

# Symmetric-key Authenticated Key Exchange (SAKE) with Perfect Forward Secrecy

Gildas Avoine<sup>1,2</sup>, Sébastien Canard<sup>3</sup>, and Loïc Ferreira<sup>3,1</sup>

<sup>1</sup> Univ Rennes, INSA Rennes, CNRS, IRISA, France

<sup>2</sup> Institut Universitaire de France

`gildas.avoine@irisa.fr`

<sup>3</sup> Orange Labs, Applied Crypto Group, Caen, France

`{sebastien.canard,loic.ferreira}@orange.com`

**Abstract.** Key exchange protocols in the asymmetric-key setting are known to provide stronger security properties than protocols in symmetric-key cryptography. In particular, they can provide *perfect forward secrecy*, as illustrated by key exchange protocols based on the Diffie-Hellman scheme. However public-key algorithms are too heavy for low-resource devices, which can then not benefit from forward secrecy. In this paper, we describe a scheme that solves this issue. Using a nifty resynchronisation technique, we propose an authenticated key exchange protocol in the symmetric-key setting that guarantees perfect forward secrecy. We prove that the protocol is sound, and provide a formal security proof.

**Keywords:** Authenticated key agreement · Symmetric-key cryptography · Perfect forward secrecy · Key-evolving.

## 1 Introduction

An authenticated key exchange (AKE) protocol executed between two parties aims at authenticating the parties, and computing a fresh shared session key. Well-known two-party authenticated key exchange protocols make use of digital signatures to provide authentication, and apply the Diffie-Hellman (DH) scheme [20] to compute a shared session key. However, such protocols are too heavy for low-resource devices. More suited protocols, solely based on symmetric-key functions, have been proposed (e.g., [11, 15, 23, 27, 32, 33, 38, 39] to cite a few), including widely deployed ones (e.g., in 3G/UMTS [2] and 4G/LTE [3]). Such symmetric-key protocols are needed in various applications, ranging from Wireless Sensor Networks (WSNs), Radio Frequency Identification (RFID) tags, smart cards, Controller Area Networks (CANs) for vehicular systems, smart home, up to industrial Internet of Things (IoT). Yet, existing symmetric-key based protocols lack a fundamental security property usually provided by the DH scheme: *perfect forward secrecy* (PFS) [21, 25]. PFS is a very strong form of long-term security which, informally, guarantees that future disclosures of some long-term secret keys do not compromise past session keys. Not all public-key

schemes are suitable to guarantee PFS. But, it is widely accepted that PFS can only be provided by asymmetric schemes (at least regarding stateless protocols). Indeed, in protocols based on symmetric-key functions, the two parties must share a long-term symmetric key (which the session keys are computed from). Therefore the disclosure of this static long-term key allows an adversary to compute all the past (and future) session keys. In this paper, we introduce an AKE protocol in the *symmetric-key* setting, and, yet, that does guarantee PFS.

### 1.1 Motivation of Our Approach

Symmetric-key based protocols do not provide the same security guarantees as those based on asymmetric algorithms. In particular, they do not guarantee forward secrecy. Nonetheless, (a few) attempts aim at proposing symmetric-key protocols that incorporate forward secrecy, as illustrated by the following related work.

*Related Work.* Dousti and Jalili [22] describe a key exchange protocol where the shared master key is updated based on time. Their protocol requires perfect synchronicity between the parties otherwise this leads to two main consequences. Firstly, in order to handle the key exchange messages, the parties may use different values of the master key corresponding to consecutive epochs, which causes the session to abort. Secondly, this allows an adversary to trivially break forward secrecy. Once a party deems the protocol run is correct and the session key can be safely used (i.e., once the party “accepts”), the adversary corrupts its partner (which still owns the previous, not updated yet, master key), and computes the current session key. Furthermore, achieving perfect time synchronisation may be quite complex in any context, in particular for low-resource devices. Contrary to Dousti et al., the protocol we propose explicitly deals with the issue of updating the master keys at both parties without requiring any additional functionality (such as a synchronised clock).

In the RFID field, the protocol proposed by Le, Burmester, and de Medeiros [31] aims at authenticating a tag to a server, and at computing a session key in order to establish a secure channel (which they do not describe). The master key is updated throughout the protocol run. To deal with the possible desynchronisation between the reader and the tag, the server keeps two consecutive values of the key: the current and the previous one. If the tag does not update its master key (which happens when the last message is dropped), the server is able to catch up during the next session. This implies that, in case of desynchronisation, the server computes the session key from the updated master key, whereas the tag still stores the previous value. Hence, an adversary that corrupts the tag can compute the previous session key with respect to the server. In fact, since the server always keeps the previous value of the master key, together with the current one, the scheme is intrinsically insecure in strong security models (i.e., models that allow the adversary to corrupt any of the partners, once the targeted party accepts). Yet, Le et al. analyse their protocol in a model where any server corruption is forbidden, and corrupting a tag is allowed only once it

accepts. In our scheme, one of the party also keeps in memory (a few) samples of a master key corresponding to different epochs (including a previous one). Yet the disclosure of all these values does *not* compromise past session keys. Furthermore, the (strong) security model we use allows the adversary to corrupt either partner as soon as the targeted party accepts.

Brier and Peyrin [16] propose a forward secret key derivation scheme in a client-server setting, that aims at improving a previous proposal [7]. In addition to forward secrecy, another constraint is that the amount of calculation to compute the master key (directly used as encryption key) on the server side must be low. Their solution implies to store, on the client side, several keys in parallel and to use a (short) counter, which is involved in the keys update. The keys belong to a tree whose each leaf (key) is derived from the previous one and the counter. The client must send the counter with the encrypted message for the server to be able to compute the corresponding key. The main drawback of this scheme is that the number of possible encryption keys is reduced. Increasing that limit implies increasing the counter size and the number of keys stored in parallel on the client side. Moreover, Brier et al. (as well as [7]) focus on forward secrecy with respect to the client only. The server is deemed as incorruptible, and is supposed to compute an encryption key only upon reception of a client’s message (the secure channel is unidirectional, and the server does not need to send encrypted messages to the client). Therefore, the scheme does not need to deal with the issue of *both* parties being in sync (with respect to the key computation), and providing forward secrecy. In addition, the purpose of Brier et al. (as well as [7]) is not to provide mutual authentication. More generally sending additional information in order to resynchronise (such as a sufficiently large counter) is a simple (and inefficient) way to build a forward secret protocol. But this yields several drawbacks. Firstly, the size of such a counter must be large enough in order to avoid any exhaustion. Secondly, sending the counter (at least periodically) is necessary for the two parties to resynchronise, which consumes bandwidth. Thirdly, resynchronisation may imply multiple updates of the master keys at once (the scheme of Brier et al. and [7] aims at limiting that amount of calculation, but it leads to a narrowed number of possible encryption keys). Our scheme avoids all these drawbacks.

The more general question of forward security in symmetric cryptography has been also investigated by Bellare and Yee [13]. They propose formal definitions and practical constructions of forward secure primitives (e.g., MAC, symmetric encryption algorithm). Their constructions protect against decryption of past messages, or antedated forgeries of messages (i.e., previously authenticated messages are made untrustworthy). Their algorithms are based on key-evolving schemes [9]. Nonetheless, Bellare et al. consider only algorithms (but not protocols) and they do not deal with the issue of synchronising the evolution of the shared key at *both* parties. That is, they propose out-of-context (non-interactive) solutions with respect to our purpose.

Abdalla and Bellare [4] investigate a related question which is “re-keying”. Their formal analysis shows that appropriate re-keying techniques “increase” the

lifetime of a key. They consider re-keying in the context of symmetric encryption (in order to thwart attacks based on the ability to get lots of encrypted messages under the same key), and forward security (in order to protect past keys). Yet, they confine their analysis to algorithms and not protocols. Hence, as Bellare et al. [13], they do not treat the synchronisation issues that arise from evolving a shared symmetric key.

The Signal messaging protocol [1] uses a key derivation scheme called “double ratchet algorithm” [34]. This scheme combines a DH based mechanism with a symmetric key-evolving mechanism (based on a one-way function). The first mechanism provides an asymmetric ratchet, whereas the second provides a symmetric ratchet. The asymmetric ratchet is applied when a fresh DH share is received (included in an application message) from the peer. The symmetric ratchet is applied when a party wants to send several successive messages without new incoming message from its partner. Thanks to the DH scheme, the asymmetric ratchet is supposed to provide forward secrecy.<sup>4</sup> Regarding the symmetric ratchet, each party is compelled to store the decryption keys of the not yet received messages. This is due to the asynchronous nature of the Signal protocol. Therefore, the symmetric ratchet in Signal does not provide forward secrecy, as stated in their security analysis by Cohn-Gordon, Cremers, Dowling, Garratt, and Stebila [18]: “*old but unused receiving keys are stored at the peer for an implementation dependent length of time, trading off forward security for transparent handling of outdated messages. This of course weakens the forward secrecy of the keys*”. Consequently, Cohn-Gordon et al. choose not to model this weakened property. In turn, Alwen, Coretti, and Dodis [6] incorporate the latter in the security analysis of their “generalised Signal protocol”. But the crucial difference in their notion of forward security is that, as soon as the receiver is compromised, no more security can be provided. On the contrary, we tackle the synchronisation issue, and solve it in our protocol. The security model we use captures forward secrecy and allows corrupting a party and its partner as soon as the targeted party “accepts” (i.e., deems the session key can be safely used). With regard to Signal, our protocol can be compared to the asymmetric ratchet (in synchronous mode), and yet does not implement asymmetric functions.

We stress that the goals of several of the aforementioned protocols are not the same as ours. Nonetheless, the small number of existing symmetric-key protocols that provide forward secrecy, and the lukewarm security level they achieve illustrate that combining symmetric-key cryptography and (a strong form of) forward secrecy is a non-trivial task.

## 1.2 Contributions

In this paper, we propose an AKE protocol, that we call SAKE, *solely* based on symmetric-key functions. Not only does our protocol provide mutual authentication and key agreement, but it guarantees perfect forward secrecy. We attain

<sup>4</sup> In Signal, the DH exchanges can be asynchronous. This impairs the forward secrecy property usually ensured by this scheme.

this very strong form of long-term security by using a key-evolving scheme. As soon as two parties make a shared (symmetric) key evolve, a synchronisation problem arises. We provide a simple and efficient solution to this issue. We do not require using neither a clock, nor an additional resynchronising procedure. Our solution is based on a second (independent) chain of master keys. These keys allow tracking the evolution of the internal state, and resynchronising the parties if necessary. The parties authenticate each other prior to updating their master keys. Hence the possible gap is bounded (as we prove it), and each party is always able to catch up in case of desynchronisation (of course, if the session is correct and complete). Mutual authentication, key exchange (with forward secrecy), and resynchronisation are done in the continuity of the protocol run. In addition, the protocol we describe has the following characteristics.

- It is self-synchronising. That is, after a correct and complete session (and whatever the internal state of the parties prior to the session), the two parties involved in the protocol run share a new session key, and their internal state is updated and synchronised.
- It allows establishing an (virtually) unlimited number of sessions (as opposite to protocols that make use of a predefined list of master keys, each being used once only).
- The amount of calculation done by both parties in a single protocol run is strictly bounded. In particular we avoid the need of sending additional information in order to resynchronise, such as a (sufficiently large) counter that keeps track of the evolution of the master keys, and the subsequent drawbacks: periodically doing a great amount of computations at once (when resynchronisation is necessary), and consuming bandwidth (to transmit the additional data).

In addition, we provide a formal security proof for SAKE. We also present a variant of SAKE (that we call SAKE-AM) which is an “aggressive mode” of the protocol.<sup>5</sup> This variant inverts the role of the initiator and the responder in terms of calculations (in SAKE, the initiator performs – at most – two additional MAC computations compared to the responder). Using SAKE and SAKE-AM together results in an implementation (gathering all the aforementioned properties, starting with the forward secrecy property) that allows any party to be either initiator or responder of a session, and such that the smallest amount of calculation is always done by the same party. This is particularly convenient in the context of a set of (low-resource) end-devices communicating with a central server. In such a case, the end-device supports the smallest amount of calculation, whereas either the server or the end-device can initiate a session.

Furthermore, we elaborate on the differences between our approach and the DH scheme (beyond the intrinsic distinction between public-key and symmetric-key cryptography).

---

<sup>5</sup> The variant (with one message less) allows computing the synchronisation gap  $\delta$  earlier (with the first message). But the responder must wait for the third message to confirm that value. In a sense, this variant is also more optimistic.

### 1.3 Outline of the Paper

In Section 2 we detail the security model used to analyse the protocol we propose. Our authenticated key exchange protocol in symmetric-key setting with forward secrecy is described in Section 3. Formal proofs of soundness and security for the protocol are presented in Section 4. The differences between our approach and the DH scheme are highlighted in Section 5. Finally, we conclude in Section 6.

## 2 Security Model

Before describing our symmetric-key protocol in Section 3 (which is self-sufficient and contains all the specifics required to understand the protocol), we present in this section the security model that we employ to formally prove its security.

In a nutshell, we use the model for authenticated key exchange protocols described by Brzuska, Jacobsen, and Stebila [17]. This model incorporates all the features that are usually considered when analysing key agreement protocols in the public-key setting (e.g., DH-based protocols with signature). In this model, the adversary has full control over the communication network. It can forward, alter, drop any message exchanged by honest parties, or insert new messages. Brzuska et al.’s model then captures adaptive corruptions but also perfect forward secrecy. This appears in the definition of the security experiment.

### 2.1 Execution Environment

In this section, we present the security model for authenticated key exchange protocols described by Brzuska et al. [17], and reuse the corresponding notations.

