




Symmetric Key Exchange with Full Forward Security and Robust Synchronization

Colin Boyd¹ Gareth T. Davies²  Bor de Kock¹ 
Kai Gellert²  Tibor Jager² Lise Millerjord¹

¹NTNU – Norwegian University of Science and Technology, Trondheim, Norway

²Bergische Universität Wuppertal, Wuppertal, Germany

May 27, 2021

Abstract

We construct lightweight authenticated key exchange protocols based on pre-shared keys, which achieve *full* forward security and rely only on simple and efficient symmetric-key primitives. All of our protocols have rigorous security proofs in a strong security model, all have low communication complexity, and are particularly suitable for resource-constrained devices.

We describe three protocols that apply *linear* key evolution to provide different performance and security properties. *Correctness* in parallel and concurrent protocol sessions is difficult to achieve for linearly key-evolving protocols, emphasizing the need for assurance of availability alongside the usual confidentiality and authentication security goals. We introduce *synchronization robustness* as a new formal security goal, which essentially guarantees that parties can re-synchronize efficiently. All of our new protocols achieve this property.

Since protocols based on linear key evolution cannot guarantee that all concurrently initiated sessions successfully derive a key, we also propose two constructions with *non-linear* key evolution based on puncturable PRFs. These are instantiable from standard hash functions and require $O(C \cdot \log(|\text{CTR}|))$ memory, where C is the number of concurrent sessions and $|\text{CTR}|$ is an upper bound on the total number of sessions per party. These are the first protocols to simultaneously achieve full forward security, synchronization robustness, and concurrent correctness.

This work was supported by Deutscher Akademischer Austauschdienst (DAAD) and Norges forskningsråd (NFR) under the PPP-Norwegen programme. Colin Boyd and Lise Millerjord have been supported by NFR project number 288545. Tibor Jager and Gareth T. Davies have been supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme, grant agreement 802823.

Contents

1	Introduction	2
2	Preliminaries	7
2.1	Message Authentication Codes	7
2.2	Pseudorandom Functions	8
3	Authenticated Key Exchange in the Symmetric Setting	9
3.1	Execution Environment.	9
3.2	AKE Security	11
3.3	Concurrent Execution Synchronization Robustness	13
4	Linear Key Evolution	14
4.1	Key Derivation via Linear Evolution	15
4.2	LP3: a Three-Message Protocol	17
4.2.1	AKE-M of LP3	17
4.2.2	Bounded Gap: Non-Concurrent Setting.	21
4.2.3	Bounded Gap: Concurrent Setting.	23
4.2.4	wSR of LP3.	24
4.3	LP2: A Two-Message Protocol with Fixed Roles	26
4.3.1	AKE-M of LP2	27
4.3.2	wSR of LP2	30
4.4	LP1: A One-Message Protocol with Fixed Roles	33
4.4.1	AKE-R of LP1	34
4.4.2	wSR of LP1	36
5	Non-Linear Key Evolution	40
5.1	Puncturable Pseudorandom Functions	40
5.2	PPRF-based Symmetric AKE	41
5.2.1	AKE-M of PP2	42
5.2.2	SR of PP2	45
5.3	Protocol PP1	50
5.3.1	AKE-R of PP1	50
5.3.2	SR of PP1	50
5.4	Instantiation	51

1 Introduction

Authenticated key exchange protocols based on pre-shared long-term symmetric keys (PSK-AKE) enable two parties to use a previously established symmetric key, agreed upon via out-of-band communication, to (mutually) authenticate and derive a shared session key. Prominent examples of such protocols are the PSK modes of TLS 1.3 and prior TLS versions, but these examples still make use of public-key techniques for key derivation, even if authentication uses symmetric keys. PSK-AKE protocols can be significantly more efficient than classical public-key AKE protocols, particularly when they can be constructed exclusively

based on symmetric key primitives (“symmetric AKE”) for both authentication and key derivation. Therefore such protocols are especially desirable for performance-constrained devices, such as battery-powered wireless IoT devices, where every computation and every transmitted bit has a negative impact on battery life. More generally, such protocols may be preferable in “closed-world” applications, such as industrial settings, where pre-sharing keys may be easier and more practical than deploying a public-key infrastructure. Furthermore, protocols based purely on symmetric-key techniques, such as hash functions and symmetric encryption, also achieve security against quantum attacks by adjusting security parameters appropriately.

Forward Security in Symmetric AKE Protocols. *Forward security* is today a standard security goal of key exchange protocols. It requires that past session keys remain secure, even if the secret long-term key material is compromised. Note that this is only achievable if past session keys are *not* efficiently computable from a current long-term key. Forward security is comparatively easily achievable if *public key* cryptography is used. For instance, a classical approach is to use *ephemeral* keys for key establishment, such as the Diffie-Hellman protocol or, more generally, a key encapsulation mechanism (KEM). *Independent* long-term keys can then be used for authentication via digital signatures or another KEM.

The only currently known way to avoid public key techniques and use only symmetric key primitives is based on the “*derive-then-evolve*” approach, where first a session key is derived from a long-term key, and then the long-term key is evolved. This key evolution prevents efficient re-computation of prior session keys which yields forward security. Both steps can be implemented with simple key derivation functions. There are two common ways to use this approach:

1. *Synchronized key evolution.* In this case, both parties evolve their long-term keys in “epochs”, e.g., once per day. Note that this approach cannot achieve “*full*” forward security, but only a weaker “*delayed*” form. This is because all session keys of the current epoch can be computed from the current long-term secret, so forward security only holds for session keys of past epochs. Moreover, this approach requires synchronized clocks between parties, even to achieve correctness. For many applications this seems impractical, in particular for cheap low-performance devices, for which symmetric AKE protocols are particularly relevant.
2. *Triggered key evolution.* In this case the protocol ensures that both parties advance their key material during protocol execution. This approach directly achieves “*full*” forward security for every session, and therefore seems preferable. However, this apparently simple approach turns out to be much less trivial to realize than might be expected, because both parties must remain “in sync”, such that correctness is guaranteed even in presence of *concurrent* sessions or message loss due to network failures or *active attacks*. This approach has similarities with *ratcheting* [ACD19], but there are significant differences in our setting as discussed under *Related Work* below.

Concurrency and Key-Evolving Protocols. The possibility of running concurrent protocol sessions in parallel is a standard correctness requirement for protocols, and reflected in all common AKE security models, such as the BR and CK models [BR94, CK01] and their countless variants and refinements. The main technical challenge of key evolution is to achieve full forward security while maintaining *correctness* in the presence of parallel and concurrent protocol sessions.

Even if we assume that all parties are *honest* and that all messages are transmitted *reliably* (i.e., without being dropped because of an unreliable network or influence from an adversary) this is already highly non-trivial and we do not know of any currently existing forward-secure symmetric AKE protocol which achieves *correctness* and full forward security in such a setting. The difficulty is essentially that one session might

Table 1: Overview of our protocols and comparison to SAKE [ACF20]. The number in the protocol name indicates the total number of messages per protocol run, “R only” means that only the responder authenticates its communication partner. The third column considers the communication complexity, where **C** is the number of counter values that are sent, **M** the number of MACs, and **N** the number of nonces. **Sync. Rob.** indicates the achieved level of synchronization robustness, **Bd. Gap** whether the gap between two parties is bounded (for non-concurrent executions), **CC** whether concurrent correctness is achieved, and **FS** whether full forward security is achieved.

Protocol	Auth.	(C, M, N)	Sync. Rob.	Bd. Gap	CC	FS
SAKE (5) [ACF20]	mutual	(0,4,2) + ID	✗	✓	✗	✓
SAKE-AM (4) [ACF20]	mutual	(0,4,2) + ID	✗	✓	✗	✓
LP3	mutual	(3,3,2)	weak	✓	✗	✓
LP2	mutual	(2,2,0)	weak	✗	✗	✓
LP1	R only	(1,1,0)	weak	✗	✗	✓
PP2	mutual	(1,2,0)	full	✓	✓	✓
PP1	R only	(1,1,0)	full	✓	✓	✓

advance a key “too early” for another concurrent session to be completed, which breaks correctness. No such difficulty appears in classical forward-secure public key protocols, since long-term keys are usually *static* and different sessions use independent randomness. So it turns out that, somewhat surprisingly, forward security and correctness is more difficult to achieve for symmetric AKE.

To complicate matters even further, note that the assumption of honest parties and reliable message transmission is very strong and may not be realistic for many applications. Therefore we actually want to achieve forward security and “synchronization robustness” in the presence of an adversary which intentionally aims to *break* synchronization, e.g., by adaptively dropping or re-ordering messages. Such an adversary is attacking *availability* properties of the AKE protocol, an important aspect of security usually omitted from key exchange security models. The development of techniques to ensure availability for stateful key exchange is an unsolved foundational problem.

Our Contributions. In this work we develop several new lightweight forward-secure symmetric AKE protocols with different efficiency and correctness properties. Table 1 summarizes the main security and efficiency properties of our new protocols. This includes the first protocols that provably achieve synchronization robustness, a formal availability security notion we introduce, and correctness in the presence of concurrent sessions. More concretely we achieve the following.

Security model. We describe a security model suited to forward-secure symmetric AKE capturing entity authentication (one-sided and mutual), indistinguishability of established keys, and forward security. Our model follows a standard approach for AKE protocols based on the Bellare-Rogaway model [BR94], adapted to the requirements of symmetric AKE with evolving keys.

Synchronization robustness. We formalize a new property called *synchronization robustness* (SR), which is trivially achieved for traditional AKE protocols with fixed long-term keys, but turns out to be a

crucial feature for key-evolving protocols such as forward-secure symmetric AKE. Essentially, SR captures whether parties in a protocol can efficiently re-synchronize their states in order to complete a successful protocol run. This should even hold if an adversary controls the network and/or some of the parties.

We define two flavours. Both consider an active adversary that may execute arbitrary protocol sessions to manipulate the state of parties, and whose goal is to manipulate the state such that a subsequent protocol execution fails.

In *weak* SR the ‘target’ protocol session must then be executed without adversarial interaction (similar to the corresponding requirement in Krawczyk’s weak forward security [Kra05]). ‘Full’ SR allows the adversary arbitrary queries between messages of the ‘target’ session, even to parties of the oracles involved in the ‘target’ session.

Linear key evolving protocols. We define the notion of *linear key evolution*, which makes the classical “derive-then-evolve” approach concrete. We argue that protocols based on linear key evolution can only achieve weak SR and cannot achieve concurrent correctness.

We construct three different protocols (LP1, LP2, LP3, cf. Table 1), all of which achieve weak SR. Most interestingly, LP3 even achieves a “bounded gap” property, which means that no active adversary in control of the network is able to force the state of two parties to differ by more than one key evolving step, so that a party is always able to catch up quickly, if necessary. For all three protocols we show that in a setting where concurrent runs between two parties are allowed, this number of steps required to catch up is bounded in the number of concurrent runs. To this end, we apply a new approach to precisely analyze the state machine of a protocol. Furthermore, we also show two extremely lightweight protocols LP1 and LP2, which provide one-sided and mutual authentication, respectively, and where the communication complexity is only one (resp. two) MAC and one (resp. two) counter value.

Full SR and concurrent correctness. This leads to the question of whether and how full synchronization robustness and concurrent correctness (CC) can be achieved. We propose the use of puncturable pseudorandom functions (PPRFs) to apply a “non-linear” key evolving strategy, and we construct two protocols PP1 and PP2, which both achieve full SR and CC.

Since PPRFs can be efficiently instantiated from cryptographic hash functions, both protocols are extremely lightweight. PP1 achieves one-sided authentication with a single counter value and a single MAC, PP2 mutual authentication with one counter and two MACs. Furthermore, while repeated puncturing PPRFs may lead to large secret keys [AGJ19, AGJ21] we take advantage of the stateful nature of symmetric AKE protocols to instantiate the PPRF such that secret key size is at most *logarithmic* in the number of sessions.

Hence, we offer a versatile catalogue of lightweight and forward-secure symmetric AKE protocols with significantly stronger correctness and security properties. This includes the first protocols to achieve concurrent correctness and full synchronization robustness, or weak SR with bounded gap. Which of these protocols is best for a particular application depends on the nature of the security and functionality requirements. Further, in LP3 the parties exchange nonces: we recognize that in some applications sufficient randomness will not be available and so we prove the protocol secure for any nonce generation procedure, which could be random selection or (stateful) use of a counter.

Related Work. Bellare and Yee [BY03] analyzed forward security for symmetric-key primitives, specifically pseudo-random generation, message authentication codes and symmetric encryption. They provide constructions using key evolution which are similar to the linear key evolution that we employ, and their protocols use some techniques from key-evolving schemes such as prior work on forward-secure signatures [BM99]. Their work does not deal with key exchange.

Brier and Peyrin [BP10] gave a tree-based protocol for key establishment, with the stated aim of improving the DUKPT scheme defined in ANSI X9.24 [ANS09]. The idea in DUKPT is that the client device (payment terminal) is highly constrained in terms of memory, yet needs to derive a unique key per transaction from an original pre-shared key, by applying a PRF (based on Triple-DES) to a counter and the base derivation key. Their work involves formalizing the specific problem faced in the payment terminal setting, and their scheme assumes an incorruptible server: a far weaker security model than the one that we consider. A similar security assumption was used by Le et al. [LBdM07], who presented a protocol for use in the context of Radio Frequency Identification (RFID), where the server keeps two values of the key derivation key to deal with potential synchronization loss.

Li et al. [LSY⁺14] analyzed the pre-shared key ciphersuites of TLS 1.2, using an adaption of the ACCE model of Jager et al. [JKSS12]. In this setting, Li et al. presented a formalization of the prior AKE-style models, but where parties could share PSK material with other parties in addition to their long-term key pairs.

Dousti and Jalili [DJ15] presented a key exchange protocol called FORSAKES, which is based on synchronized time-based key evolution. Their protocol requires 3 messages and assumes perfect synchronicity of parties to achieve correctness, and as we have already mentioned their approach can only obtain *delayed forward security*. A discussion of delayed forward security and more generally the various challenges involved in defining forward security was given Boyd and Gellert [BG20].

The concept of evolving symmetric keys is reminiscent of Signal’s double ratchet [ACD19], a well-known example of a symmetric protocol with evolving keys. Signal employs a Diffie-Hellman-ratchet, which adds new key material at every step through multiple Diffie-Hellman exchanges along the way. At every step of this main ratchet a separate linear key evolving ratchet is ‘branched off’, which is similar to how linear evolution works in our protocols — however, a critical difference is that in our scenario we evolve the key shared across different sessions as opposed to evolving a key within one session as happens in the Signal protocol. It is this difference which leads to the complexity of managing synchronization between sessions which run concurrently. In addition to this difference, which anyway makes Signal unusable for our setting, use of Diffie-Hellman in the Signal ratchet means that there is a vector for quantum attacks, while our protocol is purely based on symmetric primitives.

Another primitive conceptually similar to PPRFs is *puncturable encryption*, which was introduced by Green and Miers in 2015 [GM15], and has since led to several follow-up constructions of puncturable encryption [GHJL17, DJSS18, CRSS20, SSS⁺20, DGJ⁺21]. However, all those constructions rely on expensive public-key techniques (such as bilinear pairings) and are thus impractical in the context of this work.

Comparison with Avoine et al. [ACF20]. In Table 1 above we have mentioned two protocols named SAKE and SAKE-AM that were presented by Avoine et al. [ACF20] (henceforth ACF20). Their paper was the first to provide key exchange protocols that attain forward security via linear evolution. Their system assumptions are largely the same as ours, with the crucial difference that our models are equipped to capture parallel executions. The security model of ACF20 explicitly disallows concurrent sessions, which not only yields a weak security notion, but also sidesteps the major difficulty of achieving even correctness

in the presence of concurrent sessions in key-evolving symmetric-key protocols. Indeed, the protocols from ACF20 completely break down when executed concurrently, allowing an adversary to prevent the parties from computing any session keys in future sessions. We consider this an unrealistic and impractical restriction for many applications. Therefore we introduce the new notion of synchronization robustness, which formally defines the ability of key-evolving protocols to deal with concurrent executions, including in adversarial environments.

We embrace the use of (explicit) counters to acquire linear key evolving protocols that are conceptually simpler and require fewer messages than those provided by ACF20, in a way that additionally provides (weak) synchronization robustness. In any protocol that uses PSK evolution to achieve forward security a party must update the key state after a successful protocol run, and in embedded devices this requires writing to persistent storage. Our protocols require the updating (writing) of one key and one counter per session, while SAKE and SAKE-AM require updating two keys. Since a sequentially evolving key can also be seen as an implicit counter, conceptually the distinction between counters and evolving keys seems to be minor. The storage overhead of our protocols compared to ACF20’s protocols is the (usually 8-byte) counter, while the linear key evolving protocols in our paper and ACF20 require storage of two keys (usually 16 or 32 bytes).

We note that ACF20 remarked that the parties could use separate PSKs for concurrent executions, however this solution requires an a priori bound on the number of possible concurrent sessions that could occur and a corresponding multiplication in key storage: none of our protocols require this. Further, implementing their approach would require a modification of their protocols, since parties need to know which PSK to use, and the security of these modified protocols is not proven.

2 Preliminaries

We denote the security parameter as λ . For any $n \in \mathbb{N}$ let 1^n be the unary representation of n and let $[n] = \{1, \dots, n\}$ be the set of numbers between 1 and n . We write $x \xleftarrow{\$} \mathcal{S}$ to indicate that we choose element x uniformly at random from set \mathcal{S} . For a probabilistic polynomial-time algorithm \mathcal{A} we define $y \xleftarrow{\$} \mathcal{A}(a_1, \dots, a_n)$ as the execution of \mathcal{A} (with fresh random coins) on input a_1, \dots, a_n and assigning the output to y . The function $\text{NextOdd}(x)$ takes as input an integer and outputs the next odd integer greater than x , i.e. whichever element of $\{x + 1, x + 2\}$ is odd. Our protocols require the use of counters, and integer $|\text{CTR}|$ is the largest possible counter value. Furthermore, we write $[n] \times [n] \setminus (i^*, j^*)$ as a shorthand for $(i, j) \in [n]^2 \setminus \{(i^*, j^*)\}$ with $i < j$.

2.1 Message Authentication Codes

Throughout this paper we assume that all MACs are deterministic. This is to reduce complexity in our proofs, however most MACs used in practice are deterministic [CMA05, GMA07, HMA08, ISO11, KMA16].

Definition 1 (Message Authentication Codes). A *message authentication code* consists of three probabilistic polynomial-time algorithms $\text{MAC} = (\text{KGen}, \text{Mac}, \text{Vrfy})$ with key space \mathcal{K}_{MAC} and the following properties:

- $\text{KGen}(1^\lambda)$ takes as input a security parameter λ and outputs a symmetric key $K^{\text{MAC}} \in \mathcal{K}_{\text{MAC}}$;
- $\text{Mac}(K^{\text{MAC}}, m)$ takes as input a key $K^{\text{MAC}} \in \mathcal{K}_{\text{MAC}}$ and a message m . Output is a tag σ ;

- $\text{Vrfy}(k^{\text{MAC}}, m, \sigma)$ takes as input a key $k^{\text{MAC}} \in \mathcal{K}_{\text{MAC}}$, a message m , and a tag σ . Output is a bit $b \in \{0, 1\}$.

We call a message authentication code *correct* if for all m , we have

$$\Pr_{k^{\text{MAC}} \leftarrow \mathcal{K}\text{Gen}(1^\lambda)} \left[\text{Vrfy}(k^{\text{MAC}}, m, \text{Mac}(k^{\text{MAC}}, m)) = 1 \right] = 1.$$

We define MAC security as strong existential unforgeability under chosen message attack, where the adversary has access to a verification oracle. In the more common version of this game, which we denote SEUF-CMA-1, the adversary must stop running after it submits its first verification query: this is a subcase of our more general definition. Bellare et al. [BGM04] showed that in the strong unforgeability case these definitions are equivalent up to a loss factor Q .

$\mathcal{G}_{\text{MAC}}^{\text{SEUF-CMA-}Q}(\mathcal{A})$	$\mathcal{O}_{\text{Mac}}(m)$	$\mathcal{O}_{\text{Vrfy}}(m, \sigma)$
1 : $k^{\text{MAC}} \leftarrow \mathcal{K}\text{Gen}(1^\lambda)$	7 : $\sigma \leftarrow \text{Mac}(k^{\text{MAC}}, m)$	10 : $b \leftarrow \text{Vrfy}(m, \sigma)$
2 : $\mathcal{Q}, \mathcal{V} \leftarrow \emptyset$	8 : $\mathcal{Q} := \mathcal{Q} \cup \{(m, \sigma)\}$	11 : if $b = 1$
3 : $\mathcal{A}^{\mathcal{O}_{\text{Mac}}(\cdot), \mathcal{O}_{\text{Vrfy}}(\cdot, \cdot)}(1^\lambda)$	9 : return σ	12 : $\mathcal{V} := \mathcal{V} \cup \{(m, \sigma)\}$
4 : if $\exists (m, \sigma) \in \mathcal{V} \setminus \mathcal{Q}$		13 : return b
5 : return 1		
6 : return 0		

Figure 1: The SEUF-CMA- Q security experiment for message authentication code MAC. \mathcal{A} can make Q queries to $\mathcal{O}_{\text{Vrfy}}$.

Definition 2 (MAC Security). The advantage of an adversary \mathcal{A} in the SEUF-CMA- Q security experiment defined in Fig. 1 for message authentication code MAC is

$$\text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-}Q}(\mathcal{A}) := \Pr \left[\mathcal{G}_{\text{MAC}}^{\text{SEUF-CMA-}Q}(\mathcal{A}) = 1 \right].$$

2.2 Pseudorandom Functions

Definition 3 (Pseudorandom Functions). A *pseudorandom function* is a deterministic function $y = \text{PRF}(k, x)$ that takes as input some key $k \in \mathcal{K}_{\text{PRF}}$ and some element of a domain \mathcal{D}_{PRF} , and returns an element $y \in \mathcal{R}_{\text{PRF}}$.

Definition 4 (PRF Security). The advantage of an adversary \mathcal{A} in the PRF-sec security experiment defined in Fig. 2 for pseudorandom function PRF is

$$\text{Adv}_{\text{PRF}}^{\text{PRF-sec}}(\mathcal{A}) := \left| \Pr \left[\mathcal{G}_{\text{PRF}}^{\text{PRF-sec}}(\mathcal{A}) = 1 \right] - \frac{1}{2} \right|.$$

$G_{\text{PRF}}^{\text{PRF-sec}}(\mathcal{A})$	$\mathcal{O}_f(x)$
1 : $b \xleftarrow{\$} \{0, 1\}$	8 : if $b = 1$
2 : $k_{\text{PRF}} \xleftarrow{\$} \mathcal{K}_{\text{PRF}}$	9 : $y \leftarrow f(k_{\text{PRF}}, x)$
3 : $g \xleftarrow{\$} \{\mathcal{F} : \mathcal{D}_{\text{PRF}} \rightarrow \mathcal{R}_{\text{PRF}}\}$	10 : else
4 : $b^* \xleftarrow{\$} \mathcal{A}^{\mathcal{O}_f(\cdot)}(1^\lambda)$	11 : $y \leftarrow g(x)$
5 : if $b^* = b$	12 : return y
6 : return 1	
7 : return 0	

Figure 2: The PRF-sec security experiment for pseudorandom function PRF. $\{\mathcal{F} : \mathcal{D}_{\text{PRF}} \rightarrow \mathcal{R}_{\text{PRF}}\}$ is the set of all functions from \mathcal{D}_{PRF} to \mathcal{R}_{PRF} .

3 Authenticated Key Exchange in the Symmetric Setting

In this section we describe our model for authenticated key exchange with forward security in the symmetric setting. Our model follows the standard approach of AKE protocols based on the Bellare-Rogaway model [BR94], adopted to the requirements of symmetric AKE with evolving keys. This includes definitions for entity authentication (one-sided or mutual), key indistinguishability, and forward security. Furthermore, we define the property of synchronization robustness, which is a crucial feature for forward-secure symmetric key exchange protocols. Parts of our formalization take inspiration from the models of Jager et al. [JKSS12].

Differences to public-key AKE models. The most notable difference in the symmetric key setting is that each pair of parties is initialized with shared key material, which is specified before the actual protocol is run. This key material typically includes MAC keys or key derivation keys that have been established in an out-of-band communication (e.g., chosen during the manufacturing process of devices). In order to achieve forward-security via “key evolving techniques” in the symmetric key setting, we additionally have to provide (sessions of) parties with the ability to modify the party’s key material. As a consequence, the shared key material of two parties will not always be equal: While one party might evolve their key before preparing the first protocol message, the responder can (at the earliest) evolve after it has received that message.

3.1 Execution Environment.

We consider a set of n parties $\{P_1, \dots, P_n\}$, where each party is a potential protocol participant. We refer to parties by P_i or by their label i if context is clear. Initially, each pair of parties (P_i, P_j) with $i \neq j$ share a common secret $\text{PSK}_{i,j}$, which is the initial key material generated during protocol initialization (e.g., MAC keys or key derivation keys). Note that this key material may evolve over time and that $\text{PSK}_{i,j}$ and $\text{PSK}_{j,i}$ may not necessarily be equal at all times.

