_i to S_j . Additionally, C_i can use $k_{\text{tmp}}^{i,j}$ to encrypt some data M_i .

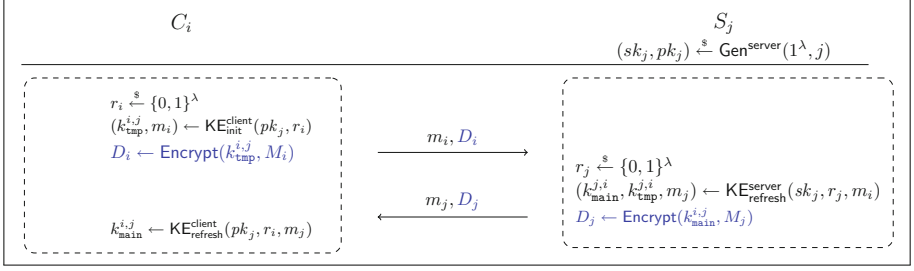


Fig. 2. Execution of a 0-RTT KE Protocol with Server-Only Authentication in Parallel to **Encrypted Application Data**. Note that the messages D_i and D_j correspond to the symmetric encryption protocol used to encrypt payload data, and are therefore *not* part of the 0-RTT KE protocol, but a separate protocol. These messages are only displayed here only to illustrate the basic, parallel application message flow to that of a 0-RTT KE protocol. While it would in principle be possible to define the symmetric encryption directly as part of the protocol, this would require a significantly more complex “ACCE-style” [17] security model, which we avoid for sake of simplicity.

2. Upon reception of message m_i , S_j initializes a new oracle $S_{j,t}$. This oracle chooses $r_j \in \{0, 1\}^\lambda$ and computes $(k_{\text{main}}^{j,i}, k_{\text{tmp}}^{j,i}, m_j) \leftarrow \text{KE}_{\text{refresh}}^{\text{server}}(sk_j, r_j, m_i)$. The server may use the ephemeral key $k_{\text{tmp}}^{j,i}$ to decrypt D_i . Then, the server sends m_j and optionally some data M_j encrypted with the key $k_{\text{main}}^{j,i}$ to the client.
3. C_i computes $k_{\text{main}}^{i,j} \leftarrow \text{KE}_{\text{refresh}}^{\text{client}}(pk_j, r_i, m_j)$ and can optionally decrypt D_j . Correctness of the 0-RTT KE scheme guarantees that $k_{\text{main}}^{i,j} = k_{\text{main}}^{j,i}$.

3.2 Execution Environment

We provide an adversary \mathcal{A} against a 0-RTT KE protocol with the following *execution environment*. Clients, which are not in possession of a long-term secret are represented by oracles C_1, \dots, C_d (without any particular “identity”). We consider ℓ servers, each server has a long-term key pair (sk_j, pk_j) ⁷, $j \in \{1, \dots, \ell\}$, and each client has access to all public keys pk_1, \dots, pk_ℓ . Each server is represented by a collection of k oracles $S_{j,1}, \dots, S_{j,k}$, where each oracle represents a process that executes one single instance of the protocol.

We use the following variables to maintain the internal state of oracles.

Clients. Each client oracle C_i , $i \in [d]$, maintains

- two variables k_i^{tmp} and k_i^{main} to store the temporal and main keys of a session,

⁷ We do not distinguish between static (i.e. long-lived) and semi-static (i.e. medium lived) key pairs. Thus the long-lived keys in this model correspond to the server configuration file keys of QUIC and TLS 1.3.

- a variable Partner_i , which contains the identity of the intended communication partner, and
- variables $\mathcal{M}_i^{\text{in}}$ and $\mathcal{M}_i^{\text{out}}$ containing messages sent and received by the oracle.

The internal state of a client oracle is initialized to $(k_i^{\text{tmp}}, k_i^{\text{main}}, \text{Partner}_i, \mathcal{M}_i^{\text{in}}, \mathcal{M}_i^{\text{out}}) := (\emptyset, \emptyset, \emptyset, \emptyset, \emptyset)$.