*Parties.* A two-party protocol is carried out by a set of parties  $\mathcal{P} = \{P_0, \dots, P_{n-1}\}$ . Each party  $P_i$  has an associated long-term key  $P_i.\text{ltk}$ . The same long-term key ltk is shared by a unique pair of parties  $(P_i, P_j)$ . That is,  $P_i.\text{ltk} = P_j.\text{ltk}$ .<sup>6</sup>

*Instances.* Each party can take part in multiple sequential executions of the protocol. We prohibit *parallel* executions of the protocol. Indeed, since the protocol we propose is based on shared *evolving* symmetric keys, running multiple instances in parallel may cause some executions to abort (we elaborate more on this in Section 5). This is the only restriction we demand compared to AKE security models used in the public-key setting.

Each run of the protocol is called a session. To each session of a party  $P_i$ , an instance  $\pi_i^s$  is associated which embodies this (local) session’s execution of the protocol, and has access to the long-term key of the party. In addition, each instance maintains the following state specific to the session.

- $\rho$ : the role  $\rho \in \{\text{init}, \text{resp}\}$  of the session in the protocol execution, being either the initiator or the responder.

<sup>6</sup> Note that ltk can be a set of master keys (e.g., each one used by the party for a different purpose).

- **pid**: the identity  $\text{pid} \in \mathcal{P}$  of the intended communication partner of  $\pi_i^s$ .
- $\alpha$ : the state  $\alpha \in \{\perp, \text{running}, \text{accepted}, \text{rejected}\}$  of the instance.
- **sk**: the session key derived by  $\pi_i^s$ .
- $\kappa$ : the status  $\kappa \in \{\perp, \text{revealed}\}$  of the session key  $\pi_i^s.\text{sk}$ .
- **sid**: the identifier of the session.
- **b**: a random bit  $b \in \{0, 1\}$  sampled at initialisation of  $\pi_i^s$ .

We put the following correctness requirements on the variables  $\alpha$ , **sk**, **sid** and **pid**. For any two instances  $\pi_i^s, \pi_j^t$ , the following must hold:

$$\pi_i^s.\alpha = \text{accepted} \Rightarrow \pi_i^s.\text{sk} \neq \perp \wedge \pi_i^s.\text{sid} \neq \perp \quad (1)$$

$$\pi_i^s.\alpha = \pi_j^t.\alpha = \text{accepted} \wedge \pi_i^s.\text{sid} = \pi_j^t.\text{sid} \Rightarrow \begin{cases} \pi_i^s.\text{sk} = \pi_j^t.\text{sk} \\ \pi_i^s.\text{pid} = P_j \\ \pi_j^t.\text{pid} = P_i \end{cases} \quad (2)$$

*Adversarial queries.* The adversary  $\mathcal{A}$  is assumed to control the network, and interacts with the instances by issuing the following queries to them.

- **NewSession**( $P_i, \rho, \text{pid}$ ): this query creates a new instance  $\pi_i^s$  at party  $P_i$ , having role  $\rho$ , and intended partner **pid**.
- **Send**( $\pi_i^s, m$ ): this query allows the adversary to send any message  $m$  to  $\pi_i^s$ . If  $\pi_i^s.\alpha \neq \text{running}$ , it returns  $\perp$ . Otherwise  $\pi_i^s$  responds according to the protocol specification.
- **Corrupt**( $P_i$ ): this query returns the long-term key  $P_i.\text{ltk}$  of  $P_i$ . If **Corrupt**( $P_i$ ) is the  $\nu$ -th query issued by the adversary, then we say that  $P_i$  is  $\nu$ -corrupted. For a party that has not been corrupted, we define  $\nu = +\infty$ .
- **Reveal**( $\pi_i^s$ ): this query returns the session key  $\pi_i^s.\text{sk}$ , and  $\pi_i^s.\kappa$  is set to **revealed**.
- **Test**( $\pi_i^s$ ): this query may be asked only once throughout the game. If  $\pi_i^s.\alpha \neq \text{accepted}$ , then it returns  $\perp$ . Otherwise it samples an independent key  $sk_0 \xleftarrow{\$} \mathcal{K}$ , and returns  $sk_b$ , where  $sk_1 = \pi_i^s.\text{sk}$ . The key  $sk_b$  is called the *Test-challenge*.

**Definition 1 (Partnership).** *Two instances  $\pi_i^s$  and  $\pi_j^t$  are partners if  $\pi_i^s.\text{sid} = \pi_j^t.\text{sid}$ .*

**Definition 2 (Freshness).** *An instance  $\pi_i^s$  is said to be fresh with intended partner  $P_j$ , if*

- (a)  $\pi_i^s.\alpha = \text{accepted}$  and  $\pi_i^s.\text{pid} = P_j$  when  $\mathcal{A}$  issues its  $\nu_0$ -th query,
- (b)  $\pi_i^s.\kappa \neq \text{revealed}$  and  $P_i$  is  $\nu$ -corrupted with  $\nu_0 < \nu$ , and
- (c) for any partner instance  $\pi_j^t$  of  $\pi_i^s$ , we have that  $\pi_j^t.\kappa \neq \text{revealed}$  and  $P_j$  is  $\nu'$ -corrupted with  $\nu_0 < \nu'$ .

Note that the notion of freshness incorporates a requirement for forward secrecy.

An *authenticated key exchange protocol* (AKE) is a two-party protocol satisfying the correctness requirements 1 and 2, and where the security is defined in terms of an AKE experiment played between a challenger and an adversary. This experiment uses the execution environment described above. The adversary can win the AKE experiment in one of two ways: (i) by making an instance accept maliciously, or (ii) by guessing the secret bit of the `Test`-instance.

**Definition 3 (Entity Authentication (EA)).** *An instance  $\pi_i^s$  of a protocol  $\Pi$  is said to have accepted maliciously in the AKE security experiment with intended partner  $P_j$ , if*

- (a)  $\pi_i^s.\alpha = \text{accepted}$  and  $\pi_i^s.\text{pid} = P_j$  when  $\mathcal{A}$  issues its  $\nu_0$ -th query,
- (b)  $P_i$  and  $P_j$  are  $\nu$ - and  $\nu'$ -corrupted with  $\nu_0 < \nu, \nu'$ , and
- (c) there is no unique instance  $\pi_j^t$  such that  $\pi_i^s$  and  $\pi_j^t$  are partners.

The adversary's advantage is defined as its winning probability:

$$\text{adv}_{\Pi}^{\text{ent-auth}}(\mathcal{A}) = \Pr[\mathcal{A} \text{ wins the EA game}].$$

**Definition 4 (Key Indistinguishability).** *An adversary  $\mathcal{A}$  against a protocol  $\Pi$ , that issues its `Test`-query to instance  $\pi_i^s$  during the AKE security experiment, answers the `Test`-challenge correctly if it terminates with output  $b'$ , such that*

- (a)  $\pi_i^s$  is fresh with some intended partner  $P_j$ , and
- (b)  $\pi_i^s.\text{b} = b'$ .

The adversary's advantage is defined as

$$\text{adv}_{\Pi}^{\text{key-ind}}(\mathcal{A}) = \left| \Pr[\pi_i^s.\text{b} = b'] - \frac{1}{2} \right|.$$

Definitions 3 and 4 allow the adversary to corrupt an instance involved in the security experiment (once the targeted instance has accepted, in order to exclude trivial attacks). Therefore, protocols secure with respect to Definition 5 below provide *perfect forward secrecy*. Note that we do not allow the targeted instance to be corrupted before it accepts. This security model does not capture key-compromise impersonation attacks (KCI) [14] since that would allow trivially breaking key exchange protocols solely based on shared symmetric keys.

**Definition 5 (AKE Security).** *We say that a two-party protocol  $\Pi$  is a secure AKE protocol if  $\Pi$  satisfies the correctness requirements 1 and 2, and for all probabilistic polynomial time adversary  $\mathcal{A}$ ,  $\text{adv}_{\Pi}^{\text{ent-auth}}(\mathcal{A})$  and  $\text{adv}_{\Pi}^{\text{key-ind}}(\mathcal{A})$  are a negligible function of the security parameter.*



## 2.2 Security Definitions of SAKE's Building Blocks

In this section, we recall the definitions of the main security notions we use in our results. The security definition of a pseudo-random function is taken from Bellare, Desai, Jokipii, and Rogaway [8], and that of a MAC strongly unforgeable under chosen-message attacks from Bellare and Namprempre [10]. We recall also the definition of matching conversations initially proposed by Bellare and Rogaway [11], and modified by Jager, Kohlar, Schäge, and Schwenk [28].

*Secure PRF.* A *pseudo-random function* (PRF)  $F$  is a deterministic algorithm which given a key  $K \in \{0, 1\}^\lambda$  and a bit string  $x \in \{0, 1\}^*$  outputs a string  $y = F(K, x) \in \{0, 1\}^\gamma$  (with  $\gamma$  being polynomial in  $\lambda$ ). Let  $Func$  be the set of all functions of domain  $\{0, 1\}^*$  and range  $\{0, 1\}^\gamma$ . The security of a PRF is defined with the following experiment between a challenger and an adversary  $\mathcal{A}$ :

1. The challenger samples  $K \xleftarrow{\$} \{0, 1\}^\lambda$ ,  $G \xleftarrow{\$} Func$ , and  $b \xleftarrow{\$} \{0, 1\}$  uniformly at random.
2. The adversary may adaptively query values  $x$  to the challenger. The challenger replies to each query with either  $y = F(K, x)$  if  $b = 1$ , or  $y = G(x)$  if  $b = 0$ .
3. Finally, the adversary outputs its guess  $b' \in \{0, 1\}$  of  $b$ .

The adversary's advantage is defined as

$$\text{adv}_F^{\text{prf}}(\mathcal{A}) = \left| \Pr[b = b'] - \frac{1}{2} \right|.$$

**Definition 6 (Secure PRF).** A function  $F: \{0, 1\}^\lambda \times \{0, 1\}^* \rightarrow \{0, 1\}^\gamma$  is said to be a secure pseudo-random function (PRF) if, for all probabilistic polynomial time adversary  $\mathcal{A}$ ,  $\text{adv}_F^{\text{prf}}(\mathcal{A})$  is a negligible function in  $\lambda$ .

*Secure MAC.* A message authentication code (MAC) consists of two algorithms  $(\text{Mac}, \text{Vrf})$ . The tagging algorithm  $\text{Mac}$  takes as input a key  $K \in \{0, 1\}^k$  and a message  $m \in \{0, 1\}^*$  and returns a tag  $\tau \in \{0, 1\}^\gamma$  (with  $\gamma$  being polynomial in  $k$ ). The verification algorithm  $\text{Vrf}$  takes as input the key  $K$ , a message  $m$ , and a candidate tag  $\tau$  for  $m$ . It outputs 1 if  $\tau$  is a valid tag on message  $m$  with respect to  $K$ . Otherwise, it returns 0. The notion of *strong unforgeability under chosen-message attacks* (SUF-CMA) for a MAC  $G = (\text{Mac}, \text{Vrf})$  is defined with the following experiment between a challenger and an adversary  $\mathcal{A}$ :

1. The challenger samples  $K \xleftarrow{\$} \{0, 1\}^k$ , and sets  $S \leftarrow \emptyset$ .
2. The adversary may adaptively query values  $m$  to the challenger. The challenger replies to each query with  $\tau = \text{Mac}(K, m)$  and records  $(m, \tau)$ :  $S \leftarrow S \cup \{(m, \tau)\}$ .
3. Finally, the adversary sends  $(m^*, \tau^*)$  to the challenger.

The adversary's advantage is defined as

$$\text{adv}_G^{\text{suf-cma}}(\mathcal{A}) = \Pr[\text{Vrf}(K, m^*, \tau^*) = 1 \wedge (m^*, \tau^*) \notin S].$$

**Definition 7 (SUF-CMA).** A message authentication code  $G = (\text{Mac}, \text{Vrf})$  with  $\text{Mac}: \{0, 1\}^k \times \{0, 1\}^* \rightarrow \{0, 1\}^\gamma$  is said to be strongly unforgeable under chosen-message attacks (SUF-CMA) if, for all probabilistic polynomial time adversary  $\mathcal{A}$ ,  $\text{adv}_G^{\text{suf-cma}}(\mathcal{A})$  is a negligible function in  $k$ .

*Matching conversations.* Let  $T_{i,s}$  be the sequence of all (valid) messages sent and received by an instance  $\pi_i^s$  in chronological order. For two transcripts  $T_{i,s}$  and  $T_{j,t}$ , we say that  $T_{i,s}$  is a prefix of  $T_{j,t}$  if  $T_{i,s}$  contains at least one message, and the messages in  $T_{i,s}$  are identical to the first  $|T_{i,s}|$  messages of  $T_{j,t}$ .

**Definition 8 (Matching Conversations).** We say that  $\pi_i^s$  has a matching conversation to  $\pi_j^t$ , if

- $\pi_i^s$  has sent all protocol messages and  $T_{j,t}$  is a prefix of  $T_{i,s}$ , or
- $\pi_j^t$  has sent all protocol messages and  $T_{i,s} = T_{j,t}$ .

*Remark 1.* Defining matching conversations as per Definition 8 means that we use a post-specified session identifier  $\text{sid}$  equal to the first five or four messages of the protocol that we describe in Section 3.

### 3 Our Symmetric-key AKE Protocol with Perfect Forward Secrecy

In this section we describe our main protocol. Although all the calculations are based on shared master keys, forward secrecy is guaranteed by using a key-evolving scheme. More precisely, we use two types of keys: one to compute the session keys, the other to authenticate messages and resynchronise when necessary. This second type of keys allows tracking the master keys evolution, and limit the gap (in terms of keys update) between both parties. Mutual authentication, key exchange, and synchronised update of the master keys are done in the same session.

#### 3.1 Key Concepts

The protocol allows two parties  $A$  (initiator) and  $B$  (responder) to mutually authenticate and compute a shared session key. It is based on two types of master keys: a derivation master key  $K$  and an authentication master key  $K'$ . The protocol makes use of symmetric-key functions only. Each pair of parties ( $A, B$ ) shares distinct master keys. The main lines of the protocol are as follows: the two parties exchange pseudo-random values  $r_A, r_B$ . These two values are used to

- authenticate each other: each party sends back the value it has received in a message that is MAC-ed with the authentication master key  $K'$ . For instance, if  $B$  receives  $r_A$  it replies with  $r_B \parallel \tau_B$  where  $\tau_B = \text{MAC}(K', B \parallel A \parallel r_B \parallel r_A)$ .