We model parallel executions of a protocol by equipping each party i with $q \in \mathbb{N}$ session oracles π_i^1, \dots, π_i^q . Each session oracle represents a process that executes one single instance of the protocol. All oracles have access to the “global key material” PSK (including the ability to modify the key material PSK). Moreover, each oracle maintains an internal state consisting of the following variables:

Variable	Description
α	execution state $\in \{\text{uninitialized}, \text{negotiating}, \text{accept}, \text{reject}\}$
pid	identity of the intended partner $\in \{P_1, \dots, P_n\}$
ρ	role $\in \{\text{Initiator}, \text{Responder}\}$
sk	session key $\in \mathcal{K}_s \cup \perp$ for some session key space \mathcal{K}_s
κ	freshness of session key $\in \{\text{exposed}, \text{fresh}\}$
sid	session identifier
b	security bit $\in \{0, 1\}$

Additionally, we assume that each oracle has an additional temporary state variable, used to store ephemeral values or the transcript of messages. As initial state of the oracle, we have $\alpha = \text{uninitialized}$ and $\kappa = \text{fresh}$ and $b \xleftarrow{\$} \{0, 1\}$. Note that pid and ρ are set when the adversary interacts with the respective oracles and that sid and sk are defined as the protocol/adversary progresses.

As usual, if an oracle derives a session key then it will enter the execution state `accept`. If an oracle reaches the execution state `reject`, then it will no longer accept any messages. Later on when we describe protocols, the event `Abort` will identify points at which this action would be triggered.

To begin any of the experiments in this section, the challenger initializes n parties $\{P_1, \dots, P_n\}$, with each pair of parties sharing symmetric key material PSK as specified by the protocol.

An adversary interacts with session oracles π_i^s by issuing the following queries. Several of these queries add output to an oracle transcript (defined below) which is available to the adversary.

- `NewSessionI`(π_i^s, pid) initializes a new initiator session for party P_i with intended partner pid . Specifically, this query assigns pid , $\rho = \text{Initiator}$ and $\alpha = \text{negotiating}$ to π_i^s , creates the first protocol message and adds this to transcript of π_i^s .
- `NewSessionR`(π_i^s, pid, m) initializes a new responder session for party P_i with $\rho = \text{Responder}$ and intended partner pid , and delivers a protocol message to this oracle. Specifically, it assigns pid and $\rho = \text{Responder}$ to π_i^s and processes message m . The message m and consequent protocol messages (if any) are added to its transcript, and the execution state is set to `negotiating`.
- `Send`(π_i^s, m) delivers message m to oracle π_i^s . This input message, and consequent protocol messages (if any), are added to this oracle's transcript.
- `RevealKey`(π_i^s) reveals session key sk_i^s and sets $\pi_i^s.\kappa$ to `exposed`.
- `Corrupt`(P_i, P_j) (issued to some oracle of P_i or P_j) returns $\text{PSK}_{i,j}$. If the query `Corrupt`(P_i, P_j) is the τ -th query issued by \mathcal{A} , we say that all oracles π_i with $\text{pid} = j$ are τ -corrupted. (i.e., party P_i becomes τ -corrupted with respect to the other party P_j). An uncorrupted oracle is considered as $+\infty$ -corrupted.
- `Test`(π_i^s) chooses $sk_0 \xleftarrow{\$} \mathcal{K}_s$, sets $sk_1 = \pi_i^s.sk$ and returns sk_b . This oracle can only be queried once, and the query making this action is labelled τ_0 .

The adversary must call `NewSessionI` or `NewSessionR` in order to specify a role and intended partner identifier for each oracle it wishes to use. Afterwards, the adversary can use the `Send` query to convey messages to these oracles.

3.2 AKE Security

To define entity authentication we use *matching conversations* [BR94] for oracle partnering, which requires a definition of an oracle’s *transcript*: T_i^s is the sequence of all messages sent and received by π_i^s in chronological order. The standard definition of matching conversations, reflects that the party that sends the last message cannot be sure that the responder received that protocol message. We use this definition for entity authentication.

Note that an oracle π_i^s only has a transcript, T_i^s , if $\pi_i^s.\alpha \neq \text{uninitialized}$. Transcript T_j^t is a *prefix* of T_i^s if T_j^t contains at least one message and messages in T_j^t are identical to and in the same order as the first $|T_j^t|$ messages of T_i^s .

Definition 5 (Partial-transcript Matching conversations [JKSS12, Def. 3]). π_i^s has a partial-transcript matching conversation to π_j^t if

- T_j^t is a prefix of T_i^s and π_i^s has sent the last message(s), or
- $T_i^s = T_j^t$ and π_j^t has sent the last message(s).

However, standard matching conversations are not strong enough to define key indistinguishability in a symmetric setting and leave room for a trivial attack (intuitively, this is due to the “asynchronous evolution” of the global key material PSK). Consider an adversary that uses the above execution environment to execute some protocol between two (sessions of two) parties. The adversary forwards all messages but the last one between both parties. At this point the party that sent the last message must have reached the accept state and applied some one-way procedure to its key material PSK in order to achieve forward security. However, the other party still needs to receive the final message in order to derive the session key and update its version of the key material. If the adversary were now to use Test on the accepting party while using Corrupt on the other party, this leads to a trivial distinguishing attack in standard key indistinguishability games (e.g., in [JKSS12]). Hence, we need to introduce a slightly stronger notion of matching conversations to precisely capture when Corrupt queries are allowed: the conversation is only deemed to be matching if all messages were delivered.

Definition 6 (Guaranteed Delivery Matching conversations). π_i^s has a guaranteed delivery matching conversation to π_j^t if $T_i^s = T_j^t$.

As usual, we say that the adversary breaks entity authentication if it forces a fresh oracle to accept maliciously, and breaks key indistinguishability if it can distinguish from random an established key that it cannot trivially access.

Definition 7 (Entity Authentication). Let Π be a protocol. Let $G_{\Pi}^{\text{Ent-Auth}}(\mathcal{A})$ be the following game:

- The challenger initializes n parties and their keys;
- \mathcal{A} may issue queries to oracles NewSessionI, NewSessionR, Send, RevealKey, Corrupt and Test as defined above;
- Once \mathcal{A} has concluded, the experiment outputs 1 if and only if there exists an accepting oracle π_i^s such that the following conditions hold:

1. both P_i (w.r.t. P_j) and intended partner P_j (w.r.t. P_i) were not corrupted before query τ_0 ;

2. there is no unique π_j^t , with $\rho_i^s \neq \rho_j^t$, such that π_i^s has a partial-transcript matching conversation to π_j^t .

Define the advantage of an adversary \mathcal{A} in the Ent-Auth security experiment $G_{\Pi}^{\text{Ent-Auth}}(\mathcal{A})$ as

$$\text{Adv}_{\Pi}^{\text{Ent-Auth}}(\mathcal{A}) := \Pr \left[G_{\Pi}^{\text{Ent-Auth}}(\mathcal{A}) = 1 \right].$$

An oracle π_i^s accepting in the above sense ‘accepts maliciously’.

Later on we separate the analysis of an initiator oracle accepting maliciously from a responder oracle accepting maliciously. Further, we will present protocols that only provide one-sided authentication: this requires separation of the AKE definition. To this end, we use the following notation:

$$\text{Adv}_{\Pi}^{\text{Ent-Auth}}(\mathcal{A}) = \text{Adv}_{\Pi}^{\text{Ent-Auth-I}}(\mathcal{A}) + \text{Adv}_{\Pi}^{\text{Ent-Auth-R}}(\mathcal{A}).$$

Definition 8 (Key Indistinguishability). Let Π be a protocol. Let $G_{\Pi}^{\text{Key-Ind}}(\mathcal{A})$ be the following game:

- The challenger initializes n parties and their keys;
- \mathcal{A} may issue queries to oracles `NewSessionI`, `NewSessionR`, `Send`, `RevealKey`, `Corrupt` and `Test` as defined above;
- Once \mathcal{A} has output (i, s, b') to indicate its conclusion, the experiment outputs 1 if and only if there exists an oracle π_i^s such that the following holds:
 1. π_i^s accepts, with a unique oracle π_j^t , such that π_i^s has a partial-transcript matching conversation to π_j^t , when \mathcal{A} issues its τ_0 -th query;
 2. \mathcal{A} did not issue `RevealKey` to π_i^s nor π_j^t (so $\kappa_i^s = \text{fresh}$) and $\rho_i^s \neq \rho_j^t$;
 3. P_i (w.r.t. P_j) is τ_i -corrupted and P_j (w.r.t. P_i) is τ_j -corrupted, with $\tau_i, \tau_j > \tau_0$;
 4. at the point of query τ_j , oracle π_j^t had a guaranteed delivery matching conversation to π_i^s , and
 5. $b' = \pi_i^s.b$.

Define the advantage of an adversary \mathcal{A} in the Key-Ind security experiment $G_{\Pi}^{\text{Key-Ind}}(\mathcal{A})$ as

$$\text{Adv}_{\Pi}^{\text{Key-Ind}}(\mathcal{A}) := \left| \Pr \left[G_{\Pi}^{\text{Key-Ind}}(\mathcal{A}) = 1 \right] - \frac{1}{2} \right|.$$

We assume that all adversaries in the Key-Ind game are valid, meaning that they terminate and provide an output in the correct format (i.e. $(i, s, b') \in [n] \times [q] \times \{0, 1\}$). Later on in our proofs we will follow the game-hopping strategy, and in doing so we will often simplify exposition by additionally assuming adversaries that do not trigger a trivial win (in the Key-Ind game or any subsequent modifications of this game).

We define AKE security in three flavors, distinguished by the level of entity authentication that is achieved by the protocol. An adversary breaks the AKE security of a protocol if it wins either the entity authentication game, or the key indistinguishability game.

Definition 9 (Authenticated Key Exchange). Let Π be a protocol. The advantage of an adversary \mathcal{A} in terms of AKE-M (mutual entity authentication), resp. AKE-I (initiator authenticates the responder), resp. AKE-R (responder authenticates the initiator) is defined as follows:

$$\begin{aligned} \text{Adv}_{\Pi}^{\text{AKE-M}}(\mathcal{A}) &:= \text{Adv}_{\Pi}^{\text{Key-Ind}}(\mathcal{A}) + \text{Adv}_{\Pi}^{\text{Ent-Auth-I}}(\mathcal{A}) + \text{Adv}_{\Pi}^{\text{Ent-Auth-R}}(\mathcal{A}). \\ \text{Adv}_{\Pi}^{\text{AKE-I}}(\mathcal{A}) &:= \text{Adv}_{\Pi}^{\text{Key-Ind}}(\mathcal{A}) + \text{Adv}_{\Pi}^{\text{Ent-Auth-I}}(\mathcal{A}). \\ \text{Adv}_{\Pi}^{\text{AKE-R}}(\mathcal{A}) &:= \text{Adv}_{\Pi}^{\text{Key-Ind}}(\mathcal{A}) + \text{Adv}_{\Pi}^{\text{Ent-Auth-R}}(\mathcal{A}). \end{aligned}$$

We do not provide any protocols that provide AKE-I in this paper, however it is defined here for completeness.

3.3 Concurrent Execution Synchronization Robustness

We now describe a novel property for key exchange protocols. The goal is to capture, in a formal manner, how *robust* a protocol is in the event of adversarial control of the network and/or some of the parties. We seek a definition that asks: after an adversary has had control of the communication network (by executing arbitrary Send and NewSessionI/NewSessionR queries), can an honest protocol run be executed successfully? Specifically, if it is possible for the parties to lose synchronization (due to dropped messages or adversarial control) such that the parties cannot, in one protocol run, regain synchronization and compute the same key, then the protocol does not meet this property.

Our formalization follows the execution environment of the Ent-Auth and Key-Ind games described above, and allows an adversary to specify the protocol run (that it is attempting to ‘interrupt’) at the end of its execution by specifying two oracles. The challenger awards success if the two parties (specifically those two oracles) did not accept with the same session key. We define two flavours: a weaker version wSR in which the ‘target’ protocol run must be executed without any other messages interleaved, and a stronger version SR which allows arbitrary queries in between messages of the ‘target’ run, even to parties of the oracles involved in the ‘target’ run (though of course not to the two oracles).

We define an *honest protocol run* (via adversarial queries) between two oracles (with initial state set to uninitialized) as follows: a NewSessionI query was made that produced a protocol message m_1 , a NewSessionR query was made to the other oracle with input message m_1 , and if this query produced a protocol message m_2 then this value was given as a Send query to the other oracle, and so on, until all protocol messages have been created and delivered, if possible. In the event that any of these queries fails (returns \perp) the honest protocol run aborts. This honest protocol run can be thought of as a genuine attempt to execute a protocol execution.

Definition 10 ((weak) Synchronization Robustness). Let Π be a protocol. Let $G_{\Pi}^{\text{wSR}}(\mathcal{A})$ with boxed text or $G_{\Pi}^{\text{SR}}(\mathcal{A})$ with dashed boxed text be the following game:

- The challenger initializes n parties and their keys;
- \mathcal{A} may issue queries NewSessionI, NewSessionR and Send as defined above;
- Once \mathcal{A} has output (i, j, s, t) to indicate its conclusion, the experiment outputs 1 if and only if the following conditions hold:
 1. $\pi_i^s.\text{pid} = P_j$ and $\pi_j^t.\text{pid} = P_i$;
 2. $\pi_i^s.\text{sk} \neq \pi_j^t.\text{sk}$ or both values are \perp ;

- 3. an honest protocol run was executed between π_i^s and π_j^t ;
- 4. no queries were made by \mathcal{A} to interrupt the protocol execution between π_i^s and π_j^t .
- 4. no protocol messages in the transcripts of π_i^s and π_j^t were sent to any *other* oracles before they were delivered in the honest run.

Define the advantage of an adversary \mathcal{A} in the XX security experiment $G_{\Pi}^{XX}(\mathcal{A})$, for $XX \in \{\text{wSR}, \text{SR}\}$, as

$$\text{Adv}_{\Pi}^{XX}(\mathcal{A}) := \Pr \left[G_{\Pi}^{XX}(\mathcal{A}) = 1 \right].$$

Notes on the definitions.

- Condition (4.) in the SR experiment states that for each genuine protocol message in the ‘target’ session, \mathcal{A} must not have provided this message to any other oracles before that message is delivered as part of the ‘target’ run. This prevents a trivial attack where \mathcal{A} delivers the final protocol message to two oracles: first to some other oracle than the ‘target’ oracle (but of the same party), then to the target oracle. When the (genuine) protocol message is delivered to the party for the second time the target oracle would abort. The parties have still created exactly one key for this genuine protocol run, and so condition (4.) essentially fixes the allowable output oracles as the ones that are processing protocol messages for the first time. (Replay attacks are not an issue in the wSR setting, since the execution must be uninterrupted and so any action made *after* that run has occurred has no impact on \mathcal{A} ’s chances of winning.)
- We do not allow Corrupt queries in this definition: in all of the protocols in this paper we assume pairwise shared key material (and specifically, no keys that are used by a party for communication with multiple other parties). This means that the adversary is not allowed to corrupt the parties in the target run with respect to each other, and that all other Corrupt queries will be of no benefit to an attacker. A similar argument follows for RevealKey queries. This simplifies the security experiment, while capturing the property that we wish to assess.
- In an alternative formulation of our definitions, the target protocol run could be performed by the challenger as an Execute query as seen in past literature [BPR00]. We avoid this approach for two reasons. First, in the SR case, in order to support interleaving, the adversary would have to call the challenger to initiate each stage of the execution (i.e. $k + 1$ times for a k -message protocol), and this is notationally awkward. Secondly and perhaps more importantly, our model allows the adversary to attempt to win its game in multiple protocol runs, and output the oracles which provides the best chance of success. Thus to retain the strength of the definition we would require *multiple* Execute queries, resulting in a model that looks very similar to what we have presented here.

4 Linear Key Evolution

In this section we present a number of protocols that use linear key evolution to derive session keys. All of these protocols achieve wSR. It is not hard to see that full robustness (SR) is not achievable with linearly evolving protocols. To win the SR game the adversary makes a new complete protocol run after the target run has started and the session key is computed at one party, but before the session key is computed at the

second party. This means that when the target session completes, the long-term key has already evolved and the key will be computed with the wrong version of the long-term key at the second party. Either the session will fail at the second party or the key will be different at the two parties. (There is a third case when the key is independent of the long-term key, but in that case the protocol fails to achieve key indistinguishability.)

The first linear key evolving protocol that we present, LP3, exchanges three short messages and has the attractive property of bounding the gap between the counters of the two parties. We present two further protocols which are even more efficient at the cost of some restrictions. LP2 is a two-message protocol but in order to maintain mutual authentication we insist that parties running LP2 have fixed their role as either initiator or responder (not an unreasonable assumption in many application scenarios). Our simplest protocol, LP1, has only a single message but, in addition to requiring fixed roles, like any other one-message protocol it can only achieve unilateral authentication. For all of our protocols we provide theorems guaranteeing authentication, key indistinguishability and weak synchronization robustness (wSR) security.

Syntax and Conventions. All protocols in this paper use message authentication codes to ensure that parties can only process messages that are meant for them. This means that party A stores a key K_{AB}^{MAC} (static) for MAC and key derivation key k_{AB}^{CTR} (evolving) to communicate with B , and K_{AC}^{MAC} and k_{AC}^{CTR} to communicate with C , etc. We describe the key derivation process in more detail in Sec. 4.1.

In LP2 and LP3, the party sending the protocol message includes its own identity in the MAC computation: this stops redirection/reflection attacks of protocol messages to the sending party. For LP1 this is not necessary since the sending party advances after sending its protocol message, meaning that its state is ahead and therefore it is unable to process messages that it has already sent.

4.1 Key Derivation via Linear Evolution

Before looking at specific protocols, we define what we mean by linear key evolution and present an abstract security definition for it. Party A holds a *key derivation key* k_{AB}^{CTR} for use in communication with party B , where the value CTR is an integer that defines the current key state, which is the number of times the key has evolved since its creation. After a party has participated in a key exchange run and computed a session key, it will apply a function Advnc to this key derivation key in order to obtain the next key derivation key and update the counter. This process is detailed in Fig. 3a. Looking ahead, forward security will be obtained if the function that computes $k_{AB}^{\text{CTR}+1}$ from k_{AB}^{CTR} is one-way: this stipulation ensures that an adversary corrupting a party has no way to move upwards in the figure.

The initial “key derivation key” (KDK) is k_{AB}^0 . Subsequent KDKs are derived using a pseudorandom function PRF with $\mathcal{K}_{\text{PRF}} = \mathcal{R}_{\text{PRF}}$ as

$$k_{AB}^{i+1} = \text{PRF}(k_{AB}^i, \text{"ad"}) \quad (1)$$


and session keys are derived as

$$sk_{AB}^i = \text{PRF}(k_{AB}^i, \text{"der"})$$

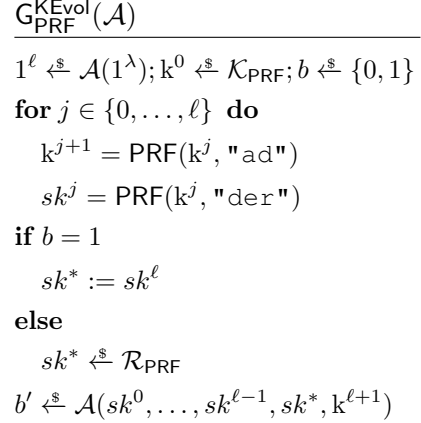
where “ad”(“advance”) and “der”(“derive”) are constant labels used for domain separation.

Furthermore, for convenience, we define a function Advnc which performs multiple key derivations, if necessary. That is, $\text{Advnc}(k_{AB}^i, i, z)$ takes an i -th key derivation key for some counter i and an integer z , and applies PRF iteratively z times to obtain the $(i+z)$ -th KDK such that (1) is satisfied, and sets $i := i+z$. For example:

$$k_{AB}^{i+z}, i+z \leftarrow \text{Advnc}(k_{AB}^i, i, z).$$



(a) Linear key evolution scheme.



(b) The $G_{\text{PRF}}^{\text{KEvol}}(\mathcal{A})$ security experiment.

Figure 3: Linear key evolution and the corresponding experiment.

Security. For the security proofs of our protocols it will be convenient to have an abstract security definition for such a key derivation scheme, which we will show to be implied by the security of the PRF. To this end, Fig. 3b represents a security experiment for the linear key evolution scheme that we describe. The adversary \mathcal{A} outputs an integer 1^ℓ (in unary, to ensure that the number ℓ is polynomially bounded for any efficient \mathcal{A}), and the adversary's task is to distinguish sk^ℓ from random, when given all prior session keys $sk^0, \dots, sk^{\ell-1}$ and the 'next' key derivation key $k^{\ell+1}$.

Definition 11. The advantage of \mathcal{A} in the KEvol security experiment defined in Fig. 3b for pseudorandom function PRF is defined as

$$\text{Adv}_{\text{PRF}}^{\text{KEvol}}(\mathcal{A}) := \left| \Pr[b = b'] - \frac{1}{2} \right|.$$

We now prove the following straightforward theorem.

Theorem 1. *Let PRF be a pseudorandom function. For any adversary \mathcal{A} against the KEvol security of PRF, there exists an adversary \mathcal{B} against the PRF-sec of PRF such that*

$$\text{Adv}_{\text{PRF}}^{\text{KEvol}}(\mathcal{A}) \leq \ell \cdot \text{Adv}_{\text{PRF}}^{\text{PRF-sec}}(\mathcal{B}).$$

PROOF. The proof is a straightforward hybrid argument. Let H_0 be the original experiment. For $i \in \{1, \dots, \ell\}$ let H_i be an experiment which proceeds exactly like H_0 , except that k^j and sk^{j-1} are chosen uniformly random for all $j \leq i$, while all other keys are generated exactly as in the original experiment.

Let X_i denote the event that $b = b'$ in H_i . Then we have

$$\Pr[X_0] = \text{Adv}_{\text{PRF}}^{\text{KEvol}}(\mathcal{A}) \quad \text{and} \quad \Pr[X_\ell] = \frac{1}{2}$$

because the key sk^* is always uniformly random in H_ℓ , independent of the hidden bit b .

We now construct an adversary \mathcal{B} such that

$$|\Pr[X_i] - \Pr[X_{i+1}]| \leq \text{Adv}_{\text{PRF}}^{\text{PRF-sec}}(\mathcal{B}) \tag{2}$$

for all $i \in \{0, \dots, \ell - 1\}$, which yields the claim. \mathcal{B} proceeds exactly like H_i , except that it defines $k^i = \mathcal{O}(\text{"ad"})$ and $sk^{i-1} = \mathcal{O}(\text{"der"})$, where \mathcal{O} is the PRF oracle provided by the PRF security experiment. If \mathcal{A} outputs b' such that $b = b'$, then \mathcal{B} outputs 1, otherwise 0. Note that if \mathcal{O} implements a “real” PRF, then \mathcal{B} perfectly simulates H_i for \mathcal{A} , whereas if the oracle implements a “random” function, then it perfectly simulates H_{i+1} , which yields (2). \square

4.2 LP3: a Three-Message Protocol

Intuition. In Fig. 4 we present a three-message protocol called LP3, which puts a bound on how far initiator and responder can be out of sync, allows either party to initiate communications, and provides mutual authentication. After the first message is sent by an initiator, the responding party advances to catch up if they are behind. Then they respond, and the initiator does the same if they are behind. A third message confirms that both parties are now in sync again, and only after that a session key is established. We make use of state analysis proofs to show that the gap between the two states will be bounded even if messages are lost on the way (Lemma 6) and extend this proof to a scenario where concurrent runs are allowed (Lemma 7). We then show that the number of concurrent runs is a bound on the gap that can occur. We show in Theorem 8 that this also implies that the protocol achieves weak synchronization robustness (wSR). The protocol uses MACs and nonces to achieve mutual authentication (AKE-M). The functions Advnc and KDF, for PSK advancement and session key derivation respectively, are implemented using a PRF as described in Fig. 3a and Sec. 4.1.

State. The protocol uses nonces on both the initiating (N_A) and responding (N_B) sides. Local session state keeps track of these, and so it is only necessary to send N_A in the first protocol message and only N_B in the second message. The nonce generation procedure is denoted GenN, and this process could be, for example, random selection of a bitstring of some fixed length, or a (per-recipient) counter maintained by the party (note however that this counter is distinct from CTR, which tracks the key derivation key’s evolution stage). This choice depends on the application scenario, and this abstraction is for cleaner proofs. In the absence of a hardware RNG, random nonces require allocated memory for a software CSPRNG code while maintaining a counter requires writing to persistent storage (though such writes must be made anyway in linear key evolving protocols). The probability of a collision in random selection from \mathcal{NS} can be bounded by $\text{coll}[q_N, \text{GenN}] \leq \frac{q_N^2}{2|\mathcal{NS}|}$, and the collision probability of a (per-recipient) counter of size $|\mathcal{NS}|$ that is called q_N times is

$$\text{coll}[q_N, \text{GenN}] = \begin{cases} 0 & \text{for } 0 \leq q_N \leq |\mathcal{NS}| - 1, \\ 1 & \text{for } q_N \geq |\mathcal{NS}|. \end{cases}$$

We do not specify the additional counters required to make LP3 deterministic, so it is specified here as a protocol with random nonces.