Servers. Each server oracle $S_{j,t}$, $(j, t) \in [\ell] \times [k]$, maintains:

- two variables k_i^{tmp} and k_i^{main} to store the temporal and main keys of a session, and
- variables $\mathcal{M}_{j,t}^{\text{in}}$ and $\mathcal{M}_{j,t}^{\text{out}}$ containing messages sent and received by the server.

The internal state of a server oracle is initialized to $(k_{j,t}^{\text{tmp}}, k_{j,t}^{\text{main}}, \mathcal{M}_{j,t}^{\text{in}}, \mathcal{M}_{j,t}^{\text{out}}) := (\emptyset, \emptyset, \emptyset, \emptyset)$.

We say that an oracle has *accepted the temporal key* if $k^{\text{tmp}} \neq \emptyset$, and *accepted the main key* if $k^{\text{main}} \neq \emptyset$.

In the security experiment, the adversary is able to interact with the oracles by issuing the following queries.

Send($C_i/S_{j,t}, m$). The adversary sends a message m to the requested oracle.

The oracle processes m according to the protocol specification. Any response generated by the oracle according to the protocol specification is returned to the adversary.

If a client oracle C_i receives m as the first message, then the oracle checks if m consists of a special initialization message ($m = (\text{init}, j)$). If true, then the oracle responds with the first protocol message generated for intended partner S_j , else it outputs \perp .

Reveal($C_i/S_{j,t}, \text{tmp/main}$). This query returns the key of the given stage if it already has been computed, or \perp otherwise.

Corrupt(j). On input of a server identity j , this query returns the long-term private key of the server. If **Corrupt**(j) is the τ -th query issued by \mathcal{A} , we say a party is τ -*corrupted*. For parties that are not corrupted we define $\tau := \infty$.

Test($C_i/S_{j,t}, \text{tmp/main}$). This query is used to test a key and is only asked once.

It is answered as follows: If the variable of the requested key is not empty, a random $b \xleftarrow{\$} \{0, 1\}$ is selected, and

- if $b = 0$ then the requested key is returned, else
- if $b = 1$ then a random key, according to the probability distribution of keys generated by the protocol, is returned.

Otherwise \perp is returned.

Security Model Security Game $\mathcal{G}_{\mathcal{A}}^{\text{ORTT-sa}}$. After receiving a security parameter λ the challenger \mathcal{C} simulates the protocol and keeps track of all variables of the execution environment: he generates the long-lived key pairs of all server parties and answers faithfully to all queries by the adversary.

The adversary receives all public keys pk_1, \dots, pk_ℓ and can interact with the challenger by issuing any combination of the queries **Send**(), **Corrupt**(), and

Reveal(). At some point the adversary queries **Test()** to an oracle and receives a key, which is either the requested key or a random value. The adversary may continue asking **Send()**, **Corrupt()**, and **Reveal()**-queries after receiving the key and finally outputs some bit b' .

Definition 5 (0-RTT KE-Security with Server-Only Authentication).

Let an adversary \mathcal{A} interact with the challenger in game $\mathcal{G}_A^{0RTT-sa}$ as it is described above. We say the challenger outputs 1, denoted by $\mathcal{G}_A^{0RTT-sa}(\lambda) = 1$, if $b = b'$ and the following conditions hold:

- if \mathcal{A} issues **Test**(C_i, tmp) all of the following hold:
 - **Reveal**(C_i, tmp) was never queried by \mathcal{A}
 - **Reveal**($S_{j,t}, \text{tmp}$) was never queried by \mathcal{A} for any oracle $S_{j,t}$ such that $\text{Partner}_i = j$ and $\mathcal{M}_{j,t}^{\text{in}} = \mathcal{M}_i^{\text{out}}$
 - the communication partner $\text{Partner}_i = j$, if it exists, is not τ -corrupted with $\tau < \infty$
- if \mathcal{A} issues **Test**(C_i, main) all of the following hold:
 - **Reveal**(C_i, main) was never queried by \mathcal{A}
 - **Reveal**($S_{j,t}, \text{main}$) was never queried by \mathcal{A} , where $\text{Partner}_i = j$, $\mathcal{M}_{j,t}^{\text{in}} = \mathcal{M}_i^{\text{out}}$, and $\mathcal{M}_i^{\text{in}} = \mathcal{M}_{j,t}^{\text{out}}$
 - the communication $\text{Partner}_i = j$ is not τ -corrupted with $\tau < \tau_0$, where **Test**(C_i, main) is the τ_0 -th query issued by \mathcal{A}
- if \mathcal{A} issues **Test**($S_{j,t}, \text{tmp}$) all of the following hold:
 - **Reveal**($S_{j,t}, \text{tmp}$) was never queried by \mathcal{A}
 - there exists an oracle C_i with $\mathcal{M}_i^{\text{out}} = \mathcal{M}_{j,t}^{\text{in}}$
 - **Reveal**(C_i, tmp) was never queried by \mathcal{A} to any oracle C_i with $\mathcal{M}_i^{\text{out}} = \mathcal{M}_{j,t}^{\text{in}}$
 - **Reveal**($S_{j,t'}, \text{tmp}$) was never queried by \mathcal{A} for any oracle $S_{j,t'}$ with $\mathcal{M}_{j,t'}^{\text{in}} = \mathcal{M}_{j,t}^{\text{in}}$
 - j is not τ -corrupted with $\tau < \infty$
- if \mathcal{A} issues **Test**($S_{j,t}, \text{main}$) all of the following hold:
 - **Reveal**($S_{j,t}, \text{main}$) was never queried by \mathcal{A}
 - there exists an oracle C_i with $\mathcal{M}_i^{\text{out}} = \mathcal{M}_{j,t}^{\text{in}}$
 - **Reveal**(C_i, main) was never queried by \mathcal{A} , if $\mathcal{M}_i^{\text{in}} = \mathcal{M}_{j,t}^{\text{out}}$

else the game outputs a random bit. We define the advantage of \mathcal{A} in the game $\mathcal{G}_A^{0RTT-sa}(\lambda)$ by

$$\text{Adv}_A^{0RTT-sa}(\lambda) := |2 \cdot \Pr[\mathcal{G}_A^{0RTT-sa}(\lambda) = 1] - 1|.$$

Definition 6. We say that a 0-RTT key exchange protocol is **test-secure**, if there exists a negligible function $\text{negl}(\lambda)$ such that for all PPT adversaries \mathcal{A} interacting according to the security game $\mathcal{G}_A^{0RTT-sa}(\lambda)$ it holds that

$$\text{Adv}_A^{0RTT-sa}(\lambda) \leq \text{negl}(\lambda).$$

Remark 2. Our security model captures forward secrecy for the `main-key`, because key indistinguishability is required to hold even if the adversary is able to corrupt the communication partner of the `test-oracle` (but only after the `test-oracle` has accepted, of course, in order to avoid trivial attacks).

Moreover, strong key independence is modeled by the fact that an adversary which attempts to distinguish a `tmp-key` from random (i.e., an adversary which asks $\text{Test}(X, \text{tmp})$ for $X \in \{C_i, S_{j,t}\}$ for some $i, j, t\}$) is allowed to learn the `main-key` of X . Similarly, an adversary which tries to distinguish a `main-key` from random by asking $\text{Test}(X, \text{main})$ is allowed to learn the `tmp-key` of X as well. Security in this sense guarantees that the `tmp-key` and the `main-key` look independent to a computationally-bounded adversary.

Remark 3. Note that the requirements of $\mathcal{M}_i^{\text{out}} = \mathcal{M}_{j,t}^{\text{in}}$ etc. in the above security model essentially adopt the notion of *matching conversations*, defined by Bellare and Rogaway [2] for general, multi-message key exchange protocols, to the special case of 0-RTT KE.

