



Tightly-Secure Authenticated Key Exchange, Revisited

Tibor Jäger¹(✉), Eike Kiltz², Doreen Riepel², and Sven Schäge²

¹ Bergische Universität Wuppertal, Wuppertal, Germany
`tibor.jaeger@uni-wuppertal.de`

² Ruhr-Universität Bochum, Bochum, Germany
{eike.kiltz,doreen.riepel,sven.schaege}@rub.de

Abstract. We introduce new tightly-secure authenticated key exchange (AKE) protocols that are extremely efficient, yet have only a *constant* security loss and can be instantiated in the random oracle model both from the standard DDH assumption and a subgroup assumption over RSA groups. These protocols can be deployed with optimal parameters, independent of the number of users or sessions, without the need to compensate a security loss with increased parameters and thus decreased computational efficiency.

We use the standard “Single-Bit-Guess” AKE security (with forward secrecy and state corruption) requiring all challenge keys to be simultaneously pseudo-random. In contrast, most previous papers on tightly secure AKE protocols (Bader et al., TCC 2015; Gjøsteen and Jäger, CRYPTO 2018; Liu et al., ASIACRYPT 2020) concentrated on a non-standard “Multi-Bit-Guess” AKE security which is known not to compose tightly with symmetric primitives to build a secure communication channel.

Our key technical contribution is a new generic approach to construct tightly-secure AKE protocols based on non-committing key encapsulation mechanisms. The resulting DDH-based protocols are considerably more efficient than all previous constructions.

Keywords: Authenticated key exchange · Tightness · Non-committing encryption · Forward security

1 Introduction

Authenticated Key Exchange (AKE) is a fundamental cryptographic primitive with immense practical importance. The goal is to securely establish a session key between two parties in a network where an adversary can read, send, modify or delete messages and may also corrupt selected parties and sessions.

TIGHTNESS OF AKE. When proving a cryptographic scheme secure, one commonly describes a security reduction which transforms an adversary \mathcal{A} that breaks the cryptographic scheme into an adversary \mathcal{B} that solves some underlying complexity assumption. For instance, if \mathcal{A} has advantage ϵ in breaking the

scheme and \mathcal{B} solves the problem with advantage $\epsilon' = \epsilon/L$, then L is called the reduction’s security loss. If L is constant (and in particular independent of the number of \mathcal{A} ’s oracle queries) and additionally the running times of \mathcal{A} and \mathcal{B} are roughly identical, then we say the reduction is *tight*. Especially when choosing protocol-specific system parameters, the tightness of a security proof plays an important role. In the security model for AKE the attacker can actively control all messages sent between the involved parties and is additionally allowed to reveal secret information such as a long-term secret key (by corrupting a party), or a session key. The adversary breaks security if it is able to distinguish non-revealed session keys from random.

MULTI-CHALLENGE SECURITY DEFINITIONS. The standard and well established security notion in the context of multiple challenges [3, 10, 18, 20] is “Single-Bit Guess” (SBG) security. The blueprint of a SBG security experiment is as follows. First, the experiment picks a secret random bit $b \in \{0, 1\}$. Next, the adversary is allowed to make multiple (up to, say, T) challenge queries. On each challenge query, the experiment returns a “real key” if $b = 0$, and an independent “random key” if $b = 1$. The adversary wins if it can guess the challenge bit b with a probability better than $1/2$.

In AKE protocols, challenge queries are usually called test queries and non-revealed session keys can be accessed by making multiple calls to a TEST oracle. If $b = 0$, a query to TEST returns the real challenge key; if $b = 1$, a query to TEST returns an independent random challenge key. This notation of multi-challenge SBG security for AKE was first formalized in 2019 by Cohn-Gordon et al. [10]. By conditioning on bit b , SBG security is known to be tightly equivalent to (single-bit) “Real-Or-Random” (ROR) security, where the adversary has to distinguish a real game (where all challenge keys output by TEST are real) from a random game (where all challenge keys are random). Using the above equivalence, SBG security precisely captures the intuition that *all challenge keys* are simultaneously pseudo-random.

Surprisingly, in the first publication on tightly secure AKE protocols in 2015, Bader et al. [1] defined a different and non-standard “Multi-Bit-Guess” (MBG) AKE security notion. In MBG security, the experiment picks multiple independent challenge bits b_1, \dots, b_T and, on the i -th TEST query, it returns a real challenge key if $b_i = 0$ and a random challenge key if $b_i = 1$. That is, each of the T challenge keys depends on an independent challenge bit b_i . The adversary wins if it can guess correctly one of the T challenge bits b_{i^*} with a probability better than $1/2$. We are not aware of any meaningful multi-bit ROR security game that is tightly equivalent to MBG security.¹ This makes it difficult to provide a good intuition of what MBG security tries to model.

¹ If one tries to apply a similar conditioning argument as in the single-bit case, MBG can be shown equivalent to a ROR-type security experiment where in the real game ($b_{i^*} = 0$) the i^* -th challenge key output by TEST is real and in the random game ($b_{i^*} = 1$) it is random. However, the remaining $T - 1$ keys still depends on the random bits b_i ($i \neq i^*$): the i -th challenge key is real if $b_i = 0$ and it is random if $b_i = 1$. Hence, about one half of the challenge keys is expected to be real (the ones

CHOOSING A MEANINGFUL SECURITY MODEL FOR AKE. SBG and MBG security are asymptotically equivalent but only imply each other with a security loss of T , the total number of TEST queries. Hence, when considering tightness, one has to carefully choose a meaningful security model.

First off, as already pointed out, SBG security is the standard and well established security notion in the context of multiple challenges [3, 10, 18, 20]. Cohn-Gordon et al. [10, Section 3] already pointed out that, in the AKE setting, SBG security tightly composes with symmetric primitives, whereas MBG security doesn't. Let us elaborate. AKE is not intended to be used as a stand-alone primitive. Rather, it is naturally composed with symmetric primitives to establish a secure channel [7, 24], for example to encrypt (e.g., using AES) a message with the session key. Since SBG security is tightly equivalent to ROR security, it offers precisely the right security interface to switch *all challenge keys at once* from real to random. This step allows to infer the privacy of the encrypted messages from the security properties of the symmetric primitive. MBG security, on the other hand, does not have a meaningful ROR-style security, which makes it difficult to argue about the privacy of the encrypted messages without relying on a hybrid argument. In summary, in the context of tightness of AKE protocols, SBG security is a meaningful notion whereas MBG isn't.

PREVIOUS RESULTS. Previous work on tight AKE protocols by Gjøsteen and Jager [21] and Liu et al. [32] exclusively concentrated on the MBG model by Bader et al. [1]. We now give a brief overview of existing AKE protocols in the context of tight SBG security.

- At CRYPTO 2019, Cohn-Gordon et al. [10] presented highly efficient two message AKE protocols with implicit authentication, in the style of HMQV [26] and similar protocols. Their schemes achieve a loss of $O(N)$ in the SBG security model with weak forward secrecy, where N is the number of users. They also extend the impossibility results from [2] to show that a loss of $O(N)$ is unavoidable for many natural protocols (including HMQV [26], NAXOS [28], Kudla-Paterson [27], KEA+ [29], and more) with respect to typical cryptographic security proofs (so-called simple reductions). Furthermore, since their protocol does not feature explicit authentication, a well-known impossibility result applies [6, 26, 34] and their protocol cannot achieve full forward security.
- Diemert and Jager [16] and independently Davis and Günther [15] considered the three message TLS 1.3 handshake AKE protocol with explicit authentication. Its design follows the standard “ $1 \times \text{KEM} + 2 \times \text{SIG}$ ” (aka. signed Diffie-Hellman) AKE approach [9, 14–16, 21, 32]. TLS 1.3, when instantiated with standardized signatures (e.g., RSA-PSS, RSA-PKCS #1 v1.5, ECDSA, or EdDSA), has rather non-tight SBG security with full forward security. But when instantiated with tightly secure signatures in the multi-user setting with adaptive corruptions [1], then SBG security of TLS 1.3 actually becomes tight. Since the TLS 1.3 protocol contains two signatures, the inefficiency of cur-

with $b_i = 0$) whereas the other half is random, and the adversary does not have any information on them.

rently known tightly secure signature schemes [1, 21] makes the resulting TLS instantiation very impractical.

1.1 The Difficulty of Constructing Tightly Secure AKE

Security models for authenticated key exchange are extremely complex, as they consider very strong adversaries that may modify, drop, or inject messages. Furthermore, usually an adversary may adaptively corrupt users' long-term secrets via CORRUPT-queries, session keys via REVEAL-queries, and sometimes even ephemeral states of sessions via REV-STATE-queries. Security is formalized with multiple TEST queries, where the adversary specifies a session, receives back a real key or a random key, and has to distinguish these. This complexity makes achieving tight security challenging, particularly because all the following difficulties must be tackled simultaneously.

THE "COMMITMENT" PROBLEM. As explained in more detail in [21], this problem is the reason why nearly all security proofs of classical key exchange protocols have a quadratic security loss. Essentially, the problem is that most AKE protocols have security proofs where a reduction can only extract a solution to a computationally hard problem if an instance of the problem is embedded into the protocol messages of the TESTED sessions, but at the same time the reduction is not able to answer REVEAL queries for such sessions. The standard way to resolve this is to let the reduction guess the TESTED session, and to embed an instance of a computationally hard problem only there. However, this incurs a significant security loss. A tight reduction has to be able to respond to *both* TEST and REVEAL queries for *every* session.

THE PROBLEM OF LONG-TERM KEY REVEALS. A CORRUPT query in typical AKE security models enables the adversary to obtain the long-term key of certain users. If we want to avoid a security loss that results from guessing corrupted and non-corrupted parties, then we must be able to construct a reduction that "knows" valid-looking long-term keys for all users throughout the security experiment. However, this is a major difficulty, for instance, in protocols where the long-term keys are key pairs for a digital signature scheme. The difficulty is that in the security proof we would have to describe a reduction that is able to extract a solution to a computationally hard problem from a forged signature, even though it "knows" the signing key and thus is able to compute a valid signature itself. Hence, in order to obtain a tightly-secure AKE protocol, one needs to devise a way such that a reduction always knows all secret keys, yet is able to argue that an adversary is, e.g., not able to forge signatures.