4.2.1 AKE-M of LP3

Theorem 2 (AKE-M of LP3). *Let Π be the two-message protocol in Fig. 4, built using $\text{MAC} = \{\text{KGen}, \text{Mac}, \text{Vrfy}\}$ and PRF, with n parties. Then for any adversary \mathcal{A} against the AKE-M security of Π that makes a maximum of q queries that initiate new sessions for each party (with $q < |\text{CTR}|$), there exists an adversary $\mathcal{B}_{2.1}$ against the SEUF-CMA-Q of MAC and an adversary $\mathcal{B}_{2.2}$ against the KEvol security of KDF such that*

$$\text{Adv}_{\Pi}^{\text{AKE-M}}(\mathcal{A}) \leq n^2 \cdot \left(4\text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{2.1}) + 4\text{coll}[q, \text{GenN}] + q \cdot \text{Adv}_{\text{PRF}}^{\text{KEvol}}(\mathcal{B}_{2.2}) \right).$$

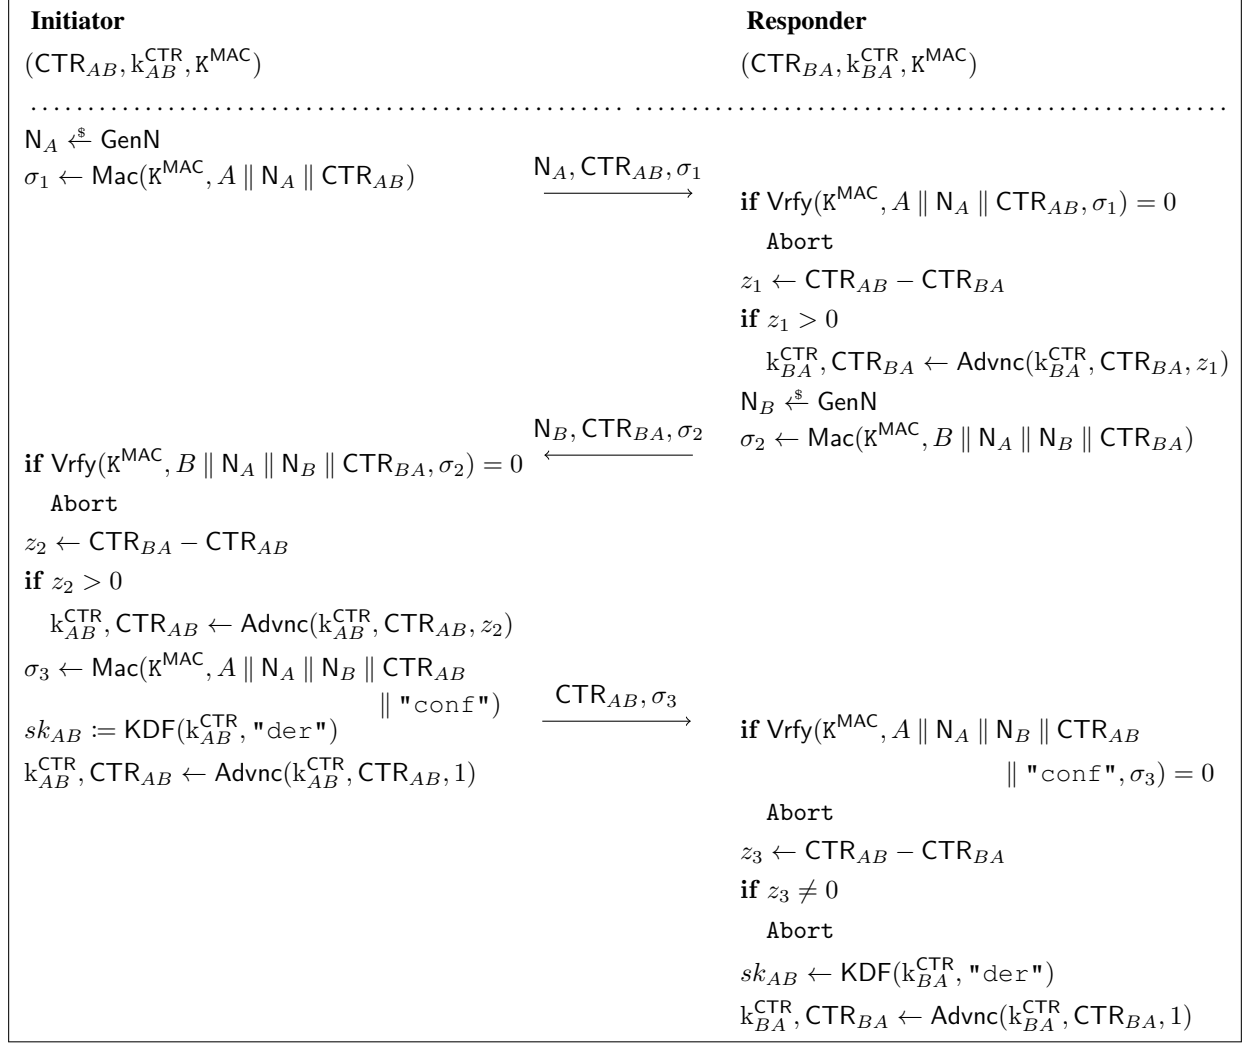


Figure 4: LP3, a three-message protocol.

We form a bound for each of the three ways in which an adversary can break AKE-M security, namely Ent-Auth-R, Ent-Auth-I and Key-Ind, and then sum these bounds. There are two MAC security terms, for entity authentication of responder and initiator, and so $2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{2.1})$ bounds these two terms by fixing $\mathcal{B}_{2.1}$ to be whichever of $\mathcal{B}_{2.1r}$ and $\mathcal{B}_{2.1i}$ has greater advantage.

We now give intuition regarding the Ent-Auth proofs. Note that since $q < |\text{CTR}|$, if there is no collision in the generation of N_A for party A then all first protocol messages are unique. Further, after receiving a first protocol message a responder always generates a nonce, thus if no nonce collisions occur, sending the same first protocol message to two (or more) different responder oracles will result in distinct transcripts. This rules out the possibility of *multiple* oracles with non-unique matching conversations accepting, in either Ent-Auth-R or Ent-Auth-I. To conclude our proofs, we need to show that the only way an adversary can force an oracle to accept and for there not to exist any other oracle with a matching transcript, the adversary must forge a MAC message.

For Lemma 3 (Ent-Auth-R), a responder oracle accepts when it receives (what it believes to be) a third

protocol message, and so to win there must not exist any oracle with the same partial transcript as this accepting oracle. This accepting oracle's transcript must consist of (i) an input first protocol message, (ii) the resulting second protocol message, and (iii) the third protocol message that resulted in accept being reached. Since the MAC is calculated on both nonce values and the counter value, the first and third input messages to the accepting oracle must have been generated by an oracle of the communication partner. So the only viable route to victory if a forgery has not occurred is if these messages came from different oracles: and since the nonces are stored as part of the session state, this victory could only occur in the event of a nonce collision.

A similar argument applies for Lemma 4 (Ent-Auth-I), except now the accepting oracle computes a session key after receiving a valid second protocol message. Again that second protocol message must have come from a genuine invocation of NewSessionR by the intended partner if no forgery has occurred, and thus that oracle's transcript is a prefix of the accepting oracle's transcript.

Lemma 3 (Ent-Auth-R of LP3). *For any adversary \mathcal{A} , the probability that there exists an oracle with $\rho = \text{Responder}$ that accepts maliciously can be bounded by*

$$\text{Adv}_{\Pi}^{\text{Ent-Auth-R}}(\mathcal{A}) \leq n^2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{2.1r}) + \text{coll}[q, \text{GenN}]$$

where all quantities are defined as stated in Theorem 2.

PROOF. We proceed using a sequence of games.

Game 0. This is the original Ent-Auth game.

$$\Pr [G_{\Pi}^{\text{Ent-Auth}}(\mathcal{A}) = 1] = \Pr [G_0^{\mathcal{A}} = 1]$$

Game 1. This game is the same as Game 0 except that we assume all nonces generated by NewSessionI and NewSessionR queries are different, and abort otherwise. The only way an adversary can notice the difference between these games is by making nonces collide, so we can write the following:

$$\Pr [G_0^{\mathcal{A}} = 1] \leq \Pr [G_1^{\mathcal{A}} = 1] + \text{coll}[q, \text{GenN}].$$

Game 2. In this game we guess which responder will be the first to accept maliciously and its partner identity, and abort if this guess is wrong. The game is the same as Game 1 except that the challenger guesses $(i^*, j^*) \xleftarrow{\$} [n] \times [n]$, and if an oracle π_i^s (for some s) accepts maliciously with $\pi_i^s.\rho \neq \text{Responder}$ or $i^* \neq i$ or $j^* \neq \pi_i^s.\text{pid}$, then the challenger aborts.

$$\Pr [G_1^{\mathcal{A}} = 1] = n^2 \cdot \Pr [G_2^{\mathcal{A}} = 1]$$

In order for an adversary to win without an abort in Game 2, it must make a responder oracle (of party i^*) accept maliciously, with no initiator oracle (of party j^*) having a matching conversation.

We construct a reduction $\mathcal{B}_{2.1r}$ that is playing against the SEUF-CMA-Q security of MAC that simulates the environment for an underlying adversary \mathcal{A} that attempts to win in game Game 2.

The reduction generates (initial) key derivation keys for all pairs of parties, and authentication keys for all pairs of parties except i^* and j^* . When responding to queries by \mathcal{A} regarding all other parties, the reduction will honestly provide messages as specified in the protocol specification and the Ent-Auth game. For any query made between oracles of parties i^* and j^* , the reduction will use its \mathcal{O}_{Mac} oracle and provide

the received value in its simulation for \mathcal{A} . For example, to initialize initiator oracles π_{j^*} , the reduction calls $N \leftarrow \text{GenN}$, checks the current state counter $\text{CTR}_{j^*i^*}$ and calls $\mathcal{O}_{\text{Mac}}(j^* \parallel N \parallel \text{CTR}_{j^*i^*})$. If \mathcal{A} provides any value that it has not been given by a prior protocol message as input to a Send query from i^* to j^* or j^* to i^* , then the reduction sends this to its $\mathcal{O}_{\text{Vrfy}}$ oracle in the SEUF-CMA-Q game. The simulation of Game 2 is perfect and any win for \mathcal{A} directly corresponds to a valid signature forgery, so we can write

$$\Pr [G_2^{\mathcal{A}} = 1] \leq \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{2.1r}).$$

Summing these terms gives the bound in the statement of Lemma 3. □

Lemma 4 (Ent-Auth-I of LP3). *For any adversary \mathcal{A} , the probability that there exists an oracle with $\rho = \text{Initiator}$ that accepts maliciously can be bounded by*

$$\text{Adv}_{\Pi}^{\text{Ent-Auth-I}}(\mathcal{A}) \leq n^2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{2.1i}) + \text{coll}[q, \text{GenN}]$$

where all quantities are defined as stated in Theorem 2.

PROOF. Games 0, 1 and 2 are as in the proof of Lemma 3. The reduction $\mathcal{B}_{2.1i}$ is as in the proof of Lemma 3, except that now we are guessing which initiator will be the first to accept maliciously (and its intended partner), and so the abort occurs if a responder oracle accepts maliciously. The loss of n^2 incurred by selection of parties is the same. □

The proof of key indistinguishability is very similar to that of Lemma 12 and the game hops proceed following the same strategy. Again we use a reduction to the KEvol security of PRF.

Lemma 5 (Key-Ind of LP3). *For any adversary \mathcal{A} and (any fixed) entity authentication adversary $\mathcal{B}_{2.1}$, the probability that \mathcal{A} answers the Test challenge correctly can be bounded by*

$$\text{Adv}_{\Pi}^{\text{Key-Ind}}(\mathcal{A}) \leq \text{Adv}_{\Pi}^{\text{Ent-Auth}}(\mathcal{B}_{2.1}) + n^2 \cdot q \cdot \text{Adv}_{\text{PRF}}^{\text{KEvol}}(\mathcal{B}_{2.2})$$

where all quantities are defined as stated in Theorem 2.

PROOF. Let b' be the bit output by \mathcal{A} in each game, and b be the bit sampled as part of the Test query.

Game 0. This is the original Key-Ind game.

$$\Pr [G_{\Pi}^{\text{Key-Ind}}(\mathcal{A}) = 1] = \Pr [G_0^{\mathcal{A}} = 1]$$

Game 1. This game is the same as Game 0, except the challenger aborts and chooses $b' \xleftarrow{\$} \{0, 1\}$ if any oracle accepts maliciously.

$$\Pr [G_0^{\mathcal{A}} = 1] \leq \Pr [G_1^{\mathcal{A}} = 1] + \text{Adv}_{\Pi}^{\text{Ent-Auth-R}}(\mathcal{B}_{2.1})$$

At this stage, the oracle to which \mathcal{A} asks its Test query has a unique partner oracle with a matching conversation.

Game 2. This game is the same as Game 1, except the challenger guesses the two parties involved in the Test query via $(i^*, j^*) \leftarrow_{\mathcal{S}} [n] \times [n]$, and additionally guesses the counter value which identifies the key derivation key state of the session key in the Test query. If \mathcal{A} issues a Test($\pi_{i^*}^s$) query with $i^* \neq i$ or $j^* \neq \pi_{i^*}^s.\text{pid}$ (for some s), or the session key computed via Test($\pi_{i^*}^s$), i.e. $sk_{i^*j^*}$ is not equal to $\text{KDF}(k_{i^*j^*}^{\text{CTR}}, \text{"der"})$, then the challenger aborts. Note that the challenger does not guess which oracle of i^* the Test query will be made to, only the counter value linked to the session key in that query.

$$\Pr [G_1^{\mathcal{A}} = 1] \leq n^2 \cdot q \cdot \Pr [G_2^{\mathcal{A}} = 1]$$

At this stage, if the challenger has guessed correctly then the Test query will be asked to an oracle after the key derivation counter has been advanced a fixed number of times, and this oracle has unique partner with a matching conversation.

Game 3. This game is the same as Game 2, except that when the challenger runs KDF on the key derivation values used in the Test query, the challenger instead responds with a random key from the session key space. Noticing this change results in an adversary that is successful in the KEvol game for PRF, so we can write

$$\Pr [G_2^{\mathcal{A}} = 1] = \Pr [G_3^{\mathcal{A}} = 1] + \text{Adv}_{\text{PRF}}^{\text{KEvol}}(\mathcal{B}_{2.2}).$$

The reduction $\mathcal{B}_{2.2}$ is detailed in Fig. 5. $\mathcal{B}_{2.2}$ initially guesses the target parties in the Test session and the counter value associated with the Test session, as per the previous game hop.

For this reduction the challenger determines how an oracle should process each Send(π_i^s) query using a label $\text{MsgNr} \in \{1, 2, 3\}$ as follows. If $\pi_i^s.\rho = \text{Responder}$ and the message is of the form N, CTR, σ then $\text{MsgNr} \leftarrow 1$. If $\pi_i^s.\rho = \text{Initiator}$ then $\text{MsgNr} \leftarrow 2$. If $\pi_i^s.\rho = \text{Responder}$ and the message is of the form CTR, σ then $\text{MsgNr} \leftarrow 3$. For all other query types or improperly formatted messages, the challenger returns \perp and sets $\pi_i^s.\alpha \leftarrow \text{reject}$

In the event that $\mathcal{B}_{2.2}$ is in the ‘real’ version of its own game, where it receives a genuine evaluation of the function KDF, $\mathcal{B}_{2.2}$ perfectly simulates Game 2 for \mathcal{A} , and otherwise it perfectly simulates Game 3.

At this stage, the Test query is asked on a key that is randomly chosen, and thus independent of the protocol and the security game. Consequently,

$$\Pr [G_3^{\mathcal{A}} = 1] = \frac{1}{2} \Rightarrow \text{Adv}_{\Pi}^3(\mathcal{A}) = 0.$$

□

4.2.2 Bounded Gap: Non-Concurrent Setting.

We will now prove that the ‘gap’ between the state of the two parties in LP3 is bounded in the non-concurrent setting, that is:

Lemma 6. *Let A and B be respectively the initiator and the responder of a single — non-concurrent — LP3-run. Let δ_{AB} be the gap between A and B with respect to the evolution of the master keys of both parties. Then $\delta_{AB} \in \{-1, 0, 1\}$, assuming MAC-security.*

The messages in LP3 are counted in a natural way, as indicated in Fig. 6a. For this non-concurrent setting the proof is similar to [ACF20, Lemma 1]. Then the notation “(CTR_{AB}, CTR_{BA})” means that, when the run ends, the last valid message received by A is message of index CTR_{AB}, and the last valid

Reduction $\mathcal{B}_{2.2}$ playing $G_{\text{KDF}}^{\text{KEvol}}(\mathcal{B}_{2.2})$

```

1 :  $i^*, j^* \xleftarrow{\$} [n]; \text{CTR}^* \xleftarrow{\$} [q]$ 
2 : for  $i, j \in [n]$  do
3 :    $\text{CTR}_{ij} = \text{CTR}_{ji} \leftarrow 0$ 
4 :    $\text{K}_{ij}^{\text{MAC}} = \text{K}_{ji}^{\text{MAC}} \xleftarrow{\$} \mathcal{K}_{\text{MAC}}$ 
5 :   for  $[n] \times [n] \setminus (i^*, j^*)$  do
6 :      $k_{ij}^0 = k_{ji}^0 \xleftarrow{\$} \mathcal{K}_{\text{PRF}}$ 
7 :   output  $\text{CTR}^*$ 
8 :   receive  $(sk^0, \dots, sk^{\text{CTR}^*-1}, sk^*, k^{\text{CTR}^*+1})$ 
9 :    $k_{i^*j^*}^{\text{CTR}^*+1} \leftarrow k^{\text{CTR}^*+1}$ 
10 :  $\mathcal{A}^{\text{oracles}}$ 
11 : When  $\mathcal{A}$  calls  $\text{Test}(\pi_i^s)$  do
12 :   if  $i \neq i^*$  or  $j^* \neq \pi_{i^*}^s \cdot \text{pid}$ 
13 :     return Abort
14 :   return  $\pi_i^s \cdot sk$ 
15 :    $(i^*, s^*, b') \xleftarrow{\$} \mathcal{A}(\pi_i^s \cdot sk)$ 
16 :   return  $b'$ 

```

NewSessionI(π_i^s, pid)

```

17 :  $\pi_i^s \cdot \rho \leftarrow \text{Initiator}$ 
18 :  $\pi_i^s \cdot \alpha \leftarrow \text{negotiating}$ 
19 :  $\pi_i^s \cdot \text{pid} \leftarrow \text{pid} / =j$ 
20 :  $N \leftarrow \text{GenN}$ 
21 :  $\sigma_1 \leftarrow \text{Mac}(\text{K}^{\text{MAC}}, i \parallel N \parallel \text{CTR}_{ij})$ 
22 :  $m' \leftarrow N, \text{CTR}_{ij}, \sigma_1$ 
23 : return  $m'$ 

```

NewSessionR(π_i^s, pid, m)

```

24 :  $\pi_i^s \cdot \rho \leftarrow \text{Responder}$ 
25 :  $\pi_i^s \cdot \alpha \leftarrow \text{negotiating}$ 
26 :  $\pi_i^s \cdot \text{pid} \leftarrow \text{pid}$ 
27 : do  $\text{Send}(\pi_i^s, m)$ 

```

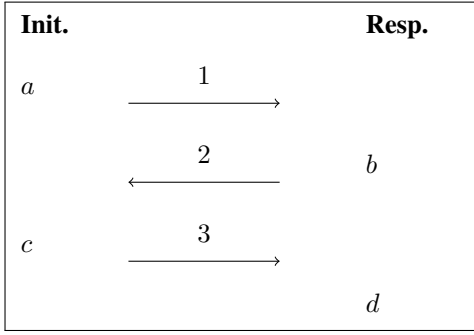
$\text{Send}(\pi_i^s, m) / \text{pid} = j$

```


28 : if  $\text{MsgNr} = 1$ 
29 :   if  $\text{Vrfy}(\text{K}^{\text{MAC}}, j \parallel N_j \parallel \text{CTR}_{ji}, \sigma_1) = 0$ 
30 :     return Abort
31 :    $z_1 \leftarrow \text{CTR}_{ji} - \text{CTR}_{ij}$ 
32 :   if  $z_1 > 0$ 
33 :      $k_{ij}, \text{CTR}_{ij} \leftarrow \text{Advnc}(k_{ij}, \text{CTR}_{ij}, z_1)$ 
34 :    $N_i \leftarrow \text{GenN}$ 
35 :    $\sigma_2 \leftarrow \text{Mac}(\text{K}^{\text{MAC}}, i \parallel N_j \parallel N_i \parallel \text{CTR}_{ij})$ 
36 :    $m' \leftarrow N_i, \text{CTR}_{ij}, \sigma_2$ 
37 :   return  $m'$ 
38 : if  $\text{MsgNr} = 2$ 
39 :   if  $\text{Vrfy}(\text{K}^{\text{MAC}}, j \parallel N_i \parallel N_j \parallel \text{CTR}_{ji}, \sigma_2) = 0$ 
40 :     return Abort
41 :    $z_2 \leftarrow \text{CTR}_{ji} - \text{CTR}_{ij}$ 
42 :   if  $z_2 > 0$ 
43 :      $k_{ij}, \text{CTR}_{ij} \leftarrow \text{Advnc}(k_{ij}, \text{CTR}_{ij}, z_2)$ 
44 :    $\sigma_3 \leftarrow \text{Mac}(\text{K}^{\text{MAC}}, i \parallel N_i \parallel N_j \parallel \text{CTR}_{ij} \parallel \text{"conf"})$ 
45 :    $m' \leftarrow \text{CTR}_{ij}, \sigma_3$ 
46 :   return  $m'$ 
47 : else /  $\text{MsgNr} = 3$ 
48 :   if  $\text{Vrfy}(\text{K}^{\text{MAC}}, j \parallel N_j \parallel N_i \parallel \text{CTR}_{ji} \parallel \text{"conf"}, \sigma_3) = 0$ 
49 :     return Abort
50 :    $z_3 \leftarrow \text{CTR}_{ji} - \text{CTR}_{ij}$ 
51 :   if  $z_3 \neq 0$ 
52 :     return Abort
53 :   if  $\text{MsgNr} \in \{2, 3\}$ 
54 :     if  $(i, j) = (i^*, j^*) \vee (j^*, i^*)$ 
55 :       if  $\text{CTR}_{ij} < \text{CTR}^*$ 
56 :          $\pi_i^s \cdot sk \leftarrow sk^{\text{CTR}_{ij}}$ 
57 :       if  $\text{CTR}_{ij} = \text{CTR}^*$ 
58 :          $\pi_i^s \cdot sk \leftarrow sk^*$ 
59 :        $s^* \leftarrow s$ 
60 :        $\text{CTR}_{ij} \leftarrow \text{CTR}_{ij} + 1$ 
61 :     else
62 :        $\pi_i^s \cdot sk \leftarrow \text{KDF}(k_{ij}^{\text{CTR}}, \text{"der"})$ 
63 :        $k_{ij}^{\text{CTR}}, \text{CTR}_{ij} \leftarrow \text{Advnc}(k_{ij}^{\text{CTR}}, \text{CTR}_{ij}, 1)$ 
64 :        $\pi_i^s \cdot \alpha \leftarrow \text{accept}$ 

```

Figure 5: Reduction $\mathcal{B}_{2.2}$ for the proof of Lemma 5. If at any time \mathcal{A} causes an oracle to accept maliciously, then $\mathcal{B}_{2.2}$ simply does Abort. $\mathcal{B}_{2.2}$ provides \mathcal{A} with access to $\text{oracles} = \text{NewSessionI}(\cdot, \cdot, \cdot), \text{NewSessionR}(\cdot, \cdot), \text{Send}(\cdot, \cdot, \cdot), \text{RevealKey}(\cdot), \text{Corrupt}(\cdot, \cdot)$, however RevealKey and Corrupt are omitted for space reasons: they are exactly as in Fig. 9 (Key-Ind proof of LP2).



(a) Numbering of states for the proofs of Lemmas 6 (1, 2, 3) and 7 (a, b, c, d).



(b) Synchronization state for LP3 in the non-concurrent setting.

Figure 6: Different states for LP3, and transitions between them.

message received by B is message of index CTR_{AB} . We call a (CTR_{AB}, CTR_{BA}) -run a run where the last message received by A is message CTR_{AB} , and the last message received by B is message CTR_{BA} . By convention $CTR_{AB} = 0$ means that no message has been received by A . In Fig. 6b, we define the states to be the different values of δ_{AB} . The transitions are the possible messages. An example: if our protocol instance is in state $\delta_{AB} = -1$, and B responds to message 1 with message 2, i.e. transition $(2, 1)$ in the state diagram, the initiator will advance twice and the state will be $\delta_{AB} = 1$. A then sends the third message: transition $(2, 3)$ takes place and we end up in state $\delta_{AB} = 0$ since this third message will cause the responder to advance.

PROOF. We prove Lemma 6. The protocol is initialized with $\delta_{AB} = 0$ and the first step is sending message 1: either the message never reaches the responder, or the message is received correctly. In either case neither party advances, so $\delta_{AB} = 0$ — i.e. transition $(0, 1)$ in Fig. 6b is fired. If the protocol now terminates we end up in state 0, while sending and receiving message 2 would cause the initiator to advance, or in terms of the state diagram, fire $(2, 1)$ and transition to $\delta_{AB} = 1$.