- Compute a session key: a pseudo-random function KDF is keyed with the derivation master key  $K$  and uses the pseudo-random values as input. That is,  $sk \leftarrow \text{KDF}(K, f(r_A, r_B))$ . Function  $f$  is deliberately left undefined. For instance,  $f(r_A, r_B)$  can be equal to the concatenation or the bitwise addition of  $r_A$  and  $r_B$ .<sup>7</sup>

*Providing forward secrecy.* The shared key  $K$  is used to compute the session keys. If this key remains unchanged throughout all sessions, its disclosure allows computing all past (and future) session keys. To solve this issue we apply a key-evolving technique. We update the master key such that a previous master key cannot be computed from an updated one. Each of the two parties involved in a session updates its own copy of the derivation master key  $K$  with a non-invertible function  $\text{update}$ :  $K \leftarrow \text{update}(K)$ . Hence this protects past sessions in case the (current value of) master key  $K$  is revealed. Each party authenticates its peer prior to updating the derivation master key. If the master key is updated throughout the session, it may happen that one of the two involved parties update its master key whereas the other do not. This leads to a *synchronisation problem*.

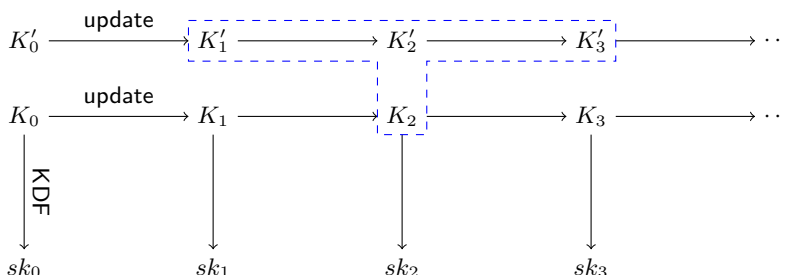


Fig. 1: Master key chains in SAKE. At epoch  $j$ , the initiator stores four keys:  $K = K_j$ , and  $K'_{j-1}$ ,  $K'_j$ ,  $K'_{j+1}$ . The responder stores two keys:  $K = K_j$  and  $K' = K'_j$ . An illustration with  $j = 2$  corresponds to the keys surrounded by the blue dashed box.

*The synchronisation problem.* If two parties use a different key  $K$ , they are obviously not able to compute a shared session key. Hence they must resynchronise first. More fundamentally, if a party initiates a session with some derivation master key  $K$ , and its partner stores a master key corresponding to an earlier epoch, then an adversary that corrupts the partner can compute past session keys with respect to the initiator, hence trivially break forward secrecy. Therefore, it is of

<sup>7</sup> The function  $f$  must be chosen such that the security of KDF is not impaired. We assume here that the cryptographic functions used are ideal (investigating this topic is beyond the scope of this paper).

paramount importance that the parties know if the master key of its partner has actually been updated. We provide a solution to both issues in the continuity of a *single* session. In particular, no extra procedure is needed in order for a desynchronised party to catch up. We avoid the need of sending additional information in order to resynchronise, such as a (sufficiently large) counter that keeps track of the evolution of the master keys, and the subsequent drawbacks: periodically doing a great amount of computations at once (when resynchronisation is necessary), and consuming bandwidth (to transmit the additional data). We base our solution on the second master key  $K'$  used to authenticate the messages exchanged during a session. The solution is to update  $K'$  at the same time as  $K$ . Therefore the evolution of  $K'$  follows that of  $K$ . The party that receives the first authenticated message uses the MAC tag to learn which epoch the sender belongs to. Of course,  $K'$  can also be desynchronised in the same way as  $K$ . This is why, whereas one party (responder  $B$ ) stores only one sample of the key  $K'$ , the other party (initiator  $A$ ) stores several samples of the authentication master key  $K'$  corresponding to several consecutive epochs. We *prove* that only three keys  $K'_{j+1}$ ,  $K'_j$ ,  $K'_{j-1}$ , corresponding respectively to the next, the current, and the previous epochs, are sufficient in order for  $A$  and  $B$  to resynchronise. The initiator ( $A$ ) is the one able to deal with the synchronisation issue, and consequently tells  $B$  how to behave. Each party “accepts” only after it has received a confirmation (final MAC-ed messages) that its partner has already updated its own master keys. In such a case, the party ending in accepting state deems that the fresh session key can be safely used. Otherwise (in particular when the parties are desynchronised), the session key is discarded.

Since two independent master keys are used (authentication and session key derivation), one can safely maintain a copy of  $K'$  corresponding to an earlier epoch ( $K'_{j-1}$ ) without risk of threatening forward secrecy. Only one sample of the derivation master key  $K$  is kept: the most up-to-date.

### 3.2 Description of the Protocol

The protocol is depicted by Figure 2. We call it *Symmetric-key Authenticated Key Exchange* (SAKE). The parameter  $\delta_{AB}$  computed by  $A$  corresponds to the gap between  $A$  and  $B$  with respect to the evolution of the master keys. We prove that  $\delta_{AB} \in \{-1, 0, 1\}$  (see Section 4.1). That is,  $A$  can only be either *one step* behind, or in sync, or *one step* ahead to  $B$ . During a session,  $A$  uses the keys  $K'_j$ ,  $K'_{j-1}$ ,  $K'_{j+1}$  (by order of likelihood) and the first message ( $m_B$ ) sent by  $B$  to learn  $\delta_{AB}$ . The message  $m_B$  is computed with the current value  $K'$  of  $B$ . Therefore  $m_B$  indicates the current synchronisation state of  $B$ . Then  $A$  informs  $B$ . One bit  $\epsilon$  is enough (message  $m_A$ ) because  $B$  takes two behaviours only: if  $\delta_{AB} \in \{-1, 0\}$  ( $\epsilon = 0$ ), and if  $\delta_{AB} = 1$  ( $\epsilon = 1$ ).  $A$  and  $B$  behaves as follows.

- If  $A$  is in sync with  $B$  ( $\delta_{AB} = 0$ ),  $A$  computes the new session key, and updates its master keys. Then, upon reception of  $m_A$ ,  $B$  does the same.
- If  $A$  is in advance ( $\delta_{AB} = 1$ ),  $A$  waits for  $B$  to resynchronise (i.e.,  $B$  updates its master keys a first time), and to proceed with the regular operations (i.e.,

$B$  computes the new session key, and updates its master keys a second time). Then, once  $A$  receives a confirmation that  $B$  is synchronised (message  $\tau'_B$ ),  $A$  performs the regular operations as well (session key computation, master keys update). Since  $A$  waits for  $B$  to resynchronise before proceeding, the gap between the parties is *bounded* (as proved in Section 4.1).

- If  $A$  is late ( $\delta_{AB} = -1$ ), it resynchronises (i.e., it updates its master keys a first time), and then performs the regular operations (session key computation, master keys update). Then (upon reception of message  $m_A$ ),  $B$  applies the regular operations.

Once a correct and complete session ends, three goals are achieved in the *same* protocol run: (i) the two parties have updated their master keys, (ii) they are synchronised (which stems in particular from the fact that the gap between  $A$  and  $B$  is bounded, i.e.,  $|\delta_{AB}| \leq 1$ ), and (iii) they share a new session key. In other words, the protocol is *self-synchronising*.

The session can be reduced from five to four messages in some cases. Indeed, regarding the synchronisation state, in two cases (when  $\delta_{AB} \in \{-1, 0\}$ , that is  $\epsilon = 0$ ),  $A$  and  $B$  are synchronised, and share a session key once  $B$  has received message  $m_A$  and executed the subsequent operations. Therefore, in such a case, the session can end upon reception of message  $\tau'_B$  by  $A$ . More precisely

- if  $\delta_{AB} = 1$  ( $\epsilon = 1$ ), then  $A$  accepts upon reception of  $\tau'_B$ , and  $B$  accepts upon reception of  $\tau'_A$ ;
- if  $\delta_{AB} \in \{-1, 0\}$  ( $\epsilon = 0$ ), then  $A$  accepts upon reception of  $\tau'_B$ , and  $B$  accepts upon reception of  $m_A$ .

Although this does not appear explicitly in Figure 2, a party aborts the session if it receives a message computed with an invalid identity. For the responder  $B$ , an invalid identity corresponds to an initiator party  $A$  it does not share master keys with. For an initiator  $A$ , the particular case  $B = A$ , among other possibilities, yields an error (i.e., each party must have a distinct identity).

*Remark 2.* Since  $K'_{j+1}$  and  $K'_j$  can be computed from  $K'_{j-1}$ , it is also possible to store only  $K'_{j-1}$ , and to compute the two other keys when necessary during the session.

*Remark 3.* Alternatively, the three authentication keys  $K'_{j-1}$ ,  $K'_j$ ,  $K'_{j+1}$  can be replaced by two local counters  $c_A$ ,  $c_B$  (respectively stored by  $A$  and  $B$ ) that keep track of the evolution of the derivation master key  $K$ , with one static authentication master key  $K'$ .<sup>8</sup> On the initiator's side, the MAC verifications are then done with consecutive values of the counter  $j-1$ ,  $j$ ,  $j+1$ . On the overall, the sequence of operations and the computations are similar to that of SAKE. This means mainly replacing function  $x \mapsto \text{MAC}(K'_j, x)$  with  $x \mapsto \text{MAC}(K', j||x)$ . This alternative implies the storage of two keys and one counter:  $K$ ,  $K'$  and  $c_A/c_B$ , instead of two keys only:  $K$  and  $K'_{j-1}/K'$  (and, on the initiator's side only, one or two additional calls to `update` in order to compute  $K'_j$  and, possibly,  $K'_{j+1}$ ).

<sup>8</sup> This alternative has been suggested by anonymous reviewers of Crypto 2019.

*Remark 4.* With respect to the security model presented in Section 2, the long-term key of  $A$  and  $B$  corresponds respectively to  $A.\text{ltk} = (K, K'_{j-1})$  and  $B.\text{ltk} = (K, K')$ . We could have allowed the authentication master key  $K'_{j-1}/K'$  to be disclosed prior to the start of the session. This would not impair the forward secrecy of the derivation master key  $K$ . Nonetheless, knowing the authentication master key an adversary could desynchronise a legitimate party so that the party could not catch up anymore. Hence our choice to include both master keys in the response to a `Corrupt`-query.

### 3.3 Notations

For the sake of clarity, we use the following notations in Figure 2:

- `kdf` corresponds to:  $sk \leftarrow \text{KDF}(K, f(r_A, r_B))$
- `updA` corresponds to
  1.  $K \leftarrow \text{update}(K)$
  2.  $K'_{j-1} \leftarrow K'_j$
  3.  $K'_j \leftarrow K'_{j+1}$
  4.  $K'_{j+1} \leftarrow \text{update}(K'_{j+1})$
- `updB` corresponds to
  1.  $K \leftarrow \text{update}(K)$
  2.  $K' \leftarrow \text{update}(K')$

Moreover, `verif`( $k, m, \tau$ ) denotes the MAC verification function that takes as input a secret key  $k$ , a message  $m$ , and a tag  $\tau$ . It outputs `true` if  $\tau$  is a valid tag on message  $m$  with respect to  $k$ . Otherwise, it returns `false`.

Before the first session between  $A$  and  $B$ , the master keys are initialised as follows<sup>9</sup>:

- $K$  and  $K'$  are uniformly chosen at random.
- $K'_{j-1} \leftarrow \perp$
- $K'_j \leftarrow K'$
- $K'_{j+1} \leftarrow \text{update}(K')$

### 3.4 SAKE-AM: a Variant of SAKE

From SAKE, we can derive an aggressive variant that allows any party to be either initiator or responder, and such that the smallest amount of calculation is always done by the *same* party.

In SAKE the initiator  $A$  owns the three keys  $K'_{j+1}$ ,  $K'_j$ ,  $K'_{j-1}$ , and the responder  $B$  does the lightest computations. In this variant  $B$  owns the three keys, and  $A$  does the smallest amount of calculation. The main idea is to skip the first SAKE message  $A||r_A$ . Hence the roles between the two parties are swapped. This leads to other minor changes in message format compared to SAKE. Despite these differences, the messages and the calculations are essentially the same

<sup>9</sup> During the first protocol run,  $A$  needs only  $K'_j$  to verify message  $m_B$ .