In order to resolve this issue, previous works [1, 21] constructed signature schemes based on non-interactive OR-proof systems, which enable a reduction to "know" one out of two signing keys. It is argued that the adversary will forge a signature with respect to the other, unknown key with sufficiently high probability. However, these signature schemes are much less efficient than classical ones, and thus impose a performance penalty on the protocols.

THE PROBLEM OF EPHEMERAL STATE REVEALS. Yet another difficulty arises when the security model allows ephemeral state reveals. Previous works on tightly-secure AKE did not consider this very strong security notion at all, therefore we face (and solve) this problem for the first time. From a high-level perspective, the issue is similar to the long-term key reveal problem, except that ephemeral states are considered. In order to achieve tightness, the reduction must be able to output valid-looking states for all sessions. Note that this includes even TESTED sessions, where ephemeral states may be revealed when parties are not corrupted.

1.2 Main Contributions

Summarizing the previous paragraphs, we can formulate the following natural questions related to tightly secure AKE:

- Q1:** Do there exist implicitly authenticated two-message AKEs with tight SBG security, state reveals, and weak forward security?
- Q2:** Do there exist explicitly authenticated two-message AKEs with tight SBG security, state reveals, and full forward security, with *one single* signature?

In this work, we answer the two questions to the positive. Following [4, 10], we consider SBG security, allowing adaptive corruptions of long-term secrets, adaptive reveals of session keys, and multiple adaptive TEST queries. Our model also captures (weak and full) forward security (FS), and prevents key-compromise impersonation and reflection attacks. In comparison to prior work on tightly-secure key exchange [1, 10, 15, 16, 21], we consider a model which additionally allows to reveal some internal state information.

OUR DDH-BASED AKE PROTOCOLS. Our two protocols instantiated from DDH are given in Fig. 1. $\text{AKE}_{\text{wFS,DDH}}$ is an implicitly-authenticated two-message protocol $\text{AKE}_{\text{wFS,DDH}}$ in the sense of [26]. It requires the exchange of only five group elements in total, and thus is the first efficient implicitly-authenticated protocol with weak FS that achieves full tightness.

Our second protocol $\text{AKE}_{\text{FS,DDH}}$ achieves full FS. Instead of using the standard “1×KEM+2×SIG” approach, it replaces one of the signatures with a more efficient MAC and an additional KEM ciphertext, which yields a “2×KEM+1×SIG+1×MAC” construction. When instantiated at “128-bit security” with the most efficient tightly-secure signatures of [21],² the communication complexity is 448 bytes, again with ephemeral state reveals. In comparison, the previously most efficient tightly and fully forward-secure protocol with SBG security TLS* (which is TLS 1.3 instantiated with the tightly-secure signature of [21]) requires three messages, the transmission of 704 bytes and does not allow state reveals. See Fig. 2 for a comparison of our protocols with previous works. Note that the communication bottleneck in all full FS protocols is the number of signatures. For completeness the figure also list previous protocols with tight MBG security [21, 32].

² The signatures of [21] consist of 2 group elements, 4 elements in \mathbb{Z}_p and 2 hashes in $\{0, 1\}^k$. At “128-bit security” this corresponds to 256 bytes per signature.

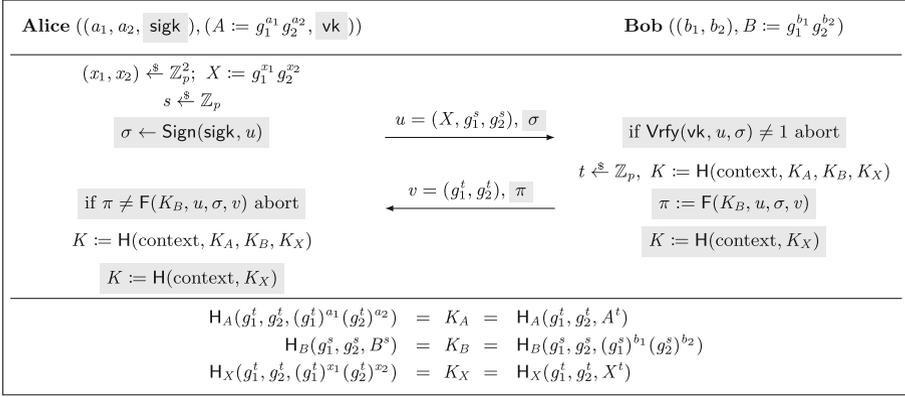


Fig. 1. The two message protocols $\text{AKE}_{\text{wFS}, \text{DDH}}$ (without the gray boxes) and $\text{AKE}_{\text{FS}, \text{DDH}}$ (including the gray boxes), where K is the resulting session key. We define $\text{context} := (A, B, X, \text{vk}, g_1^s, g_2^s, g_1^t, g_2^t, \sigma, \pi)$. H, H_A, H_B, H_X and F are hash functions.

GENERIC CONSTRUCTIONS OF AKE FROM NCKE. Our main technical tool is a new approach to achieve a tight reduction for authenticated key exchange protocols. Our starting point is an extension of (receiver) non-committing encryption (NCE) [8, 33] to *non-committing key encapsulation (NCKE) in the multi-user setting with corruptions*. We construct an NCKE scheme in the random oracle model from any smooth projective hash proof system (HPS) [11]. If the HPS' subset membership problem (SMP) is hard in the multi-instance setting, then the NCKE scheme is also tightly secure in our multi-user setting. We provide two such HPS, one from the DDH assumption, and another one from a subgroup assumption over groups of unknown order. The construction allows us to address the commitment problem described above.

We give a generic construction of an implicitly authenticated two-message AKE protocol AKE_{wFS} with weak forward security from any NCKE scheme, whose security is tightly based on the multi-user security of the underlying NCKE scheme. Furthermore, we give a generic construction of an explicitly authenticated two-message AKE protocol AKE_{FS} with perfect forward security by adding a tightly-secure signature scheme and a message authentication code (MAC) to our first construction, see Fig. 3. Thus, we require only a single signature which is particularly useful for tightly-secure key exchange, because known constructions of suitable tightly-secure signature schemes [1, 21] have relatively large signatures and replacing one signature with a MAC significantly improves the computational efficiency and communication complexity of the protocol.³

All these generic constructions leverage NCKE in order to resolve the technical difficulties in constructing tightly-secure AKE protocols described before.

³ [31] showed how to generically avoid signatures in forward-secure AKE protocols, but at the cost of additional messages.

Protocol	Comm. ($\mathbb{G}, \{0, 1\}^\kappa, \text{Sig}$)	Bytes	#Msg.	Assump.	Auth.	Model	State Reveal	Sec. Loss
Protocols with full forward security								
TLS* [16,15]	(2, 4, 2)	704	3	Strong-DH + DDH	expl.	SBG	no	$O(1)$
GJ [21]	(2, 1, 2)	608	3	DDH	expl.	MBG	no	$O(1)$
LLGW [32]	(3, 0, 2)	608	2	DDH	expl.	MBG	no	$O(1)$
AKE _{FS,DDH} (Fig. 1)	(5, 1, 1)	448	2	DDH	expl.	SBG	yes	$O(1)$
Protocols with weak forward security								
HMQV [26]	(2, 0, 0)	64	2	CDH	impl.	SBG	yes	$O(TN^2\ell^2)$
CCGJJ [10]	(2, 0, 0)	64	2	Strong-DH	impl.	SBG	no	$O(N)$
CCGJJ _{Twin} [10]	(3, 0, 0)	96	2	CDH	impl.	SBG	no	$O(N)$
AKE _{wFS,DDH} (Fig. 1)	(5, 0, 0)	160	2	DDH	impl.	SBG	yes	$O(1)$

Fig. 2. Comparison of AKE protocols over a group \mathbb{G} , where N refers to the number of parties, ℓ to the number of sessions per party and T is the number of test queries. TLS* refers to the TLS 1.3 handshake, instantiated with the tightly-secure signatures of [21]. The column **Comm.** counts the communication complexity of the protocols in terms of the number of group elements, hashes, and signatures. The column **Model** lists the AKE security model and distinguishes between multi-bit guessing (MBG) and the single-bit-guessing (SBG) security.

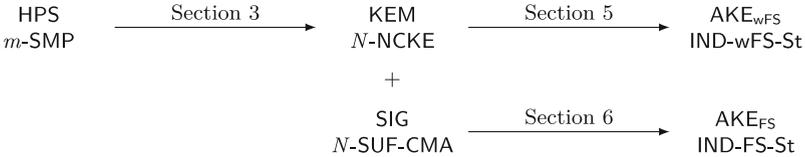


Fig. 3. Overview of our transformations, where N is the maximum number of users in the NCKE security game and in the SUF-CMA security game. The subset membership problem of HPS is m -fold for $m = N \cdot q$, where q is the maximum number of challenge queries in the NCKE security game.

HANDLING EPHEMERAL STATE REVEALS. Our protocols are secure against ephemeral state reveals. We construct the first tightly-secure protocols to achieve this. Note that this requires us to deal with the situation that the reduction must “know” valid ephemeral states for *all* sessions, even tested sessions. To this end, we encrypt the state information with a symmetric long-term key. An adversary now needs to query both long-term secret key and ephemeral state to reveal the secret state information, similarly to the approach used in the NAXOS protocol [28]. While the idea of achieving security against ephemeral state reveals by relying on the security of long-term keys was used before [5, 19, 28, 36], the approach to simply encrypt the state is new. It avoids the expensive re-computation of protocol messages required in prior generic approaches, which makes it particularly efficient. Also, previous work did not focus on tightness and it is unclear if a tight proof can be achieved in an even stronger security model which requires to reveal the randomness.

Our approach does not work generically, e.g., it cannot be applied to the protocols in [10, 21], so we have to design our protocols such that they are compatible. This is due to the fact that in both works, the state is a secret DH exponent which is implicitly determined by rerandomizing the CDH (or DDH) challenge and then is embedded in multiple sessions. Thus, the reduction is able to extract the solution independently of which session is the test session, but it also does not know any of the secret exponents, which the adversary could reveal for non-test sessions.