Because we restrict ourselves to non-concurrent executions, the only possible option no matter the state is to advance with one message or terminate and start from $(0, 1)$. Adding all possible transitions to the state diagram, we observe that there are no reachable states other than 0 and 1. Since the protocol does not have fixed roles we can reach a state -1 by changing roles after we reached state 1. From there, there are two transitions that bring us back to states 0 and 1. Since we assume that MACs cannot be forged, these are the only reachable states, thus $\delta_{AB} \in \{-1, 0, 1\}$ always holds. \square

4.2.3 Bounded Gap: Concurrent Setting.

We will now extend Lemma 6 to the concurrent setting.

Lemma 7. *Let A and B be respectively the initiator and the responder of C concurrent LP3-runs. Let δ_{AB} be the gap between A and B with respect to the evolution of the master keys of both parties. Then $-C \leq \delta_{AB} \leq 1 + C$, assuming MAC-security.*

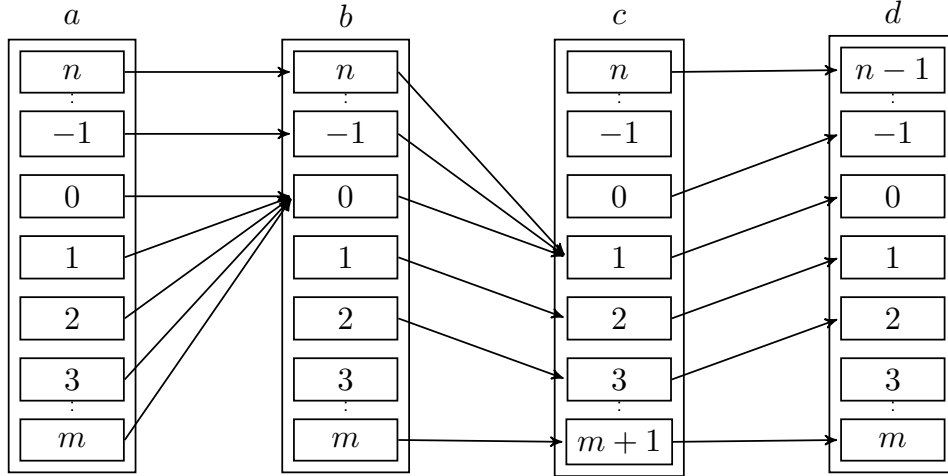


Figure 7: Synchronization state for LP3 in the non-concurrent setting.

To illustrate the (in a sense) multidimensional effect of concurrent runs on the protocol, we will now use a different message labelling convention. Figure 6a defines the different states the protocol execution can be in. The state diagram in Fig. 7 now uses these four possible protocol states as diagram states — a message between state a and b is thus necessarily message 1. The internal state of the four ‘macro states’ in the diagram now represents the value of δ_{AB} .

Observe that for the transitions from a to b and from b to c , i.e. the sending of messages 1 and 2, respectively, the evolution of δ_{AB} depends on the actual value of a . For all transitions caused by message 3, the change is systematic:

1. Any transition from c to d will decrease δ_{AB} by 1;
2. any transition from b to c will increase δ_{AB} by at least 1.

Additionally there are two ‘resets’, since

3. any transition from a to b will set δ_{AB} to 0, if the gap is 1 or more;
4. any transition from b to c will set δ_{AB} to 1, if the gap is 0 or less.

PROOF. We prove Lemma 7. In Lemma 6, the normal range is defined as $\delta_{AB} \in \{-1, 0, 1\}$. Extremities beyond this range are possible when the condition in 1. or 2. above occurs during a run, so each consecutive run can influence δ_{AB} with -1 or $+1$ at most. Since we assume MAC-security, the adversary cannot influence the protocol with messages other than those authentically sent during one of the runs. Inductively, we conclude $-C \leq \delta_{AB} \leq 1 + C$. \square

4.2.4 wSR of LP3.

We now argue that LP3 obtains weak synchronization robustness (wSR), the property that captures how well a protocol can recover from network errors and interleaving of protocol runs. In the wSR game the adversary can make arbitrary NewSessionI, NewSessionR and Send queries, and at its conclusion it outputs

the identifiers of two oracles: it is said to win the wSR game if these oracles engaged in an uninterrupted protocol run but did not compute the same session key. As such, a proof of wSR must argue that whatever values of party state exist before the target protocol run occurs, neither of the parties will abort and both will arrive at the same session key.

Our general approach for proving robustness of all of the protocols in this paper is to separate adversaries that win the wSR game via forging a MAC value, and those that do not produce a forgery during their execution. LP1 (Fig. 11) and LP2 (Fig. 8) have fixed roles and as a result the initiator's counter value must always be at least the size of the responder's counter value for the protocol to have correctness. Thus a MAC forgery can force the responding party's counter value to be arbitrarily large, and the target protocol run will cause at least one party to abort, and the adversary wins the wSR game. LP3, on the other hand, is actually not vulnerable in the sense of synchronization robustness if a MAC forgery does occur. This is due to LP3 being designed to have correctness for all starting (integer) counter values, since in any session, both parties can catch up from being arbitrarily far behind.

We formally prove this below, however to see this visually, consider Fig. 7 for the execution of a single protocol run, i.e. from a to d . For any initial state difference a , the state c after the second protocol message has been processed is always 1 (the initiator computes a session key and advances once), leading to state difference 0 after the responder processes the final protocol message (deriving a session key and advancing once).

Theorem 8 (wSR of LP3). *Let Π be the three-message protocol in Fig. 4, built using $\text{MAC} = \{\text{KGen}, \text{Mac}, \text{Vrfy}\}$ and PRF with n parties. Then for any adversary A against the wSR security of Π , $\text{Adv}_{\Pi}^{\text{wSR}}(A) = 0$.*

PROOF. The only places where **Abort** occurs in the protocol description (Fig. 4) are after MAC verification failures: in the target protocol session all messages are honestly generated so this cannot occur (assuming perfect correctness of MAC). As a result, the only route to victory in the wSR game for an adversary is to make the parties compute different session keys. This occurs if the parties compute session keys but have different counter values once all three protocol messages have been delivered and processed: following the notation and arguments in Lemma 7, this is the same as showing that $\delta = 0$ after a $(2, 3)$ session for any starting delta value. More precisely, let A and B be the parties involved in the target session where A sends the first protocol message, let δ_{AB}^{pre} be the gap between A and B with respect to the evolution of the master keys of both parties and the point *before* the target session begins (i.e. before the adversary calls `NewSessionI` for the target session), and let $\delta_{AB}^{\text{post}}$ be the gap *after* the target session has occurred. Figure 6b shows that $\delta_{AB}^{\text{post}} = 0$ for $\delta_{AB}^{\text{pre}} \in \{-1, 0, 1\}$, so to complete the proof we need to show that this also holds for arbitrary δ_{AB}^{pre} .

If $\delta_{AB}^{\text{pre}} \in \{1, 2, \dots\}$, i.e. CTR_{AB} is ahead of CTR_{BA} by $\delta_{AB}^{\text{pre}} = z_1$ steps, then the first protocol message processing by B results in B advancing its counter CTR_{BA} by δ_{AB}^{pre} steps, leading to state difference 0. This means that A will not advance on receiving the second protocol message and both parties will compute a session key for state CTR_{AB} and then advance once, and so $\delta_{AB}^{\text{post}} = 0$.

If $\delta_{AB}^{\text{pre}} \in \{-1, -2, \dots\}$, i.e. CTR_{BA} is ahead of CTR_{AB} by $-\delta_{AB}^{\text{pre}} = z_2$ steps, B does not advance in processing the first message, however A does advance by $-\delta_{AB}^{\text{pre}} = z_2$ steps on receiving the second protocol message. Again this leads to state difference 0 and here a session key is computed for state CTR_{BA} and then both parties advance once, so $\delta_{AB}^{\text{post}} = 0$.

This concludes the proof, since any initial state will lead to the target protocol run computing the same session key for the involved parties. \square

4.3 LP2: A Two-Message Protocol with Fixed Roles

In Fig. 8 we present a two-message protocol, LP2, with linear key evolution. The roles of initiator and responder are fixed, so the same party initiates every session: this is enforced by $CTR_{AB} \geq CTR_{BA}$ (for A initiating).

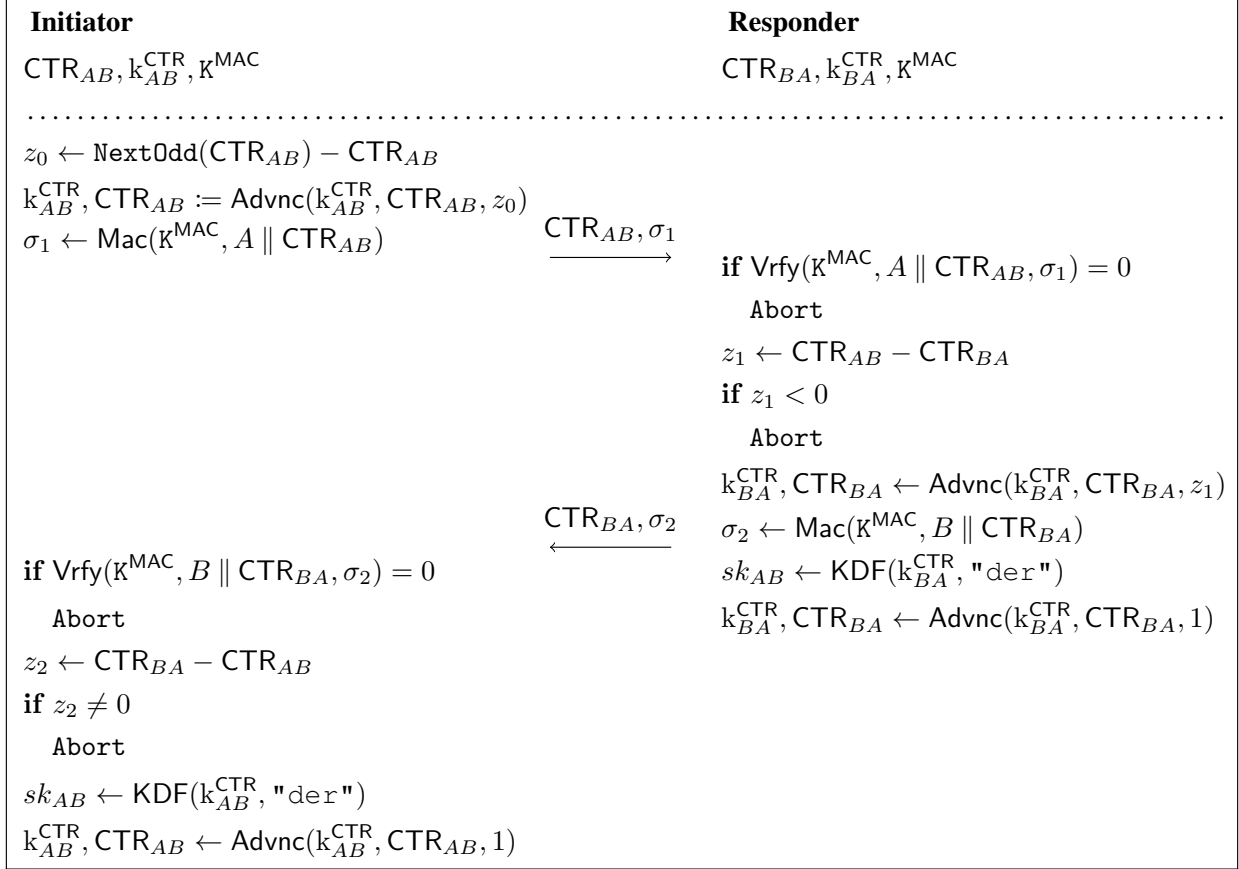


Figure 8: LP2, a two-message protocol with fixed roles.

Achieving weak synchronization robustness (wSR) is slightly more complicated in LP2 than it was in LP3. If we were to adapt LP3 to a two-message protocol by simply dropping the last message and having the responder accept (thus, deriving a session key and advancing its state), we could end up in a situation where we break the requirement that the responder should never advance past the state of the initiator. In this hypothetical protocol, the initiator will initiate the key exchange, but will not derive a session key until it has authenticated the responder. The responder, however, will authenticate the initiator upon receiving the first protocol message (rather than waiting for a key confirmation message as in LP3) and produce the second protocol message, after which it will immediately derive a session key and advance its state. Thus, if this second protocol message is not delivered, the responder will have advanced its state, but the initiator has not, contradicting our requirement that $CTR_{AB} \geq CTR_{BA}$.

In order to avoid this in LP2, the initiator A will always advance to the next *odd* value of its counter at the beginning of each session. How many steps the initiator advances depends on what has happened earlier. If a complete session has been executed as A 's previous action, A starts by advancing once, so that its state

counter is ahead of B . If in the previous session A never processed the second protocol message, A will advance twice at the beginning of the next session, in order to catch up to B and move ahead. The reasoning behind this is the separation of A 's counter set: if the counter is an even integer then A has most recently received a message (and derived a key), whereas if it is an odd integer then A most recently sent a (session opening) protocol message. In both cases, advancing to $\text{NextOdd}(\text{CTR}_{AB})$ will have the desired effect.

With this simpler protocol we are able to achieve most of the desired properties from SP3, but with a more lightweight protocol. Fixing the roles makes this possible, and this demonstrates the fine balance between forward security and (weak) synchronization robustness. In the event that the reduced communication complexity of LP2 compared to LP3 is desirable when choosing a protocol, but if the application demands that either party can initiate, it is possible to run LP2 in *duplex mode*. In duplex mode, both parties keep separate key derivation keys and counters for initiating and responding such that both parties can have both roles without violating the condition $\text{CTR}_{AB} \geq \text{CTR}_{BA}$.

4.3.1 AKE-M of LP2

Theorem 9 (AKE-M of LP2). *Let Π be the two-message protocol in Fig. 8, built using $\text{MAC} = \{\text{KGen}, \text{Mac}, \text{Vrfy}\}$, and PRF, with n parties. Then for any adversary \mathcal{A} against the AKE-M security of Π that makes a maximum of q queries that initiate new sessions for each party (with $q < \frac{|\text{CTR}|}{2}$), there exists an adversary $\mathcal{B}_{9,1}$ against the SEUF-CMA- Q of MAC and an adversary $\mathcal{B}_{9,2}$ against the KEvol security of PRF such that*

$$\text{Adv}_{\Pi}^{\text{AKE-M}}(\mathcal{A}) \leq n^2 \cdot \left(4 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{9,1}) + q \cdot \text{Adv}_{\text{PRF}}^{\text{KEvol}}(\mathcal{B}_{9,2}) \right).$$

We form a bound for each of the three ways in which an adversary can break AKE-M security, namely Ent-Auth-R, Ent-Auth-I and Key-Ind, and then sum these bounds. There are two MAC security terms, for entity authentication of responder and initiator, and so $2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{9,1})$ bounds these two terms by fixing $\mathcal{B}_{9,1}$ to be whichever of $\mathcal{B}_{9,1r}$ and $\mathcal{B}_{9,1i}$ has greater advantage.

We now give intuition regarding the Ent-Auth proofs. Note that since $q < \frac{|\text{CTR}|}{2}$, all first protocol messages are unique, and thus the first message in every transcript is unique. Further, after receiving a first protocol message a responder always advances its state (at least) once, and so delivering that same first protocol message to any oracle of the same responder party will trigger the *if* $z_1 < 0$ branch and result in Abort. This rules out the possibility of *multiple* oracles with non-unique matching conversations accepting. To conclude our proofs, we need to show that the only way an adversary can force an oracle to accept and for there not to exist any other oracle with a matching transcript, the adversary must forge a MAC message.

For Lemma 10 (Ent-Auth-R), a responder oracle accepts when it receives a first protocol message, and so to win there must not exist any oracle with the same partial transcript as this accepting oracle. Since the MAC is calculated on the the counter value and the communicating parties' identities, the input message to the accepting oracle must have been generated by an oracle of the communication partner (in which case there is a matching conversation and the adversary has not won) or the adversary has produced a MAC forgery (in the SEUF-CMA- Q sense).

A similar argument applies for Lemma 11 (Ent-Auth-I), except now the accepting oracle computes a session key after receiving a valid second protocol message. Again if no forgery has occurred that second protocol message must have come from a genuine invocation of `NewSessionR` by the intended partner (with the first protocol message of the accepting oracle provided as input), and thus that oracle's transcript is a prefix of the accepting oracle's transcript. Specifically, if an initiator session was instantiated after the first protocol message of the accepting oracle was produced then $z_2 \neq 0$ and we have a contradiction since this oracle could not then reach accept.

Lemma 10 (Ent-Auth-R of LP2). *For any adversary \mathcal{A} , the probability that there exists an oracle with $\rho = \text{Responder}$ that accepts maliciously can be bounded by*

$$\text{Adv}_{\Pi}^{\text{Ent-Auth-R}}(\mathcal{A}) \leq n^2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{9.1r})$$

where all quantities are defined as stated in Theorem 9.

PROOF. We proceed using a sequence of games.

Game 0. This is the original Ent-Auth game.

$$\Pr [G_{\Pi}^{\text{Ent-Auth}}(\mathcal{A}) = 1] = \Pr [G_0^{\mathcal{A}} = 1]$$

Game 1. In this game we guess which responder will be the first to accept maliciously and its partner identity, and abort if this guess is wrong. The game is the same as Game 0 except that the challenger guesses $(i^*, j^*) \xleftarrow{\$} [n] \times [n]$, and if an oracle π_i^s (for some s) accepts maliciously with $\pi_i^s.\rho \neq \text{Responder}$ or $i^* \neq i$ or $j^* \neq \pi_i^s.\text{pid}$, then the challenger aborts.

$$\Pr [G_0^{\mathcal{A}} = 1] = n^2 \cdot \Pr [G_1^{\mathcal{A}} = 1]$$

We construct a reduction $\mathcal{B}_{9.1r}$ that is playing against the SEUF-CMA- Q security of MAC that simulates the environment for an underlying adversary \mathcal{A} that attempts to win in game Game 1. The reduction generates (initial) key derivation keys for all pairs of parties, and authentication keys for all pairs of parties except i^* and j^* . When responding to queries by \mathcal{A} regarding all other parties, the reduction will honestly provide messages as specified in the protocol specification and the Ent-Auth game. For any query made between oracles of parties i^* and j^* , the reduction will use its \mathcal{O}_{Mac} oracle and provide the received value in its simulation for \mathcal{A} . For example, to initialize initiator oracles π_{j^*} , the reduction checks the current state counter $\text{CTR}_{j^*i^*}$ and calls $\mathcal{O}_{\text{Mac}}(j^* \parallel \text{CTR}_{j^*i^*})$. If \mathcal{A} provides any value that it has not been given as an initialization query as input to a NewSessionR query from i^* to j^* , then the reduction sends this to its $\mathcal{O}_{\text{Vrfy}}$ oracle in the SEUF-CMA- Q game. The simulation of Game 1 is perfect and any win for \mathcal{A} directly corresponds to a valid signature forgery, so we can write

$$\Pr [G_1^{\mathcal{A}} = 1] \leq \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{9.1r}).$$

Summing these terms gives the bound in the statement of Lemma 10. □

The second proof is very similar, and considers malicious acceptance by an initiator, i.e. as a result of a full protocol run of two messages. We only detail significant changes and note that our term collection is exactly the same.

Lemma 11 (Ent-Auth-I of LP2). *For any adversary \mathcal{A} , the probability that there exists an oracle with $\rho = \text{Initiator}$ that accepts maliciously can be bounded by*

$$\text{Adv}_{\Pi}^{\text{Ent-Auth-I}}(\mathcal{A}) \leq n^2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{9.1i})$$

where all quantities are defined as stated in Theorem 9.

PROOF. Games 0 and 1 are exactly as in the proof of Lemma 10. The reduction is as in the proof of Lemma 10, except that now we are guessing which initiating party will be the first to accept maliciously (and its intended partner), and so the abort occurs if a responder oracle accepts maliciously. The loss of n^2 incurred by selection of parties is the same.

Again, the next step is a reduction that plays against SEUF-CMA- Q of the MAC scheme MAC. However, the reduction of course must behave slightly differently, since it must send to its own $\mathcal{O}_{\text{Vrfy}}$ oracle any message that was called as a Send query for the targeted initiator oracle, but which was not given as an output protocol message by a NewSessionR query (to responder oracle j^*). Further, we need to ensure that the forgery attempt is on an oracle that does not have a matching conversation with any others: in the proof of Lemma 10 this was straightforward since there was only one unique message in question, but here the transcripts that we are interested in consist of (up to) two flows between each oracle. This is just a matter of bookkeeping, and as before $\mathcal{B}_{9.1i}$ forwards all attempted forgeries to its own verification oracle, so any query that would have caused \mathcal{A} to win the entity authentication game (in the simulation that \mathcal{A} is experiencing) also corresponds to success in the game that $\mathcal{B}_{9.1i}$ is playing. \square

Lemma 12 (Key-Ind of LP2). *For any adversary \mathcal{A} and (any fixed) entity authentication adversary $\mathcal{B}_{9.1}$, the probability that \mathcal{A} answers the Test challenge correctly can be bounded by*

$$\text{Adv}_{\Pi}^{\text{Key-Ind}}(\mathcal{A}) \leq \text{Adv}_{\Pi}^{\text{Ent-Auth}}(\mathcal{B}_{9.1}) + n^2 q \cdot \text{Adv}_{\text{PRF}}^{\text{KEvol}}(\mathcal{B}_{9.2}).$$

where all quantities are defined as stated in Theorem 9.

PROOF. Let b' be the bit output by \mathcal{A} in each game, and b be the bit sampled as part of the Test query.

Game 0. This is the original Key-Ind game.

$$\Pr [G_{\Pi}^{\text{Key-Ind}}(\mathcal{A}) = 1] = \Pr [G_0^{\mathcal{A}} = 1]$$

Game 1. This game is the same as Game 0, except the challenger aborts and chooses $b' \xleftarrow{\$} \{0, 1\}$ if any oracle accepts maliciously.

$$\Pr [G_0^{\mathcal{A}} = 1] \leq \Pr [G_1^{\mathcal{A}} = 1] + \text{Adv}_{\Pi}^{\text{Ent-Auth}}(\mathcal{B}_{9.1})$$

At this stage, the oracle to which \mathcal{A} asks its Test query has a unique partner oracle with a matching conversation.

Game 2. This game is the same as Game 1, except the challenger guesses the two parties involved in the Test query via $(i^*, j^*) \xleftarrow{\$} [n] \times [n]$, and additionally guesses the counter value which identifies the key derivation key state of the session key in the Test query. Note that session keys are only derived for counters equal to odd integers in $\{1, 3, \dots, q\}$, so the challenger chooses CTR_{i^*, j^*}^* from this set. If \mathcal{A} issues a Test($\pi_{i^*}^s$) query with $i^* \neq i$ or $j^* \neq \pi_{i^*}^s.\text{pid}$ (for some s), or the session key computed via Test($\pi_{i^*}^s$), i.e. sk_{i^*, j^*} is not equal to $\text{KDF}(k_{i^*, j^*}^{\text{CTR}}, \text{"der"})$, then the challenger aborts. Note that the challenger does not guess which oracle of i^* the Test query will be made to, only the counter value linked to the session key in that query.

$$\Pr [G_1^{\mathcal{A}} = 1] \leq n^2 \cdot q \cdot \Pr [G_2^{\mathcal{A}} = 1]$$

At this stage, if the challenger has guessed correctly then the Test query will be asked to an oracle after the key derivation counter has been advanced a fixed number of times, and this oracle has unique partner with a matching conversation.

Game 3. This game is the same as Game 2, except that when the challenger runs KDF on the key derivation values used in the Test query, the challenger instead responds with a random key from the session key space. Noticing this change results in an adversary that is successful in the KEvol game for PRF, so we can write

$$\Pr [G_2^A = 1] = \Pr [G_3^A = 1] + \text{Adv}_{\text{PRF}}^{\text{KEvol}}(\mathcal{B}_{9.2}).$$

The reduction $\mathcal{B}_{9.2}$ is detailed in Fig. 9. $\mathcal{B}_{9.2}$ initially guesses the target parties in the Test session and the counter value associated with the Test session, as per the previous game hop. In this reduction, instances of **return Abort** indicate that the adversary has done something that the reduction cannot respond to (since it has been ruled out by earlier game hops) and thus the reduction must abort. Instances of **return \perp** indicate that the adversary has done something that it is not allowed to do, for example doing **RevealKey** on a non-existent session key or delivering an invalid message to a(n oracle of a) party. In this case the reduction stops the action of the query, and if instructed by the model in Section 3, sets the execution state α of the queried oracle to **reject**. Finally, **return x** indicates that $\mathcal{B}_{9.2}$ gives value x to \mathcal{A} as a result of one of \mathcal{A} 's queries.

In the event that $\mathcal{B}_{9.2}$ is in the ‘real’ version of its own game, where it receives a genuine evaluation of the function KDF, $\mathcal{B}_{9.2}$ perfectly simulates Game 2 for \mathcal{A} , and otherwise it perfectly simulates Game 3.