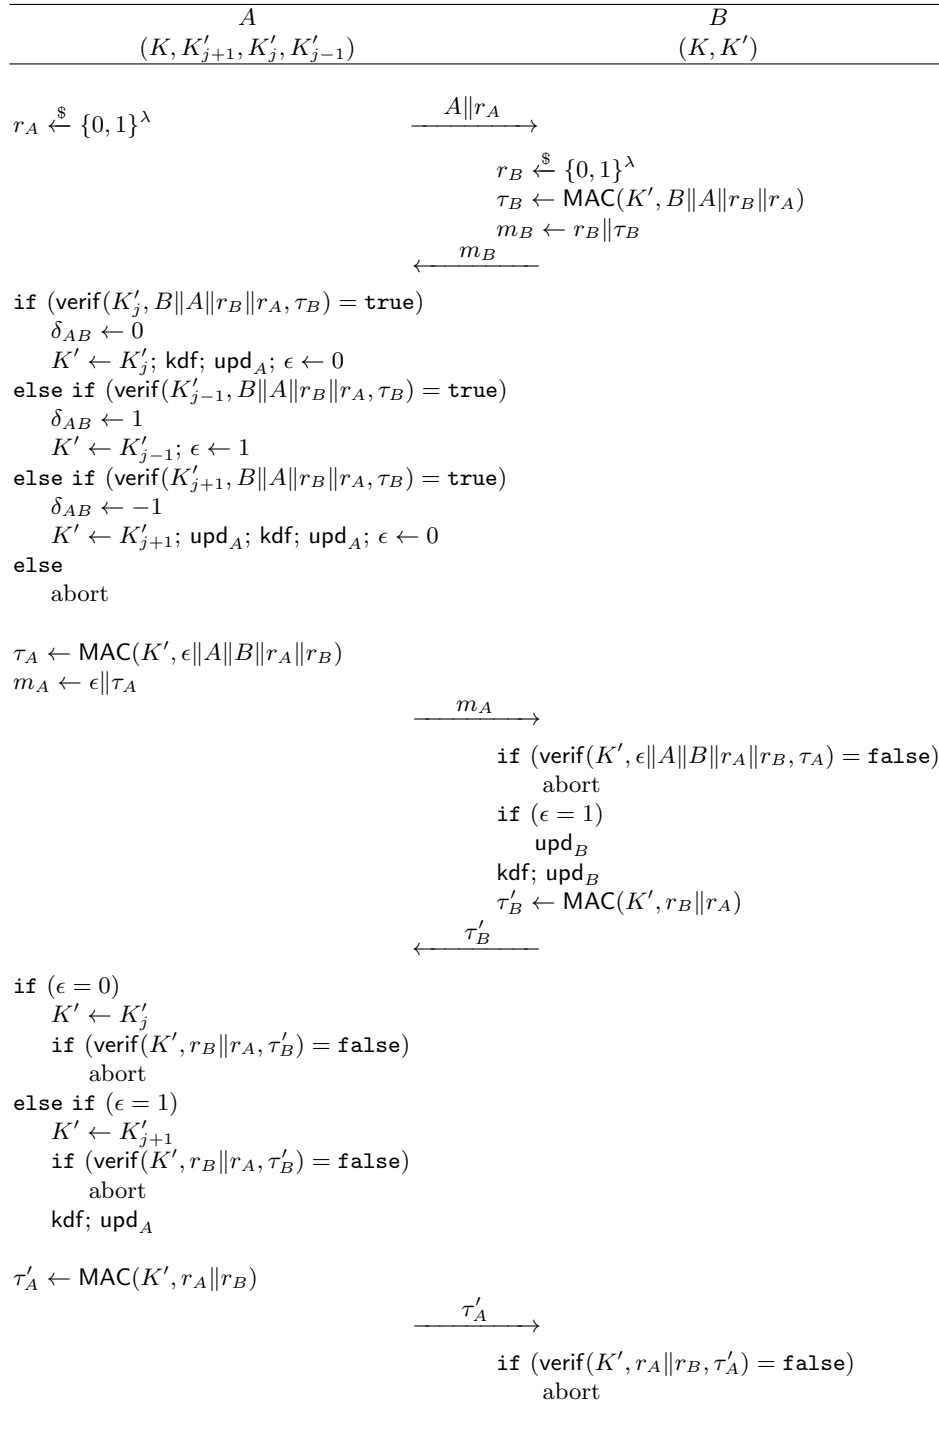


Fig. 2: SAKE protocol

as in SAKE. This variant remains a sound and secure AKE protocol (according to Definition 5).<sup>10</sup> We call this variant *SAKE in aggressive mode* (SAKE-AM).

This can be applied in the context of industrial IoT when a set of end-devices (e.g., sensors, actuators) communicate with a central server. When the end-device wants to initiate a communication, protocol SAKE-AM is launched. Otherwise (the server is initiator), SAKE is used (see Figure 3). Therefore, the end-device always does the lightest computations.

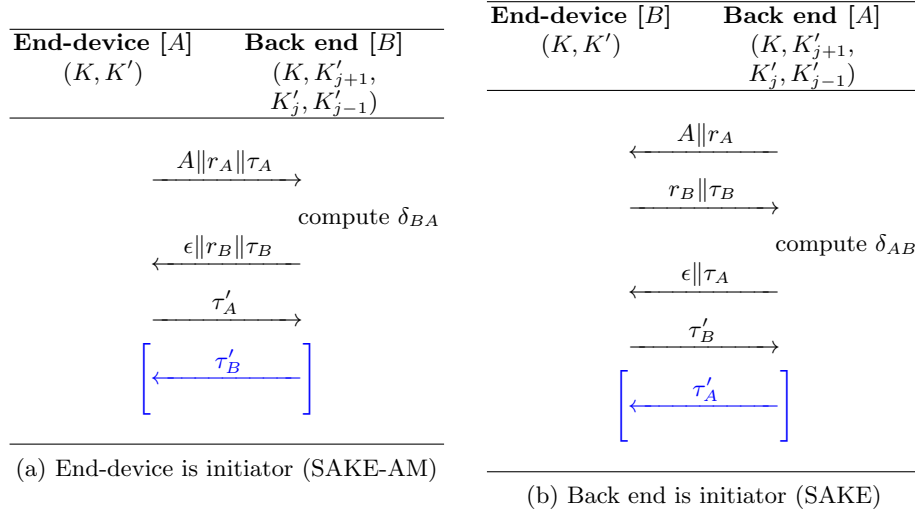


Fig. 3: Symmetric-key authenticated key exchange with forward secrecy between a low-resource end-device and a back-end server. Both parties may initiate the session. In some cases, the last message can be skipped.

## 4 Security and Soundness for SAKE

In this section we prove that (i) SAKE is *sound*, and (ii) it is a *secure AKE* protocol according to Definition 5.

### 4.1 Soundness of SAKE

We want to show that SAKE is *sound*, which essentially means that, once a correct session is complete, both parties have updated their respective internal state, are synchronised, and share the same (new) session key. We call a “benign”

<sup>10</sup> The proofs of soundness and security for SAKE-AM are essentially the same as for SAKE (see Section 4.2). They are given in Appendix B.




adversary an adversary that faithfully forwards all messages between an initiator  $A$  and a responder  $B$ .

**Lemma 1.** *Let  $A$  and  $B$  be respectively the initiator and the responder of a SAKE session. Let  $\delta_{AB}$  be the gap between  $A$  and  $B$  with respect to the evolution of the master keys of both parties. The following conditions always hold:*

1.  $\delta_{AB} \in \{-1, 0, 1\}$ , and
2. *whatever the synchronisation state between  $A$  and  $B$  at the beginning of a session (i.e., whatever  $A$  and  $B$  are synchronised or not), when that session completes in presence of a benign adversary, then*
  - (a)  *$A$  and  $B$  have updated their master keys at least once, and*
  - (b)  *$A$  and  $B$  are synchronised (with respect to their master keys), and*
  - (c)  *$A$  and  $B$  share the same session key.*

In order to prove Lemma 1, we use the following notations. The messages exchanged during a session are numbered in a natural way:



The notation “ $(i_A, i_B)$ ” means that, when the session ends, the last valid message received by  $A$  is message of index  $i_A$ , and the last valid message received by  $B$  is message of index  $i_B$ . We call a  $(i_A, i_B)$ -*session* a session where the last message received by  $A$  is message  $i_A$ , and the last message received by  $B$  is message  $i_B$ . By convention  $i_A = 0$  means that no message has been received by  $A$ .

It may happen that  $A$  send a first message which is not received by  $B$ .  $B$  cannot know if it has missed a first message. But this is of no consequence regarding the synchronisation between  $A$  and  $B$  ( $A$  may simply run the protocol anew). Therefore we do not use the value  $i_B = 0$  (it is equivalent to  $i_B = 5$ ). At initialisation (i.e., before the first run of the protocol),  $(i_A, i_B)$  is set to  $(4, 5)$ . Since  $A$  sends message  $i \in \{3, 5\}$  only upon reception of a valid message  $i - 1$ , and  $B$  sends message  $j \in \{2, 4\}$  only upon reception of a valid message  $j - 1$ , the only possible values for  $(i_A, i_B)$  are as listed in Table 1.

*Proof.* We prove Lemma 1. We first prove item 1.

Let  $c_A$  (resp.  $c_B$ ) be a (virtual) monotonically increasing counter initialised to 0 that follows the evolution of the master keys held by  $A$  (resp.  $B$ ). That is,  $c_A$  (resp.  $c_B$ ) is increased each time the master keys  $K, K'_{j+1}, K'_j, K'_{j-1}$  (resp.  $K, K'$ ) are updated. The parameter  $\delta_{AB}$  corresponds to the gap between  $A$  and  $B$  with respect to the evolution of their master keys, hence  $\delta_{AB} = c_A - c_B$ .

Table 1: Possible values for  $(i_A, i_B)$  in SAKE

$i_A \backslash i_B$	1	3	5
0	✓	✗	✗
2	✓	✓	✗
4	✗	✓	✓

The different possible sessions are listed in Table 1. We prove item 1 by constructing iteratively Table 2b.

Before the first session,  $A$  and  $B$  are synchronised. That is  $\delta_{AB} = c_A - c_B = 0$ , and  $(c_A, c_B) = (i, i)$  (with  $i = 0$ ). Therefore,  $A$  can validate  $\tau_B$  (in message  $m_B$ ) with the same key  $K'_j = K'$  as  $B$ . Hence  $A$  computes  $\delta_{AB} = 0$ , and  $\epsilon = 0$ . Consequently, if one carries out the protocol run starting with  $\delta_{AB} = 0$  and  $\epsilon = 0$ , for each possible value  $(i_A, i_B)$ , one eventually gets the following:

- $(c_A, c_B) = (i, i)$  and  $\delta_{AB} = 0$  after a  $(0, 1)$ -session,
- $(c_A, c_B) = (i + 1, i)$  and  $\delta_{AB} = 1$  after a  $(2, 1)$ -session,
- $(c_A, c_B) = (i + 1, i + 1)$  and  $\delta_{AB} = 0$  after a  $(2, 3)$ -session,
- $(c_A, c_B) = (i + 1, i + 1)$  and  $\delta_{AB} = 0$  after a  $(4, 3)$ -session,
- $(c_A, c_B) = (i + 1, i + 1)$  and  $\delta_{AB} = 0$  after a  $(4, 5)$ -session.

This corresponds to the first column of Tables 2b and 2a. As we can see, the only possible values for  $\delta_{AB}$  after any session are 0 and 1.  $\delta_{AB} = 0$  has already been investigated. Hence, starting with  $\delta_{AB} = 1$  (i.e.,  $(c_A, c_B) = (i + 1, i)$ ), we look for all the values  $\delta_{AB}$  may have when the session ends, considering any possible session.

$(c_A, c_B) = (i + 1, i)$  means that  $A$  is in advance with respect to  $B$ . In such a case,  $A$  succeeds in validating  $\tau_B$  with  $K'_{j-1}$  (and, indeed, finds  $\delta_{AB} = 1$ ). Then  $A$  uses  $\delta_{AB} = 1$  and  $\epsilon = 1$ . If one carries out the protocol run using these two values, one gets:

- $(c_A, c_B) = (i + 1, i)$  and  $\delta_{AB} = 1$  after a  $(0, 1)$ -session,
- $(c_A, c_B) = (i + 1, i)$  and  $\delta_{AB} = 1$  after a  $(2, 1)$ -session,
- $(c_A, c_B) = (i + 1, i + 2)$  and  $\delta_{AB} = -1$  after a  $(2, 3)$ -session,
- $(c_A, c_B) = (i + 2, i + 2)$  and  $\delta_{AB} = 0$  after a  $(4, 3)$ -session,
- $(c_A, c_B) = (i + 2, i + 2)$  and  $\delta_{AB} = 0$  after a  $(4, 5)$ -session.

This corresponds to the second column of Table 2b. This shows that a third value is possible for  $\delta_{AB}$ , which is  $-1$  (i.e.,  $(c_A, c_B) = (i, i + 1)$ ).

Then we restart the protocol with all possible sessions, assuming that  $(c_A, c_B) = (i, i + 1)$  at the beginning of the run. This means that  $A$  is one step late with respect to  $B$ . In such a case,  $A$  succeeds in validating  $\tau_B$  with key  $K'_{j+1}$  (and, indeed, finds  $\delta_{AB} = -1$ ). Then  $A$  uses  $\delta_{AB} = -1$  and  $\epsilon = 0$ . If one carries out the protocol run using these two values, one gets:

- $(c_A, c_B) = (i, i + 1)$  and  $\delta_{AB} = -1$  after a  $(0, 1)$ -session,

- $(c_A, c_B) = (i + 2, i + 1)$  and  $\delta_{AB} = 1$  after a  $(2, 1)$ -session,
- $(c_A, c_B) = (i + 2, i + 2)$  and  $\delta_{AB} = 0$  after a  $(2, 3)$ -session,
- $(c_A, c_B) = (i + 2, i + 2)$  and  $\delta_{AB} = 0$  after a  $(4, 3)$ -session,
- $(c_A, c_B) = (i + 2, i + 2)$  and  $\delta_{AB} = 0$  after a  $(4, 5)$ -session.

We end with three possible values for  $\delta_{AB}$  (third column of Table 2b):  $-1, 0$  and  $1$ , that have already been explored. This proves that, whatever the sequences of sessions, the only possible values for  $\delta_{AB}$  are in  $\{-1, 0, 1\}$ .

Table 2: Possible values for  $\delta_{AB}$  and  $(c_A, c_B)$  among all sequences of sessions in SAKE

(a) Possible values for $(c_A, c_B)$				(b) Possible values for $\delta_{AB}$			
session \ $(c_A, c_B)$	$(i, i)$	$(i + 1, i)$	$(i, i + 1)$	session \ $\delta_{AB}$	$0$	$1$	$-1$
$(0, 1)$	$(i, i)$	$(i + 1, i)$	$(i, i + 1)$	$(0, 1)$	$0$	$1$	$-1$
$(2, 1)$	$(i + 1, i)$	$(i + 1, i)$	$(i + 2, i + 1)$	$(2, 1)$	$1$	$1$	$1$
$(2, 3)$	$(i + 1, i + 1)$	$(i + 1, i + 2)$	$(i + 2, i + 2)$	$(2, 3)$	$0$	$-1$	$0$
$(4, 3)$	$(i + 1, i + 1)$	$(i + 2, i + 2)$	$(i + 2, i + 2)$	$(4, 3)$	$0$	$0$	$0$
$(4, 5)$	$(i + 1, i + 1)$	$(i + 2, i + 2)$	$(i + 2, i + 2)$	$(4, 5)$	$0$	$0$	$0$

Now we prove item 2 of Lemma 1.