1.3 Related Work and Open Problems

Concurrent and independent work of Liu et al. [32] also proposed a tightly secure 2-message AKE with full forward security. Compared to our protocols, they do not consider state reveal attacks and their proofs only hold in the MBG security model. Their AKE construction LLGW follows the well known $1 \times \text{KEM} + 2 \times \text{SIG}$ approach, meaning that even neglecting the issues with the MBG security model, it is still considerably less efficient than ours (c.f. Fig. 2). The main novelty of [32] is the new KEM security notion of (multi-bit) “IND-mCPA with adaptive reveals” that gives them the handle to prove tight MBG security. It is a natural question whether this KEM security notion can be adapted to a single-bit notion such that the resulting AKE protocol achieves tight SBG (rather than MBG) security. This is in particular interesting since IND-mCPA KEMs with adaptive reveals can be instantiated in the standard model, whereas our NCKE notion seem to inherently rely on random oracles. More concretely this raises the question whether (variants of) [32] can also be proved in the SBG model, without relying on random oracles.

2 Preliminaries

For an integer n , $[n]$ denotes the set $\{1, \dots, n\}$. For a set S , $s \xleftarrow{\$} S$ denotes that s is sampled uniformly and independently at random from S . $y \leftarrow \mathcal{A}(x_1, x_2, \dots)$ denotes that on input x_1, x_2, \dots the probabilistic algorithm \mathcal{A} returns y . $\mathcal{A}^{\mathcal{O}}$ denotes that algorithm \mathcal{A} has access to oracle \mathcal{O} . We will use code-based games as introduced in [35]. An adversary is a probabilistic algorithm. $\Pr[G^{\mathcal{A}} \Rightarrow 1]$ denotes the probability that the final output $G^{\mathcal{A}}$ of game G running adversary \mathcal{A} is 1.

3 Multi-receiver Non-committing Key Encapsulation

In this section, we introduce Multi-Receiver Non-Committing Key Encapsulation (NCKE). We will use this concept to resolve the “commitment problem” described in the introduction, which often makes proofs for multi-party protocols with adaptive corruptions non-tight, as for example AKE protocols.

SYNTAX. A key encapsulation mechanism $\text{KEM} = (\text{Gen}, \text{Encaps}, \text{Decaps})$ consists of three algorithms. The key generation algorithm Gen outputs a key pair (pk, sk) ,

where \mathbf{pk} is the public key and \mathbf{sk} the secret key. The encapsulation algorithm inputs a public key \mathbf{pk} and outputs a ciphertext c and a key K from the key space \mathcal{K} , where c is called an encapsulation of K . The deterministic decapsulation algorithm inputs the secret key \mathbf{sk} and a ciphertext c and outputs K .

By μ we denote the *collision probability* of the key generation algorithm. In particular,

$$\Pr[(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen}, (\mathbf{pk}', \mathbf{sk}') \leftarrow \text{Gen} : \mathbf{pk} = \mathbf{pk}'] \leq 2^{-\mu} .$$

We denote the *min-entropy* of the encapsulation algorithm Encaps by $\gamma(\mathbf{pk}) := -\log \max_{c \in \mathcal{C}} \Pr[c = \text{Encaps}(\mathbf{pk})]$. We say KEM is γ -spread if for all $(\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Gen} : \gamma(\mathbf{pk}) \geq \gamma$. This implies that for all $c \in \mathcal{C}$:

$$\Pr[c = \text{Encaps}(\mathbf{pk})] \leq 2^{-\gamma} .$$

SECURITY. Following [33], we introduce a security definition of Multi-Receiver Non-Committing Key Encapsulation (NCKE) for a key encapsulation mechanism KEM in the random oracle model, i. e., the KEM algorithms have access to a random oracle $\mathsf{H} : \{0, 1\}^* \rightarrow \{0, 1\}^\kappa$, indicated by $\text{Encaps}^{\mathsf{H}}$. Our definition is relative to a simulator $\text{Sim} = (\text{SimGen}, \text{SimEncaps}, \text{SimHash})$. The simulated key generation algorithm SimGen generates a key pair $(\mathbf{pk}, \mathbf{sk})$. The simulated encapsulation algorithm SimEncaps takes both the public and private key and outputs a ciphertext c . The simulated hash algorithm SimHash inputs the key pair as well as three sets (used for bookkeeping) and deterministically computes a simulated hash value.

We define the two games $\text{NCKE}_{\text{real}}$ and NCKE_{sim} in Fig. 4 where we consider N receivers each holding a key pair $(\mathbf{pk}_n, \mathbf{sk}_n)$. In the $\text{NCKE}_{\text{real}}$ game, the original Encaps algorithm is used. We give each user an individual hash function H_n such that keys are computed independently. (In general, this can be implemented by using the user's public key and identity as input to the hash function as well, where collisions have to be considered.) In the NCKE_{sim} game, the SimEncaps algorithm is used to compute the ciphertexts. Keys are chosen uniformly at random. The adversary may also adaptively corrupt some receivers. We require that ciphertexts of corrupted receivers always decapsulate to the key output by ENCAPS , which is modeled by the SimHash algorithm. Therefore, if the receiver is corrupted, the algorithm takes sets \mathcal{CK} , \mathcal{D} and \mathcal{H} , where the first one stores all challenge ciphertexts and keys output to the adversary, the second one stores all decapsulation queries and the third one stores all hash queries which have been issued so far. Thus, the SimHash algorithm can answer future queries based on everything that is known to the adversary. If the receiver is not corrupted, set \mathcal{C} is used instead of \mathcal{CK} . This set stores only challenge ciphertexts and thus a hash value is computed independently of previous challenge keys.

The goal of an adversary \mathcal{A} is to distinguish between the real KEM algorithms used in game $\text{NCKE}_{\text{real}}$ and the simulated algorithms used in game NCKE_{sim} . This is captured in Definition 1. Note that the non-committing property is due to the SimHash algorithm. In particular, the SimHash algorithm ensures that a (uniformly random) challenge key can be explained by the corresponding ciphertext generated by SimEncaps as soon as the receiver is corrupted.

<p>NCKE_{real} and NCKE_{sim}</p> <pre> 00 for $n \in [N]$ 01 $(pk_n, sk_n) \leftarrow \text{Gen}$ 02 $(pk_n, sk_n) \leftarrow \text{SimGen}$ 03 $\text{opened}[n] := \text{false}$ 04 $\mathcal{CK}_n := \emptyset, \mathcal{C}_n := \emptyset, \mathcal{D}_n := \emptyset, \mathcal{H}_n := \emptyset$ 05 $b' \leftarrow \mathcal{A}^{\text{ENCAPS, DECAPS, OPEN, H}_1, \dots, \text{H}_N}(pk_1, \dots, pk_N)$ 06 return b' $\mathcal{H}_n(M)$ 07 if $\exists h$ s. t. $(M, h) \in \mathcal{H}_n$ return h 08 $h \xleftarrow{\\$} \{0, 1\}^k$ 09 if $\text{opened}[n]$ 10 $h \leftarrow \text{SimHash}(pk_n, sk_n, \mathcal{CK}_n, \mathcal{D}_n, \mathcal{H}_n, M)$ 11 else 12 $h \leftarrow \text{SimHash}(pk_n, sk_n, \mathcal{C}_n, \mathcal{D}_n, \mathcal{H}_n, M)$ 13 $\mathcal{H}_n := \mathcal{H}_n \cup \{(M, h)\}$ 14 return h </pre>	<pre> ENCAPS($n \in [N]$) 15 $(c, K) \leftarrow \text{Encaps}^{\text{H}_n}(pk_n)$ 16 $c \leftarrow \text{SimEncaps}(pk_n, sk_n)$ 17 $K \xleftarrow{\\$} \mathcal{K}$ 18 $\mathcal{CK}_n := \mathcal{CK}_n \cup \{(c, K)\}$ 19 $\mathcal{C}_n := \mathcal{C}_n \cup \{(c, \perp)\}$ 20 return (c, K) DECAPS($n \in [N], c$) 21 if $\exists K$ s. t. $(c, K) \in \mathcal{CK}_n$ 22 return \perp 23 $K := \text{Decaps}^{\text{H}_n}(sk_n, c)$ 24 $\mathcal{D}_n := \mathcal{D}_n \cup \{c\}$ 25 return K OPEN($n \in [N]$) 26 $\text{opened}[n] := \text{true}$ 27 return sk_n </pre>
---	---

Fig. 4. Real and simulated game for N -receiver non-committing key encapsulation in the random oracle model.

Definition 1 (N -Receiver Non-Committing Key Encapsulation). We define games $\text{NCKE}_{\text{real}}$ and NCKE_{sim} as in Fig. 4, where N is the number of users. The simulator $\text{Sim} = (\text{SimGen}, \text{SimEncaps}, \text{SimHash})$ is defined relative to KEM and is used in NCKE_{sim} . The advantage of an adversary \mathcal{A} against KEM and Sim is defined as

$$\text{Adv}_{\text{KEM}, \text{Sim}}^{N\text{-NCKE}}(\mathcal{A}) := \left| \Pr[\text{NCKE}_{\text{real}}^{\mathcal{A}} \Rightarrow 1] - \Pr[\text{NCKE}_{\text{sim}}^{\mathcal{A}} \Rightarrow 1] \right|.$$

When we write NCKE , we mean $\text{NCKE}\text{-CCA}$, where the adversary is allowed to access a decapsulation oracle. Sometimes we will explicitly write $\text{NCKE}\text{-CCA}$ to differentiate from $\text{NCKE}\text{-CPA}$, where the adversary cannot issue decapsulation queries.

We stress that compared to the standard definition of non-committing encryption in the random oracle model (e.g., [33]), Definition 1 is for KEMs (rather than encryption), only considers receiver corruptions (rather than sender and receiver corruptions), and considers multiple receivers (rather than one single receiver).

INSTANTIATIONS FROM HASH PROOF SYSTEMS. We recall the definition of hash proof systems by Cramer and Shoup [11] and properties defined in [25].