At this stage, the Test query is asked on a key that is randomly chosen, and thus independent of the protocol and the security game. Consequently,

$$\Pr [G_3^A = 1] = \frac{1}{2} \Rightarrow \text{Adv}_{\Pi}^3(\mathcal{A}) = 0.$$

□

4.3.2 wSR of LP2

We now argue that LP2 obtains weak synchronization robustness (wSR). A proof of wSR must argue that whatever the adversary does before the target protocol run occurs, during the target protocol run itself neither of the parties will abort and both will arrive at the same session key. Protocol LP2, like LP1 (Fig. 11), only has correctness when the initiator’s counter is at least the size of the responder’s counter, i.e. $\text{CTR}_{AB} \geq \text{CTR}_{BA}$ – this inequality is guaranteed in our protocol by MAC security. By inspection, if an adversary forges a MAC on $A \parallel \text{CTR}_{AB}$ for some CTR_{AB} larger than the current value of CTR_{BA} and delivers this MAC as part of a protocol message to an oracle of B , then any subsequent protocol run will cause B to **Abort** and thus will be a winning target protocol run for this adversary.

The formal proof is below, but here we outline the proof idea. We first define an event E_{13} that is triggered if the adversary in the wSR game forges a MAC, i.e. produces a message-tag pair that verifies correctly that it has not seen before, and the challenger aborts if this occurs: bounding this event is of course straightforward. Then, we must argue that if a MAC forgery has not occurred then there are in fact no viable routes to victory in the wSR game. To see this, note that for the (uninterrupted) target session, if $z_1 = \text{CTR}_{AB} - \text{CTR}_{BA} \geq 0$ then B will always catch up to the counter value of A (i.e. advance by z_1 steps) and both parties will compute a session key for counter value CTR_{AB} . Note also that in the target session, it is not possible for the **if** $z_2 \neq 0$ to be triggered after A receives the second protocol message since the counter value CTR_{BA} that B sends will always have caught up to CTR_{AB} in the processing of the first protocol message. Thus to conclude, we just need to show that, if a forgery has not occurred, it is not possible for the adversary to force $\text{CTR}_{AB} < \text{CTR}_{BA}$. Every time the adversary creates a new initiator session, the initiator’s counter is incremented by either 1 or 2 steps, whereas B can advance an arbitrary number of times to catch up to (what B thinks is) A 's current counter state. Since the MAC includes party

Reduction $\mathcal{B}_{9,2}$ playing $G_{\text{KDF}}^{\text{KEvol}}(\mathcal{B}_{9,2})$

```

1 :  $i^*, j^* \xleftarrow{\$} [n]$ ;  $\text{CTR}^* \xleftarrow{\$} \{1, 3, \dots, q\}$ 
2 : for  $i, j \in [n]$  do
3 :    $\text{CTR}_{ij} = \text{CTR}_{ji} \leftarrow 0$ 
4 :    $K_{ij}^{\text{MAC}} = K_{ji}^{\text{MAC}} \xleftarrow{\$} \mathcal{K}_{\text{MAC}}$ 
5 :   for  $[n] \times [n] \setminus (i^*, j^*)$  do
6 :      $k_{ij}^0 = k_{ji}^0 \xleftarrow{\$} \mathcal{K}_{\text{PRF}}$ 
7 :   output  $\text{CTR}^*$ 
8 :   receive( $sk^0, \dots, sk^{\text{CTR}^* - 1}, sk^*, k^{\text{CTR}^* + 1}$ )
9 :    $k_{i^* j^*}^{\text{CTR}^* + 1} \leftarrow k^{\text{CTR}^* + 1}$ 
10 :  $\mathcal{A}^{\text{oracles}}$ 
11 : When  $\mathcal{A}$  calls  $\text{Test}(\pi_i^s)$  do
12 :   if  $i \neq i^*$  or  $j^* \neq \pi_{i^*}^{s^*}.\text{pid}$ 
13 :     return Abort
14 :   return  $\pi_i^s.sk$ 
15 :    $(i^*, s^*, b') \xleftarrow{\$} \mathcal{A}(\pi_i^s.sk)$ 
16 :   return  $b'$ 

```

NewSessionI(π_i^s, pid)

```

17 :  $\pi_i^s.\rho \leftarrow \text{Initiator}$ 
18 :  $\pi_i^s.\alpha \leftarrow \text{negotiating}$ 
19 :  $\pi_i^s.\text{pid} \leftarrow \text{pid} \neq j$ 
20 :  $z_0 \leftarrow \text{NextOdd}(\text{CTR}_{ij}) - \text{CTR}_{ij}$ 
21 :  $k_{ij}^{\text{CTR}}, \text{CTR}_{ij} \leftarrow \text{Advnc}(k_{ij}^{\text{CTR}}, \text{CTR}_{ij}, z_0)$ 
22 :  $\sigma_1 \leftarrow \text{Mac}(K_{ij}^{\text{MAC}}, i \parallel \text{CTR}_{ij})$ 
23 :  $m' \leftarrow \text{CTR}_{ij}, \sigma_1$ 
24 : return  $m'$ 

```

NewSessionR(π_i^s, pid, m)

```

25 :  $\pi_i^s.\rho \leftarrow \text{Responder}$ 
26 :  $\pi_i^s.\alpha \leftarrow \text{negotiating}$ 
27 :  $\pi_i^s.\text{pid} \leftarrow \text{pid}$ 
28 : do  $\text{Send}(\pi_i^s, m)$ 

```

Corrupt(P_i, P_j)

```

29 : if  $(i, j) = (i^*, j^*) \vee (j^*, i^*)$ 
30 :   if  $\text{CTR}_{ij} \leq \text{CTR}^*$ 
31 :     return Abort
32 :   return  $k_{ij}$ 

```

RevealKey(π_i^s)

```

33 : if  $\pi_i^s.\alpha \neq \text{accept}$ 
34 :   return  $\perp$ 
35 : if  $(i, s) = (i^*, s^*)$ 
36 :   return Abort
37 :  $\pi_i^s.\kappa \leftarrow \text{exposed}$ 
38 : return  $\pi_i^s.sk$ 

```

Send(π_i^s, m) / $\text{pid} = j$

```

39 : Parse  $m$  as  $\text{CTR}_{ji}, \sigma$ 
40 : if  $\text{Vrfy}(K_{ij}^{\text{MAC}}, j \parallel \text{CTR}_{ji}, \sigma) = 0$ 
41 :   return  $\perp$ 
42 : if  $\pi_i^s.\rho = \text{Responder}$ 
43 :    $z_1 \leftarrow \text{CTR}_{ji} - \text{CTR}_{ij}$ 
44 :   if  $z_1 < 0$ 
45 :     return  $\perp$ 
46 :    $k_{ij}^{\text{CTR}}, \text{CTR}_{ij} \leftarrow \text{Advnc}(k_{ij}^{\text{CTR}}, \text{CTR}_{ij}, z_1)$ 
47 :    $\sigma_2 \leftarrow \text{Mac}(K_{ij}^{\text{MAC}}, i \parallel \text{CTR}_{ij})$ 
48 :    $m' \leftarrow \text{CTR}_{ij}, \sigma_2$ 
49 :   return  $m'$ 
50 : else
51 :    $z_2 \leftarrow \text{CTR}_{ij} - \text{CTR}_{ji}$ 
52 :   if  $z_2 \neq 0$ 
53 :     return  $\perp$ 
54 :   if  $(i, j) = (i^*, j^*) \vee (j^*, i^*)$ 
55 :     if  $\text{CTR}_{ij} < \text{CTR}^*$ 
56 :        $\pi_i^s.sk \leftarrow sk^{\text{CTR}_{ij}}$ 
57 :     if  $\text{CTR}_{ij} = \text{CTR}^*$ 
58 :        $\pi_i^s.sk \leftarrow sk^*$ 
59 :      $s^* \leftarrow s$ 
60 :      $\text{CTR}_{ij} \leftarrow \text{CTR}_{ij} + 1$ 
61 :   else
62 :      $\pi_i^s.sk \leftarrow \text{KDF}(k_{ij}^{\text{CTR}}, \text{"der"})$ 
63 :      $k_{ij}^{\text{CTR}}, \text{CTR}_{ij} \leftarrow \text{Advnc}(k_{ij}^{\text{CTR}}, \text{CTR}_{ij}, 1)$ 
64 :    $\pi_i^s.\alpha \leftarrow \text{accept}$ 

```

Figure 9: Reduction $\mathcal{B}_{9,2}$ for the proof of Lemma 12. If at any time \mathcal{A} causes an oracle to accept maliciously, then $\mathcal{B}_{9,2}$ simply does Abort. $\mathcal{B}_{9,2}$ provides \mathcal{A} with access to $\text{oracles} = \text{NewSessionI}(\cdot, \cdot, \cdot), \text{NewSessionR}(\cdot, \cdot), \text{Send}(\cdot, \cdot, \cdot), \text{RevealKey}(\cdot), \text{Corrupt}(\cdot, \cdot)$.

identification information and the initiator's counter value, in the absence of MAC forgeries the adversary cannot produce a valid protocol message with verifying MAC for any counter larger than the ones that it has seen as a result of genuine invocations of new protocol sessions.

Theorem 13 (wSR of LP2). *Let Π be the two-message protocol in Fig. 8, built using $\text{MAC} = \{\text{KGen}, \text{Mac}, \text{Vrfy}\}$ and PRF, with n parties. Then for any adversary \mathcal{A} against the wSR security of Π that makes a maximum of q queries that initiate new sessions for each party (with $q < \frac{|\text{CTR}|}{2}$), there exists an adversary \mathcal{B}_{13} against the SEUF-CMA- Q of MAC such that*

$$\text{Adv}_{\Pi}^{\text{wSR}}(\mathcal{A}) \leq n^2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-}Q}(\mathcal{B}_{13}).$$

PROOF. We proceed using a sequence of games.

Game 0. This is the original wSR game.

$$\Pr [G_{\Pi}^{\text{wSR}}(\mathcal{A}) = 1] = \Pr [G_0^{\mathcal{A}} = 1]$$

Game 1. This game is the same as Game 0 except that we define an event E_{13} , that is said to occur if the adversary successfully forges a MAC while it is running, and the challenger aborts if E_{13} occurs.

$$\Pr [G_0^{\mathcal{A}} = 1] = \Pr [G_1^{\mathcal{A}} = 1] + \Pr [E_{13}]$$

We now bound the probability that E_{13} occurs. We construct a reduction \mathcal{B}_{13} to the SEUF-CMA- Q security of MAC with a verification oracle. First, the reduction guesses which parties will be involved in the forgery that triggers E_{13} , and then simulates the environment of Game 0. To do this, \mathcal{B}_{13} must select (initial) key derivation keys from the appropriate keyspace and randomly pick MAC keys for all pairs of parties except for the guessed parties i^* and j^* . It is then simple to simulate all queries to oracles of parties except communication between the guessed pair. Note that this choice of parties involved in the forgery is independent of the parties that the underlying adversary \mathcal{A} will eventually output for the target protocol run.

For any query to an oracle of party i^* with $\text{pid} = j^*$ or an oracle of j^* with $\text{pid} = i^*$, the reduction needs to forward queries to its own MAC and verification oracles. However note that \mathcal{B}_{13} knows the key derivation keys even for these parties, so responding to queries is straightforward except for its calls to \mathcal{O}_{Mac} and $\mathcal{O}_{\text{Vrfy}}$. The full reduction is detailed in Fig. 10.

$$\Pr [E_{13}] \leq n^2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-}Q}(\mathcal{B}_{13})$$

At this point, the adversary cannot win via a MAC forgery.

A win can be obtained if either party in the target session aborts, or if the parties compute different keys. If the target session begins with $\text{CTR}_{AB} \geq \text{CTR}_{BA}$, then the parties will always compute the same key via $sk_{AB} \leftarrow \text{KDF}(k_{AB}^{\text{CTR}}, \text{"der"})$ so we only need to argue that the adversary cannot force the state $\text{CTR}_{BA} > \text{CTR}_{AB}$ without forging the MAC. Note that every initiator session advances its counter (and thus key state) once either once or twice, while responder sessions (oracles) can advance an arbitrary number of times. However, the responder only advances if it has received and accepted a protocol message. For the responder to advance without the initiator having done so, the adversary must deliver a valid message to the responder without having called `NewSessionI`. This message must also have a counter value CTR'_{AB} that is greater than CTR_{AB} (the actual counter value of A with respect to B), and a valid MAC. Producing such a message that will be processed by B requires it to have a valid MAC on $A \parallel \text{CTR}'_{AB}$. Since we assume that the adversary does not produce a MAC forgery it must have seen this message before, which means it must have been output by a `NewSessionI` query, and we have a contradiction. This concludes the proof. \square

Reduction \mathcal{B}_{13} playing $G_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{13})$	Send(π_i^s, m) / pid = j
<pre> 1 : $i^*, j^* \xleftarrow{\\$} [n]$ 2 : for $i, j \in [n]$ do 3 : $k_{ij}^{\text{CTR}} = k_{ji}^{\text{CTR}} \xleftarrow{\\$} \mathcal{K}_{\text{PRF}}$ 4 : for $[n] \times [n] \setminus (i^*, j^*)$ do 5 : $K_{ij}^{\text{MAC}} = k_{\text{MAC}}^{j,i} \xleftarrow{\\$} \mathcal{K}_{\text{MAC}}$ 6 : $\text{CTR}_{ij} = \text{CTR}_{ji} \leftarrow 0$ 7 : $\mathcal{A}^{\text{oracles}}$ </pre>	<pre> 23 : Parse m as CTR_{ji}, σ 24 : if $(i, j) = (i^*, j^*) \vee (j^*, i^*)$ 25 : $b \leftarrow \text{call } \mathcal{O}_{\text{Vrfy}}(j \parallel \text{CTR}_{ji}), \sigma$ 26 : if $b = 0$ 27 : return \perp 28 : else 29 : if $\text{Vrfy}(K^{\text{MAC}}, j \parallel \text{CTR}_{ij}, \sigma) = 0$ 30 : return \perp 31 : if $\pi_i^s.\rho = \text{Responder}$ 32 : $z_1 \leftarrow \text{CTR}_{ji} - \text{CTR}_{ij}$ 33 : if $z_1 < 0$ 34 : return \perp 35 : $k_{ij}^{\text{CTR}}, \text{CTR}_{ij} \leftarrow \text{Advnc}(k_{ij}^{\text{CTR}}, \text{CTR}_{ij}, z_1)$ 36 : if $(i, j) = (i^*, j^*) \vee (j^*, i^*)$ 37 : $\sigma_2 \leftarrow \text{call } \mathcal{O}_{\text{Mac}}(i \parallel \text{CTR}_{ij})$ 38 : else 39 : $\sigma_2 \leftarrow \text{Mac}(K_{ij}^{\text{MAC}}, i \parallel \text{CTR}_{ij})$ 40 : return $m' \leftarrow \text{CTR}_{ij}, \sigma_2$ 41 : else 42 : $z_2 \leftarrow \text{CTR}_{ij} - \text{CTR}_{ji}$ 43 : if $z_2 \neq 0$ 44 : return \perp 45 : $\pi_i^s.sk \leftarrow \text{KDF}(k_{ij}^{\text{CTR}}, \text{"der"})$ 46 : $\pi_i^s.\alpha \leftarrow \text{accept}$ 47 : $k_{ij}^{\text{CTR}}, \text{CTR}_{ij} \leftarrow \text{Advnc}(k_{ij}^{\text{CTR}}, \text{CTR}_{ij}, 1)$ </pre>
NewSessionI(π_i^s, pid)	
<pre> 8 : $\pi_i^s.\rho \leftarrow \text{Initiator}$ 9 : $\pi_i^s.\alpha \leftarrow \text{negotiating}$ 10 : $\pi_i^s.\text{pid} \leftarrow \text{pid} / = j$ 11 : $z_0 \leftarrow \text{NextOdd}(\text{CTR}_{ij}) - \text{CTR}_{ij}$ 12 : $k_{ij}^{\text{CTR}}, \text{CTR}_{ij} \leftarrow \text{Advnc}(k_{ij}^{\text{CTR}}, \text{CTR}_{ij}, z_0)$ 13 : if $(i, j) = (i^*, j^*) \vee (j^*, i^*)$ 14 : $\sigma_1 \leftarrow \text{call } \mathcal{O}_{\text{Mac}}(i \parallel \text{CTR}_{ij})$ 15 : else 16 : $\sigma_1 \leftarrow \text{Mac}(K_{ij}^{\text{MAC}}, i \parallel \text{CTR}_{ij})$ 17 : $m' \leftarrow \text{CTR}_{ij}, \sigma_1$ 18 : return m' </pre>	
NewSessionR(π_i^s, pid, m)	
<pre> 19 : $\pi_i^s.\rho \leftarrow \text{Responder}$ 20 : $\pi_i^s.\alpha \leftarrow \text{negotiating}$ 21 : $\pi_i^s.\text{pid} \leftarrow \text{pid}$ 22 : do Send(π_i^s, m) </pre>	

Figure 10: Reduction \mathcal{B}_{13} for the proof of Theorem 13. \mathcal{B}_{13} provides \mathcal{A} with access to `oracles = NewSessionI(\cdot, \cdot, \cdot), NewSessionR(\cdot, \cdot), Send(\cdot, \cdot, \cdot)`.

4.4 LP1: A One-Message Protocol with Fixed Roles

In Fig. 11 we present a one-message protocol, LP1, with linear key evolution. Like in LP2, the roles of initiator and responder are fixed, so the same party initiates every session: i.e. $\text{CTR}_{AB} \geq \text{CTR}_{BA}$ (for A initiating). Ensuring that the counter states of the communicating parties is slightly simpler in LP1 than LP2, since we do not have to worry about the responder advancing before the initiator. The initiator simply advances once every time it participates in a session, and both parties advance exactly once after computing a session key. If protocol messages are dropped then it may be necessary for the responder to advance before it can compute the session key.

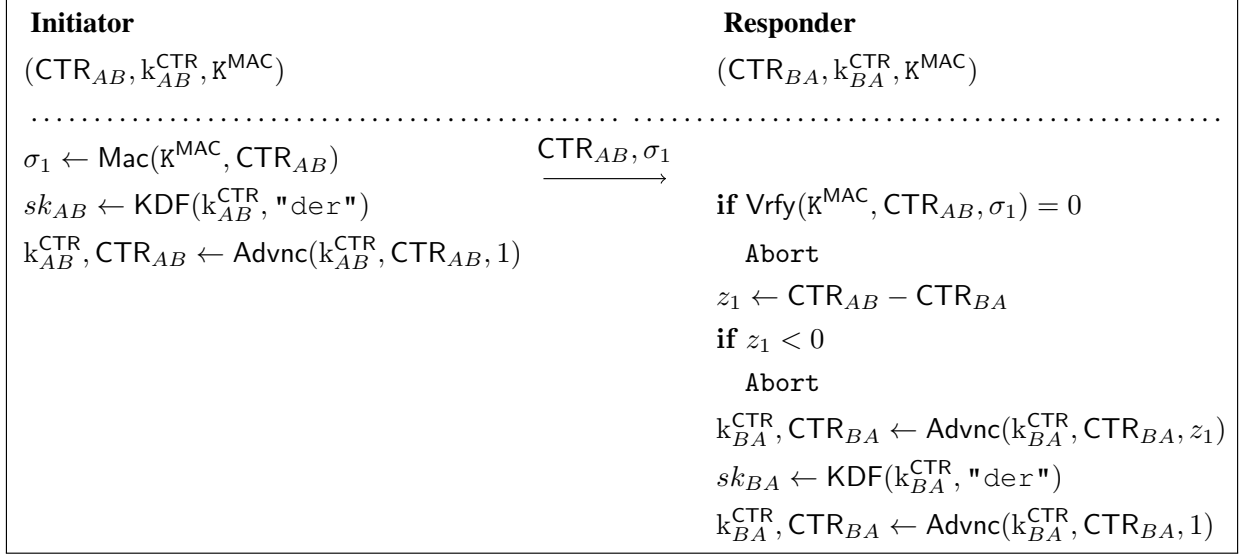


Figure 11: LP1, a one-message protocol with fixed roles.

In Theorem 14 we show that LP1 achieves one-sided authentication (responder authenticates initiator). Achieving weak synchronization robustness (wSR, Theorem 17) is similar in LP1 and LP2, and is guaranteed by MAC security. Like with LP2, if both parties need to be able to initiate then LP1 can be run in *duplex mode*.

4.4.1 AKE-R of LP1

Theorem 14 (AKE-R of LP1). *Let Π be the one-message protocol in Fig. 11, built using $\text{MAC} = \{\text{KGen}, \text{Mac}, \text{Vrfy}\}$ and PRF, with n parties. Then for any adversary \mathcal{A} against the AKE-R security of Π that makes a maximum of q queries that initiate new sessions for each party (with $q < |\text{CTR}|$), there exists an adversary $\mathcal{B}_{14.1}$ against the SEUF-CMA- Q security of MAC and an adversary $\mathcal{B}_{14.2}$ against the KEvol security of PRF such that*

$$\text{Adv}_{\Pi}^{\text{AKE-R}}(\mathcal{A}) \leq n^2 \cdot \left(2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-}Q}(\mathcal{B}_{14.1}) + q \cdot \text{Adv}_{\text{PRF}}^{\text{KEvol}}(\mathcal{B}_{14.2}) \right).$$

We form a bound for each of the two ways in which an adversary can break AKE-R security, namely Ent-Auth-R and Key-Ind, and then sum these bounds.

We now give intuition regarding the Ent-Auth-R proof. Note that since $q < \frac{|\text{CTR}|}{2}$, all first protocol messages are unique, and thus the first message in every transcript is unique. This rules out the possibility of *multiple* oracles with non-unique matching conversations accepting. To conclude our proofs, we need to show that the only way an adversary can force an oracle to accept and for there not to exist any other oracle with a matching transcript, the adversary must forge a MAC message.

A responder oracle accepts when it receives a first protocol message, and so to win there must not exist any oracle with the same partial transcript as this accepting oracle. Since the MAC is calculated on the counter value and the communicating parties' identities, the input message to the accepting oracle must have been generated by an oracle of the communication partner (in which case there is a matching conversation and the adversary has not won) or the adversary has produced a MAC forgery (in the SEUF-CMA- Q sense).

Lemma 15 (Ent-Auth-R of LP1). *For any adversary \mathcal{A} , the probability that there exists an oracle with $\rho = \text{Responder}$ that accepts maliciously can be bounded by*

$$\text{Adv}_{\Pi}^{\text{Ent-Auth-R}}(\mathcal{A}) \leq n^2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{14.1}).$$

PROOF. We proceed using a sequence of games.

Game 0. This is the original Ent-Auth game.

$$\Pr \left[G_{\Pi}^{\text{Ent-Auth}}(\mathcal{A}) = 1 \right] = \Pr \left[G_0^{\mathcal{A}} = 1 \right]$$

Game 1. In this game we guess which responder will be the first to accept maliciously and its partner identity, and abort if this guess is wrong. The game is the same as Game 0 except that the challenger guesses $(i^*, j^*) \xleftarrow{\$} [n] \times [n]$, and if an oracle π_i^s (for some s) accepts maliciously with $\pi_i^s.\rho \neq \text{Responder}$ or $i^* \neq i$ or $j^* \neq \pi_i^s.\text{pid}$, then the challenger aborts.

$$\Pr \left[G_0^{\mathcal{A}} = 1 \right] = n^2 \cdot \Pr \left[G_1^{\mathcal{A}} = 1 \right]$$

We construct a reduction $\mathcal{B}_{14.1}$ that is playing against the SEUF-CMA- Q security of MAC that simulates the environment for an underlying adversary \mathcal{A} that attempts to win in game Game 1. The reduction generates (initial) key derivation keys for all pairs of parties, and authentication keys for all pairs of parties except i^* and j^* . When responding to queries by \mathcal{A} regarding all other parties, the reduction will honestly provide messages as specified in the protocol specification and the Ent-Auth game. For any query made between oracles of parties i^* and j^* , the reduction will use its \mathcal{O}_{Mac} oracle and provide the received value in its simulation for \mathcal{A} . For example, to initialize initiator oracles π_{j^*} , the reduction checks the current state counter $\text{CTR}_{j^*i^*}$ and calls $\mathcal{O}_{\text{Mac}}(\text{CTR}_{j^*i^*})$. If \mathcal{A} provides any value that it has not been given as an initialization query as input to a `NewSessionR` query from i^* to j^* , then the reduction sends this to its $\mathcal{O}_{\text{Vrfy}}$ oracle in the SEUF-CMA- Q game. The simulation of Game 1 is perfect and any win for \mathcal{A} directly corresponds to a valid signature forgery, so we can write

$$\Pr \left[G_1^{\mathcal{A}} = 1 \right] \leq \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{14.1}).$$

Summing these terms gives the bound in the statement of Lemma 15. □

The proof of key indistinguishability is very similar to that of Lemma 12 and the game hops proceed following the same strategy. Again we use a reduction to the KEvol security of PRF.

Lemma 16 (Key-Ind of LP1). *For any adversary \mathcal{A} and (any fixed) entity authentication adversary $\mathcal{B}_{14.1}$, the probability that \mathcal{A} answers the Test challenge correctly can be bounded by*

$$\text{Adv}_{\Pi}^{\text{Key-Ind}}(\mathcal{A}) \leq \text{Adv}_{\Pi}^{\text{Ent-Auth}}(\mathcal{B}_{14.1}) + n^2 \cdot q \cdot \text{Adv}_{\text{PRF}}^{\text{KEvol}}(\mathcal{B}_{14.2})$$

where all quantities are defined as stated in Theorem 14.

PROOF. Let b' be the bit output by \mathcal{A} in each game, and b be the bit sampled as part of the Test query.