We know that  $\delta_{AB} \in \{-1, 0, 1\}$ . For each possible value of  $\delta_{AB}$  at the beginning of the session, the last line of Table 2b indicates the value of that parameter after a correct and complete session (i.e., a  $(4, 5)$ -session). As we can see,  $A$  and  $B$  are always synchronised (i.e.,  $\delta_{AB} = 0$ ) in such a case whatever the value of  $\delta_{AB}$  when the session starts. Furthermore, the session key computation immediately precedes the last update of the derivation master key  $K$ . Hence, when a correct and complete session ends,  $A$  and  $B$  use the same derivation master key  $K$  to compute the session key. Therefore, using the same values  $r_A, r_B$ ,  $A$  and  $B$  compute the same session key.

In addition, Table 2a shows that, whatever the synchronisation state between  $A$  and  $B$  (i.e.,  $c_A$  and  $c_B$ ) at the beginning of the session, after a correct and complete session,  $A$  and  $B$  have updated their internal state at least once (as the last line of the table, corresponding to a  $(4, 5)$ -session, indicates).  $\square$

The diagram depicted by Figure 4 represents all the possible sequences of sessions with SAKE.

## 4.2 Security of SAKE

In order to prove that the protocol SAKE is a secure AKE protocol, we use the execution environment described in Section 2.1. We define the partnering between two instances with the notion of *matching conversations*. That is, we




Fig. 4: Diagram of SAKE. The circled values correspond to the gap  $\delta_{AB}$ , and each edge to a  $(i_A, i_B)$ -session.

define  $\text{sid}$  to be the transcript, in chronological order, of all the (valid) messages sent and received by an instance during the key exchange, but, possibly, the last one (see Definition 8). Furthermore, we choose the function  $\text{update}$  to be a PRF, that is  $\text{update} : K \mapsto \text{PRF}(K, x)$  for some (constant) value  $x$ .

**Theorem 1.** *The protocol SAKE is a secure AKE protocol, and for any probabilistic polynomial time adversary  $\mathcal{A}$  in the AKE security experiment against SAKE*

$$\begin{aligned} \text{adv}_{SAKE}^{\text{ent-auth}}(\mathcal{A}) &\leq nq \left( (nq - 1)2^{-\lambda} + (q + 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + 2\text{adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{C}) \right) \\ \text{adv}_{SAKE}^{\text{key-ind}}(\mathcal{A}) &\leq nq \left( (q - 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + \text{adv}_{\text{KDF}}^{\text{prf}}(\mathcal{D}) \right) + \text{adv}_{SAKE}^{\text{ent-auth}}(\mathcal{A}) \end{aligned}$$

where  $n$  is the number of parties,  $q$  the number of instances (sessions) per party,  $\lambda$  the size of the pseudo-random values  $(r_A, r_B)$ , and  $\mathcal{B}$  is an adversary against the PRF-security of  $\text{update}$ ,  $\mathcal{C}$  an adversary against the SUF-CMA-security of MAC, and  $\mathcal{D}$  an adversary against the PRF-security of KDF.

We give a sketch proof of Theorem 1. The complete proof is given in Appendix A.

*Proof.* We proceed through a sequence of games between a challenger and an adversary  $\mathcal{A}$ . First we consider the entity authentication experiment. We use the following hops.

*Game 0* corresponds to the entity authentication security experiment described in Section 2.1.

In *Game 1*, the challenger aborts if there exists an instance that chooses a

random value  $r_A$  or  $r_B$  that is not unique. There is at most  $n \times q$  random values, each uniformly drawn at random in  $\{0, 1\}^\lambda$ . Therefore, the two games are equivalent up to a collision term  $\frac{nq(nq-1)}{2^\lambda}$ .

In *Game 2*, the adversary targets a single instance  $\pi$ . This is equivalent to guessing the targeted instance, hence implies a security loss equal to  $\frac{1}{nq}$ .

In *Game 3*, the challenger aborts if the initiator (resp. responder) instance  $\pi$  ever receives a message  $m_B$  (resp.  $m_A$ ), but no instance having a matching conversation to  $\pi$  has output that message. Here, we reduce the probability of this event to the security of the functions `MAC` (used to compute the MAC tags) and `update` (used to update the MAC key). Hence we evaluate the overall security loss.

By assumption, the genuine value of the MAC key  $K'$  (i.e., the value used during the first session between two same parties) is uniformly chosen at random. Then  $K'$  (and  $K$ ) is updated throughout the session with the function `update`. If  $K'$  is random, we can rely on the pseudo-randomness of `update`( $\cdot$ ) = `PRF`( $\cdot$ ,  $\cdot$ ). In turn, since `PRF`( $K'$ ,  $\cdot$ ) can be replaced with a truly random function, its output (updated  $K'$ ) is random. Therefore, one can rely upon the pseudo-randomness of the function `update` keyed with this new value  $K'$ , and so forth. Each transition (i.e., each update of  $K'$ ) implies a loss equal to  $\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B})$  corresponding to the ability of an adversary  $\mathcal{B}$  to distinguish `update` from a random function.

On the initiator side, if the initiator party  $P_i$  is synchronised with the responder  $P_j$  ( $\delta_{AB} = 0$ ),  $P_i$  updates its master keys once (upon reception of  $m_B$ ). If  $P_i$  is in advance ( $\delta_{AB} = 1$ ), it updates its keys at most once (if a valid message  $\tau'_B$  is received). If  $P_i$  is late ( $\delta_{AB} = -1$ ), it updates its keys twice. Yet, in that case,  $P_i$  did not update its keys during the previous session. Therefore, on average,  $P_i$  updates its keys at most once per session. Hence, when the  $u$ -th session starts,  $P_i$  has updated its keys at most  $u - 1$  times on average, and, upon reception of  $\tau'_B$ ,  $P_i$  updates the keys at most two times.

This is similar regarding the responder  $P_j$ . A responder instance at  $P_j$  accepts only if the two messages  $m_A$  and  $\tau'_A$  are valid. Upon reception of a valid message  $m_A$ , the keys are updated once ( $\epsilon = 0$ ) or twice ( $\epsilon = 1$ ). In the latter case, the keys have not been updated during the previous session. This means that the keys are updated on average at most once per session. Therefore, when the  $u$ -th session starts,  $P_j$  has updated its keys at most  $u - 1$  times on average, and, upon reception of  $m_A$ , the keys are updated at most two times.

Since there is at most  $q$  sessions per party, this implies a security loss at most  $(q - 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B})$ . Then we reduce the probability of the adversary  $\mathcal{A}$  to win this game to the ability of an adversary  $\mathcal{C}$  to forge a valid tag  $\tau_B$  (resp.  $\tau_A$ ). Hence an additional factor  $\text{adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{C})$ .

In *Game 4*, the challenger aborts if the targeted instance  $\pi$  ever receives a valid message  $\tau'_B$  (resp.  $\tau'_A$ ), but no instance having a matching conversation to  $\pi$  has output that message. As explained above, between the reception of message  $m_B$  (resp.  $m_A$ ) and message  $\tau'_B$  (resp.  $\tau'_A$ ), the MAC key  $K'$  is updated at most twice. Here, we reduce the probability of the adversary to win this game to the security of the MAC function used to compute the message  $\tau'_B$  (resp.  $\tau'_A$ ). In turn

we rely on the randomness of the MAC key, hence on the security of the function `update` used to update the MAC key  $K'$ . Therefore this implies a security loss at most  $2\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B})$ . Then we reduce the probability of the adversary  $\mathcal{A}$  to win this game to the ability of an adversary  $\mathcal{C}$  to forge a valid tag  $\tau'_B$  (resp.  $\tau'_A$ ). Hence an additional factor  $\text{adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{C})$ .

To that point, the only way for the adversary to make  $\pi$  accept maliciously is to send a valid message  $\tau'_B$  (resp.  $\tau'_A$ ) different from all the messages sent by all the instances. However, in such a case, the challenger aborts. Hence the adversary has no chance to win.

Now we prove the key indistinguishability security.

*Game 0* corresponds to the key indistinguishability experiment described in Section 2.1.

In *Game 1*, the challenger aborts the experiment and chooses  $b' \in \{0, 1\}$  uniformly at random if there exists an instance that accepts maliciously. In other words, in this game we make the same modifications as in the games performed during the entity authentication proof. Hence a loss  $\text{adv}_{\text{SAKE}}^{\text{ent-auth}}(\mathcal{A})$ .

In *Game 2*, the adversary interacts with a single instance. This is equivalent to guessing the targeted instance, hence implies a security loss equal to  $\frac{1}{nq}$ .

In *Game 3*, we reduce the advantage of the adversary to win this game to the security of the function KDF used to compute the session key. That is, we rely upon the pseudo-randomness of the KDF function. This is possible if the key  $K$  used to compute the session key is random. The genuine value of  $K$  is uniformly chosen at random by assumption. Then  $K$  is updated with `update` at most once per session on average. Therefore, when the  $u$ -th session starts,  $K$  has been updated at most  $u - 1$  times already. Therefore we must take into account the successive losses due to the key update with respect to the pseudo-randomness of `update`. Since there is at most  $q$  sessions per party, this loss is at most  $(q - 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B})$ . Then we reduce the ability of  $\mathcal{A}$  to win to the security of KDF. Hence an additional factor  $\text{adv}_{\text{KDF}}^{\text{prf}}(\mathcal{D})$ , corresponding to the advantage of an adversary  $\mathcal{D}$  to distinguish KDF from a random function. To that point, the adversary can do no better than guess.  $\square$

## 5 Comparison with the DH paradigm

The protocol SAKE is based on shared master keys and apply symmetric-key functions only. In particular it does not require the application of any kind of DH-like scheme. Yet it provides a strong form of forward secrecy. Despite this result, our protocol differs from a DH scheme in several ways beyond the intrinsic distinction between public-key and symmetric-key cryptography.<sup>11</sup>

*Concurrent executions.* Our protocol does not allow parallel executions. Indeed, since it is based on shared evolving symmetric keys, running multiple instances in

<sup>11</sup> The same holds regarding the variant SAKE-AM.

parallel may cause some sessions to abort. A way to relax this restriction is that each party use separate master keys for concurrent executions. On the contrary, the DH scheme allows an (virtually) unlimited number of parallel executions.

*KCI attacks.* The ephemeral DH scheme (when using safe parameters) is resistant against KCI attacks [14], whereas our protocol is not (due to the dependency between the (updated) master keys).<sup>12</sup> Moreover if an adversary succeeds in getting the key  $K'$  (or  $K'_j$ ), she can compute the subsequent key (corresponding to  $K'_{j+1}$ ). Hence the adversary can forge a message  $m_B$  in SAKE that brings the initiator to update its master keys twice consecutively. Therefore, that party is desynchronised with respect to an honest partner, with no possibility to resynchronise.

Note that KCI attacks affect also the static DH scheme (when a party uses a fixed DH share, whereas the other generates a fresh ephemeral one [26]).

Another consequence of the dependency of the master keys in SAKE, is that once the keys are revealed, an adversary can passively compromise all subsequent session keys. This is not the case in general with ephemeral DH. Yet, this is also true regarding non-DH public-key protocols (e.g., TLS-RSA), but also ephemeral DH (in some pathological cases) when small, fixed public parameters are used [5].

*Post-quantum setting.* Now a probable benefit of our protocol compared to the DH scheme is that, since it is based on symmetric-key functions, it can likely survive in a post-quantum world (with a suitable choice of the primitives, in particular whose key length is large enough [24]). On the contrary, the DH scheme is known to be insecure in such a context [30, 35, 36]. Yet, we observe that there exists a post-quantum variant of the original DH scheme [19, 29], but it is based on larger parameters and heavier computations than SAKE. Moreover this post-quantum variant does not provide entity authentication.

*Computations.* The DH scheme implies heavier computations (modular exponentiations, elliptic curve point multiplication) than SAKE which is solely built on symmetric-key functions. In practice, SAKE is likely more suitable to be implemented on constrained devices which have limited computational (and communication) capabilities.

## 6 Conclusion

We have described SAKE, an authenticated key exchange protocol in the symmetric-key setting. Although this protocol is solely based on symmetric-key algorithms, it provides perfect forward secrecy without requiring any additional procedure (e.g., resynchronisation phase) or functionality (e.g., shared clock).

<sup>12</sup> When a party  $P_i$ 's long-term secret key is disclosed, an adversary can impersonate  $P_i$  to other parties. In the same context, a key compromise impersonation (KCI) allows the adversary to impersonate other parties to  $P_i$ .

The underlying idea is to make the shared master keys evolve. We solve the synchronisation problem that stems from this evolving principle with an elegant and efficient solution.

SAKE guarantees that, whatever the synchronisation state of the involved parties prior to the session, both parties share a new session key, and their internal state is updated and synchronised, once a correct session is complete: SAKE is self-synchronising. As in the public-key setting, our protocol allows an (virtually) unlimited number of sessions. Furthermore, we prove that SAKE is sound, and provide a formal proof of its security in a strong model.

Finally, we describe SAKE-AM, a variant of our protocol, which, used in conjunction with SAKE, results in an implementation that gathers all the aforementioned properties (starting with forward secrecy). This implementation allows any party to be initiator or responder of a session, such that the smallest amount of calculation is always done by the same party. This is particularly convenient in the context of IoT where a set of (low-resource) end-devices communicates with a back-end server.

To the best of our knowledge, this is the first protocol with perfect forward secrecy in the symmetric-key setting that is comparable to the DH scheme, beyond the intrinsic distinction between public-key and symmetric-key cryptography.