3.3 Composing a 0-RTT KE Protocol with Symmetric Encryption

The security model described above considers only the 0-RTT KE protocol, *without* symmetric encryption of payload data (that is, without the messages D_i and D_j displayed in Fig. 2). A protocol secure in this sense guarantees the indistinguishability of keys in a hypothetical setting, where the key is not used for symmetric encryption of payload messages potentially known to the adversary. One may think that this is not sufficient for 0-RTT KE, because the key *will* be used to encrypt payload data, and this will enable an adversary to trivially distinguish a “real” key from a “random” key (this holds for both the “temporal” key $k_{\text{tmp}}^{i,j}$ and the actual “main” session key $k_{\text{main}}^{i,j}$). Note that this argument applies not only to the above 0-RTT KE security model, but actually to any security model for (authenticated) key exchange which is based on the indistinguishability of keys, such as the classical model of Bellare and Rogaway and many similar models [2, 5, 7, 10, 20, 27]. In practice, this key will usually be used in a cryptographic protocol, e.g. to encrypt messages, and therefore trivially allow for distinguishing “real” from “random” keys. The security of the composition of a protocol secure in the sense of [2, 5, 7, 10, 20, 27] with a symmetric encryption protocol follows from a standard two-step hybrid argument, which essentially proceeds as follows:

1. In the original security experiment, the adversary interacts with a composed protocol, where the KE protocol is first used to derive a key k , which is then used to encrypt payload data with the symmetric encryption protocol.
2. In the next hybrid experiment, the adversary interacts with a composed protocol, where the symmetric encryption does not use the key k computed by the KE protocol, but an independent random key. Note that an adversary that distinguishes this hybrid from the original game can be used to distinguish a “real” key of the KE protocol from a “random” one.

Now the adversary interacts with an encryption protocol that uses a key which is *independent* of the KE protocol. This allows for a reduction of the security of the composed protocol to the security of the symmetric protocol.

A similarly straightforward hybrid argument applies to the composition of 0-RTT KE with symmetric encryption, which works as follows:

1. In the original security experiment, the adversary interacts with a composed protocol, where the 0-RTT KE protocol is first used to derive a key $k_{\text{tmp}}^{i,j}$, which is then used to encrypt the payload data sent along with the first protocol message. Then the 0-RTT KE protocol is used to derive the main key $k_{\text{main}}^{i,j}$, which in turn is used to encrypt all further payload data.
2. In the first hybrid experiment, the adversary interacts with a composed protocol, where only $k_{\text{tmp}}^{i,j}$ is replaced with an independent random value. An adversary that distinguishes this hybrid from the original game can be used to distinguish a “real” $k_{\text{tmp}}^{i,j}$ from a “random” one.
Now the adversary interacts with an encryption protocol that encrypts the first payload message with a key which is *independent* of the 0-RTT KE protocol. This allows for a reduction of the security of the first payload message to the security of the symmetric protocol.
3. In the second hybrid experiment, the adversary interacts with a composed protocol, where $k_{\text{main}}^{i,j}$ is now also replaced with an independent random value. An adversary that distinguishes this hybrid from the previous one can be used to distinguish a “real” $k_{\text{main}}^{i,j}$ from a “random” one. This allows for a reduction of the security of all further payload messages to the security of the symmetric protocol.

Following the long tradition of previous works on indistinguishability-based key exchange security models [2, 5, 7, 10, 20, 27], we can thus consider an indistinguishability-based security model for 0-RTT KE even though in practice key exchange messages will be interleaved with messages of the symmetric encryption protocol. This allows for simple security models, and enables a modular analysis of the building blocks of a composed protocol.

4 Generic Construction of 0-RTT KE from NIKE

Now we are ready to describe our generic NIKE-based 0-RTT KE protocol and its security analysis.

4.1 Generic Construction

Let $\text{NIKE} = (\text{NIKE.Gen}, \text{NIKE.Key})$ be a NIKE scheme according to Definition 2 and let $\text{SIG} = (\text{SIG.Gen}, \text{SIG.Sign}, \text{SIG.Vfy})$ be a signature scheme. Then we construct a 0-RTT KE scheme $\text{0-RTT} = (\text{Gen}^{\text{server}}, \text{KE}_{\text{init}}^{\text{client}}, \text{KE}_{\text{refresh}}^{\text{client}}, \text{KE}_{\text{refresh}}^{\text{server}})$, per Definition 4, in the following way (cf. Fig. 3).




Fig. 3. 0-RTT KE from NIKE. Again, it is possible to include the parallel execution of a symmetric encryption protocol which would behave as in Fig. 2 for encrypted application data. As such a protocol is *not* part of the 0-RTT KE protocol, we omit it here for simplicity.