Game 0. This is the original Key-Ind game.

$$\Pr [G_{\Pi}^{\text{Key-Ind}}(\mathcal{A}) = 1] = \Pr [G_0^{\mathcal{A}} = 1]$$

Game 1. This game is the same as Game 0, except the challenger aborts and chooses $b' \xleftarrow{\$} \{0, 1\}$ if any oracle accepts maliciously.

$$\Pr [G_0^{\mathcal{A}} = 1] \leq \Pr [G_1^{\mathcal{A}} = 1] + \text{Adv}_{\Pi}^{\text{Ent-Auth-R}}(\mathcal{B}_{14.1})$$

At this stage, the oracle to which \mathcal{A} asks its Test query has a unique partner oracle with a matching conversation.

Game 2. This game is the same as Game 1, except the challenger guesses the two parties involved in the Test query via $(i^*, j^*) \xleftarrow{\$} [n] \times [n]$, and additionally guesses the counter value which identifies the key derivation key state of the session key in the Test query. If \mathcal{A} issues a Test($\pi_{i^*}^s$) query with $i^* \neq i$ or $j^* \neq \pi_{i^*}^s.\text{pid}$ (for some s), or the session key computed via Test($\pi_{i^*}^s$), i.e. $sk_{i^*j^*}$ is not equal to $\text{KDF}(k_{i^*j^*}^{\text{CTR}}, \text{"der"})$, then the challenger aborts. Note that the challenger does not guess which oracle of i^* the Test query will be made to, only the counter value linked to the session key in that query.

$$\Pr [G_1^{\mathcal{A}} = 1] \leq n^2 \cdot q \cdot \Pr [G_2^{\mathcal{A}} = 1]$$

At this stage, if the challenger has guessed correctly then the Test query will be asked to an oracle after the key derivation counter has been advanced a fixed number of times, and this oracle has a unique partner with a matching conversation.

Game 3. This game is the same as Game 2, except that when the challenger runs KDF on the key derivation values used in the Test query, the challenger instead responds with a random key from the session key space. Noticing this change results in an adversary that is successful in the KEvol game for PRF, so we can write

$$\Pr [G_2^{\mathcal{A}} = 1] = \Pr [G_3^{\mathcal{A}} = 1] + \text{Adv}_{\text{PRF}}^{\text{KEvol}}(\mathcal{B}_{9.2}).$$

The reduction $\mathcal{B}_{14.2}$ is detailed in Fig. 12. $\mathcal{B}_{14.2}$ initially guesses the target parties in the Test session and the counter value associated with the Test session, as per the previous game hop.

In the event that $\mathcal{B}_{14.2}$ is in the ‘real’ version of its own game, where it receives a genuine evaluation of the function KDF, $\mathcal{B}_{14.2}$ perfectly simulates Game 2 for \mathcal{A} , and otherwise it perfectly simulates Game 3.

At this stage, the Test query is asked on a key that is randomly chosen, and thus independent of the protocol and the security game. Consequently,

$$\Pr [G_3^{\mathcal{A}} = 1] = \frac{1}{2} \Rightarrow \text{Adv}_{\Pi}^3(\mathcal{A}) = 0.$$

□

4.4.2 wSR of LP1

The wSR security of LP1 is achieved in a similar manner to LP2 and we proceed to prove the theorem.

<p>Reduction $\mathcal{B}_{14.2}$ playing $G_{\text{KDF}}^{\text{KEvol}}(\mathcal{B}_{14.2})$</p> <pre> 1 : $i^*, j^* \xleftarrow{\\$} [n]$; $\text{CTR}^* \xleftarrow{\\$} [q]$ 2 : for $i, j \in [n]$ do 3 : $\text{CTR}_{ij} = \text{CTR}_{ji} \leftarrow 0$ 4 : $K_{ij}^{\text{MAC}} = K_{ji}^{\text{MAC}} \xleftarrow{\\$} \mathcal{K}_{\text{MAC}}$ 5 : for $[n] \times [n] \setminus (i^*, j^*)$ do 6 : $k_{ij}^0 = k_{ji}^0 \xleftarrow{\\$} \mathcal{K}_{\text{PRF}}$ 7 : output CTR^* 8 : receive $(sk^0, \dots, sk^{\text{CTR}^*-1}, sk^*, k^{\text{CTR}^*+1})$ 9 : $k_{i^*j^*}^{\text{CTR}^*+1} \leftarrow k^{\text{CTR}^*+1}$ 10 : $\mathcal{A}^{\text{oracles}}$ 11 : When \mathcal{A} calls $\text{Test}(\pi_i^s)$ do 12 : if $i \neq i^*$ or $j \neq j^*$ do 13 : return Abort 14 : return $\pi_i^s.sk$ 15 : $(i^*, s^*, b') \xleftarrow{\\$} \mathcal{A}(\pi_i^s.sk)$ 16 : return b' </pre> <p>NewSessionI(π_i^s, pid)</p> <pre> 17 : $\pi_i^s.\rho \leftarrow \text{Initiator}$ 18 : $\pi_i^s.\alpha \leftarrow \text{negotiating}$ 19 : $\pi_i^s.\text{pid} \leftarrow \text{pid} \neq j$ 20 : $\sigma_1 \leftarrow \text{Mac}(K^{\text{MAC}}, \text{CTR}_{ij})$ 21 : $m' \leftarrow \text{CTR}_{ij}, \sigma_1$ 22 : $k_{ij}^{\text{CTR}}, \text{CTR}_{ij} \leftarrow \text{Advnc}(k_{ij}^{\text{CTR}}, \text{CTR}_{ij}, 1)$ 23 : return m' </pre> <p>NewSessionR(π_i^s, pid, m)</p> <pre> 24 : $\pi_i^s.\rho \leftarrow \text{Responder}$ 25 : $\pi_i^s.\alpha \leftarrow \text{negotiating}$ 26 : $\pi_i^s.\text{pid} \leftarrow \text{pid}$ 27 : do $\text{Send}(\pi_i^s, m)$ </pre>	<p>Corrupt(P_i, P_j)</p> <pre> 28 : if $(i, j) = (i^*, j^*) \vee (j^*, i^*)$ 29 : if $\text{CTR}_{ij} \leq \text{CTR}^*$ 30 : return Abort 31 : return k_{ij} </pre> <p>RevealKey(π_i^s)</p> <pre> 32 : if $\pi_i^s.\alpha \neq \text{accept}$ 33 : return \perp 34 : if $(i, s) = (i^*, s^*)$ 35 : return Abort 36 : $\pi_i^s.\kappa \leftarrow \text{exposed}$ 37 : return $\pi_i^s.sk$ </pre> <p>Send(π_i^s, m) / $\text{pid} = j$</p> <pre> 38 : Parse m as $\text{CTR}_{ji}, \sigma_1$ 39 : if $\text{Vrfy}(K_{ij}^{\text{MAC}} \parallel \text{CTR}_{ji}, \sigma_1) = 0$ 40 : return \perp 41 : $z_1 \leftarrow \text{CTR}_{ji} - \text{CTR}_{ij}$ 42 : if $z_1 < 0$ 43 : return \perp 44 : $k_{ij}^{\text{CTR}}, \text{CTR}_{ij} \leftarrow \text{Advnc}(k_{ij}^{\text{CTR}}, \text{CTR}_{ij}, z_1)$ 45 : if $(i, j) = (i^*, j^*) \vee (j^*, i^*)$ 46 : if $\text{CTR}_{ij} < \text{CTR}^*$ 47 : $\pi_i^s.sk \leftarrow sk^{\text{CTR}_{ij}}$ 48 : if $\text{CTR}_{ij} = \text{CTR}^*$ 49 : $\pi_i^s.sk \leftarrow sk^*$ 50 : $s^* \leftarrow s$ 51 : $\text{CTR}_{ij} \leftarrow \text{CTR}_{ij} + 1$ 52 : else 53 : $\pi_i^s.sk \leftarrow \text{KDF}(k_{ij}^{\text{CTR}}, \text{"der"})$ 54 : $k_{ij}^{\text{CTR}}, \text{CTR}_{ij} \leftarrow \text{Advnc}(k_{ij}^{\text{CTR}}, \text{CTR}_{ij}, 1)$ 55 : $\pi_i^s.\alpha \leftarrow \text{accept}$ </pre>
--	---

Figure 12: Reduction $\mathcal{B}_{14.2}$ for the proof of Lemma 16. If at any time \mathcal{A} causes an oracle to accept maliciously, then $\mathcal{B}_{14.2}$ simply does Abort. $\mathcal{B}_{14.2}$ provides \mathcal{A} with access to $\text{oracles} = \text{NewSessionI}(\cdot, \cdot, \cdot), \text{NewSessionR}(\cdot, \cdot), \text{Send}(\cdot, \cdot, \cdot), \text{RevealKey}(\cdot), \text{Corrupt}(\cdot, \cdot)$.

Theorem 17 (wSR of LP1). *Let Π be the one-message protocol in Fig. 11, built using $\text{MAC} = \{\text{KGen}, \text{Mac}, \text{Vrfy}\}$ and PRF, with n parties. Then for any adversary \mathcal{A} against the wSR security of Π that makes a maximum of q queries that initiate new sessions for each party (with $q < |\text{CTR}|$), there exists an adversary \mathcal{B}_{17} against the SEUF-CMA- Q security of MAC such that*

$$\text{Adv}_{\Pi}^{\text{wSR}}(\mathcal{A}) \leq n^2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-}Q}(\mathcal{B}_{17}).$$

We now argue that LP1 obtains weak synchronization robustness (wSR). A proof of wSR must argue that whatever the adversary does before the target protocol run occurs, neither of the parties will abort during the target protocol run itself and both will arrive at the same session key. Protocol LP1, like LP2 (Fig. 8), only has correctness when the initiator’s counter is at least the size of the responder’s counter, i.e. $\text{CTR}_{AB} \geq \text{CTR}_{BA}$ — this inequality is guaranteed in our protocol by MAC security. By inspection, if an adversary forges a MAC on CTR_{AB} for some CTR_{AB} larger than the current value of CTR_{BA} and delivers this MAC as part of a protocol message to an oracle of B , then any subsequent protocol run will cause B to Abort and thus will be a winning target protocol run for this adversary.

The formal proof is below, but here we outline the proof idea. We first define an event E_{17} that is triggered if the adversary in the wSR game forges a MAC, i.e. produces a message-tag pair that verifies correctly and that it has not seen before, and the challenger aborts if this occurs: bounding this event is of course straightforward. Then, we must argue that if a MAC forgery has not occurred then there are in fact no viable routes to victory in the wSR game. To see this, note that for the (uninterrupted) target session, if $z_1 = \text{CTR}_{AB} - \text{CTR}_{BA} \geq 0$ then B will always catch up to the counter value of A (i.e. advance by z_1 steps) and both parties will compute a session key for counter value CTR_{AB} . Thus to conclude, we just need to show that, if a forgery has not occurred, it is not possible for the adversary to force $\text{CTR}_{AB} < \text{CTR}_{BA}$. Every time the adversary creates a new initiator session, the initiator’s counter is incremented by 1, whereas B can advance an arbitrary number of times to catch up to (what B thinks is) A ’s current counter state. Since the MAC includes party identification information and the initiator’s counter value, in the absence of MAC forgeries the adversary cannot produce a valid protocol message with verifying MAC for any counter larger than the ones that it has seen as a result of genuine invocations of new protocol sessions.

PROOF. We proceed using a sequence of games.

Game 0. This is the original wSR game.

$$\Pr [G_{\Pi}^{\text{wSR}}(\mathcal{A}) = 1] = \Pr [G_0^A = 1]$$

Game 1. This game is the same as Game 0 except that we define an event E_{17} , that is said to occur if the adversary successfully forges a MAC while it is running, and the challenger aborts if E_{17} occurs.

$$\Pr [G_0^A = 1] = \Pr [G_1^A = 1] + \Pr [E_{17}]$$

We now bound the probability that E_{17} occurs. We construct a reduction \mathcal{B}_{17} to the SEUF-CMA- Q security of MAC with a verification oracle. First, the reduction guesses which parties will be involved in the forgery that triggers E_{17} , and then simulates the environment of Game 0. To do this, \mathcal{B}_{17} must select (initial) key derivation keys from the appropriate keyspace and randomly pick MAC keys for all pairs of parties except for the guessed parties i^* and j^* . It is then easy to simulate all queries to the parties’ oracles, except communication between the guessed pair. Note that this choice of parties involved in the forgery is independent of the parties that the underlying adversary \mathcal{A} will eventually output for the target protocol run.

For any query to an oracle of party i^* with $\text{pid} = j^*$ or an oracle of j^* with $\text{pid} = i^*$, the reduction needs to forward queries to its own MAC and verification oracles. However note that \mathcal{B}_{17} knows the key derivation keys even for these parties, so responding to queries is straightforward except for its calls to \mathcal{O}_{Mac} and $\mathcal{O}_{\text{Vrfy}}$. The full reduction is detailed in Fig. 13.

$$\Pr [E_{17}] \leq n^2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{17})$$

Reduction \mathcal{B}_{17} playing $G_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{17})$	NewSessionR(π_i^s, pid, m)
1 : $i^*, j^* \xleftarrow{s} [n]$ 2 : for $i, j \in [n]$ do 3 : $k_{ij}^{\text{CTR}} = k_{ji}^{\text{CTR}} \xleftarrow{s} \mathcal{K}_{\text{PRF}}$ 4 : for $[n] \times [n] \setminus (i^*, j^*)$ do 5 : $K_{ij}^{\text{MAC}} = k_{\text{MAC}}^{j,i} \xleftarrow{s} \mathcal{K}_{\text{MAC}}$ 6 : $\text{CTR}_{ij} = \text{CTR}_{ji} \leftarrow 0$ 7 : $\mathcal{A}^{\text{oracles}}$	20 : $\pi_i^s.\rho \leftarrow \text{Responder}$ 21 : $\pi_i^s.\alpha \leftarrow \text{negotiating}$ 22 : $\pi_i^s.\text{pid} \leftarrow \text{pid}$ 23 : do Send(π_i^s, m)
NewSessionI(π_i^s, pid)	Send(π_i^s, m) / $\text{pid} = j$
8 : $\pi_i^s.\rho \leftarrow \text{Initiator}$ 9 : $\pi_i^s.\alpha \leftarrow \text{negotiating}$ 10 : $\pi_i^s.\text{pid} \leftarrow \text{pid} / = j$ 11 : if $(i, j) = (i^*, j^*) \vee (j^*, i^*)$ 12 : $\sigma_1 \leftarrow \text{call } \mathcal{O}_{\text{Mac}}(\text{CTR}_{ij})$ 13 : else 14 : $\sigma_1 \leftarrow \text{Mac}(K_{ij}^{\text{MAC}}, \text{CTR}_{ij})$ 15 : $m' \leftarrow \text{CTR}_{ij}, \sigma_1$ 16 : return m' 17 : $\pi_i^s.sk \leftarrow \text{KDF}(k_{ij}^{\text{CTR}}, \text{"der"})$ 18 : $\pi_i^s.\alpha \leftarrow \text{accept}$ 19 : $k_{ij}^{\text{CTR}}, \text{CTR}_{ij} \leftarrow \text{Advnc}(k_{ij}^{\text{CTR}}, \text{CTR}_{ij}, 1)$	24 : Parse m as $\text{CTR}_{ji}, \sigma_1$ 25 : if $(i, j) = (i^*, j^*) \vee (j^*, i^*)$ 26 : $b \leftarrow \text{call } \mathcal{O}_{\text{Vrfy}}(\text{CTR}_{ji}, \sigma_1)$ 27 : if $b = 0$ 28 : return \perp 29 : else 30 : if $\text{Vrfy}(K_{ij}^{\text{MAC}}, \text{CTR}_{ij}, \sigma_1) = 0$ 31 : return \perp 32 : if $\pi_i^s.\rho = \text{Initiator}$ 33 : return \perp 34 : $z_1 \leftarrow \text{CTR}_{ji} - \text{CTR}_{ij}$ 35 : if $z_1 < 0$ 36 : return \perp 37 : $k_{ij}^{\text{CTR}}, \text{CTR}_{ij} \leftarrow \text{Advnc}(k_{ij}^{\text{CTR}}, \text{CTR}_{ij}, z_1)$ 38 : $\pi_i^s.sk \leftarrow \text{KDF}(k_{ij}^{\text{CTR}}, \text{"der"})$ 39 : $\pi_i^s.\alpha \leftarrow \text{accept}$ 40 : $k_{ij}^{\text{CTR}}, \text{CTR}_{ij} \leftarrow \text{Advnc}(k_{ij}^{\text{CTR}}, \text{CTR}_{ij}, 1)$

Figure 13: Reduction \mathcal{B}_{17} for the proof of Theorem 17. \mathcal{B}_{17} provides \mathcal{A} with access to oracles = NewSessionI(\cdot, \cdot, \cdot), NewSessionR(\cdot, \cdot), Send(\cdot, \cdot, \cdot).

At this point, the adversary cannot win via MAC forgery.

A win can be obtained if either party in the target session aborts, or if the parties compute different keys. If the target session begins with $\text{CTR}_{AB} \geq \text{CTR}_{BA}$, then the parties will always compute the same key via $sk_{AB} \leftarrow \text{KDF}(k_{AB}^{\text{CTR}}, \text{"der"})$, so we only need to argue that the adversary cannot force the state $\text{CTR}_{BA} > \text{CTR}_{AB}$ without forging the MAC. Note that every initiator session advances its counter (and thus key state) exactly once, while responder sessions (oracles) can advance an arbitrary number of times.

However, the responder only advances if it has received and accepted a protocol message. For the responder to advance without the initiator having done so, the adversary must deliver a valid message to the responder without having called `NewSessionI`. This message must also have a counter value CTR'_{AB} that is greater than CTR_{AB} (the actual counter value of A with respect to B , and a valid MAC. Producing such a message that will be processed by B requires it to have a valid MAC on CTR'_{AB} . Since we assume that the adversary does not produce a MAC forgery it must have seen this message before, which means it must have been output by a `NewSessionI` query, and we have a contradiction. This concludes the proof. \square

5 Non-Linear Key Evolution

In the previous section, we have considered protocols that deploy a linear key evolving mechanism. We have seen that the linearity of these mechanisms has significant downsides when the protocol runs multiple times in parallel between the same two parties. Especially interleaving of messages might cause all but one protocol execution to abort, which is an undesirable behavior.

In this section, we present a protocol that uses puncturable pseudorandom functions (PPRFs) as a “non-linear” key evolution mechanism. We show that this protocol can establish many parallel sessions between two parties, while only requiring some additional storage (logarithmic in the supported maximum number sessions) and computations (in practice hash function evaluations logarithmic in the supported maximum number of sessions).

5.1 Puncturable Pseudorandom Functions

We briefly recall the basic definition of puncturable pseudorandom functions (PPRF). A PPRF is a special case of a pseudorandom function, where it is possible to compute punctured keys, which do not allow evaluation on inputs that have been punctured. We recall the definition of a PPRF and its security [SW14].

Definition 12 (PPRF). A *puncturable pseudorandom function* with key space $\mathcal{K}_{\text{PPRF}}$, domain $\mathcal{D}_{\text{PPRF}}$, and range $\mathcal{R}_{\text{PPRF}}$ consists of three probabilistic polynomial-time algorithms $\text{PPRF} = (\text{Setup}, \text{Eval}, \text{Punct})$, which are described as follows:

- $\text{Setup}(1^\lambda)$: This algorithm takes as input the security parameter λ and outputs a description of a key $k \in \mathcal{K}_{\text{PPRF}}$.
- $\text{Eval}(k, x)$: This algorithm takes as input a key $k \in \mathcal{K}_{\text{PPRF}}$ and a value $x \in \mathcal{D}_{\text{PPRF}}$, and outputs a value $y \in \mathcal{R}_{\text{PPRF}}$, or a failure symbol \perp .
- $\text{Punct}(k, x)$: This algorithm takes as input a key $k \in \mathcal{K}_{\text{PPRF}}$ and a value $x \in \mathcal{D}_{\text{PPRF}}$, and returns a punctured key $k' \in \mathcal{K}_{\text{PPRF}}$.

Note that the puncturing procedure can also output an unmodified key (i.e. $k' = k$). This is for example reasonable if the procedure is called on an already-punctured value.

Definition 13 (PPRF Correctness). A PPRF is *correct* if for every subset $\{x_1, \dots, x_t\} = \mathcal{S} \subseteq \mathcal{D}_{\text{PPRF}}$ and all $x \in \mathcal{D}_{\text{PPRF}} \setminus \mathcal{S}$, it holds that

$$\Pr \left[\text{Eval}(k_0, x) = \text{Eval}(k_t, x) : \begin{array}{l} k_0 \stackrel{\$}{\leftarrow} \text{Setup}(1^\lambda); \\ k_i = \text{Punct}(k_{i-1}, x_i) \text{ for } i \in [t]; \end{array} \right] = 1.$$

$G_{\text{PPRF}}^{\text{rand}}(\mathcal{A})$	$\mathcal{O}_{\text{Eval}}(x)$
$k \xleftarrow{\$} \text{Setup}(1^\lambda)$	$y \leftarrow \text{Eval}(k, x)$
$b \xleftarrow{\$} \{0, 1\}; \mathcal{Q} := \emptyset$	$\mathcal{Q} := \mathcal{Q} \cup \{x\}$
$x^* \xleftarrow{\$} \mathcal{A}^{\mathcal{O}_{\text{Eval}}(\cdot)}(1^\lambda)$	$k \leftarrow \text{Punct}(k, x)$
$y_0 \xleftarrow{\$} \mathcal{R}_{\text{PPRF}}; y_1 \leftarrow \text{Eval}(k, x^*)$	return y
$k \leftarrow \text{Punct}(k, x^*)$	
$b^* \xleftarrow{\$} \mathcal{A}(k, y_b)$	
return 1 if $b = b^*$ and $x^* \notin \mathcal{Q}$	
return 0	

Figure 14: The rand security experiment for puncturable PRF PPRF.

The security experiment asks that an adversary cannot distinguish an evaluation of a real input (provided by the adversary) from a random output range element, even if the adversary has access to an evaluation oracle and the key that results from puncturing on the challenge input.

Definition 14 (PPRF Security). The advantage of an adversary \mathcal{A} in the rand security experiment $G_{\text{PPRF}}^{\text{rand}}(\mathcal{A})$ defined in Fig. 14 is

$$\text{Adv}_{\text{PPRF}}^{\text{rand}}(\mathcal{A}) := \left| \Pr \left[G_{\text{PPRF}}^{\text{rand}}(\mathcal{A}) = 1 \right] - \frac{1}{2} \right|.$$

5.2 PPRF-based Symmetric AKE

Intuition. The main idea of our PPRF-based protocol is to derive the session key via an evaluation of the PPRF. That is, both parties share a PPRF evaluation key k , which is used to derive session keys by computing $\text{Eval}(k, N_A)$ for some value N_A (in our protocols this will be a counter). After derivation of a session key, the PPRF key will also be punctured at the value N_A by computing $k \leftarrow \text{Punct}(k, N_A)$. Note that the new key k cannot recompute $\text{Eval}(k, N_A)$ as it has been punctured for N_A . This will be our leverage to achieve forward security.

Additionally, the PPRF is an essential building block to achieve full synchronization robustness in our protocols. Intuitively, the puncturing procedure of a PPRF does not evolve its key “linearly” but rather enables fine-grained removal of evaluation capabilities. This guarantees that every protocol run with some fresh value N_A for $\text{Eval}(k, N_A)$ will be completed successfully, even if other protocol runs with some value $N'_A \neq N_A$ are executed in-between.

Our protocols. We present a one-message and a two-message protocol, based on PPRFs. Both protocols have fixed roles, meaning the same party will always initiate (and only this party is required to store the counter). The two-message protocol implicitly authenticates both parties (and thus achieves mutual authentication), while the one-message protocol inherently only achieves responder-only authentication (responder authenticates initiator).

Another important aspect of our protocols is that they use counters to systematically “exhaust” the PPRF. We will later discuss that this approach assists the efficiency of tree-based PPRFs as discussed in Aviram et al. [AGJ19]. The number of session keys that can be derived is equal to the size of the counter space.