## References

1. Signal, <https://signal.org/>
2. 3rd Generation Partnership Project: Technical Specifications 33, available via <http://www.3gpp.org/DynaReport/33-series.htm>
3. 3rd Generation Partnership Project: Technical Specifications 35, available via <http://www.3gpp.org/DynaReport/35-series.htm>
4. Abdalla, M., Bellare, M.: Increasing the lifetime of a key: a comparative analysis of the security of re-keying techniques. In: Okamoto, T. (ed.) *Advances in Cryptology – ASIACRYPT 2000*. Lecture Notes in Computer Science, vol. 1976, pp. 546–559. Springer, Heidelberg, Germany, Kyoto, Japan (Dec 3–7, 2000). [https://doi.org/10.1007/3-540-44448-3\\_42](https://doi.org/10.1007/3-540-44448-3_42)
5. Adrian, D., Bhargavan, K., Durumeric, Z., Gaudry, P., Green, M., Halderman, J.A., Heninger, N., Springall, D., Thomé, E., Valenta, L., VanderSloot, B., Wustrow, E., Zanella-Béguelin, S., Zimmermann, P.: Imperfect forward secrecy: How Diffie-Hellman fails in practice. In: Ray, I., Li, N., Kruegel, C. (eds.) *ACM CCS 2015: 22nd Conference on Computer and Communications Security*. pp. 5–17. ACM Press, Denver, CO, USA (Oct 12–16, 2015). <https://doi.org/10.1145/2810103.2813707>
6. Alwen, J., Coretti, S., Dodis, Y.: The Double Ratchet: Security Notions, Proofs, and Modularization for the Signal Protocol. *Cryptology ePrint Archive*, Report 2018/1037 (2018), <https://eprint.iacr.org/2018/1037>
7. American National Standards Institute: ANSI X9.24-1:2009 Retail Financial Services Symmetric Key Management Part 1: Using Symmetric Techniques (2009)
8. Bellare, M., Desai, A., Jokipii, E., Rogaway, P.: A concrete security treatment of symmetric encryption. In: *38th Annual Symposium on Foundations of Computer Science*. pp. 394–403. IEEE Computer Society Press, Miami Beach, Florida (Oct 19–22, 1997). <https://doi.org/10.1109/SFCS.1997.646128>



9. Bellare, M., Miner, S.K.: A forward-secure digital signature scheme. In: Wiener, M.J. (ed.) *Advances in Cryptology – CRYPTO’99*. Lecture Notes in Computer Science, vol. 1666, pp. 431–448. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 15–19, 1999). [https://doi.org/10.1007/3-540-48405-1\\_28](https://doi.org/10.1007/3-540-48405-1_28)
10. Bellare, M., Namprempre, C.: Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. *Journal of Cryptology* **21**(4), 469–491 (Oct 2008). <https://doi.org/10.1007/s00145-008-9026-x>
11. Bellare, M., Rogaway, P.: Entity authentication and key distribution. In: Stinson, D.R. (ed.) *Advances in Cryptology – CRYPTO’93*. Lecture Notes in Computer Science, vol. 773, pp. 232–249. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 22–26, 1994). [https://doi.org/10.1007/3-540-48329-2\\_21](https://doi.org/10.1007/3-540-48329-2_21)
12. Bellare, M., Rogaway, P.: Code-based game-playing proofs and the security of triple encryption. *Cryptology ePrint Archive*, Report 2004/331 (2004), <http://eprint.iacr.org/2004/331>
13. Bellare, M., Yee, B.S.: Forward-security in private-key cryptography. In: Joye, M. (ed.) *Topics in Cryptology – CT-RSA 2003*. Lecture Notes in Computer Science, vol. 2612, pp. 1–18. Springer, Heidelberg, Germany, San Francisco, CA, USA (Apr 13–17, 2003). [https://doi.org/10.1007/3-540-36563-X\\_1](https://doi.org/10.1007/3-540-36563-X_1)
14. Blake-Wilson, S., Johnson, D., Menezes, A.: Key agreement protocols and their security analysis. In: Darnell, M. (ed.) *6th IMA International Conference on Cryptography and Coding*. Lecture Notes in Computer Science, vol. 1355, pp. 30–45. Springer, Heidelberg, Germany, Cirencester, UK (Dec 17–19, 1997)
15. Boyd, C., Mathuria, A.: *Protocols for Authentication and Key Establishment*. Information Security and Cryptography, Springer (2003)
16. Brier, E., Peyrin, T.: A forward-secure symmetric-key derivation protocol - how to improve classical DUKPT. In: Abe, M. (ed.) *Advances in Cryptology – ASIACRYPT 2010*. Lecture Notes in Computer Science, vol. 6477, pp. 250–267. Springer, Heidelberg, Germany, Singapore (Dec 5–9, 2010). [https://doi.org/10.1007/978-3-642-17373-8\\_15](https://doi.org/10.1007/978-3-642-17373-8_15)
17. Brzuska, C., Jacobsen, H., Stebila, D.: Safely exporting keys from secure channels: On the security of EAP-TLS and TLS key exporters. In: Fischlin, M., Coron, J.S. (eds.) *Advances in Cryptology – EUROCRYPT 2016, Part I*. Lecture Notes in Computer Science, vol. 9665, pp. 670–698. Springer, Heidelberg, Germany, Vienna, Austria (May 8–12, 2016). [https://doi.org/10.1007/978-3-662-49890-3\\_26](https://doi.org/10.1007/978-3-662-49890-3_26)
18. Cohn-Gordon, K., Cremers, C., Dowling, B., Garratt, L., Stebila, D.: A Formal Security Analysis of the Signal Messaging Protocol. In: *2017 IEEE European Symposium on Security and Privacy (EuroS&P)*. pp. 451–466. IEEE (April 2017). <https://doi.org/10.1109/EuroSP.2017.27>
19. Costello, C., Longa, P., Naehrig, M.: Efficient algorithms for supersingular isogeny Diffie-Hellman. In: Robshaw, M., Katz, J. (eds.) *Advances in Cryptology – CRYPTO 2016, Part I*. Lecture Notes in Computer Science, vol. 9814, pp. 572–601. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 14–18, 2016). [https://doi.org/10.1007/978-3-662-53018-4\\_21](https://doi.org/10.1007/978-3-662-53018-4_21)
20. Diffie, W., Hellman, M.E.: New directions in cryptography. *IEEE Transactions on Information Theory* **22**(6), 644–654 (1976)
21. Diffie, W., van Oorschot, P.C., Wiener, M.J.: Authentication and authenticated key exchanges. *Designs, Codes and Cryptography* **2**(2), 107–125 (Jun 1992)
22. Dousti, M.S., Jalili, R.: FORSAKES: A forward-secure authenticated key exchange protocol based on symmetric key-evolving schemes. *Cryptology ePrint Archive*, Report 2014/123 (2014), <http://eprint.iacr.org/2014/123>

23. GlobalPlatform: GlobalPlatform – Card Specification – Version 2.3.1 (March 2018), reference GPC.SPE\_034. Available via <https://www.globalplatform.org/specificationscard.asp>
24. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: 28th Annual ACM Symposium on Theory of Computing. pp. 212–219. ACM Press, Philadelphia, PA, USA (May 22–24, 1996). <https://doi.org/10.1145/237814.237866>
25. Günther, C.G.: An identity-based key-exchange protocol. In: Quisquater, J.J., Vandewalle, J. (eds.) Advances in Cryptology – EUROCRYPT’89. Lecture Notes in Computer Science, vol. 434, pp. 29–37. Springer, Heidelberg, Germany, Houthalen, Belgium (Apr 10–13, 1990). [https://doi.org/10.1007/3-540-46885-4\\_5](https://doi.org/10.1007/3-540-46885-4_5)
26. Hlauschek, C., Gruber, M., Fankhauser, F., Schanes, C.: Prying Open Pandora’s Box: KCI Attacks Against TLS. In: Proceedings of the 9th USENIX Conference on Offensive Technologies. WOOT’15, USENIX Association (2015)
27. International Organization for Standardization: ISO/IEC 11770-2 – Information technology – Security techniques – Key Management – Part 2: Mechanisms using Symmetric Techniques (2008)
28. Jager, T., Kohlar, F., Schäge, S., Schwenk, J.: On the security of TLS-DHE in the standard model. Cryptology ePrint Archive, Report 2011/219 (2011), <http://eprint.iacr.org/2011/219>
29. Jao, D., De Feo, L.: Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. In: Yang, B.Y. (ed.) Post-Quantum Cryptography - 4th International Workshop, PQCrypto 2011. pp. 19–34. Springer, Heidelberg, Germany, Taipei, Taiwan (Nov 29 – Dec 2 2011). [https://doi.org/10.1007/978-3-642-25405-5\\_2](https://doi.org/10.1007/978-3-642-25405-5_2)
30. Kaliski Jr., B.S.: A Quantum “Magic Box” for the Discrete Logarithm Problem. Cryptology ePrint Archive, Report 2017/745 (2017), <https://eprint.iacr.org/2017/745>
31. Le, T.V., Burmester, M., de Medeiros, B.: Universally composable and forward-secure RFID authentication and authenticated key exchange. In: Bao, F., Miller, S. (eds.) ASIACCS 07: 2nd ACM Symposium on Information, Computer and Communications Security. pp. 242–252. ACM Press, Singapore (Mar 20–22, 2007)
32. Park, T., Shin, K.G.: LiSP: A Lightweight Security Protocol for Wireless Sensor Networks. ACM Trans. Embed. Comput. Syst. **3**(3), 634–660 (August 2004)
33. Perrig, A., Szewczyk, R., Tygar, J., Wen, V., Culler, D.E.: SPINS: Security Protocols for Sensor Networks. Wireless Networks **8**(5), 521–534 (September 2002)
34. Perrin, T., Marlinspike, M.: The Double Ratchet Algorithm (2016), <https://signal.org/docs/specifications/doublerratchet/>, Revision 1, 20/11/2016
35. Proos, J., Zalka, C.: Shor’s Discrete Logarithm Quantum Algorithm for Elliptic Curves. Quantum Info. Comput. **3**(4), 317–344 (Jul 2003)
36. Shor, P.W.: Algorithms for quantum computation: Discrete logarithms and factoring. In: 35th Annual Symposium on Foundations of Computer Science. pp. 124–134. IEEE Computer Society Press, Santa Fe, NM, USA (Nov 20–22, 1994). <https://doi.org/10.1109/SFCS.1994.365700>
37. Shoup, V.: Sequences of games: a tool for taming complexity in security proofs. Cryptology ePrint Archive, Report 2004/332 (2004), <http://eprint.iacr.org/2004/332>
38. Sornin, N., Luis, M., Eirich, T., Kramp, T.: LoRaWAN Specification (July 2016), LoRa Alliance, version 1.0
39. ZigBee Alliance: ZigBee specification, available via <http://www.zigbee.org/download/standards-zigbee-specification/>

## A Extended Security Proof for SAKE

In this section, we give a proof of Theorem 1. That is, SAKE is a secure AKE protocol according to Definition 5.

*Proof.* In order for an initiator instance  $\pi_i^s$  at some party  $P_i$  to accept, two valid messages (i.e., with valid MAC tags) must be received by  $\pi_i^s$  ( $m_B$  and  $\tau'_B$ ). We reduce the security of the MAC function to the (in)ability to forge a valid output. Therefore we use the fact that the key  $K'$  is random. By assumption, the genuine value of  $K'$  (i.e., the value used during the first session between two same parties) is uniformly chosen at random. Yet  $K'$  (and  $K$ ) is updated throughout the session with the function `update`. If  $K'$  is random, we can rely on the pseudo-randomness of `update` =  $\text{PRF}(\cdot, \cdot)$ . In turn, since  $\text{PRF}(K', \cdot)$  can be replaced with a truly random function, its output (updated  $K'$ ) is random. Therefore, one can rely upon the pseudo-randomness of the function `update` keyed with this new value  $K'$ , and so forth. Each transition (i.e., each update of  $K'$ ) implies a loss equal to  $\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B})$  corresponding to the ability of an adversary  $\mathcal{B}$  to distinguish `update` from a random function.

If  $P_i$  is synchronised with the responder ( $\delta_{AB} = 0$ ),  $P_i$  updates its master keys once (upon reception of  $m_B$ ). If  $P_i$  is in advance ( $\delta_{AB} = 1$ ), it updates its keys at most once (if a valid message  $\tau'_B$  is received). If  $P_i$  is late ( $\delta_{AB} = -1$ ), it updates its keys twice. Yet, in that case,  $P_i$  did not update its keys during the previous session. Therefore, on average,  $P_i$  updates its keys at most once per session. Hence, when the  $u$ -th session starts,  $P_i$  has updated its keys at most  $u - 1$  times on average, and, upon reception of  $\tau'_B$ ,  $P_i$  updates the keys at most two times.

This is similar regarding the responder. A responder instance  $\pi_j^t$  at some party  $P_j$  accepts only if the two messages  $m_A$  and  $\tau'_A$  are valid. Upon reception of a valid message  $m_A$ , the keys are updated once ( $\epsilon = 0$ ) or twice ( $\epsilon = 1$ ). In the latter case, the keys have not been updated during the previous session. This means that the keys are updated on average at most once per session. Therefore, when the  $u$ -th session starts,  $P_j$  has updated its keys at most  $u - 1$  times on average, and, upon reception of  $m_A$ , the keys are updated at most two times.

We can now proceed with the proof. We proceed through a sequence of games [12, 37], where each consecutive game aims at reducing the challenger's dependency on the functions MAC, `update` and KDF. We first prove the entity authentication security. Let  $E_i$  be the event that the adversary win the entity authentication experiment in Game  $i$ .