SMOOTH PROJECTIVE HASHING. Let \mathcal{Y} and \mathcal{Z} be sets and $\mathcal{X} \subset \mathcal{Y}$ a language. Let $\Lambda_{\text{sk}} : \mathcal{Y} \rightarrow \mathcal{Z}$ be a hash function indexed with $\text{sk} \in \mathcal{SK}$, where \mathcal{SK} is a set. A hash function Λ_{sk} is projective if there exists a projection $\mu : \mathcal{SK} \rightarrow \mathcal{PK}$ such that $\mu(\text{sk}) \in \mathcal{PK}$ defines the action of Λ_{sk} over \mathcal{X} . In particular, for every $c \in \mathcal{X}$, $Z = \Lambda_{\text{sk}}(c)$ is uniquely determined by $\mu(\text{sk})$ and c . However, there is no guarantee for $c \in \mathcal{Y} \setminus \mathcal{X}$ and it may not be possible to compute $\Lambda_{\text{sk}}(c)$ from $\mu(\text{sk})$ and C . A projective hash function is k -entropic if for all $c \in \mathcal{Y} \setminus \mathcal{X}$ it holds that $H_{\infty}(\Lambda_{\text{sk}}(c) \mid \text{pk}) \geq k$, where $\text{pk} = \mu(\text{sk})$ for $\text{sk} \xleftarrow{\$} \mathcal{SK}$.

Gen(par)	Encaps ^H (pk)	Decaps ^H (sk, c)
00 $sk \xleftarrow{\$} \mathcal{SK}$	03 $c \xleftarrow{\$} \mathcal{X}$ with witness r	06 $K := H(c, \text{Priv}(sk, c))$
01 $pk := \mu(sk)$	04 $K := H(c, \text{Pub}(pk, c, r))$	07 return K
02 return (pk, sk)	05 return (c, K)	

Fig. 5. Key encapsulation mechanism $\text{KEM} = (\text{Gen}, \text{Encaps}, \text{Decaps})$.

SimEncaps(pk, sk)	SimHash(pk, sk, $\mathcal{E}, \mathcal{D}, \mathcal{H}, M$)
00 $c \xleftarrow{\$} \mathcal{Y} \setminus \mathcal{X}$	02 $(c, Z) := M$
01 return c	03 if $\exists K$ s. t. $(c, K) \in \mathcal{E}$ and $\text{Priv}(sk, c) = Z$
	04 $h := K$
	05 else
	06 $h \xleftarrow{\$} \{0, 1\}^\kappa$
	07 return h

Fig. 6. Simulator $\text{Sim} = (\text{SimGen}, \text{SimEncaps}, \text{SimHash})$ for KEM, where $\text{SimGen} = \text{Gen}$. List \mathcal{E} is either \mathcal{CK} or \mathcal{C} .

HASH PROOF SYSTEM. A hash proof system $\text{HPS} = (\text{Par}, \text{Priv}, \text{Pub})$ consists of three algorithms. The randomized algorithm Par generates parametrized instances of $\text{par} = (\text{group}, \mathcal{Z}, \mathcal{Y}, \mathcal{X}, \mathcal{PK}, \mathcal{SK}, \Lambda_{(\cdot)} : \mathcal{Y} \rightarrow \mathcal{Z}, \mu : \mathcal{SK} \rightarrow \mathcal{PK})$, where group may contain additional structural parameters. The deterministic public evaluation algorithm Pub inputs the projection key $pk = \mu(sk)$, $c \in \mathcal{X}$ and a witness r of the fact that $c \in \mathcal{X}$ and returns $Z = \Lambda_{sk}(c)$. The deterministic private evaluation algorithm Priv takes $sk \in \mathcal{SK}$ and returns $\Lambda_{sk}(c)$ without knowing a witness. Furthermore, we assume that μ is efficiently computable and that there are efficient algorithms for sampling $c \in \mathcal{X}$ uniformly together with a witness r , sampling $c \in \mathcal{Y}$ uniformly and checking membership in \mathcal{Y} .

(m -FOLD) SUBSET MEMBERSHIP PROBLEM. We define the m -fold subset membership problem for HPS which requires to distinguish m ciphertexts uniformly drawn from \mathcal{X} from m ciphertexts uniformly drawn from $\mathcal{Y} \setminus \mathcal{X}$. The advantage of an adversary \mathcal{A} is defined as

$$\text{Adv}_{\text{HPS}}^{m\text{-SM}}(\mathcal{A}) := |\Pr[\mathcal{A}(\mathcal{Y}, \mathcal{X}, c_1, \dots, c_m) \Rightarrow 1] - \Pr[\mathcal{A}(\mathcal{Y}, \mathcal{X}, c'_1, \dots, c'_m) \Rightarrow 1]| \quad ,$$

where $c_1, \dots, c_m \xleftarrow{\$} \mathcal{X}$ and $c'_1, \dots, c'_m \xleftarrow{\$} \mathcal{Y} \setminus \mathcal{X}$.

N -RECEIVER NCKE FROM HPS. We use a k -entropic hash proof system $\text{HPS} = (\text{Par}, \text{Pub}, \text{Priv})$ with m -fold subset membership problem and a random oracle $H : \{0, 1\}^* \rightarrow \{0, 1\}^\kappa$ in order to construct a key encapsulation algorithm KEM and a simulator Sim as shown in Figs. 5 and 6. The encapsulation algorithm Encaps samples an element c from \mathcal{X} and a witness r . It runs the public evaluation algorithm and computes the key K as $H(c, \text{Pub}(pk, c, r))$. The decapsulation algorithm Decaps uses the result of the private evaluation algorithm Priv as input to H to compute K . Instead of sampling an element from \mathcal{X} , the SimEncaps algorithm samples an element c uniformly at random from $\mathcal{Y} \setminus \mathcal{X}$ and only returns c . The SimHash algorithm takes as input three sets $\mathcal{E}, \mathcal{D}, \mathcal{H}$, where $\mathcal{E} \in \{\mathcal{C}, \mathcal{CK}\}$, and the value $M = (c, Z)$ chosen by the adversary. If there exists a key K such that $(c, K) \in \mathcal{E}$ (note that for $\mathcal{E} = \mathcal{C}$ this will never be true) and the adversary's input to H satisfies $\text{Priv}(sk, c) = Z$, then the output value h is set to K .

Theorem 1 (*k*-entropic HPS with $(N \cdot q_E)$ -fold SMP \Rightarrow N -NCKE). For any N -NCKE adversary \mathcal{A} against KEM and Sim that issues at most q_E queries to ENCAPS, q_D queries to DECAPS and at most q_H queries to each random oracle H_n for $n \in [N]$, there exists an adversary \mathcal{B} against the $(N \cdot q_E)$ -fold subset membership problem of HPS such that

$$\text{Adv}_{\text{KEM, Sim}}^{N\text{-NCKE}}(\mathcal{A}) \leq \text{Adv}_{\text{HPS}}^{(N \cdot q_E)\text{-SM}}(\mathcal{B}) + \frac{N \cdot q_E \cdot q_H}{2^k} + \frac{N \cdot q_E \cdot q_D}{|\mathcal{Y} \setminus \mathcal{X}|},$$

where HPS is k -entropic, \mathcal{Y} is the set of all ciphertexts and \mathcal{X} is the set of valid ciphertexts.

We will give an instantiation based on the DDH assumption in Sect. 7.1. For the proof of Theorem 1 and an instantiation based on the higher residuosity assumption, we refer to the full version [23].

4 Security Model for Two-Message Authenticated Key Exchange

A two-message key exchange protocol $\text{AKE} = (\text{Gen}_{\text{AKE}}, \text{Init}_I, \text{Der}_R, \text{Der}_I)$ consists of four algorithms which are executed interactively by two parties as shown in Fig. 7. We denote the party which initiates the session by P_i and the party which responds to the session by P_r . The key generation algorithm Gen_{AKE} outputs a key pair (pk, sk) for one party. The initialization algorithm Init_I inputs the initiator’s long-term secret key sk_i and the responder’s long-term public key pk_r and outputs a message I and a state st . The responder’s derivation algorithm Der_R takes as input the responder’s long-term secret key sk_r , the initiator’s long-term public key pk_i and a message I . It computes a message R and a session key K . The initiator’s derivation algorithm Der_I inputs the initiator’s long-term secret key sk_i , the responder’s long-term public key pk_r , a message R and a state st . It outputs a session key K . Note that in contrast to the initiating party P_i , the responding party P_r will not be required to save any (secret) state information besides the session key K . The session key can be derived immediately after receiving the initiator’s message.

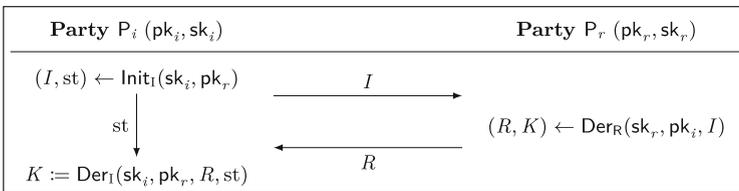


Fig. 7. Running a key exchange protocol between two parties.