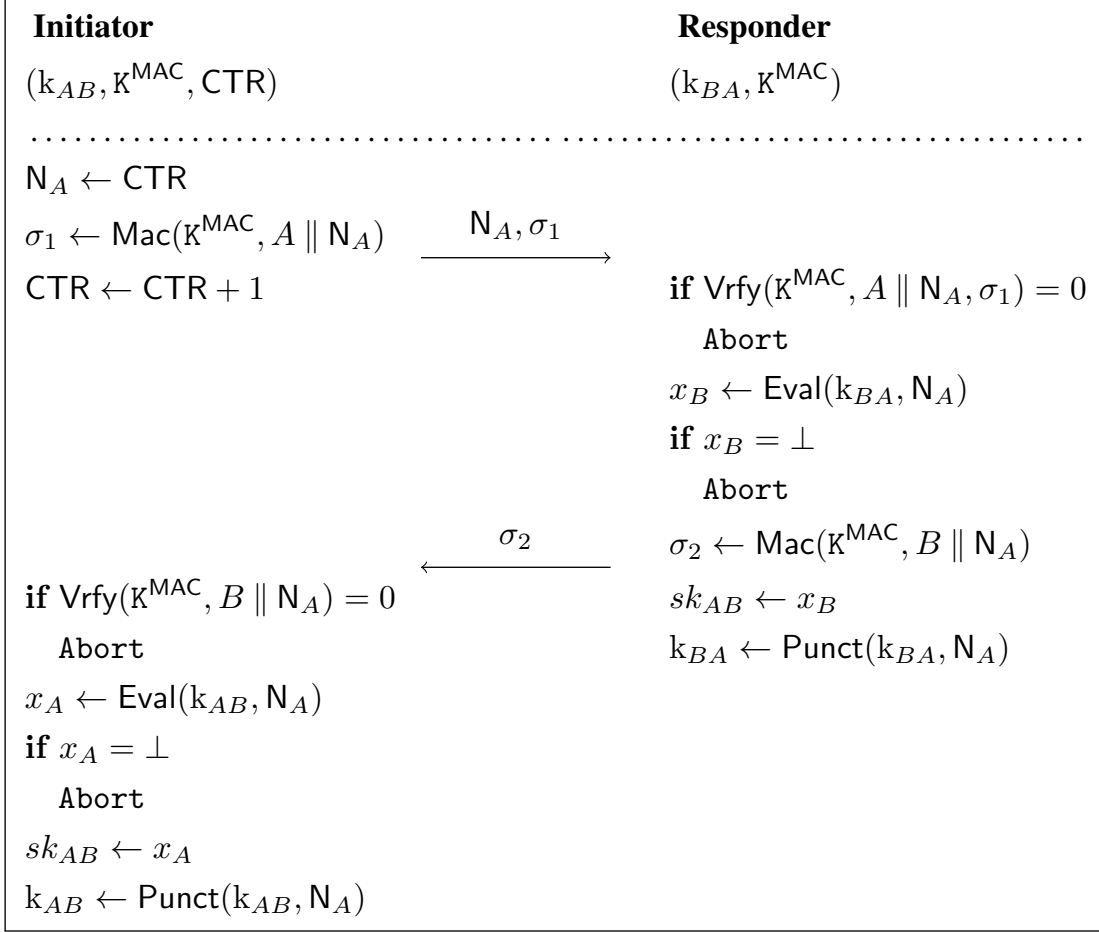


Figure 15: A symmetric AKE protocol PP2 that tolerates concurrent sessions, using a puncturable PRF PPRF = (Setup, Eval, Punct).

5.2.1 AKE-M of PP2

Theorem 18 (AKE-M of PP2). *Let Π be the two-message protocol in Fig. 15, built using $\text{MAC} = \{\text{KGen}, \text{Mac}, \text{Vrfy}\}$ and $\text{PPRF} = (\text{Setup}, \text{Eval}, \text{Punct})$ with n parties. Then for any adversary \mathcal{A} against the AKE-M security of Π that makes a maximum of q queries that initiate new sessions for each party (with $q < |\text{CTR}|$), there exists an adversary $\mathcal{B}_{18.1}$ against the SEUF-CMA- Q of MAC and an adversary $\mathcal{B}_{18.2}$ against the rand security of PPRF such that*

$$\text{Adv}_{\Pi}^{\text{AKE-M}}(\mathcal{A}) \leq 4n^2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-}Q}(\mathcal{B}_{18.1}) + n^2 \cdot q \cdot \text{Adv}_{\text{PPRF}}^{\text{rand}}(\mathcal{B}_{18.2}).$$

We form a bound for each of the three ways in which an adversary can break AKE security, namely Ent-Auth-R, Ent-Auth-I and Key-Ind, and then sum these bounds.

There are two MAC security terms, for entity authentication of responder and initiator, and so $2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-}Q}(\mathcal{B}_{18.1})$ bounds these two terms by fixing $\mathcal{B}_{18.1}$ to be whichever of $\mathcal{B}_{18.1r}$ and $\mathcal{B}_{18.1i}$ has greater advantage.

Lemma 19 (Ent-Auth-R of PP2). *For any adversary \mathcal{A} , the probability that there exists an oracle with $\rho = \text{Responder}$ that accepts maliciously can be bounded by*

$$\text{Adv}_{\Pi}^{\text{Ent-Auth-R}}(\mathcal{A}) \leq n^2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{18.1r})$$

where all quantities are defined as stated in Theorem 18.

PROOF. We proceed using a sequence of games.

Game 0. This is the original Ent-Auth game. We have

$$\Pr \left[G_{\Pi}^{\text{Ent-Auth}}(\mathcal{A}) = 1 \right] = \Pr \left[G_0^{\mathcal{A}} = 1 \right].$$

Note that that Theorem 18 requires that the adversary only initiates $q < |\text{CTR}|$ sessions (which can be guaranteed by choosing $|\text{CTR}|$ exponential in the security parameter), implying that the counter is never exhausted and thus, the initial protocol message sent from initiator to responder is unique. This rules out the possibility of *multiple* oracles with non-unique matching conversations accepting.

Game 1. In this game we guess which responder will be the first to accept maliciously and its partner identity, and abort if this guess is wrong. The game is the same as Game 0 except that the challenger guesses $(i^*, j^*) \xleftarrow{\$} [n] \times [n]$, and if an oracle π_i^s (for some s) accepts maliciously with $\pi_i^s.\rho \neq \text{Responder}$ or $i^* \neq i$ or $j^* \neq \pi_i^s.\text{pid}$, then the challenger aborts. We have

$$\Pr \left[G_0^{\mathcal{A}} = 1 \right] = n^2 \cdot \Pr \left[G_1^{\mathcal{A}} = 1 \right].$$

We conclude our reduction by constructing a reduction $\mathcal{B}_{18.1r}$ that is playing against the SEUF-CMA- Q security of MAC that simulates the environment for an underlying adversary \mathcal{A} that attempts to win in game Game 1. Note that in Game 1, the only way that an adversary can win without an abort occurring is if it makes a responder oracle (of party i^*) accept maliciously, and no initiator oracle (of party j^*) has a matching conversation. To do this, it must initialize an initiator oracle, resulting in some initial protocol message, and then provide a different message to the responder oracle that causes the responder oracle to accept. (If the messages were not different and the adversary forwarded the genuine initial message, then the conversations would match.)

The reduction generates (initial) key derivation keys for all pairs of parties, and authentication keys for all pairs of parties except i^* and j^* . When responding to queries by \mathcal{A} regarding all other parties, the reduction will honestly provide messages as specified in the protocol specification and the Ent-Auth game. For any query made between oracles of parties i^* and j^* , the reduction will use its \mathcal{O}_{Mac} oracle and provide the received value in its simulation for \mathcal{A} . For example, to initialize initiator oracles π_{j^*} , the reduction sets $N := \text{CTR}_{j^*i^*}$, increments $\text{CTR}_{j^*i^*} := \text{CTR}_{j^*i^*} + 1$, and calls $\mathcal{O}_{\text{Mac}}(j^* || N)$. If \mathcal{A} provides a value that it has not been given as an initialization query as input to a `NewSessionR` query from i^* to j^* , then the reduction sends this to its $\mathcal{O}_{\text{Vrfy}}$ oracle in the SEUF-CMA- Q game. The simulation of Game 1 is perfect and any win for \mathcal{A} directly corresponds to a valid signature forgery, so we can write

$$\Pr \left[G_1^{\mathcal{A}} = 1 \right] \leq \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{18.1r}).$$

Summing these terms gives the bound in the statement of Lemma 19. □

The second proof is very similar, and considers malicious acceptance by an initiator, i.e. as a result of a full protocol run of two messages. We only detail significant changes and note that our term collection is exactly the same.

Lemma 20 (Ent-Auth-I of PP2). *For any adversary \mathcal{A} , the probability that there exists an oracle with $\rho = \text{Initiator}$ that accepts maliciously can be bounded by*

$$\text{Adv}_{\Pi}^{\text{Ent-Auth-I}}(\mathcal{A}) \leq n^2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{18.1i})$$

where all quantities are defined as stated in Theorem 18.

PROOF. Game 0 is exactly as in the proof of Lemma 19. Game 1 is as in the proof of Lemma 19, except that now we are guessing which initiator oracle will be the first to accept maliciously (and its intended partner), and so the abort occurs if a responder oracle accepts maliciously. The loss of n^2 incurred by selection of parties is the same.

Again, the next step is a reduction that plays against SEUF-CMA-Q of the MAC scheme MAC. However, the reduction of course must behave slightly differently, since it must send to its own $\mathcal{O}_{\text{Vrfy}}$ oracle any message that was called as a Send query for the targeted initiator oracle, but which was not given as an output protocol message by a NewSessionR query (to responder oracle j^*). Further, we need to ensure that the forgery attempt is on an oracle that does not have a matching conversation with any others: in the proof of Lemma 19 this was straightforward since there was only one unique message in question, but here the transcripts that we are interested in consist of (up to) two flows between each oracle. This is just a matter of bookkeeping, and as before $\mathcal{B}_{18.1i}$ forwards all attempted forgeries to its own verification oracle, so any query that would have caused \mathcal{A} to win the entity authentication game (in the simulation that \mathcal{A} is experiencing) also corresponds to success in the game that $\mathcal{B}_{18.1i}$ is playing. \square

Lemma 21 (Key-Ind of PP2). *For any adversary \mathcal{A} and (any fixed) entity authentication adversary $\mathcal{A}_{18.2}$, the probability that \mathcal{A} answers the Test challenge correctly can be bounded by*

$$\text{Adv}_{\Pi}^{\text{Key-Ind}}(\mathcal{A}) \leq \text{Adv}_{\Pi}^{\text{Ent-Auth}}(\mathcal{A}_{18.2}) + n^2 \cdot q \cdot \text{Adv}_{\text{PPRF}}^{\text{rand}}(\mathcal{B}_{18.2})$$

where all quantities are defined as stated in Theorem 18.

PROOF. Let b' be the bit output by \mathcal{A} in each game, and b be the bit sampled as part of the Test query.

Game 0. This is the original Key-Ind game. We have

$$\Pr [G_{\Pi}^{\text{Key-Ind}}(\mathcal{A}) = 1] = \Pr [G_0^{\mathcal{A}} = 1].$$

Game 1. This game is the same as Game 0, except the challenger aborts and chooses $b' \xleftarrow{\$} \{0, 1\}$ if any oracle accepts maliciously. We have

$$\Pr [G_0^{\mathcal{A}} = 1] \leq \Pr [G_1^{\mathcal{A}} = 1] + \text{Adv}_{\Pi}^{\text{Ent-Auth}}(\mathcal{A}_{21}).$$

At this stage, the oracle to which \mathcal{A} asks its Test query has a unique partner oracle with a matching conversation. (Recall that we only consider adversaries that terminate with valid outputs, and further if \mathcal{A} does anything that would trigger a trivial loss in the original Key-Ind game then it loses in all games in this proof.)

Game 2. This game is the same as Game 1, except the challenger guesses $(i^*, j^*, s^*) \xleftarrow{\$} [n] \times [n] \times [q]$, and if \mathcal{A} issues a $\text{Test}(\pi_i^s)$ query with $(i^*, s^*) \neq (i, s)$ or $j^* \neq \pi_{i^*}^{s^*}.\text{pid}$ then the challenger aborts. We have

$$\Pr [G_1^{\mathcal{A}} = 1] = n^2 \cdot q \cdot \Pr [G_2^{\mathcal{A}} = 1].$$

Now $\pi_{i^*}^{s^*}$ is the oracle to which the Test query will be asked, and this oracle has unique partner $\pi_{j^*}^{t^*}$ with a matching conversation.

Game 3. In this game, the challenger responds to the Test query with a random element of $\mathcal{K}_{\text{PPRF}}$, the output space of PPRF.

We construct a reduction $\mathcal{B}_{18.2}$, detailed in Fig. 16 that runs an adversary that attempts to distinguish Game 3 from Game 2, while $\mathcal{B}_{18.2}$ is playing the rand game.

$$\Pr [G_2^{\mathcal{A}} = 1] = \Pr [G_3^{\mathcal{A}} = 1] + \text{Adv}_{\text{PPRF}}^{\text{rand}}(\mathcal{B}_{18.2})$$

If $b = 1$ in the rand game then $\mathcal{B}_{18.2}$ perfectly simulates Game 2 for \mathcal{A} , while if $b = 0$ in the rand game then $\mathcal{B}_{18.2}$ perfectly simulates Game 3 for \mathcal{A} .

The presentation in Fig. 16 is given for clarity, and omits a number of bookkeeping tasks performed by $\mathcal{B}_{18.2}$, such as managing and updating execution state values, partner identifiers, session key freshness values, security bit values and transcripts for all oracles. Further, if \mathcal{A} makes an invalid query, such as a Send query to an oracle that has already entered $\alpha \in \{\text{accept}, \text{reject}\}$ then $\mathcal{B}_{18.2}$ replies with \perp . Likewise if the adversary submits a Test query to an oracle that has not entered into an accept state or either it or its partner were corrupted before acceptance occurred, then $\mathcal{B}_{18.2}$ will do Abort. Recall that since the MAC keys do not update, $K_{ij}^{\text{MAC}} = K_{ji}^{\text{MAC}}$ throughout, while the key derivation keys k_{ij} and k_{ji} initially start as the same value but may differ as the protocol progresses.

In Game 3 the adversary's advantage of winning is zero, since a Test query will always return a random key that is independent of the protocol,

$$\Pr [G_3^{\mathcal{A}} = 1] = 0.$$

□

5.2.2 SR of PP2

We will now prove that PP2 achieves full synchronization robustness (SR). Intuitively we want to show that any adversary, making arbitrary message delivery queries between any of the parties (and their session oracles), cannot cause an adversarially chosen but honestly executed target protocol run to break down.

The robustness proof essentially needs three arguments: 1) the adversary cannot forge protocol messages without breaking the security of the MAC, 2) replaying messages from the target protocol run to other oracles is not beneficial to the adversary, and 3) the correctness of the PPRF ensures that interleaving queries with nonce values different to the one used in the target session will not influence the successful computation of a session key in the target session.

Theorem 22 (SR of PP2). *Let Π be the two-message protocol in Fig. 15, built using $\text{MAC} = \{\text{KGen}, \text{Mac}, \text{Vrfy}\}$ and PPRF, with n parties. Then for any adversary \mathcal{A} against the SR security of Π that makes a maximum of q queries that initiate new sessions for each party (with $q < |\text{CTR}|$), there exists an adversary \mathcal{B}_{22} against the SEUF-CMA- Q of MAC such that*

$$\text{Adv}_{\Pi}^{\text{SR}}(\mathcal{A}) \leq n^2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-}Q}(\mathcal{B}_{22}).$$

PROOF. We proceed using a sequence of games.

Reduction $\mathcal{B}_{18.2}$ playing $\text{G}_{\text{PPRF}}^{\text{rand}}(\mathcal{B}_{18.2})$	Corrupt(P_i, P_j)
1 : $i^*, j^* \xleftarrow{\$} [n]; s^* \xleftarrow{\$} [q]$ 2 : $\text{CTR}_i = 0$ for all $i \in \{1, \dots, n\}$ 3 : for $i, j \in [n]$ do 4 : $K_{ij}^{\text{MAC}} = K_{ji}^{\text{MAC}} \xleftarrow{\$} \mathcal{K}_{\text{MAC}}$ 5 : $\text{CTR}_{ij} \leftarrow 0$ 6 : for $[n] \times [n] \setminus (i^*, j^*)$ do 7 : $k_{ij} = k_{ji} \xleftarrow{\$} \mathcal{K}_{\text{PPRF}}$ 8 : $\mathcal{A}^{\text{oracles}}$ 9 : When \mathcal{A} calls Test (π_i^s) do 10 : if $(i, s) \neq (i^*, s^*)$ or $j^* \neq \pi_{i^*}^{s^*}.\text{pid}$ 11 : return Abort 12 : submit N^* , receive (k^*, y^*) 13 : $(i^*, s^*, b') \xleftarrow{\$} \mathcal{A}(y^*)$ 14 : return b'	32 : if $i \in \{i^*, j^*\}$ or $j \in \{i^*, j^*\}$ 33 : if $\pi_{i^*}^{s^*}.\alpha \neq \text{accept}$ 34 : return Abort 35 : return k_{ij}
NewSessionI(π_i^s, pid)	Send($\pi_i^s, m, \text{pid} = j$)
14 : $\pi_i^s.\rho \leftarrow \text{Initiator}$ 15 : $\pi_i^s.\alpha \leftarrow \text{negotiating}$ 16 : $\pi_i^s.\text{pid} \leftarrow \text{pid}$ 17 : $N \leftarrow \text{CTR}_{ij}$ 18 : $\text{CTR}_{ij} \leftarrow \text{CTR}_{ij} + 1$ 19 : $\sigma_1 \leftarrow \text{Mac}(K_{ij}^{\text{MAC}}, i \parallel N)$ 20 : $m' \leftarrow N, \sigma_1$ 21 : return m'	36 : Parse m as N, σ_1 37 : if $\text{Vrfy}(K_{ij}^{\text{MAC}}, j \parallel N, \sigma_1) = 0$ 38 : return \perp 39 : if $(i, j) \neq (i^*, j^*) \vee (j^*, i^*)$ 40 : $x_i \leftarrow \text{Eval}(k_{ij}, N)$ 41 : else <i>embed</i> 42 : if Test has not occurred 43 : $x_i \leftarrow \text{call } \mathcal{O}_{\text{Eval}}(N)$ 44 : else 45 : $x_i \leftarrow \text{Eval}(k_{ij}, N)$ 46 : $k_{ij} \leftarrow \text{Punct}(k_{ij}, N)$ 47 : endif 48 : if $x_i = \perp$ 49 : return \perp 50 : $sk_{ij} \leftarrow x_i$ 51 : $\pi_i^s.\alpha \leftarrow \text{accept}$ 52 : if $\pi_i^s.\rho = \text{Responder}$ 53 : $\sigma_2 \leftarrow \text{Mac}(K_{ij}^{\text{MAC}}, i \parallel N)$ 54 : $m' \leftarrow N, \sigma_2$ 55 : return m' 56 : if $(i, j) \neq (i^*, j^*) \vee (j^*, i^*)$ 57 : $k_{ij} \leftarrow \text{Punct}(k_{ij}, N)$ 58 : endif
NewSessionR(π_i^s, pid, m)	
22 : $\pi_i^s.\rho \leftarrow \text{Responder}$ 23 : $\pi_i^s.\alpha \leftarrow \text{negotiating}$ 24 : $\pi_i^s.\text{pid} \leftarrow \text{pid}$ 25 : do Send (π_i^s, m)	
RevealKey(π_i^s)	
26 : if $\pi_i^s.\alpha \neq \text{accept}$ 27 : return \perp 28 : if $(i, s) = (i^*, s^*)$ 29 : return Abort 30 : $\pi_i^s.\kappa \leftarrow \text{exposed}$ 31 : return $\pi_i^s.sk$	

Figure 16: Reduction $\mathcal{B}_{18.2}$ for the proof of Lemma 21. N^* is the initial nonce sent in the transcript between $\pi_{i^*}^{s^*}$ and $\pi_{j^*}^{t^*}$. If at any time \mathcal{A} causes an oracle to accept maliciously, then $\mathcal{B}_{18.2}$ simply does **Abort**. $\mathcal{B}_{18.2}$ provides \mathcal{A} with access to $\text{oracles} = \text{NewSessionI}(\cdot, \cdot, \cdot), \text{NewSessionR}(\cdot, \cdot, \cdot), \text{Send}(\cdot, \cdot, \cdot), \text{RevealKey}(\cdot), \text{Corrupt}(\cdot, \cdot)$.

Game 0. This is the original SR game. We have

$$\Pr [G_{\Pi}^{\text{SR}}(\mathcal{A}) = 1] = \Pr [G_0^{\mathcal{A}} = 1].$$

Note that that Theorem 22 requires that the adversary only initiates $q < |\text{CTR}|$ sessions (which can be guaranteed by choosing $|\text{CTR}|$ exponential in the security parameter), implying that the counter is never exhausted and thus, the value N^* contained in the first protocol message is *unique*.

Game 1. This game is the same as Game 0 except that we define an event E_{22} , that is said to occur if the adversary successfully forges a MAC while it is running, and the challenger aborts if E_{22} occurs. We have

$$\Pr [G_0^{\mathcal{A}} = 1] = \Pr [G_1^{\mathcal{A}} = 1] + \Pr [E_{22}].$$

We now bound the probability that E_{22} occurs. We construct a reduction \mathcal{B}_{22} to the SEUF-CMA- Q security of MAC with a verification oracle. First, the reduction guesses which parties will be involved in the forgery that triggers E_{22} , and then simulates the environment of Game 0. To do this, \mathcal{B}_{22} must select (initial) PPRF keys from the appropriate keyspace and randomly pick MAC keys for all pairs of parties except for the guessed parties i^* and j^* . It is then simple to simulate all queries to oracles of parties except communication between the guessed pair. For any query to an oracle of party i^* with $\text{pid} = j^*$ or an oracle of j^* with $\text{pid} = i^*$, the reduction needs to forward queries to its own MAC and verification oracles. The full reduction is detailed in Fig. 17.

We have

$$\Pr [E_{22}] \leq n^2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-}Q}(\mathcal{B}_{22}).$$

Game 2. This game is the same as Game 1 except that we add an additional requirement to its winning condition: the adversary may not issue queries to interrupt the target protocol execution between π_i^s and π_j^s . Note that this is the additional winning condition required by the weak synchronization robustness experiment. We claim

$$\Pr [G_1^{\mathcal{A}} = 1] = \Pr [G_2^{\mathcal{A}} = 1].$$

To prove the above equality, we need to take a closer look at the sequence of queries made by \mathcal{A} . Let π_i^s and π_j^s be the oracles of the targeted protocol execution. That is, the adversary needs to query

$$\text{NewSessionI}(\pi_i^s, j) \rightarrow m_1, \quad \text{NewSessionR}(\pi_j^t, i, m_1) \rightarrow m_2, \quad \text{Send}(\pi_i^s, m_2)$$

during its runtime, where m_1 contains some nonce N^* . While the adversary has to make those three queries in order, it may interleave the queries with queries of different protocol executions. We make the following three observations:

- Sending the message m_1 to any oracle apart from π_j^t will either make the oracle abort without any modification to k_{ij} , k_{ji} or CTR_{ij} , or make the adversary immediately lose. We distinguish three cases. (i) m_1 is sent to an oracle $\pi_j^{t'}$ with $j \neq j'$. In this case the MAC cannot be verified and that oracle will abort without any modification to k_{ij} , k_{ji} or CTR_{ij} . (ii) m_1 is sent to an oracle $\pi_j^{t'}$ with $t' \neq t$ before it is sent to π_j^t . In this case the oracle $\pi_j^{t'}$ would accept the message, however, the adversary is now unable to output its intended oracles, since the transcripts involve a message that was sent earlier. (iii) m_1 is sent to an oracle $\pi_j^{t'}$ with $t' \neq t$ after it was sent to π_j^t . In this case $\pi_j^{t'}$ will abort after receiving m_1 ($x_B = \perp$ since π_j^t already computed $k_{ji} \leftarrow \text{Punct}(k_{ji}, N^*)$), without any modification to k_{ji} .

Reduction \mathcal{B}_{22} playing $G_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{22})$	NewSessionR(π_i^s, pid, m)
<pre> 1 : $i^*, j^* \xleftarrow{\\$} [n]$ 2 : for $i, j \in [n]$ do 3 : $k_{ij} = k_{ji} \xleftarrow{\\$} \mathcal{K}_{\text{PPRF}}$ 4 : $\text{CTR}_{ij} \leftarrow 0$ 5 : for $[n] \times [n] \setminus (i^*, j^*)$ do 6 : $K_{ij}^{\text{MAC}} = K_{ji}^{\text{MAC}} \xleftarrow{\\$} \mathcal{K}_{\text{MAC}}$ 7 : $\mathcal{A}^{\text{oracles}}$ </pre>	<pre> 19 : $\pi_i^s.\rho \leftarrow \text{Responder}$ 20 : $\pi_i^s.\alpha \leftarrow \text{negotiating}$ 21 : $\pi_i^s.\text{pid} \leftarrow \text{pid}$ 22 : do Send(π_i^s, m) Send(π_i^s, m) / pid = j ----- 23 : Parse m as \mathbb{N}, σ_1 24 : if $(i, j) = (i^*, j^*) \vee (j^*, i^*)$ 25 : $b \leftarrow \text{call } \mathcal{O}_{\text{Vrfy}}(j \parallel \mathbb{N}, \sigma_1)$ 26 : if $b = 0$ 27 : return \perp 28 : if $\pi_i^s.\rho = \text{Responder}$ 29 : $\sigma_2 \leftarrow \text{call } \mathcal{O}_{\text{Mac}}(i \parallel \mathbb{N})$ 30 : return $m' \leftarrow \mathbb{N}, \sigma_2$ 31 : else / simulate 32 : if $\text{Vrfy}(K_{ij}^{\text{MAC}}, j \parallel \mathbb{N}, \sigma_1) = 0$ 33 : return \perp 34 : if $\pi_i^s.\rho = \text{Responder}$ 35 : $\sigma_2 \leftarrow \text{Mac}(K_{ij}^{\text{MAC}}, i \parallel \mathbb{N})$ 36 : return $m' \leftarrow \mathbb{N}, i, \sigma_2$ 37 : $sk_{ji} \leftarrow \text{Eval}(k_{ij}, \mathbb{N}_j)$ 38 : $k_{ij} \leftarrow \text{Punct}(k_{ij}, \mathbb{N})$ </pre>
<pre> NewSessionI(π_i^s, pid) ----- 8 : $\pi_i^s.\rho \leftarrow \text{Initiator}$ 9 : $\pi_i^s.\alpha \leftarrow \text{negotiating}$ 10 : $\pi_i^s.\text{pid} \leftarrow \text{pid} / = j$ 11 : $\mathbb{N} \leftarrow \text{CTR}_{ij}$ 12 : $\text{CTR}_{ij} \leftarrow \text{CTR}_{ij} + 1$ 13 : if $(i, j) = (i^*, j^*) \vee (j^*, i^*)$ 14 : $\sigma_1 \leftarrow \text{call } \mathcal{O}_{\text{Mac}}(i \parallel \mathbb{N})$ 15 : else 16 : $\sigma_1 \leftarrow \text{Mac}(K_{ij}^{\text{MAC}}, i \parallel \mathbb{N})$ 17 : $m' \leftarrow \mathbb{N}, \sigma_1$ 18 : return m' </pre>	

Figure 17: Reduction \mathcal{B}_{22} for the proof of Theorem 22. \mathcal{B}_{22} provides \mathcal{A} with access to $\text{oracles} = \text{NewSessionI}(\cdot, \cdot, \cdot), \text{NewSessionR}(\cdot, \cdot), \text{Send}(\cdot, \cdot, \cdot)$.