*Game 0.* This game corresponds to the entity authentication security experiment described in Section 2.1. Therefore

$$\Pr[E_0] = \text{adv}_{SAKE}^{\text{ent-auth}}(\mathcal{A})$$

*Game 1.* In this game, we add an abort rule. The challenger aborts if there exists any instance that chooses a random value  $r_A$  or  $r_B$  that is not unique. There

is at most  $n \times q$  random values, each uniformly drawn at random in  $\{0, 1\}^\lambda$ . Therefore the probability that at least two random values be equal is at most  $\frac{nq(nq-1)}{2^\lambda}$ . Hence

$$\Pr[E_0] \leq \Pr[E_1] + \frac{nq(nq-1)}{2^\lambda}$$

*Game 2.* In this game, we add an abort rule. The challenger tries to guess which instance will be the first to accept maliciously. If the guess is wrong, the game is aborted. The number of instances is at most  $nq$ . Therefore

$$\Pr[E_2] = \Pr[E_1] \times \frac{1}{nq}$$

*Game 3.* Let  $\pi$  be the instance targeted by the adversary. In this game, we add an abort rule. The challenger aborts the experiment if  $\pi$ , behaving as an initiator (resp. responder) instance, ever receives a valid message  $m_B$  (resp.  $m_A$ ) but no instance having a matching conversation to  $\pi$  has output that message. We reduce the probability of this event to the security of the functions **MAC** and **update**. As explained above, when the  $u$ -th session starts, the master keys have been updated at most  $u - 1$  times already. The genuine value of  $K'$  is uniformly chosen at random. In order to be able to replace, during the current session, the key used to compute the MAC tag in  $m_A$  (resp.  $m_B$ ) with a random value, one must rely upon the pseudo-randomness of the function **update** that outputs (the new value of)  $K'$ . In turn, this relies upon the (previous) key  $K'$  being random (and on the pseudo-randomness of **update**). Therefore, in order to replace  $K'$  with a random value one must take into account the successive losses  $\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B})$ , each corresponding to the ability of an adversary  $\mathcal{B}$  to distinguish the function **update** (keyed with a different key  $K'$ ) from a random function. Since there is at most  $q$  sessions, this loss is at most  $(q - 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B})$ . Then we reduce the probability of the adversary  $\mathcal{A}$  to win this game to the ability of an adversary  $\mathcal{C}$  to forge a valid tag  $\tau_B$  (resp.  $\tau_A$ ).

Therefore, we replace each function  $\text{update}(K') = \text{PRF}(K', x)$  (keyed with a different key  $K'$  throughout the, at most,  $q - 1$  successive sessions established, prior to that current session, by the same party that owns  $\pi$ ) with truly random functions  $F_0^{\text{update}}, \dots, F_{q-2}^{\text{update}}$ . Moreover, if an instance uses the same key  $K' = K'_i$ ,  $0 \leq i < q - 1$ , to key **update**, then we replace **update** with the corresponding random function  $F_i^{\text{update}}$ . Since, to that point, the key  $K' = K'_{q-1}$  used to compute the authentication tag  $\tau_B$  (resp.  $\tau_A$ ) is random, we reduce the ability of  $\mathcal{A}$  to win to the security of the MAC function. Hence

$$\Pr[E_2] \leq \Pr[E_3] + (q - 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + \text{adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{C})$$

*Game 4.* In this game, we add an abort rule. The challenger aborts the experiment if  $\pi$  ever receives a valid message  $\tau'_B$  (resp.  $\tau'_A$ ), but no instance having a matching conversation to  $\pi$  has output that message. Between the message  $m_B$  (resp.  $m_A$ ) being received by  $\pi$ , and the message  $\tau'_B$  (resp.  $\tau'_A$ ) being received by

$\pi$ , the master keys are updated at most twice. We reduce the probability of the adversary to win this game to the security of the MAC function used to compute the message  $\tau'_B$  (resp.  $\tau'_A$ ). In turn we must rely on the randomness of the MAC key, hence on the security of the function `update` used to update the MAC key  $K'$  (recall that, due to Game 3, the current key  $K'$  is random). Therefore

$$\Pr[E_3] \leq \Pr[E_4] + 2\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + \text{adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{C})$$

To that point, the only way for the adversary to make  $\pi$  accept maliciously is to send a valid message  $\tau'_B$  (resp.  $\tau'_A$ ) different from all the messages sent by all the instances. However, in such a case, the challenger aborts. Therefore

$$\Pr[E_4] = 0$$

Collecting all the probabilities from Game 0 to Game 4, we have that

$$\begin{aligned} \text{adv}_{SAKE}^{\text{ent-auth}}(\mathcal{A}) &= \Pr[E_0] \\ &\leq \frac{nq(nq-1)}{2^\lambda} + \Pr[E_1] \\ &= \frac{nq(nq-1)}{2^\lambda} + nq \times \Pr[E_2] \\ &\leq \frac{nq(nq-1)}{2^\lambda} + nq \left( \Pr[E_3] + (q-1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + \text{adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{C}) \right) \\ &\leq \frac{nq(nq-1)}{2^\lambda} + nq \left( \Pr[E_4] + (q+1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + 2\text{adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{C}) \right) \\ &= \frac{nq(nq-1)}{2^\lambda} + nq \left( (q+1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + 2\text{adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{C}) \right) \\ &= nq \left( (nq-1)2^{-\lambda} + (q+1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + 2\text{adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{C}) \right) \end{aligned}$$

Now we prove the key indistinguishability security. Let  $E'_i$  be the event that an adversary win the key indistinguishability experiment in Game  $i$ , and  $\text{adv}_i = \Pr[E'_i] - \frac{1}{2}$ .

*Game 0.* This game corresponds to the key indistinguishability experiment described in Section 2.1. Therefore

$$\Pr[E'_0] = \frac{1}{2} + \text{adv}_{SAKE}^{\text{key-ind}}(\mathcal{A}) = \frac{1}{2} + \text{adv}_0$$

*Game 1.* In this game, we add an abort rule. The challenger aborts the experiment and chooses  $b' \in \{0, 1\}$  uniformly at random if there exists an instance that accepts maliciously. In other words, in this game we make the same modifications as in the games performed during the entity authentication proof. Hence

$$\text{adv}_0 \leq \text{adv}_1 + \text{adv}_{SAKE}^{\text{ent-auth}}(\mathcal{A})$$

*Game 2.* In this game, we add an abort rule. The challenger tries to guess which instance is targeted by the adversary. If the guess is wrong, the game is aborted. The number of instances is at most  $nq$ . Therefore

$$\text{adv}_2 = \text{adv}_1 \times \frac{1}{nq}$$

*Game 3.* Let  $\pi$  be the instance targeted by the adversary. We reduce the advantage of the adversary to win this game to the security of the function KDF used to compute the session key. That is, we rely upon the pseudo-randomness of the KDF function. This is possible if the key  $K$  is random. The genuine value of  $K$  is uniformly chosen at random by assumption. Then  $K$  is updated with `update` at most once per session on average. Therefore, when the  $u$ -th session starts,  $K$  has been updated at most  $u - 1$  times already. Therefore we must take into account the successive losses due to the key update with respect to the pseudo-randomness of `update`. Since there is at most  $q$  sessions per party (i.e., per original key  $K$ ), this loss is at most  $(q - 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B})$ . Hence we replace each function `update`( $K$ ) =  $\text{PRF}(K, x)$  (keyed with a different key  $K$  throughout the, at most,  $q - 1$  successive sessions established, prior to that current session, by the same party that owns  $\pi$ ) with truly random functions  $\mathbf{G}_0^{\text{update}}, \dots, \mathbf{G}_{q-2}^{\text{update}}$ . Moreover, if an instance uses the same key  $K = K_i$ ,  $0 \leq i < q - 1$ , to key `update`, then we replace `update` with the corresponding random function  $\mathbf{G}_i^{\text{update}}$ . Since, to that point, the key  $K = K_{q-1}$  used to compute the session key is random, we reduce the ability of  $\mathcal{A}$  to win to the security of KDF. Therefore

$$\text{adv}_2 \leq \text{adv}_3 + (q - 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + \text{adv}_{\text{KDF}}^{\text{prf}}(\mathcal{D})$$

To that point the session key is random, therefore the adversary has no advantage in guessing whether  $\pi.b = b'$ . That is

$$\text{adv}_3 = 0$$

Collecting all the probabilities from Game 0 to Game 3, we have that

$$\begin{aligned} \text{adv}_{\text{SAKE}}^{\text{key-ind}}(\mathcal{A}) &= \text{adv}_0 \\ &\leq \text{adv}_{\text{SAKE}}^{\text{ent-auth}}(\mathcal{A}) + \text{adv}_1 \\ &= \text{adv}_{\text{SAKE}}^{\text{ent-auth}}(\mathcal{A}) + nq \times \text{adv}_2 \\ &\leq \text{adv}_{\text{SAKE}}^{\text{ent-auth}}(\mathcal{A}) + nq \left( \text{adv}_3 + (q - 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + \text{adv}_{\text{KDF}}^{\text{prf}}(\mathcal{D}) \right) \\ &= \text{adv}_{\text{SAKE}}^{\text{ent-auth}}(\mathcal{A}) + nq \left( (q - 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + \text{adv}_{\text{KDF}}^{\text{prf}}(\mathcal{D}) \right) \end{aligned}$$

□

## B Soundness and Security for SAKE-AM

In this section we prove the soundness and security of the protocol SAKE-AM. Figure 5 depicts the protocol. For the sake of clarity, we use the following notations:

- kdf corresponds to:  $sk \leftarrow \text{KDF}(K, f(r_A, r_B))$
- $\text{upd}'_A$  corresponds to
  1.  $K \leftarrow \text{update}(K)$
  2.  $K' \leftarrow \text{update}(K')$
- $\text{upd}'_B$  corresponds to
  1.  $K \leftarrow \text{update}(K)$
  2.  $K'_{j-1} \leftarrow K'_j$
  3.  $K'_j \leftarrow K'_{j+1}$
  4.  $K'_{j+1} \leftarrow \text{update}(K'_{j+1})$

### B.1 Soundness of SAKE-AM

**Lemma 2.** *Let  $A$  and  $B$  be respectively the initiator and the responder of a SAKE-AM session. Let  $\delta_{BA}$  be the gap between  $B$  and  $A$  with respect to the evolution of the master keys of both parties. The following conditions always hold:*

1.  $\delta_{BA} \in \{-1, 0, 1\}$ , and
2. *whatever the synchronisation state of  $A$  and  $B$  at the beginning of a session (i.e., whatever  $A$  and  $B$  are synchronised or not), when that session completes in presence of a benign adversary, then*
  - (a)  *$A$  and  $B$  have updated their master keys at least once, and*
  - (b)  *$A$  and  $B$  are synchronised (with respect to their master keys), and*
  - (c)  *$A$  and  $B$  share the same session key.*

We prove Lemma 2. We use again the (virtual) monotonically increasing counters  $c_A$  and  $c_B$  maintained respectively by  $A$  and  $B$ . Counter  $c_A$  (resp.  $c_B$ ) is incremented each time  $A$  (resp.  $B$ ) updates its master keys.

*Proof.* We first prove item 1 of Lemma 2.

The different possible sessions are listed in Table 3. We prove item 1 by constructing iteratively Table 4b.

Table 3: Possible values for  $(i_A, i_B)$  in SAKE-AM

	$i_B$	1	3
$i_A$		✓	✗
	0	✓	✗
	2	✓	✓
	4	✗	✓

Before the first session,  $A$  and  $B$  are synchronised. That is  $\delta_{BA} = c_B - c_A = 0$ , or  $(c_B, c_A) = (i, i)$  (with  $i = 0$ ). Therefore,  $B$  can validate  $\tau_A$  with the same key  $K'_j = K'$  as  $A$ . Hence  $\delta_{BA} = 0$ , and  $\epsilon = 0$ . Consequently, if one carries out the protocol run with  $\delta_{BA} = 0$  and  $\epsilon = 0$ , for each possible value  $(i_A, i_B)$ , one eventually gets the following:

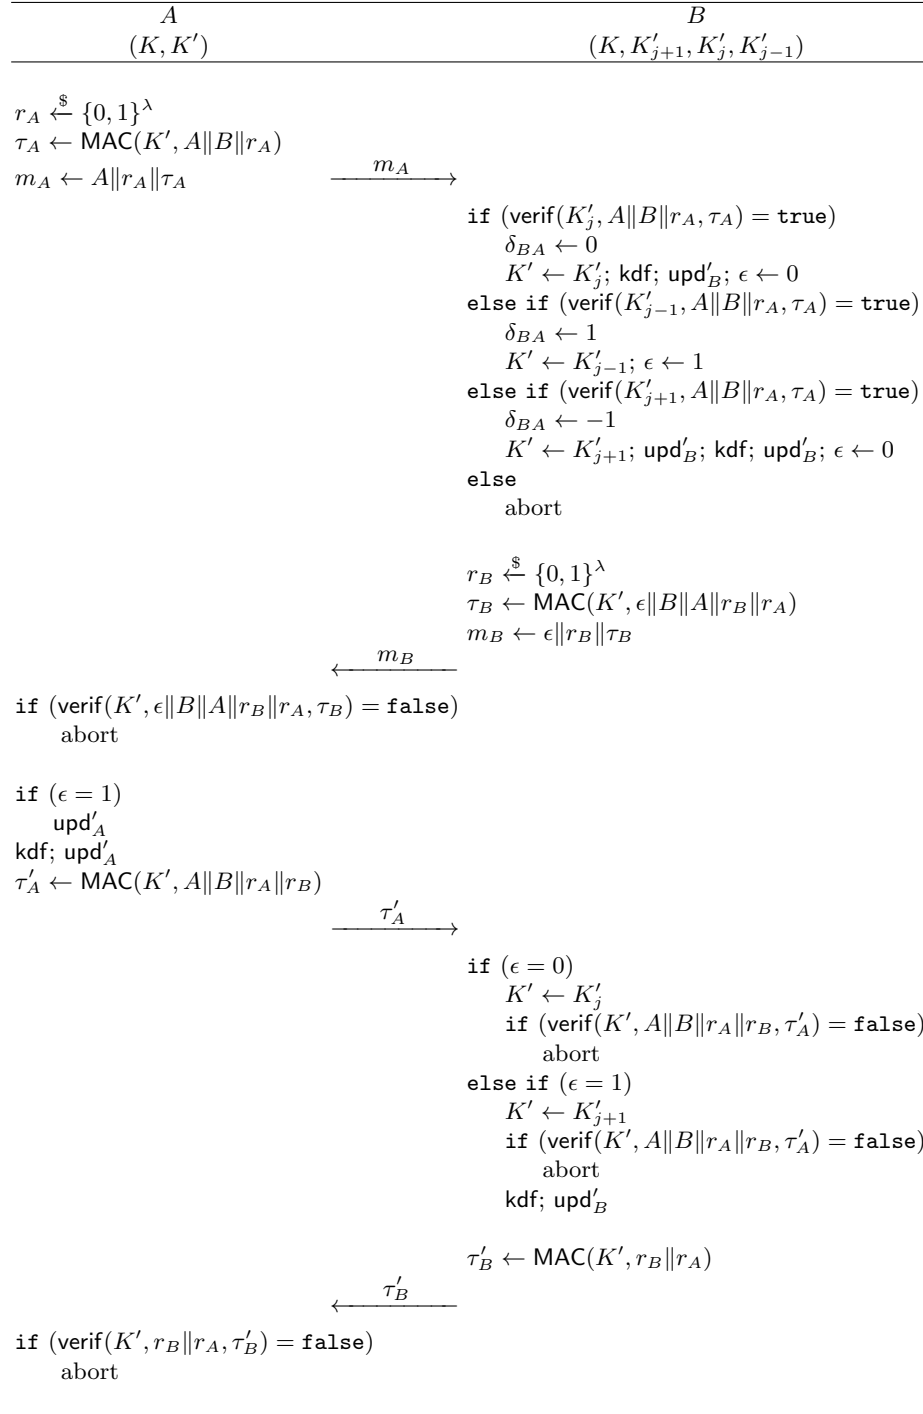


Fig. 5: SAKE-AM protocol



- $(c_B, c_A) = (i + 1, i)$  and  $\delta_{BA} = 1$  after a  $(0, 1)$ -session,
- $(c_B, c_A) = (i + 1, i + 1)$  and  $\delta_{BA} = 0$  after a  $(2, 1)$ -session,
- $(c_B, c_A) = (i + 1, i + 1)$  and  $\delta_{BA} = 0$  after a  $(2, 3)$ -session,
- $(c_B, c_A) = (i + 1, i + 1)$  and  $\delta_{BA} = 0$  after a  $(4, 3)$ -session.

This corresponds to the first column of Tables 4b and 4a. As we can see, the only possible values for  $\delta_{BA}$  after any session are 0 and 1. Hence, we explore all the values  $\delta_{BA}$  may have with any possible session if  $\delta_{BA} = 1$  (i.e.,  $(c_B, c_A) = (i + 1, i)$ ) when the session starts.

$(c_B, c_A) = (i + 1, i)$  means that  $B$  is in advance with respect to  $A$ . In such a case,  $B$  succeeds in validating  $\tau_A$  with  $K'_{j-1}$  (and, indeed, finds  $\delta_{BA} = 1$ ). Then  $B$  uses  $\delta_{BA} = 1$  and  $\epsilon = 1$ . If one carries out the protocol run using these two values, one gets:

- $(c_B, c_A) = (i + 1, i)$  and  $\delta_{BA} = 1$  after a  $(0, 1)$ -session,
- $(c_B, c_A) = (i + 1, i + 2)$  and  $\delta_{BA} = -1$  after a  $(2, 1)$ -session,
- $(c_B, c_A) = (i + 2, i + 2)$  and  $\delta_{BA} = 0$  after a  $(2, 3)$ -session,
- $(c_B, c_A) = (i + 2, i + 2)$  and  $\delta_{BA} = 0$  after a  $(4, 3)$ -session.

This corresponds to the second column of Table 4b. This shows that a third value is possible for  $\delta_{BA}$ , which is  $-1$ .

Then we restart the protocol with all possible sessions, assuming that  $(c_B, c_A) = (i, i + 1)$  at the beginning of the run. This means that  $B$  is one step late with respect to  $A$ . In such a case,  $B$  succeeds in validating  $\tau_A$  with key  $K'_{j+1}$  (and, indeed, finds  $\delta_{BA} = -1$ ). Then  $B$  uses  $\delta_{BA} = -1$  and  $\epsilon = 0$ . If one carries out the protocol run using these two values, one gets:

- $(c_B, c_A) = (i + 2, i + 1)$  and  $\delta_{BA} = 1$  after a  $(0, 1)$ -session,
- $(c_B, c_A) = (i + 2, i + 2)$  and  $\delta_{BA} = 0$  after a  $(2, 1)$ -session,
- $(c_B, c_A) = (i + 2, i + 2)$  and  $\delta_{BA} = 0$  after a  $(2, 3)$ -session,
- $(c_B, c_A) = (i + 2, i + 2)$  and  $\delta_{BA} = 0$  after a  $(4, 3)$ -session.

We end with two possible values for  $\delta_{BA}$  (third column of Table 4b): 0 and 1, that have already been explored. This proves that, whatever the sequences of sessions, the only possible values for  $\delta_{BA}$  are in  $\{-1, 0, 1\}$ .

Now we prove item 2 of Lemma 2.

We know that  $\delta_{BA} \in \{-1, 0, 1\}$ . For each possible value for  $\delta_{BA}$  at the beginning of the session, the last line of Table 4b indicates the value of that parameter after a correct session (i.e., a  $(4, 3)$ -session). As we can see,  $A$  and  $B$  are always synchronised (i.e.,  $\delta_{BA} = 0$ ) in such a case whatever the value of  $\delta_{BA}$  when the session starts. Furthermore, the session key computation immediately precedes the last update of the derivation master key  $K$ . Hence, when a correct and complete session ends,  $A$  and  $B$  use the same key  $K$  to compute the session key. Therefore, under the assumption that  $A$  and  $B$  both use the same values  $r_A, r_B$ , they do compute the same session key.

In addition, Table 4a shows that, whatever the synchronisation state of  $A$  and  $B$  (i.e.,  $c_A$  and  $c_B$ ) at the beginning of the protocol, after a correct session,  $A$  and  $B$  have updated their internal state at least once (as the last line of the table, corresponding to a  $(4, 3)$ -session, indicates).  $\square$

Table 4: Possible values for  $\delta_{BA}$  and  $(c_B, c_A)$  among all possible sequences of sessions in SAKE-AM

(a) Possible values for $(c_B, c_A)$				(b) Possible values for $\delta_{BA}$			
session \ (c_B, c_A)	(i, i)	(i + 1, i)	(i, i + 1)	session \ $\delta_{BA}$	0	1	-1
(0, 1)	(i + 1, i)	(i + 1, i)	(i + 2, i + 1)	(0, 1)	1	1	1
(2, 1)	(i + 1, i + 1)	(i + 1, i + 2)	(i + 2, i + 2)	(2, 1)	0	-1	0
(2, 3)	(i + 1, i + 1)	(i + 2, i + 2)	(i + 2, i + 2)	(2, 3)	0	0	0
(4, 3)	(i + 1, i + 1)	(i + 2, i + 2)	(i + 2, i + 2)	(4, 3)	0	0	0

The diagram depicted by Figure 6 represents all the possible sequences of sessions for SAKE-AM.




Fig. 6: Diagram of SAKE-AM. The circled values correspond to the gap  $\delta_{BA}$ , and each edge to a  $(i_A, i_B)$ -session.

## B.2 Security of SAKE-AM

With the following theorem, we claim that SAKE-AM is a secure AKE protocol according to Definition 5.

**Theorem 2.** *The protocol SAKE-AM is a secure AKE protocol, and for any probabilistic polynomial time adversary  $\mathcal{A}$  in the AKE security experiment against SAKE-AM*

$$\text{adv}_{SAKE-AM}^{\text{ent-auth}}(\mathcal{A}) \leq nq \left( (nq - 1)2^{-\lambda} + (q + 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + 2\text{adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{C}) \right)$$

$$\text{adv}_{SAKE-AM}^{\text{key-ind}}(\mathcal{A}) \leq nq \left( (q - 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + \text{adv}_{\text{KDF}}^{\text{prf}}(\mathcal{D}) \right) + \text{adv}_{SAKE-AM}^{\text{ent-auth}}(\mathcal{A})$$

where  $n$  is the number of parties,  $q$  the number of instances (sessions) per party,  $\lambda$  the size of the pseudo-random values ( $r_A, r_B$ ), and  $\mathcal{B}$  is an adversary against the PRF-security of `update`,  $\mathcal{C}$  an adversary against the SUF-CMA-security of MAC, and  $\mathcal{D}$  an adversary against the PRF-security of KDF.

*Proof.* We first prove the entity authentication security.

*Game 0.* This game corresponds to the entity authentication security experiment described in Section 2.1. Therefore

$$\Pr[E_0] = \text{adv}_{\text{SAKE-AM}}^{\text{ent-auth}}(\mathcal{A})$$

*Game 1.* In this game, we add an abort rule. The challenger aborts if there exists any instance that chooses a random value  $r_A$  or  $r_B$  that is not unique. There is at most  $n \times q$  random values, each uniformly drawn at random in  $\{0, 1\}^\lambda$ . Therefore the probability that at least two random values be equal is at most  $\frac{nq(nq-1)}{2^\lambda}$ . Hence

$$\Pr[E_0] \leq \Pr[E_1] + \frac{nq(nq-1)}{2^\lambda}$$

*Game 2.* In this game, we add an abort rule. The challenger tries to guess which instance will be the first to accept maliciously. If the guess is wrong, the game is aborted. The number of instances is at most  $nq$ . Therefore

$$\Pr[E_2] = \Pr[E_1] \times \frac{1}{nq}$$

*Game 3.* Let  $\pi$  be the instance targeted by the adversary. In this game, we add an abort rule. The challenger aborts the experiment if  $\pi$ , behaving as a responder (resp. initiator), ever receives a valid message  $m_B$  (resp.  $m_A$ ) but no instance having a matching conversation to  $\pi$  has output that message. We reduce the probability of this event to the security of the functions MAC and `update`. When the  $u$ -th session starts, the master keys have been updated at most  $u-1$  times already. The genuine value of  $K'$  is uniformly chosen at random. In order to be able to replace, during the current session, the key used to compute the MAC tag in  $m_A$  (resp.  $m_B$ ) with a random value, one must rely upon the pseudo-randomness of the function `update` that outputs (the new value of)  $K'$ . In turn, this relies upon the (previous) key  $K'$  being random (and on the pseudo-randomness of `update`). Therefore, in order to replace  $K'$  with a random value one must take into account the successive losses  $\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B})$ , each corresponding to the ability of an adversary  $\mathcal{B}$  to distinguish the function `update` (keyed with a different key  $K'$ ) from a random function. Since there is at most  $q$  sessions, this loss is at most  $(q-1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B})$ . Then we reduce the probability of the adversary  $\mathcal{A}$  to win this game to the ability of an adversary  $\mathcal{C}$  to forge a valid tag  $\tau_B$  (resp.  $\tau_A$ ).

Therefore, we replace each function `update`( $\cdot$ ) = PRF( $\cdot, \cdot$ ) (keyed with a different key  $K'$  throughout the, at most,  $q-1$  successive sessions established, prior

to that current session, by the same party that owns  $\pi$ ) with truly random functions  $F_0^{\text{update}}, \dots, F_{q-2}^{\text{update}}$ . Moreover, if an instance uses the same key  $K' = K'_i$ ,  $0 \leq i < q - 1$ , to key update, then we replace update with the corresponding random function  $F_i^{\text{update}}$ . Since, to that point, the key  $K' = K'_{q-1}$  used to compute the authentication tag  $\tau_B$  (resp.  $\tau_A$ ) is random, we reduce the ability of  $\mathcal{A}$  to win to the security of the MAC function. Hence

$$\Pr[E_2] \leq \Pr[E_3] + (q - 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + \text{adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{C})$$

*Game 4.* In this game, we add an abort rule. The challenger aborts the experiment if  $\pi$  ever receives a valid message  $\tau'_B$  (resp.  $\tau'_A$ ), but no instance having a matching conversation to  $\pi$  has output that message. Between the message  $m_B$  (resp.  $m_A$ ) being received by  $\pi$ , and the message  $\tau'_B$  (resp.  $\tau'_A$ ) being received by  $\pi$ , the master keys are updated at most twice. We reduce the probability of the adversary to win this game to the security of the MAC function used to compute the message  $\tau'_B$  (resp.  $\tau'_A$ ). In turn we must rely on the randomness of the MAC key, hence on the security of the function update used to update the MAC key  $K'$  (recall that, due to Game 3, the current key  $K'$  is random). Therefore

$$\Pr[E_3] \leq \Pr[E_4] + 2\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + \text{adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{C})$$

To that point, the only way for the adversary to make  $\pi$  accept maliciously is to send a valid message  $\tau'_B$  (resp.  $\tau'_A$ ) different from all the messages sent by all the instances. However, in such a case, the challenger aborts. Therefore

$$\Pr[E_4] = 0$$

Collecting all the probabilities from Game 0 to Game 4, we have that

$$\begin{aligned} \text{adv}_{\text{SAKE-AM}}^{\text{ent-auth}}(\mathcal{A}) &= \Pr[E_0] \\ &\leq \frac{nq(nq - 1)}{2^\lambda} + \Pr[E_1] \\ &= \frac{nq(nq - 1)}{2^\lambda} + nq \times \Pr[E_2] \\ &\leq \frac{nq(nq - 1)}{2^\lambda} + nq \left( \Pr[E_3] + (q - 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + \text{adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{C}) \right) \\ &\leq \frac{nq(nq - 1)}{2^\lambda} + nq \left( \Pr[E_4] + (q + 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + 2\text{adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{C}) \right) \\ &= \frac{nq(nq - 1)}{2^\lambda} + nq \left( (q + 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + 2\text{adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{C}) \right) \\ &= nq \left( (nq - 1)2^{-\lambda} + (q + 1)\text{adv}_{\text{update}}^{\text{prf}}(\mathcal{B}) + 2\text{adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{C}) \right) \end{aligned}$$

The reasoning and the bound for the key indistinguishability security are the same as that of SAKE (see Appendix A).

□