- $\text{Gen}^{\text{server}}(1^\lambda, r)$ computes key pairs using the NIKE key generation algorithm $(pk^{\text{nike-static}}, sk^{\text{nike-static}}) \xleftarrow{\$} \text{NIKE.Gen}(1^\lambda)$ and signature keys using the SIG algorithm $(pk^{\text{sg}}, sk^{\text{sg}}) \xleftarrow{\$} \text{SIG.Gen}$, and outputs

$$(pk, sk) := ((pk^{\text{nike-static}}, pk^{\text{sg}}), (sk^{\text{nike-static}}, sk^{\text{sg}})).$$

- $\text{KE}_{\text{init}}^{\text{client}}(pk_j, r_i)$ samples $r_i \xleftarrow{\$} \{0, 1\}^\lambda$, parses $pk_j = (pk_j^{\text{nike-static}}, pk_j^{\text{sg}})$, runs $(pk_i^{\text{nike}}, sk_i^{\text{nike}}) \leftarrow \text{NIKE.Gen}(1^\lambda, r_i)$ and $k_{i,j}^{\text{nike}} \leftarrow \text{NIKE.Key}(sk_i^{\text{nike}}, pk_j^{\text{nike-static}})$, and outputs

$$(k_{\text{tmp}}^{i,j}, m_i) := (k_{i,j}^{\text{nike}}, pk_i^{\text{nike}}).$$

- $\text{KE}_{\text{refresh}}^{\text{server}}(sk_j, r_j, m_i)$ takes in $m_i = pk_i^{\text{nike}}$, parses $sk_j = (sk_j^{\text{nike-static}}, sk_j^{\text{sg}})$, and samples $r_j \xleftarrow{\$} \{0, 1\}^\lambda$. It then computes $k_{i,j}^{\text{nike}} \leftarrow \text{NIKE.Key}(sk_j^{\text{nike-static}}, pk_i^{\text{nike}})$, $(pk_j^{\text{nike}}, sk_j^{\text{nike}}) \leftarrow \text{NIKE.Gen}(1^\lambda, r_j)$, and $\sigma_j \leftarrow \text{SIG.Sign}(sk_j^{\text{sg}}, pk_j^{\text{nike}})$. If $m_i = pk_j^{\text{nike-static}}$ then it samples $k_{\text{main}}^{\text{nike}}$ uniformly random, else it computes $k_{\text{main}}^{\text{nike}} \leftarrow \text{NIKE.Key}(sk_j^{\text{nike}}, pk_i^{\text{nike}})$, outputting

$$(k_{\text{main}}^{j,i}, k_{\text{tmp}}^{j,i}, m_j) := (k_{\text{main}}^{\text{nike}}, k_{i,j}^{\text{nike}}, (pk_j^{\text{nike}}, \sigma_j)).$$

- $\text{KE}_{\text{refresh}}^{\text{client}}(pk_j, r_i, m_j)$ parses $pk_j = (pk_j^{\text{nike-static}}, pk_j^{\text{sg}})$ and $m_j = (pk_j^{\text{nike}}, \sigma_j)$. It then checks $\text{true} \leftarrow \text{SIG.Vfy}(pk_j^{\text{sg}}, \sigma_j, pk_j^{\text{nike}})$ and computes $k_{\text{main}}^{\text{nike}} \leftarrow \text{NIKE.Key}(sk_i^{\text{nike}}, pk_j^{\text{nike}})$, outputting $k_{\text{main}}^{i,j} := k_{\text{main}}^{\text{nike}}$.

Ultimately, the construction follows by applying the NIKE NIKE.Gen algorithm and the signature SIG.Gen algorithm to generate a server configuration

file which is comprised of the server public key and a server public signature key which a client can then employ for generating the first protocol flow. In order for the 0-RTT KE construction to abstract the security guarantees of the underlying NIKE, the appropriate client $(pk_i^{\text{nike}}, sk_i^{\text{nike}})$ must be available for use in the NIKE.Key algorithm. Consequently, the $(pk_i^{\text{nike}}, sk_i^{\text{nike}})$ values are generated locally by the client, with pk_i^{nike} passed to the server as a message. Note that this construction naturally forgoes client-side authentication. Figure 3 demonstrates the construction.