Following [22], we define a game-based security model for authenticated key exchange using pseudocode. Our models for two different levels of security

GAMES IND-wFS-St_b and $\boxed{\text{IND-FS-St}_b}$	$\text{SESSION}_I((i, r) \in [N]^2)$
00 cnt := 0 //session counter	25 cnt ++
01 $\mathcal{S} := \emptyset$ //set of test sessions	26 sID := cnt
02 for $n \in [N]$	27 (init[sID], resp[sID]) := (i, r)
03 $(\text{pk}_n, \text{sk}_n) \leftarrow \text{Gen}_{\text{AKE}}$	28 type[sID] := "In"
04 $b' \leftarrow \mathcal{A}^O(\text{pk}_1, \dots, \text{pk}_N)$	29 (I, st) $\leftarrow \text{Init}_i(\text{sk}_i, \text{pk}_r)$
05 for sID* $\in \mathcal{S}$	30 ($I[\text{sID}], \text{state}[\text{sID}]) := (I, \text{st})$
06 if FRESH(sID*) = false	31 return (sID, I)
07 return 0 //session not fresh	
08 if VALID(sID*) = false	$\text{DER}_I(\text{sID}, R)$
09 return 0 //no valid attack	32 if state[sID] = \perp
10 return b'	33 return \perp //not initialized
	34 if sKey[sID] $\neq \perp$
	35 return \perp //no re-use
$\text{SESSION}_R((i, r) \in [N]^2, I)$	36 (i, r) := (init[sID], resp[sID])
11 cnt ++	37 st := state[sID]
12 sID := cnt	38 peerCorrupted[sID] := corrupted[r]
13 (init[sID], resp[sID]) := (i, r)	39 $K := \text{Der}_i(\text{sk}_i, \text{pk}_r, R, \text{st})$
14 type[sID] := "Re"	40 ($R[\text{sID}], \text{sKey}[\text{sID}]) := (R, K)$
15 peerCorrupted[sID] := corrupted[i]	41 return ε
16 (R, K) $\leftarrow \text{Derr}(\text{sk}_r, \text{pk}_i, I)$	
17 ($I[\text{sID}], R[\text{sID}], \text{sKey}[\text{sID}]) := (I, R, K)$	$\text{REVEAL}(\text{sID})$
18 return (sID, R)	42 revealed[sID] := true
	43 return sKey[sID]
$\text{TEST}(\text{sID})$	$\text{REV-STATE}(\text{sID})$
19 if sID $\in \mathcal{S}$ return \perp //already tested	44 if type[sID] \neq "In" return \perp
20 if sKey[sID] = \perp return \perp	45 revState[sID] := true
21 $\mathcal{S} := \mathcal{S} \cup \{\text{sID}\}$	46 return state[sID]
22 $K_0^* := \text{sKey}[\text{sID}]$	
23 $K_1^* \stackrel{\$}{\leftarrow} \mathcal{K}$	$\text{CORRUPT}(n \in [N])$
24 return K_b^*	47 corrupted[n] := true
	48 return sk_n

Fig. 8. Games IND-wFS-St_b and IND-FS-St_b for AKE, where $b \in \{0, 1\}$. \mathcal{A} has access to oracles $\text{O} := \{\text{SESSION}_I, \text{SESSION}_R, \text{DER}_I, \text{REVEAL}, \text{REV-STATE}, \text{CORRUPT}, \text{TEST}\}$. Helper procedures FRESH and VALID are defined in Fig. 9. If there exists any test session which is neither fresh nor valid, the game will return 0.

are given in Fig. 8. We consider N parties $\text{P}_1, \dots, \text{P}_N$ with long-term key pairs $(\text{pk}_n, \text{sk}_n)$, $n \in [N]$. Each session between two parties has a unique identification number sID and variables which are defined relative to sID:

- $\text{init}[\text{sID}] \in [N]$ denotes the initiator of the session.
- $\text{resp}[\text{sID}] \in [N]$ denotes the responder of the session.
- $\text{type}[\text{sID}] \in \{\text{"In"}, \text{"Re"}\}$ denotes the session's view, i. e. whether the initiator or the responder computes the session key.
- $I[\text{sID}]$ denotes the message that was computed by the initiator.
- $R[\text{sID}]$ denotes the message that was computed by the responder.
- $\text{state}[\text{sID}]$ denotes the state information that is stored by the initiator.
- $\text{sKey}[\text{sID}]$ denotes the session key.

To establish a session between two parties, the adversary is given access to oracles SESSION_I and SESSION_R , where the first one starts a session of type "In"

and received are the same ($I[\text{sID}] = I[\text{sID}']$ and $R[\text{sID}] = R[\text{sID}']$) and they are of different types ($\text{type}[\text{sID}] \neq \text{type}[\text{sID}']$).

- **Partially Matching Session:** A session sID' of type “In” is *partially matching* to session sID of type “Re” if the initial messages are the same ($I[\text{sID}] = I[\text{sID}']$).

Finally, the adversary is given access to oracle TEST which will return either the session key of the specified session or a uniformly random key. In our security models, we allow multiple test queries. We store test sessions in a set \mathcal{S} . In general, the adversary can disclose the complete interaction between two parties by querying the long-term secret keys, the state information and the session key. However, for each test session, we require that the adversary does not issue queries such that the session key can be trivially computed. We define the properties of freshness and validity which all test sessions have to satisfy:

- **Freshness:** A (test) session is called *fresh* if the session key was not revealed. Furthermore, if there exists a matching session, we require that this session’s key is not revealed and that this session is not also a test session.
- **Validity:** A (test) session is called *valid* if it is fresh and the adversary performed any attack which is defined in the security model. We capture this with attack tables (cf. Tables 1 and 2). A description of how to read the tables is given below.

Attack Tables. All attacks are defined using variables to indicate which queries the adversary may (not) make. We consider three dimensions covering all possible combinations of reveal queries the adversary can make:

- whether the test session is on the initiator’s ($\text{type}[\text{sID}^*] = \text{“In”}$) or the responder’s side ($\text{type}[\text{sID}^*] = \text{“Re”}$),
- all combinations of long-term secret key and state reveals (corrupted and revState variables), also taking into account when a corruption happened (peerCorrupted),
- whether the adversary acted passively (matching session), partially active (partially matching session) or actively (no matching session).

This yields a full table of 24 attacks, in particular capturing *key compromise impersonation* (KCI) and *maximal exposure* (MEX) attacks. An attack was performed if the variables are set to the corresponding values in the table. However, when considering two-message protocols, where the responder’s side does not have a state, and we only consider *weak forward security*, some of the attacks are redundant. Thus, we obtain *distilled* tables. We exclude trivial attacks, e.g., the generic attack on two-message AKE protocols with state-reveals described in [30]. Therefore, the adversary is not allowed to obtain the state of a partially matching session. Also note that by definition, a partially matching session for a two-message protocol can only be of type “Re”. Table 1 is the distilled table used for the IND-wFS-St security game and Table 2 is used for the IND-FS-St security game. Note that the numbering of attacks in the distilled tables is inherited from the full table given in the full version [23].

Table 1. Distilled table of attacks for wFS adversaries against two-message protocols. This table is obtained from the full table of attacks by using that responders do not have a state and that we are considering weak forward security. The numbering of attacks is inherited from the full table. An attack is regarded as an AND conjunction of variables with specified values as shown in the each line, where “-” means that this variable can take arbitrary value. **F** means “false” and “n/a” indicates that there is no state which can be revealed as no (partially) matching session exists.

\mathcal{A} gets (Initiator, Responder)	corrupted[i^*]	corrupted[r^*]	type[sID*]	revState[sID*]	$\exists \text{sID} \in \mathfrak{M}(\text{sID}^*) :$ revState[sID]	$ \mathfrak{M}(\text{sID}^*) $	$\exists \text{sID} \in \mathfrak{P}(\text{sID}^*) :$ revState[sID]	$ \mathfrak{P}(\text{sID}^*) $
(0) multiple partially matching sessions	-	-	-	-	-	-	-	> 1
(1 \vee 2) (long-term, long-term)	-	-	-	F	F	1	-	-
(7 \vee 8) (state, long-term)	F	-	-	-	-	1	-	-
(10) (long-term, long-term)	-	-	“Re”	F	n/a	0	F	1
(16) (state, long-term)	F	-	“Re”	F	n/a	0	-	1
(19) (state, state)	F	F	“In”	-	n/a	0	n/a	0
(21) (long-term, state)	-	F	“In”	F	n/a	0	n/a	0
(24) (state, long-term)	F	-	“Re”	F	n/a	0	n/a	0

Table 2. Distilled table of attacks for full FS adversaries against two-message protocols. This table is obtained from the full table of attacks by removing redundant rows and using that responders do not have a state. The numbering of attacks is inherited from the full table. An attack is regarded as an AND conjunction of variables with specified values as shown in the each line, where “-” means that this variable can take arbitrary value. **F** means “false” and “n/a” indicates that there is no state which can be revealed as no (partially) matching session exists.

\mathcal{A} gets (Initiator, Responder)	corrupted[i^*]	corrupted[r^*]	peerCorrupted[sID*]	type[sID*]	revState[sID*]	$\exists \text{sID} \in \mathfrak{M}(\text{sID}^*) :$ revState[sID]	$ \mathfrak{M}(\text{sID}^*) $	$\exists \text{sID} \in \mathfrak{P}(\text{sID}^*) :$ revState[sID]	$ \mathfrak{P}(\text{sID}^*) $
(0) multiple partially matching sessions	-	-	-	-	-	-	-	-	> 1
(1 \vee 2) (long-term, long-term)	-	-	-	-	F	F	1	-	-
(7 \vee 8) (state, long-term)	F	-	-	-	-	-	1	-	-
(10) (long-term, long-term)	-	-	F	“Re”	F	n/a	0	F	1
(16) (state, long-term)	F	-	-	“Re”	F	n/a	0	-	1
(17) (long-term, long-term)	-	-	F	“In”	F	n/a	0	n/a	0
(18) (long-term, long-term)	-	-	F	“Re”	F	n/a	0	n/a	0
(23) (state, long-term)	F	-	F	“In”	-	n/a	0	n/a	0

However, if the protocol does not use appropriate randomness, it should not be considered secure in our model. Thus, if the adversary is able to create more than one (partially) matching session to a test session, it may also run a trivial attack. We model this in row (0) of Tables 1 and 2.

Example. If the test session is an initiating session ($\text{type[sID}^*] = \text{“In”}$), the state was not revealed ($\text{revState[sID}^*] = \text{false}$) and there is a matching session ($|\mathfrak{M}(\text{sID}^*)| = 1$), then row (1√2) will evaluate to true. In this scenario, the adversary is allowed to query both long-term secret keys.

For all test sessions, at least one attack has to evaluate to true. Then, the adversary wins if it distinguishes the session keys from uniformly random keys which it obtains through queries to the TEST oracle.

Definition 2 (Key Indistinguishability of AKE). *We define games IND-wFS-St_b and IND-FS-St_b for $b \in \{0, 1\}$ as in Figs. 8 and 9. The advantage of an adversary \mathcal{A} against AKE in these games is defined as*

$$\text{Adv}_{\text{AKE}}^{\text{IND-wFS-St}}(\mathcal{A}) := \left| \Pr[\text{IND-wFS-St}_1^{\mathcal{A}} \Rightarrow 1] - \Pr[\text{IND-wFS-St}_0^{\mathcal{A}} \Rightarrow 1] \right| \quad \text{and}$$

$$\text{Adv}_{\text{AKE}}^{\text{IND-FS-St}}(\mathcal{A}) := \left| \Pr[\text{IND-FS-St}_1^{\mathcal{A}} \Rightarrow 1] - \Pr[\text{IND-FS-St}_0^{\mathcal{A}} \Rightarrow 1] \right| .$$

When proving the security of a protocol, the success probability for each attack strategy listed in the corresponding table will have to be analyzed, thus showing that independently of which queries the adversary makes, it cannot distinguish the session key from a uniformly random key.