- Sending the message m_2 to any oracle apart from π_i^s will make the oracle abort without any modification to k_{ij}, k_{ji} or CTR_{ij} . We distinguish two cases. (i) m_2 is sent to an oracle $\pi_{i'}^{s'}$ with $i \neq i'$. In this case the MAC cannot be verified and that oracle will abort without any modification to k_{ij}, k_{ji} or CTR_{ij} . (ii) m_2 is sent to an oracle $\pi_{i'}^{s'}$ with $s' \neq s$. In this case the uniqueness of \mathbb{N}^* guarantees that the MAC verification will fail and the oracle will abort without any modification to k_{ij} or CTR_{ij} .
- Note that the adversary may not ‘replace’ m_1 or m_2 in the target protocol run as this would immediately result in a loss for \mathcal{A} , as the target protocol run would not consist of an honest protocol run anymore.

We conclude that sending m_1 or m_2 to any oracles apart from \mathcal{A} ’s respective target oracles will either have no impact on the keys of the target parties or the counter CTR_{ij} and thus have no impact on the target protocol run, or make the target run ineligible in the SR game.

It now remains to show that the adversary cannot use any queries of different protocol runs to win the robustness experiment. We do this by ‘isolating’ queries made during the queries related to the target

protocol run. We start with the observation that any queries related to protocol runs with parties $i' \neq i$ and $j' \neq j$ does not have any influence over the target protocol run. Hence, the adversary does not gain any advantage issuing those queries in a way that interleaves with the target protocol run.

At this point, we need to consider the remaining possible queries for any oracle $\pi_i^{s'}$ with $s' \neq s$ or $\pi_j^{t'}$ with $t' \neq t$, which can be interleaved with the target protocol run. The adversary gains an advantage if it is capable of modifying the global session variables, here $(k_{ij}, k_{ji}, K_{ij}^{\text{MAC}} = K_{ji}^{\text{MAC}}, \text{CTR}_{ij})$, in such a way that the target protocol run aborts. We show that there is no such possible query that influences the outcome of the target session by a case distinction:

NewSessionI queries. Any NewSessionI query will cause the counter CTR of the initiator i to be incremented. However, since the counter is only relevant during the generation of the first protocol message, and since the initial message m_1 of the target run has already been generated, this has no impact on the target protocol run.

NewSessionR queries. Any NewSessionR query for $\pi_j^{t'}$ with $t' \neq t$ will either result in the receiving oracle aborting the protocol run (either due to an invalid MAC or a replayed first protocol message), or in puncturing k_{BA} at some position N . Since all values of N are unique and the adversary cannot re-use the first message of the target run m_1 , the adversary cannot cause a puncturing operation on the value N^* contained in the target protocol run. The correctness of the PPRF then guarantees that a consistent evaluation for N^* is possible as long as only values $N \neq N^*$ have been punctured.

Send queries. Any Send query for $\pi_i^{s'}$ with $s' \neq s$, if it does not abort due to an invalid MAC, will puncture k_{ij} (or another key belonging to i , though this does not assist the adversary) at some position N . Any oracle $\pi_i^{s'}$ with $s' \neq s$ will always puncture the PPRF key for some value N , which is not equal to the value N^* contained in the target session. This is ensured by the uniqueness of the counter during the adversary's runtime and by the session storing its respective value N during initialization. Hence, only the target initiator oracle π_i^s is able to puncture the PPRF key for the value N^* . The correctness of the PPRF then guarantees that a consistent evaluation for N^* is possible as long as only values $N \neq N^*$ have been punctured.

We have now exhausted all possible options for the adversary to cause a disturbance of the target protocol run, either via re-using values of the target protocol run before it is concluded, or via interleaving any other protocol message during the target protocol run. Both of which yield no advantage for the adversary. We hence have

$$\Pr [G_1^{\mathcal{A}} = 1] = \Pr [G_2^{\mathcal{A}} = 1].$$

Bounding the advantage of \mathcal{A} . It remains to bound the adversary's advantage in Game 2. Recall that 1) the adversary can now only execute a complete protocol run between the target session, which has full matching transcripts and is not interrupted by any other queries. Furthermore, for any protocol run between two fixed parties, the value N^* is unique, which ensures that for any protocol run, the PPRF key remains not punctured at N^* . In this case, the correctness of PP2 ensures that the target protocol runs will not abort and, in particular, will successfully derive the same session key. We get

$$\Pr [G_2^{\mathcal{A}} = 1] = 0.$$

□

5.3 Protocol PP1

In Figure 18 we give a one-message protocol that uses a PPRF in a very similar way to our two-message protocol. Here we only get one-sided entity authentication, but the proofs are very similar to the ones for the two-message protocol.

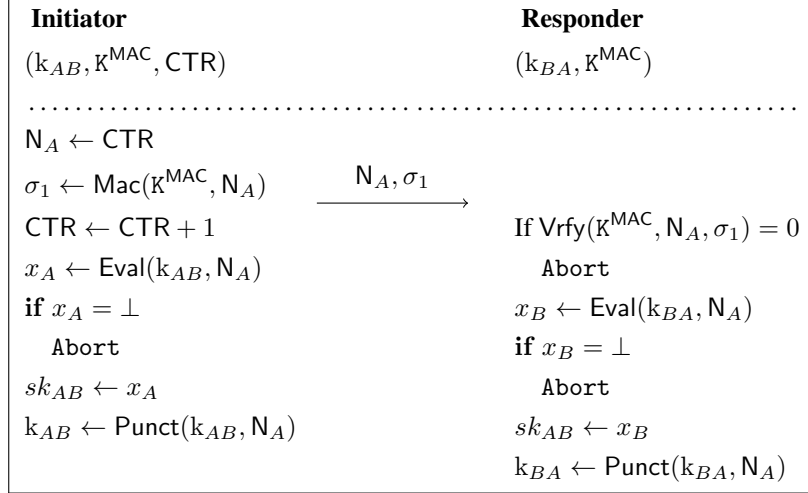


Figure 18: One-message symmetric AKE protocol that tolerates concurrent sessions, using a Puncturable PRF $\text{PPRF} = (\text{Setup}, \text{Eval}, \text{Punct})$.

5.3.1 AKE-R of PP1

Theorem 23 (AKE-R of PP1). *Let Π be the one-message protocol in Fig. 18, built using $\text{MAC} = \{\text{KGen}, \text{Mac}, \text{Vrfy}\}$ and $\text{PPRF} = (\text{Setup}, \text{Eval}, \text{Punct})$ with n parties. Then for any adversary \mathcal{A} against the AKE-R security of Π that makes a maximum of q queries that initiate new sessions for each party (with $q < |\text{CTR}|$), there exists an adversary $\mathcal{B}_{23.1}$ against the SEUF-CMA- Q of MAC and an adversary $\mathcal{B}_{23.2}$ against the rand security of PPRF such that*

$$\text{Adv}_{\Pi}^{\text{AKE-R}}(\mathcal{A}) \leq 2n^2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-}Q}(\mathcal{B}_{23.1}) + n^2 \cdot q \cdot \text{Adv}_{\text{PPRF}}^{\text{rand}}(\mathcal{B}_{23.2}).$$

The Ent-Auth-R analysis is essentially the same as in Lemma 19 and leads to the same term collection, since in PP1 the responder performs the same notable actions as in PP2 (the creation of the second protocol message is not necessary).

Proving Key-Ind is also similar to that of Lemma 21, leading to the same term collection. The final reduction is slightly more straightforward than $\mathcal{B}_{18.2}$ since the only Send queries that $\mathcal{B}_{23.2}$ needs to deal with are those that are part of NewSessionR queries.

5.3.2 SR of PP1

Theorem 24 (SR of PP1). *Let Π be the two-message protocol in Fig. 18, built using $\text{MAC} = \{\text{KGen}, \text{Mac}, \text{Vrfy}\}$ and PPRF, with n parties. Then for any adversary \mathcal{A} against the SR security of Π that makes a maximum of q queries that initiate new sessions for each party (with $q < |\text{CTR}|$), there exists an adversary \mathcal{B}_{24} against the SEUF-CMA- Q of MAC such that*

$$\text{Adv}_{\Pi}^{\text{SR}}(\mathcal{A}) \leq n^2 \cdot \text{Adv}_{\text{MAC}}^{\text{SEUF-CMA-}Q}(\mathcal{B}_{24}).$$

PROOF. The strategy for this proof is very similar to that of Theorem 22, and the term collection is the same: only the reduction logic is different (since there are no second protocol messages to deal with). The reduction \mathcal{B}_{24} is detailed in Fig. 19. Note that for one-message protocols, there are no Send queries that were not initially called by NewSessionR, since those queries will always result in the oracle going to state accept or reject (in other words, an oracle cannot ever be negotiating at the point an adversary makes a query). However, formatting is maintained to provide easier comparison with reduction \mathcal{B}_{22} .

Reduction \mathcal{B}_{24} playing $G_{\text{MAC}}^{\text{SEUF-CMA-Q}}(\mathcal{B}_{24})$	NewSessionR(π_i^s, pid, m)
1 : $i^*, j^* \xleftarrow{\$} [n]$	21 : $\pi_i^s.\rho \leftarrow \text{Responder}$
2 : for $i, j \in [n]$ do	22 : $\pi_i^s.\alpha \leftarrow \text{negotiating}$
3 : $k_{ij} = k_{ji} \xleftarrow{\$} \mathcal{K}_{\text{PPRF}}$	23 : $\pi_i^s.\text{pid} \leftarrow \text{pid}$
4 : $\text{CTR}_{ij} \leftarrow 0$	24 : do Send(π_i^s, m)
5 : for $[n] \times [n] \setminus (i^*, j^*)$ do	Send($\pi_i^s, m, \text{pid} = j$)
6 : $K_{ij}^{\text{MAC}} = K_{ji}^{\text{MAC}} \xleftarrow{\$} \mathcal{K}_{\text{MAC}}$	25 : Parse m as \mathbb{N}, σ_1
7 : $\mathcal{A}^{\text{oracles}}$	26 : if $\pi_i^s.\rho = \text{Initiator}$
NewSessionI(π_i^s, pid)	27 : return Abort
8 : $\pi_i^s.\rho \leftarrow \text{Initiator}$	28 : if $(i, j) = (i^*, j^*) \vee (j^*, i^*)$ embed
9 : $\pi_i^s.\alpha \leftarrow \text{negotiating}$	29 : $b \leftarrow \text{call } \mathcal{O}_{\text{Vrfy}}(\mathbb{N}, \sigma_1)$
10 : $\pi_i^s.\text{pid} \leftarrow \text{pid} \neq j$	30 : if $b = 0$
11 : $\mathbb{N} \leftarrow \text{CTR}_{ij}$	31 : return \perp
12 : if $(i, j) = (i^*, j^*) \vee (j^*, i^*)$	32 : else <i>simulate</i>
13 : $\sigma_1 \leftarrow \text{call } \mathcal{O}_{\text{Mac}}(\mathbb{N})$	33 : if $\text{Vrfy}(K_{ij}^{\text{MAC}}, \mathbb{N}, \sigma_1) = 0$
14 : else	34 : return \perp
15 : $\sigma_1 \leftarrow \text{Mac}(K_{ij}^{\text{MAC}}, \mathbb{N})$	35 : $sk_{ji} \leftarrow \text{Eval}(k_{ij}, \mathbb{N}_j)$
16 : $m' \leftarrow \mathbb{N}, \sigma_1$	36 : $k_{ij} \leftarrow \text{Punct}(k_{ij}, \mathbb{N})$
17 : $\text{CTR}_{ij} \leftarrow \text{CTR}_{ij} + 1$	
18 : $sk_{ji} \leftarrow \text{Eval}(k_{ij}, \mathbb{N}_j)$	
19 : $k_{ij} \leftarrow \text{Punct}(k_{ij}, \mathbb{N})$	
20 : return m'	

Figure 19: Reduction \mathcal{B}_{24} for the proof of Theorem 24. \mathcal{B}_{24} provides \mathcal{A} with access to oracles = NewSessionI(\cdot, \cdot, \cdot), NewSessionR(\cdot, \cdot), Send(\cdot, \cdot, \cdot).

□

5.4 Instantiation

It remains to discuss how PP2 can be instantiated with a PPRF and what impact the PPRF has on its efficiency. A promising candidate is the Goldreich–Goldwasser–Micali PRF [GGM86], which can be transformed to a PPRF as noted by [BW13, KPTZ13, BGI14]. We give an intuitive explanation of the construction and refer the reader to [AGJ19] for a more detailed description and analysis. This construction is

especially suitable, as both the PPRF evaluation and puncturing are solely based on hash function evaluations in practice.

Intuition. The tree-based PPRF uses two functions H_0 and H_1 both mapping from $\{0, 1\}^\lambda$ to $\{0, 1\}^\lambda$. For every input $x \in \{0, 1\}^\lambda$ of the PPRF, the binary representation of x prescribes the sequence in which H_0 and H_1 have to be repeatedly applied to x . For example, $\text{Eval}(01) = H_1(H_0(x))$. Note that the evaluation of x corresponds to a path through a binary tree, where each bit in x tells you whether to take a “left” or “right” path. The result of an evaluation always corresponds to a leaf in the binary tree.

The initial PPRF key consists of the root node, which is initialized during key generation as a randomly chosen string. To puncture values (i.e., to puncture leaves of the tree), we precompute and store all nodes on the co-path between the root and the leaf, before deleting all parent nodes (including the root node) of the leaf. Note that this procedure can be repeated for any of the leaves and note that it satisfies all puncturing-relevant properties (i.e., re-computation of $\text{Eval}(x)$ is not possible but the correctness of the PPRF remains intact).

Memory Consumption. We briefly discuss the memory consumed by the PPRF during the lifetime of PP2 (and PP1). First, note that the PPRF-based protocols deploy counters, which (if all messages are delivered in sequence) ensure a systematic puncturing from the leftmost leaf to the rightmost leaf of the binary tree. This yields the need to store at most $\log(|\text{CTR}|)$ tree nodes (i.e., at most one node per layer of the tree) at any point in time. For C concurrent sessions, this bound increases to a maximum of $C \cdot \log(|\text{CTR}|)$ tree nodes.

The analysis gets slightly more difficult if an adversary *actively drops* protocols messages. Each dropped message will either cause the initiator or both parties to not puncture at some position. One approach to tame the memory consumption in this case, would be to always puncture on all values which are smaller than $\text{CTR} - C$.¹ As we never expect more than C in parallel, this reduces additional memory caused by lost messages. In this case, the memory consumption is again upper-bounded by $C \cdot \log(|\text{CTR}|)$ tree nodes.

Finally, note that in the one-message protocol PP1 (Sec. 5.3) the initiator always punctures strictly in order and as thus has to store at most $\log(|\text{CTR}|)$ tree nodes. Note that this may be particularly useful in an application where many low-end devices communicate with a central server.

Acknowledgements. In addition to the funding bodies acknowledged on page 1, we would also like to thank Luke Mather for numerous helpful comments.

References

- [ACD19] Joël Alwen, Sandro Coretti, and Yevgeniy Dodis. The double ratchet: Security notions, proofs, and modularization for the Signal protocol. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019, Part I*, volume 11476 of *Lecture Notes in Computer Science*, pages 129–158, Darmstadt, Germany, May 19–23, 2019. Springer, Heidelberg, Germany.
- [ACF20] Gildas Avoine, Sébastien Canard, and Loïc Ferreira. Symmetric-key authenticated key exchange (SAKE) with perfect forward secrecy. In Stanislaw Jarecki, editor, *Topics in Cryptology – CT-*

¹Interestingly, the tree-based PPRF can puncture multiple values in one go by “chopping off” whole branches of the tree, instead of puncturing all values one after another.

- RSA 2020*, volume 12006 of *Lecture Notes in Computer Science*, pages 199–224, San Francisco, CA, USA, February 24–28, 2020. Springer, Heidelberg, Germany.
- [AGJ19] Nimrod Aviram, Kai Gellert, and Tibor Jager. Session resumption protocols and efficient forward security for TLS 1.3 0-RTT. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019, Part II*, volume 11477 of *Lecture Notes in Computer Science*, pages 117–150, Darmstadt, Germany, May 19–23, 2019. Springer, Heidelberg, Germany.
- [AGJ21] Nimrod Aviram, Kai Gellert, and Tibor Jager. Session resumption protocols and efficient forward security for TLS 1.3 0-RTT. *Journal of Cryptology*, 34(3):1–57, 2021.
- [ANS09] Retail Financial Services Symmetric Key Management Part 1: Using Symmetric Techniques (ANSI x9.24). Standard, American National Standards Institute, New York, USA, 2009.
- [BG20] Colin Boyd and Kai Gellert. A Modern View on Forward Security. *The Computer Journal*, 08 2020. <https://doi.org/10.1093/comjnl/bxaa104>.
- [BGI14] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In Hugo Krawczyk, editor, *PKC 2014: 17th International Conference on Theory and Practice of Public Key Cryptography*, volume 8383 of *Lecture Notes in Computer Science*, pages 501–519, Buenos Aires, Argentina, March 26–28, 2014. Springer, Heidelberg, Germany.
- [BGM04] Mihir Bellare, Oded Goldreich, and Anton Mityagin. The power of verification queries in message authentication and authenticated encryption. *Cryptology ePrint Archive*, Report 2004/309, 2004. <http://eprint.iacr.org/2004/309>.
- [BM99] Mihir Bellare and Sara K. Miner. A forward-secure digital signature scheme. In Michael J. Wiener, editor, *Advances in Cryptology – CRYPTO’99*, volume 1666 of *Lecture Notes in Computer Science*, pages 431–448, Santa Barbara, CA, USA, August 15–19, 1999. Springer, Heidelberg, Germany.
- [BP10] Eric Brier and Thomas Peyrin. A forward-secure symmetric-key derivation protocol - how to improve classical DUKPT. In Masayuki Abe, editor, *Advances in Cryptology – ASIACRYPT 2010*, volume 6477 of *Lecture Notes in Computer Science*, pages 250–267, Singapore, December 5–9, 2010. Springer, Heidelberg, Germany.
- [BPR00] Mihir Bellare, David Pointcheval, and Phillip Rogaway. Authenticated key exchange secure against dictionary attacks. In Bart Preneel, editor, *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 139–155, Bruges, Belgium, May 14–18, 2000. Springer, Heidelberg, Germany.
- [BR94] Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In Douglas R. Stinson, editor, *Advances in Cryptology – CRYPTO’93*, volume 773 of *Lecture Notes in Computer Science*, pages 232–249, Santa Barbara, CA, USA, August 22–26, 1994. Springer, Heidelberg, Germany.
- [BW13] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *Advances in Cryptology – ASIACRYPT 2013, Part II*, volume 8270 of *Lecture Notes in Computer Science*, pages 280–300, Bangalore, India, December 1–5, 2013. Springer, Heidelberg, Germany.

- [BY03] Mihir Bellare and Bennet S. Yee. Forward-security in private-key cryptography. In Marc Joye, editor, *Topics in Cryptology – CT-RSA 2003*, volume 2612 of *Lecture Notes in Computer Science*, pages 1–18, San Francisco, CA, USA, April 13–17, 2003. Springer, Heidelberg, Germany.
- [CK01] Ran Canetti and Hugo Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In Birgit Pfitzmann, editor, *Advances in Cryptology – EUROCRYPT 2001*, volume 2045 of *Lecture Notes in Computer Science*, pages 453–474, Innsbruck, Austria, May 6–10, 2001. Springer, Heidelberg, Germany.
- [CMA05] (NIST SP)-800-38B. recommendation for block cipher modes of operation: The CMAC mode for authentication. Special Publication. Standard, NIST, 2005.
- [CRSS20] Valerio Cini, Sebastian Ramacher, Daniel Slamanig, and Christoph Striecks. CCA-secure (puncturable) KEMs from encryption with non-negligible decryption errors. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2020, Part I*, volume 12491 of *Lecture Notes in Computer Science*, pages 159–190, Daejeon, South Korea, December 7–11, 2020. Springer, Heidelberg, Germany.
- [DGJ⁺21] David Derler, Kai Gellert, Tibor Jager, Daniel Slamanig, and Christoph Striecks. Bloom filter encryption and applications to efficient forward-secret 0-RTT key exchange. *Journal of Cryptology*, 34(2):1–59, 2021.
- [DJ15] Mohammad Sadeq Dousti and Rasool Jalili. FORSAKES: A forward-secure authenticated key exchange protocol based on symmetric key-evolving schemes. *Adv. Math. Commun.*, 9(4):471–514, 2015.
- [DJSS18] David Derler, Tibor Jager, Daniel Slamanig, and Christoph Striecks. Bloom filter encryption and applications to efficient forward-secret 0-RTT key exchange. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part III*, volume 10822 of *Lecture Notes in Computer Science*, pages 425–455, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Heidelberg, Germany.
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, October 1986.
- [GHJL17] Felix Günther, Britta Hale, Tibor Jager, and Sebastian Lauer. 0-RTT key exchange with full forward secrecy. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology – EUROCRYPT 2017, Part III*, volume 10212 of *Lecture Notes in Computer Science*, pages 519–548, Paris, France, April 30 – May 4, 2017. Springer, Heidelberg, Germany.
- [GM15] Matthew D. Green and Ian Miers. Forward secure asynchronous messaging from puncturable encryption. In *2015 IEEE Symposium on Security and Privacy*, pages 305–320, San Jose, CA, USA, May 17–21, 2015. IEEE Computer Society Press.
- [GMA07] (NIST SP)-800-38D. Recommendation for block cipher modes of operation: Galois/Counter Mode (GCM) and GMAC. Special Publication. Standard, NIST, 2007.
- [HMA08] FIPS 198-1. the Keyed-Hash Message Authentication Code (HMAC). Standard, NIST, 2008.

- [ISO11] ISO/IEC 9797-1:2011. Message Authentication Codes (MACs) – part 1: Mechanisms using a block cipher. Standard, International Organization for Standardization / International Electrotechnical Commission, 2011.
- [JKSS12] Tibor Jager, Florian Kohlar, Sven Schäge, and Jörg Schwenk. On the security of TLS-DHE in the standard model. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology – CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 273–293, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Heidelberg, Germany.
- [KMA16] (NIST SP)-800-185. SHA-3 derived functions: cSHAKE, KMAC, TupleHash and ParallelHash. Special Publication. Standard, NIST, 2016.
- [KPTZ13] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 2013: 20th Conference on Computer and Communications Security*, pages 669–684, Berlin, Germany, November 4–8, 2013. ACM Press.
- [Kra05] Hugo Krawczyk. HMQV: A high-performance secure Diffie-Hellman protocol. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 546–566, Santa Barbara, CA, USA, August 14–18, 2005. Springer, Heidelberg, Germany.
- [LBdM07] Tri Van Le, Mike Burmester, and Breno de Medeiros. Universally composable and forward-secure RFID authentication and authenticated key exchange. In Feng Bao and Steven Miller, editors, *ASIACCS 07: 2nd ACM Symposium on Information, Computer and Communications Security*, pages 242–252, Singapore, March 20–22, 2007. ACM Press.
- [LSY⁺14] Yong Li, Sven Schäge, Zheng Yang, Florian Kohlar, and Jörg Schwenk. On the security of the pre-shared key ciphersuites of TLS. In Hugo Krawczyk, editor, *PKC 2014: 17th International Conference on Theory and Practice of Public Key Cryptography*, volume 8383 of *Lecture Notes in Computer Science*, pages 669–684, Buenos Aires, Argentina, March 26–28, 2014. Springer, Heidelberg, Germany.
- [SSS⁺20] Shifeng Sun, Amin Sakzad, Ron Steinfeld, Joseph K. Liu, and Dawu Gu. Public-key puncturable encryption: Modular and compact constructions. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020: 23rd International Conference on Theory and Practice of Public Key Cryptography, Part I*, volume 12110 of *Lecture Notes in Computer Science*, pages 309–338, Edinburgh, UK, May 4–7, 2020. Springer, Heidelberg, Germany.
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th Annual ACM Symposium on Theory of Computing*, pages 475–484, New York, NY, USA, May 31 – June 3, 2014. ACM Press.