Remark 4. One may wonder why we define $\text{KE}_{\text{refresh}}^{\text{server}}(sk_j, r_j, m_i)$ such that it samples a random key when it takes as input a client message m_i which is equal to its own static NIKE key, that is, if $m_i = pk_j^{\text{nike-static}}$. We note that this is necessary for the security the constructed 0-RTT KE scheme to be reducible to that of the NIKE scheme, because in some cases we will not be able to simulate the key computed by a server oracle that receives as input a message which is equal to the “static” NIKE public key contained in its 0-RTT KE public key. Note that this incurs a negligible correctness error. However, it is straightforward to verify the correctness of the protocol according to Definition 4.

Theorem 1. *Let 0-RTT be executed with d clients, ℓ servers with long-term keys, and k server oracles modeling each server. From each attacker \mathcal{A} , we can construct attackers \mathcal{B}_{sig} , according to Definition 1, and $\mathcal{B}_{\text{nike}}$, according to Definition 3, such that*

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^{\text{0RTT-sa}}(\lambda) &\leq 2kd\ell \cdot \left(\text{Adv}_{\text{NIKE}, \mathcal{B}_{\text{nike}}}^{\text{CKS-light}}(\lambda) + \text{Adv}_{\text{SIG}, \mathcal{B}_{\text{sig}}}^{\text{sEUF-CMA}}(\lambda) \right) \\ &\quad + d\ell \cdot \left(k \cdot \text{Adv}_{\text{NIKE}, \mathcal{B}_{\text{nike}}}^{\text{CKS-light}}(\lambda) + \text{Adv}_{\text{SIG}, \mathcal{B}_{\text{sig}}}^{\text{sEUF-CMA}}(\lambda) \right) \\ &\quad + d\ell \cdot \left(\text{Adv}_{\text{NIKE}, \mathcal{B}_{\text{nike}}}^{\text{CKS-light}}(\lambda) + \text{Adv}_{\text{SIG}, \mathcal{B}_{\text{sig}}}^{\text{sEUF-CMA}}(\lambda) \right) + 4 \cdot \text{Adv}_{\text{NIKE}, \mathcal{B}_{\text{nike}}}^{\text{CKS-light}}(\lambda). \end{aligned}$$

The running time of \mathcal{B}_{sig} and $\mathcal{B}_{\text{nike}}$ is approximately equal to the time required to execute the security experiment with \mathcal{A} once.

Intuition for the Proof of Theorem 1. In order to prove Theorem 1, we distinguish between four types of attackers:

- adversary \mathcal{A}_1 asks $\text{Test}()$ to a client oracle and the temporary key (*CT-attacker*)
- adversary \mathcal{A}_2 asks $\text{Test}()$ to a client oracle and the main key (*CM-attacker*)
- adversary \mathcal{A}_3 asks $\text{Test}()$ to a server oracle and the temporary key (*ST-attacker*)
- adversary \mathcal{A}_4 asks $\text{Test}()$ to a server oracle and the main key (*SM-attacker*).

Let us give some intuition why this classification of attackers will be useful for the security proof. In the 0-RTT KE scheme 0-RTT each party computes 2 different keys k_{tmp}^i and k_{main}^i , where k_{tmp}^i depends on the ephemeral keys of the

client and the static keys of the server, and k'_{main} depends on the ephemeral keys of both parties. In our proof we want to be able to reduce the indistinguishability of the 0-RTT-key to the indistinguishability of the NIKE-key.

In the NIKE security experiment the attacker receives two challenge public keys $\{pk_i^{\text{nike}}, pk_j^{\text{nike}}\}$. In the reduction, we want to embed these keys in the 0-RTT security experiment, according to Sect. 3.2, and still be able to answer all **Reveal()**- and **Corrupt()**-queries correctly. In the case of adversaries that test the temporary key of the client or the server we can embed the NIKE-keys as $pk_j^{\text{nike-static}} = pk_j^{\text{nike}}$ and $m_i = pk_i^{\text{nike}}$. However, this does not work for adversaries against the main key, because k'_{main} depends on the ephemeral keys of the parties. In this case we have to embed the keys as $m_i = pk_i^{\text{nike}}$ and $m_j = pk_j^{\text{nike}}$. The **Test()**-query of the attacker in the 0-RTT experiment can then be answered with the challenge the attacker in the NIKE experiment receives.