4.1 Relation to Other Definitions

In this section, we will refer to the most widely used security definitions for authenticated key exchange protocols. In the first place, these include the CK model [9] and the stronger definition used for the HMQV protocol (CK+) in [26], the eCK model [28] and the strengthened version of [14], the definitions given in [24] and [1] which are both extensions of the BR model [4], and the definition of IND-Å security in [22]. In [12, 13], Cremers showed that the CK, CK+ und eCK model are incomparable. Thus, we will not do a formal comparison of security models, but only point out similarities and differences between our definition and the definitions listed above.

PARTY CORRUPTION. We allow the adversary to corrupt a party which means that it will obtain that party’s long-term secret key as in the eCK model and the models given in [1, 22, 24]. In contrast, a corrupt query in the CK and CK+ model will reveal all information in the memory of that party, i.e. long-term secrets and session-specific information.

STATE-REVEALS. Our model only allows state-reveal queries on initiating sessions because the initiator has to wait for the response to compute the session key. Thus, the state contains all that information that is needed to derive the session key

as soon as the responder’s message is received. The responder can directly compute the session key and does not have to store other information. The eCK model explicitly defines the state as the randomness that is used in the protocol. In the CK model, it is not clear which information is included in the state, but it is left to be specified by the AKE protocol itself. Other models such as [24], its extension given in [1] and the one used in [10] do not allow state-reveals at all. Here, we want to emphasize that in particular all previous work on tight AKE does not consider state reveals and we are the first ones to address this problem.

(WEAK) FORWARD SECURITY. Following Krawczyk [26], we specify two levels of forward security. IND-wFS-St models weak forward security, whereas IND-FS-St models full forward security. The first one is intended for 2-message protocols with implicit authentication, as those cannot achieve full forward security [26]. The second one is intended for protocols with explicit authentication. With those definitions, we capture the same properties as the most common security models given in [1, 9, 24, 26, 28], where some of them only define either weak or full forward security depending on whether they consider implicitly or explicitly authenticated protocols.

MATCHING SESSIONS AND PARTNERING. As most security models, ours use the concept of matching sessions to define a relation between two sessions. Following Cremer and Feltz [14], we additionally use the term of origin (or partially matching) sessions, which refers to a relaxation of the definition of matching sessions. The concept of origin sessions is used for full forward security, in particular we need this to handle the no-match attack described by Li and Schäge [30], where two sessions compute the same session key but do not have matching conversations. Recent works such as [10, 21] take up the approach of origin sessions and oracle partnering based on session keys as additional requirement.

ON REGISTERING CORRUPT KEYS. Some security models for AKE allow the adversary also to *register* adversarially-generated keys, this holds in particular for previous works considering tightly-secure key exchange [1, 10, 21]. Technically this makes the security model strictly stronger, as one can easily construct contrived protocols that are insecure with adversarially-registered keys, but secure without.

However, in the actual security proofs in [1, 10, 21], adversarially-registered keys are treated no differently than corrupted keys. We chose to keep model, security proofs and notation as simple as possible (it is already complex enough, anyway), and thus omitted this query. However, it is straightforward to extend our model with it, and the proofs need not to be changed. Whenever the adversary registers a new key, it would immediately be marked as “corrupted” (just like in [1, 10, 21]). Apart from that, no additional changes to the proofs are required, since the proofs deal with all corrupted keys in the same way, regardless of their distribution or whether they are generated by the experiment or an external entity. We also do not require a proof of knowledge of the corresponding secret key for the registration, or a proof that the registered public key is valid in any sense.

5 AKE with Weak Forward Security

In this section, we show how to build an implicitly authenticated AKE protocol using the concept of non-committing key encapsulation.

In particular, from two key encapsulation mechanisms $\text{KEM}_{\text{CPA}} = (\text{Gen}_{\text{CPA}}, \text{Encaps}_{\text{CPA}}, \text{Decaps}_{\text{CPA}})$ and $\text{KEM}_{\text{CCA}} = (\text{Gen}_{\text{CCA}}, \text{Encaps}_{\text{CCA}}, \text{Decaps}_{\text{CCA}})$, we construct a two-message authenticated key exchange protocol $\text{AKE}_{\text{wFS}} = (\text{Gen}_{\text{AKE}}, \text{Init}_I, \text{Der}_R, \text{Der}_I)$ as shown in Figs. 10 and 11. W.l.o.g. $\text{KEM}_{\text{CPA}}, \text{KEM}_{\text{CCA}}, \text{AKE}_{\text{wFS}}$ have identical key space \mathcal{K} . Each party holds a long-term key pair (pk, sk) for KEM_{CCA} and a symmetric key k to encrypt the secret state information which has to be stored by the initiating party. State encryption protects against state attacks and is implemented using a symmetric encryption scheme defined as $\text{E}_k(st') := (IV, \text{G}(k, IV) \oplus st')$ for a random nonce IV . Here $\text{G} : \{0, 1\}^* \rightarrow \{0, 1\}^d$ is a random oracle and d is an integer denoting the maximum bit length of the unencrypted state st' . The protocol uses an additional cryptographic hash function $\text{H} : \{0, 1\}^* \rightarrow \mathcal{K}$ to output the session key.

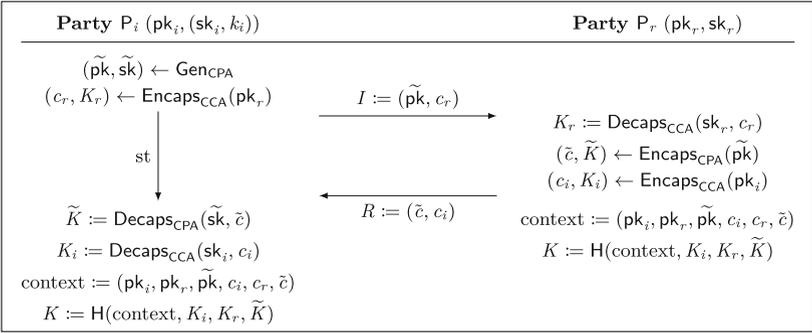


Fig. 10. Visualization: Running protocol AKE_{wFS} between two parties.

The initiating party generates an ephemeral key pair for KEM_{CPA} , then runs the $\text{Encaps}_{\text{CCA}}$ algorithm on the peer's public key to output a ciphertext c_r and a key K_r and sends the ephemeral public key and c_r to the intended receiver. All values are stored temporarily and encrypted as described above, as they will later be needed to compute the session key. The responding party uses its secret key sk_r to compute key K_r from c_r . Next, it runs the $\text{Encaps}_{\text{CPA}}$ algorithm on the received ephemeral public key to compute a ciphertext \tilde{c} and a key \tilde{K} and then the $\text{Encaps}_{\text{CCA}}$ algorithm on the initiator's public key to output c_i and K_i . It sends both ciphertexts to the initiating party and computes the session key evaluating the hash function H on all public context and computes the three shared keys K_r, K_i and \tilde{K} . The initiator retrieves the secret state information and computes K_i and \tilde{K} from c_i and \tilde{c} . Now, it can also establish the session key.

Gen_{AKE}	Der_R ((sk _r , k _r), pk _i , (pk̃, c _r))
00 (pk, sk) ← Gen _{CCA}	09 K _r := Decaps _{CCA} (sk _r , c _r)
01 k $\xleftarrow{\$}$ {0, 1} ^κ	10 (c̃, K̃) ← Encaps _{CPA} (pk)
02 return (pk', sk') := (pk, (sk, k))	11 (c _i , K _i) ← Encaps _{CCA} (pk _i)
Init_I ((sk _i , k _i), pk _r)	12 context := (pk _i , pk _r , pk, c _i , c _r , c̃)
03 (pk, sk) ← Gen _{CPA}	13 K := H(context, K _i , K _r , K̃)
04 (c _r , K _r) ← Encaps _{CCA} (pk _r)	14 return ((c̃, c _i), K)
05 IV $\xleftarrow{\$}$ {0, 1} ^κ	Der_I ((sk _i , k _i), pk _r , (c̃, c _i), st)
06 st' := (pk, sk, c _r , K _r)	15 (IV, ψ) := st
07 st := (IV, G(k _i , IV) ⊕ st')	16 (pk̃, sk, c _r , K _r) := G(k _i , IV) ⊕ ψ
08 return ((pk̃, c _r), st)	17 K̃ := Decaps _{CPA} (sk̃, c̃)
	18 K _i := Decaps _{CCA} (sk _i , c _i)
	19 context := (pk _i , pk _r , pk, c _i , c _r , c̃)
	20 K := H(context, K _i , K _r , K̃)
	21 return K

Fig. 11. Authenticated key exchange protocol AKE_{wFS} from KEM_{CPA} and KEM_{CCA} . Lines written in purple color are only used to encrypt the state.