References

1. Bellare, M., Rogaway, P.: Random oracles are practical: a paradigm for designing efficient protocols. In: Ashby, V. (ed.) ACM CCS 1993, 3–5 November 1993, Fairfax, Virginia, USA, pp. 62–73. ACM Press (1993)
2. Bellare, M., Rogaway, P.: Entity authentication and key distribution. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 232–249. Springer, Heidelberg (1994). doi:[10.1007/3-540-48329-2_21](https://doi.org/10.1007/3-540-48329-2_21)
3. Bergsma, F., Jager, T., Schwenk, J.: One-round key exchange with strong security: an efficient and generic construction in the standard model. In: Katz, J. (ed.) PKC 2015. LNCS, vol. 9020, pp. 477–494. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-46447-2_21](https://doi.org/10.1007/978-3-662-46447-2_21)
4. Bhargavan, K., Fournet, C., Kohlweiss, M., Pironti, A., Strub, P.-Y., Zanella-Béguélin, S.: Proving the TLS handshake secure (as it is). In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014. LNCS, vol. 8617, pp. 235–255. Springer, Heidelberg (2014). doi:[10.1007/978-3-662-44381-1_14](https://doi.org/10.1007/978-3-662-44381-1_14)
5. Blake-Wilson, S., Johnson, D., Menezes, A.: Key agreement protocols and their security analysis. In: Darnell, M. (ed.) Cryptography and Coding 1997. LNCS, vol. 1355, pp. 30–45. Springer, Heidelberg (1997). doi:[10.1007/BFb0024447](https://doi.org/10.1007/BFb0024447)
6. Boneh, D., Shen, E., Waters, B.: Strongly unforgeable signatures based on computational Diffie-Hellman. In: Yung et al. [30], pp. 229–240
7. Canetti, R., Krawczyk, H.: Analysis of key-exchange protocols and their use for building secure channels. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 453–474. Springer, Heidelberg (2001). doi:[10.1007/3-540-44987-6_28](https://doi.org/10.1007/3-540-44987-6_28)
8. Cash, D., Kiltz, E., Shoup, V.: The twin Diffie-Hellman problem and applications. In: Smart, N. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 127–145. Springer, Heidelberg (2008). doi:[10.1007/978-3-540-78967-3_8](https://doi.org/10.1007/978-3-540-78967-3_8)
9. Cremers, C.: Examining indistinguishability-based security models for key exchange protocols: the case of CK, CK-HMQV, and eCK. In: Cheung, B.S.N., Hui, L.C.K., Sandhu, R.S., Wong, D.S. (eds.) ASIACCS 2011, 22–24 March 2011, Hong Kong, China, pp. 80–91. ACM Press (2011)
10. Cremers, C.J.F.: Session-state reveal is stronger than ephemeral key reveal: attacking the NAXOS authenticated key exchange protocol. In: Abdalla, M., Pointcheval, D., Fouque, P.-A., Vergnaud, D. (eds.) ACNS 2009. LNCS, vol. 5536, pp. 20–33. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-01957-9_2](https://doi.org/10.1007/978-3-642-01957-9_2)