Theorem 2 (KEM_{CPA} NCKE-CPA + KEM_{CCA} NCKE-CCA $\stackrel{\text{ROM}}{\Rightarrow}$ AKE_{wFS} IND-wFS-St). *For any IND-wFS-St adversary \mathcal{A} against AKE_{wFS} with N parties that establishes at most S sessions and issues at most T queries to the test oracle TEST, q_G queries to random oracle G and at most q_H queries to random oracle H , there exists an N -NCKE-CCA adversary \mathcal{B} against KEM_{CCA} and Sim_{CCA} and an S -NCKE-CPA adversary \mathcal{C} against KEM_{CPA} and Sim_{CPA} such that*

$$\begin{aligned} \text{Adv}_{\text{AKE}_{\text{wFS}}}^{\text{IND-wFS-St}}(\mathcal{A}) &\leq 2 \cdot \left(\text{Adv}_{\text{KEM}_{\text{CCA}}, \text{Sim}_{\text{CCA}}}^{N\text{-NCKE-CCA}}(\mathcal{B}) + \text{Adv}_{\text{KEM}_{\text{CPA}}, \text{Sim}_{\text{CPA}}}^{S\text{-NCKE-CPA}}(\mathcal{C}) \right) + T \cdot \left(\frac{q_G}{2^\kappa} + \frac{q_H}{|\mathcal{K}|} \right) \\ &\quad + N^2 \cdot \left(\frac{1}{2^{\mu_{\text{CCA}}}} + \frac{1}{2^\kappa} \right) + S^2 \cdot \left(\frac{1}{2^{\mu_{\text{CPA}}}} + \frac{1}{2^{\gamma_{\text{CCA}}}} + \frac{1}{2^{\gamma_{\text{CPA}}}} + \frac{1}{2^\kappa} \right) + 2S \cdot \frac{q_G}{2^{2\kappa}}, \end{aligned}$$

where Sim_{CCA} and Sim_{CPA} are the simulators from the NCKE experiments, μ_{CCA} and μ_{CPA} are the collision probability of the key generation algorithms Gen_{CCA} and Gen_{CPA} , γ_{CCA} and γ_{CPA} are the spreadness parameters of the encapsulation algorithms $\text{Encaps}_{\text{CCA}}$ and $\text{Encaps}_{\text{CPA}}$ and κ is a security parameter. The running times of \mathcal{B} and \mathcal{C} consist essentially of the time required to execute the security experiment with the adversary once, plus a minor number of additional operations (including bookkeeping, lookups etc.).

Proof (Sketch). Let \mathcal{A} be an adversary against IND-wFS-St security of AKE_{wFS} . For $b \in \{0, 1\}$, game $G_{0,b}$ is the IND-wFS-St_b game, where we additionally exclude that collisions between long-term key pairs, ephemeral key pairs, ciphertexts and nonces occur.

In game $G_{1,b}$, we replace the computations for KEM_{CCA} by the simulator Sim_{CCA} , which allows to draw keys K_i and K_r uniformly at random. This change affects all sessions which makes the proof tight. If the adversary reveals a long-term key pair of any user, the property of receiver non-committing key encapsulation ensures that the correct keys K_i and K_r can be computed by the adversary.

Next, we want to replace the computations for KEM_{CPA} by the simulator Sim_{CPA} , which allows to draw keys \tilde{K} uniformly at random. However, the ephemeral secret key sk is part of the state and will not be available to the NCKE-CPA reduction in the first place. Thus, we introduce an intermediate game $G_{2,b}$ and do not compute the state when the session is initiated but only when the adversary queries the REV-STATE oracle. In game $G_{3,b}$, we can then use the simulator for KEM_{CPA} and draw keys \tilde{K} uniformly at random, whenever the ephemeral public key pk comes from the experiment (i.e. the adversary creates a partially matching session). Again, the non-committing property of KEM_{CPA} ensures consistency in case the adversary reveals both the state of a session and the long-term key of the initiator, which reveals the ephemeral secret key sk .

Depending on whether there exists a (partially) matching session and which queries to REV-STATE and CORRUPT the adversary makes, we can argue that at least one key K_i , K_r or \tilde{K} in each test session is chosen uniformly at random and unknown to \mathcal{A} and thus it cannot distinguish the session key from a uniformly random key in the last game $G_{4,b}$. \square

The full proof of Theorem 2 can be found in the full version [23]. Note that the non-committing property is essential to embed random KEM keys in each session and thus to achieve tightness. This way, we only need to make a case distinction at the end and can argue that for all test sessions at least one KEM key is independent of the adversary's view no matter which queries it has made (provided it did not make a trivial attack). Relying on a weaker assumption requires to make a case distinction earlier in the proof and may involve guessing as in some cases it is not clear which KEM key will be revealed (through corruption and/or reveal or state reveal) at a later point in time.

6 AKE with Full Forward Security

We show how to build an explicitly authenticated AKE protocol using the concept of non-committing key encapsulation. As we also need a signature scheme, we will first give the relevant definitions.

6.1 Digital Signatures

A digital signature scheme $\text{SIG} = (\text{Gen}_{\text{SIG}}, \text{Sign}, \text{Vrfy})$ consists of three algorithms. The key generation algorithm Gen_{SIG} outputs a key pair (vk, sigk) , where vk is the verification key and sigk the signing key. The signing algorithm Sign inputs a signing key sigk and a message m and outputs a signature σ . The deterministic verification algorithm Vrfy inputs the verification key vk , a message m and a signature σ and outputs 1 if σ is a valid signature for m , otherwise it outputs 0.

In Fig. 12, we define the security game N user Strong UnForgeability under Chosen Message Attacks with corruptions (N -SUF-CMA). The definition is similar to the one given in [1], except that we require *strong* unforgeability, i. e. the

adversary may also find a new signature for a message it queried to the SIGN oracle before. The advantage of an adversary \mathcal{A} is defined as

$$\text{Adv}_{\text{SIG}}^{N\text{-SUF-CMA}}(\mathcal{A}) := \Pr[N\text{-SUF-CMA}^{\mathcal{A}} \Rightarrow 1] .$$

GAME $N\text{-SUF-CMA}$	$\text{SIGN}(n \in [N], m)$
00 $\mathcal{S}^{\text{corr}} := \emptyset$	09 $\sigma \leftarrow \text{Sign}(\text{sig}_n^k, m)$
01 for $n \in [N]$	10 $\mathcal{S}_n := \mathcal{S}_n \cup \{(m, \sigma)\}$
02 $(\text{vk}_n, \text{sig}_n) \leftarrow \text{Gen}_{\text{SIG}}$	11 return σ
03 $\mathcal{S}_n := \emptyset$	
04 $(n^*, m^*, \sigma^*) \leftarrow \mathcal{A}^{\text{SIGN, CORRUPT}}(\text{vk}_1, \dots, \text{vk}_N)$	$\text{CORRUPT}(n \in [N])$
05 if $\text{Vrfy}(\text{vk}_{n^*}, m^*, \sigma^*) = 1$ and $n^* \notin \mathcal{S}^{\text{corr}}$	12 $\mathcal{S}^{\text{corr}} := \mathcal{S}^{\text{corr}} \cup \{n\}$
and $(m^*, \sigma^*) \notin \mathcal{S}_{n^*}$	13 return $\text{sig}_{k_{n^*}}$
06 return 1	
07 else	
08 return 0	

Fig. 12. Game $N\text{-SUF-CMA}$ for SIG.

6.2 Transformation Using NCKE and a Signature Scheme

From two key encapsulation mechanisms $\text{KEM}_{\text{CPA}} = (\text{Gen}_{\text{CPA}}, \text{Encaps}_{\text{CPA}}, \text{Decaps}_{\text{CPA}})$ and $\text{KEM}_{\text{CCA}} = (\text{Gen}_{\text{CCA}}, \text{Encaps}_{\text{CCA}}, \text{Decaps}_{\text{CCA}})$ with key space \mathcal{K} and a digital signature scheme $\text{SIG} = (\text{Gen}_{\text{SIG}}, \text{Sign}, \text{Vrfy})$, we construct a two-message authenticated key exchange protocol $\text{AKE}_{\text{FS}} = (\text{Gen}_{\text{AKE}}, \text{Init}_{\text{I}}, \text{Der}_{\text{R}}, \text{Der}_{\text{I}})$ with key space \mathcal{K} as shown in Figs. 13 and 14. Each party has a key pair $(\text{vk}, \text{sig}_k)$ for SIG, a key pair (pk, sk) for KEM_{CCA} and a symmetric key k to encrypt the secret state information which has to be stored by the initiating party (cf. Sect. 5). The protocol uses additional cryptographic hash functions $\text{F} : \{0, 1\}^* \rightarrow \{0, 1\}^k$ to compute value π and $\text{H} : \{0, 1\}^* \rightarrow \mathcal{K}$ to output the session key.

The initiating party computes an ephemeral key pair for KEM_{CPA} , runs the $\text{Encaps}_{\text{CCA}}$ algorithm on the intended receiver's public key pk_r to obtain a ciphertext c_r and a key K_r and signs both the ephemeral public key and c_r , which are sent to the receiver along with the signature. The receiver verifies the signature and then runs the $\text{Encaps}_{\text{CPA}}$ algorithm on the ephemeral public key to output a ciphertext \tilde{c} and a key \tilde{K} . It computes K_r using its secret key sk_r . It then tags the received message together with \tilde{c} and K_r by evaluating hash function F and sends the output together with \tilde{c} to the initiator. The initiator retrieves K_r from the secret state and also evaluates F . If the output is the same, it computes \tilde{K} using the ephemeral secret key. The session key is computed evaluating hash function H on all public context and key \tilde{K} . We establish the following theorem and give a proof sketch. The full proof can be found in the full version [23].