11. Fischlin, M., Günther, F.: Multi-stage key exchange and the case of Google's QUIC protocol. In: Ahn, G.-J., Yung, M., Li, N. (eds.) ACM CCS 2014, 3–7 November 2014, Scottsdale, AZ, USA, pp. 1193–1204. ACM Press (2014)
12. Freire, E.S.V., Hofheinz, D., Kiltz, E., Paterson, K.G.: Non-interactive key exchange. In: Kurosawa, K., Hanaoka, G. (eds.) PKC 2013. LNCS, vol. 7778, pp. 254–271. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-36362-7_17](https://doi.org/10.1007/978-3-642-36362-7_17)
13. Green, M.D., Miers, I.: Forward secure asynchronous messaging from puncturable encryption. In: IEEE S&P 2015 [16], pp. 305–320
14. Günther, F., Hale, B., Jager, T., Lauer, S.: 0-RTT key exchange with full forward secrecy. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10212, pp. 519–548. Springer, Cham (2017). doi:[10.1007/978-3-319-56617-7_18](https://doi.org/10.1007/978-3-319-56617-7_18)
15. Hale, B., Jager, T., Lauer, S., Schwenk, J.: Simple security definitions for and constructions of 0-RTT key exchange. Cryptology ePrint Archive, Report 2015/1214 (2015). <http://eprint.iacr.org/2015/1214>
16. IEEE Symposium on Security and Privacy, 17–21 May 2015, San Jose, CA, USA. IEEE Computer Society Press (2015)
17. Jager, T., Kohlar, F., Schäge, S., Schwenk, J.: On the security of TLS-DHE in the standard model. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 273–293. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-32009-5_17](https://doi.org/10.1007/978-3-642-32009-5_17)
18. Krawczyk, H.: HMQV: a high-performance secure Diffie-Hellman protocol. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 546–566. Springer, Heidelberg (2005). doi:[10.1007/11535218_33](https://doi.org/10.1007/11535218_33)
19. Krawczyk, H., Paterson, K.G., Wee, H.: On the security of the TLS protocol: a systematic analysis. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8042, pp. 429–448. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-40041-4_24](https://doi.org/10.1007/978-3-642-40041-4_24)
20. LaMacchia, B., Lauter, K., Mityagin, A.: Stronger security of authenticated key exchange. In: Susilo, W., Liu, J.K., Mu, Y. (eds.) ProvSec 2007. LNCS, vol. 4784, pp. 1–16. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-75670-5_1](https://doi.org/10.1007/978-3-540-75670-5_1)
21. Lauter, K., Mityagin, A.: Security analysis of KEA authenticated key exchange protocol. In: Yung et al. [30], pp. 378–394
22. Law, L., Menezes, A., Minghua, Q., Solinas, J., Vanstone, S.: An efficient protocol for authenticated key agreement. *Des. Codes Crypt.* **28**(2), 119–134 (2003)
23. Lychev, R., Jero, S., Boldyreva, A., Nita-Rotaru, C.: How secure and quick is QUIC? Provable security and performance analyses. In: IEEE S&P 2015 [16], pp. 214–231
24. NIST: SKIPJACK and KEA algorithm specifications (1998). <http://csrc.nist.gov/groups/STM/cavp/documents/skipjack/skipjack.pdf>
25. Rescorla, E.: The Transport Layer Security (TLS) Protocol Version 1.3: draft-ietf-tls-tls13-08. Technical report, August 2015. Expires 29 Feb 2016
26. Rescorla, E.: The Transport Layer Security (TLS) Protocol Version 1.3: draft-ietf-tls-tls13-18. Technical report, October 2016. Expires 29 April 2017
27. Sarr, A.P., Elbaz-Vincent, P., Bajard, J.-C.: A new security model for authenticated key agreement. In: Garay, J.A., Prisco, R. (eds.) SCN 2010. LNCS, vol. 6280, pp. 219–234. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-15317-4_15](https://doi.org/10.1007/978-3-642-15317-4_15)
28. Steinfeld, R., Pieprzyk, J., Wang, H.: How to strengthen any weakly unforgeable signature into a strongly unforgeable signature. In: Abe, M. (ed.) CT-RSA 2007. LNCS, vol. 4377, pp. 357–371. Springer, Heidelberg (2006). doi:[10.1007/11967668_23](https://doi.org/10.1007/11967668_23)

29. Yoneyama, K., Zhao, Y.: Taxonomical security consideration of authenticated key exchange resilient to intermediate computation leakage. In: Boyen, X., Chen, X. (eds.) ProvSec 2011. LNCS, vol. 6980, pp. 348–365. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-24316-5_25](https://doi.org/10.1007/978-3-642-24316-5_25)
30. Yung, M., Dodis, Y., Kiayias, A., Malkin, T. (eds.): PKC 2006. LNCS, vol. 3958. Springer, Heidelberg (2006)
31. Zhao, Y.: Identity-concealed authenticated encryption and key exchange. In: Weippl, E.R., Katzenbeisser, S., Kruegel, C., Myers, A.C., Halevi, S. (eds.) ACM CCS 2016, 24–28 October 2016, Vienna, Austria, pp. 1464–1479. ACM Press (2016)