Theorem 3 (KEM_{CPA} NCKE-CPA + KEM_{CCA} NCKE-CCA +
 SIG $N\text{-SUF-CMA} \xrightarrow{\text{ROM}}$ AKE_{FS} IND-FS-St). *For any IND-FS-St adversary \mathcal{A}*

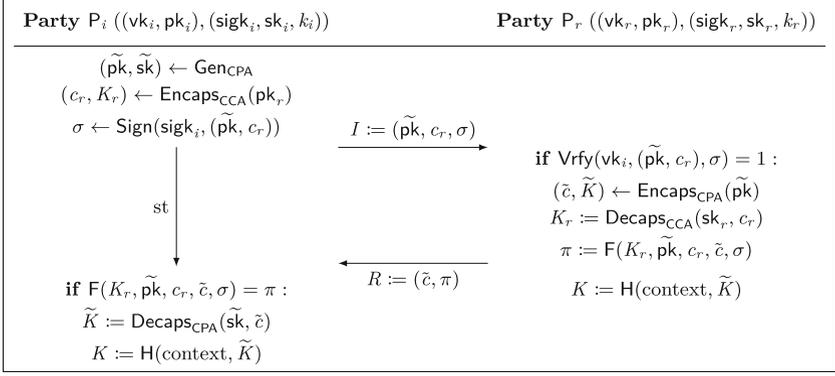


Fig. 13. Visualization: Running AKE_{FS} between two parties, where K is the resulting session key and $\text{context} := (vk_i, pk_i, vk_r, pk_r, \tilde{pk}, c_r, \tilde{c}, \sigma, \pi)$

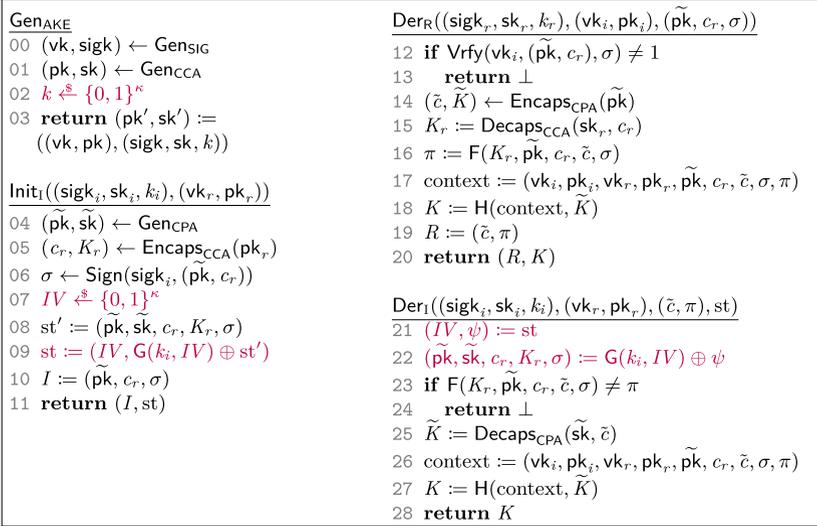


Fig. 14. Authenticated key exchange protocol AKE_{FS} from KEM_{CPA} , KEM_{CCA} and SIG . Lines written in purple color are only used to encrypt the state.

against AKE_{FS} with N parties that establishes at most S sessions and issues at most T queries to test oracle TEST , at most q_{H} , q_{G} and q_{F} queries to random oracles H , G and F , there exists an N -SUF-CMA adversary \mathcal{B} against SIG , an S -NCKE-CPA adversary \mathcal{C} against KEM_{CPA} and Sim_{CPA} and an N -NCKE-CCA adversary \mathcal{D} against KEM_{CCA} and Sim_{CCA} such that

$$\begin{aligned}
 \text{Adv}_{\text{AKE}_{\text{FS}}}^{\text{IND-FS-St}}(\mathcal{A}) &\leq 2 \cdot \left(\text{Adv}_{\text{SIG}}^{\text{N-SUF-CMA}}(\mathcal{B}) + \text{Adv}_{\text{KEM}_{\text{CPA}}, \text{Sim}_{\text{CPA}}}^{\text{S-NCKE-CPA}}(\mathcal{C}) + \text{Adv}_{\text{KEM}_{\text{CCA}}, \text{Sim}_{\text{CCA}}}^{\text{N-NCKE-CCA}}(\mathcal{D}) \right) \\
 &+ T \cdot \left(\frac{q_{\text{G}}}{2^{\kappa}} + \frac{q_{\text{H}}}{|\mathcal{K}|} \right) + N^2 \cdot \left(\frac{1}{2^{\mu_{\text{SIG}}}} + \frac{1}{2^{\mu_{\text{CCA}}}} + \frac{1}{2^{\kappa}} \right) \\
 &+ S^2 \cdot \left(\frac{1}{2^{\mu_{\text{CPA}}}} + \frac{1}{2^{\gamma_{\text{CCA}}}} + \frac{1}{2^{\gamma_{\text{CPA}}}} + \frac{1}{2^{\kappa}} \right) + 2S \cdot \frac{q_{\text{G}}}{2^{2\kappa}},
 \end{aligned}$$

where Sim_{CPA} and Sim_{CCA} are the simulators from the NCKE-CPA and NCKE-CCA experiment, $\mu_{\text{SIG}}, \mu_{\text{CPA}}, \mu_{\text{CCA}}$ are collision probabilities of the key generation algorithms $\text{Gen}_{\text{SIG}}, \text{Gen}_{\text{CPA}}$ and Gen_{CCA} and $\gamma_{\text{CPA}}, \gamma_{\text{CCA}}$ are the spreadness parameters of the encapsulation algorithms. The running times of \mathcal{B}, \mathcal{C} and \mathcal{D} consist essentially of the time required to execute the security experiment with the adversary once, plus a minor number of additional operations (including bookkeeping, lookups etc.).

Proof (Sketch). Let \mathcal{A} be an adversary against IND-FS-St security of AKE_{FS} . For $b \in \{0, 1\}$, $G_{0,b}$ is the IND-FS-St _{b} game, where we exclude that collisions between long-term key pairs, ephemeral key pairs, ciphertexts and nonces occur.

In game $G_{1,b}$, we abort when \mathcal{A} computes a valid signature for an uncorrupted user that was not output by the experiment, reducing to N -SUF-CMA security of the signature scheme.

In game $G_{2,b}$, we replace the computations for KEM_{CCA} by the simulator Sim_{CCA} in all sessions using the non-committing property of KEM_{CCA} , which allows to draw key K_r which serves as key for the MAC uniformly at random. Thus, the adversary cannot compute a valid MAC for an uncorrupted user.

In game $G_{3,b}$ (as in the proof of Theorem 2), we do not compute the state when the session is initiated but only when the adversary queries the REV-STATE oracle. After that, we can switch KEM_{CPA} to the corresponding simulator Sim_{CPA} in game $G_{4,b}$ and draw keys \tilde{K} uniformly at random, whenever the ephemeral public key pk comes from the experiment (i.e. the adversary creates a partially matching session). As the adversary can only complete a (partially) matching sessions (otherwise it would have forged a signature or MAC), we can argue that \tilde{K} in each test session is chosen uniformly at random and unknown to \mathcal{A} and thus he cannot distinguish the session key from a uniformly random key in the last game $G_{5,b}$. \square

7 Concrete Instantiation of AKE Protocols

7.1 NCKE from the DDH Assumption

Let us first describe the hash proof system we will use. Therefore, let GGen be a group generation algorithm which takes the security parameter 1^κ as input and returns (\mathbb{G}, p, g_1) , where g_1 is a generator of the cyclic group \mathbb{G} with prime order p . Define $\text{group} = (\mathbb{G}, p, g_1, g_2)$, where $g_2 = g_1^w$ for $w \xleftarrow{\$} \mathbb{Z}_p$. Define $\mathcal{Y} = \mathbb{Z}_p^2$ and $\mathcal{X} = \{(g_1^r, g_2^r) : r \in \mathbb{Z}_p\}$. A value r is a witness that $(c_1, c_2) \in \mathcal{X}$. Define $\text{SK} = \mathbb{Z}_p^2$, $\text{PK} = \mathbb{Z}_p$ and $\mathcal{Z} = \mathbb{Z}_p$. For $\text{sk} = (x_1, x_2) \in \mathbb{Z}_p^2$, define $\mu(\text{sk}) = X = g_1^{x_1} g_2^{x_2}$. This defines the output of the parameter generation algorithm Par .

For $(c_1, c_2) \in \mathcal{Y}$ define $\Lambda_{\text{sk}}(c_1, c_2) := Z = (c_1^{x_1} c_2^{x_2})$. This defines the private evaluation algorithm $\text{Priv}(\text{sk}, (c_1, c_2))$. Given $\text{pk} = \mu(\text{sk}) = X$, $(c_1, c_2) \in \mathcal{X}$ and a witness $r \in \mathbb{Z}_p$ such that $(c_1, c_2) = (g_1^r, g_2^r)$, the public evaluation algorithm $\text{Pub}(\text{pk}, (c_1, c_2), r)$ computes $Z = \Lambda_{\text{sk}}(c_1, c_2)$ as $Z = X^r$.

We define $\text{KEM}_{\text{DDH}} = (\text{Gen}_{\text{DDH}}, \text{Encaps}_{\text{DDH}}, \text{Decaps}_{\text{DDH}})$ with global parameters $\text{par} := (\mathbb{G}, p, g_1, g_2)$ as shown in Fig. 15.

$\text{Gen}_{\text{DDH}}(\text{par})$	$\text{Encaps}_{\text{DDH}}^{\text{H}}(\text{pk}, m)$	$\text{Decaps}_{\text{DDH}}^{\text{H}}(\text{sk}, (c_1, c_2))$
00 $(x_1, x_2) \xleftarrow{\$} \mathbb{Z}_p^2$	03 $r \xleftarrow{\$} \mathbb{Z}_p$	07 $K := \text{H}(c_1, c_2, c_1^{x_1} c_2^{x_2})$
01 $X := g_1^{x_1} g_2^{x_2}$	04 $(c_1, c_2) := (g_1^r, g_2^r)$	08 return K
02 return $(\text{pk} := X,$ $\text{sk} := (x_1, x_2))$	05 $K := \text{H}(c_1, c_2, X^r)$	
	06 return $((c_1, c_2), K)$	

Fig. 15. Key encapsulation mechanism $\text{KEM}_{\text{DDH}} = (\text{Gen}_{\text{DDH}}, \text{Encaps}_{\text{DDH}}, \text{Decaps}_{\text{DDH}})$.

Definition 3 (m -fold DDH Problem). Let GGen be a PPT algorithm that on input 1^k outputs a cyclic group \mathbb{G} of prime order $2^{k-1} \leq p \leq 2^k$ with generator g_1 . Furthermore let $g_2 = g_1^\omega$ for $\omega \xleftarrow{\$} \mathbb{Z}_p$. The m -DDH problem requires to distinguish m DDH tuples from m uniformly random tuples:

$$\text{Adv}_{\text{GGen}}^{m\text{-DDH}}(\mathcal{A}) := \left| \Pr[\mathcal{A}(\mathbb{G}, p, g_1, g_2, (g_1^{r_i}, g_2^{r_i})_{i \in [m]}) \Rightarrow 1] - \Pr[\mathcal{A}(\mathbb{G}, p, g_1, g_2, (g_1^{r'_i}, g_2^{r'_i})_{i \in [m]}) \Rightarrow 1] \right|,$$

where probability is taken over $(\mathbb{G}, p, g) \leftarrow \text{GGen}$, $r_i, r'_i \xleftarrow{\$} \mathbb{Z}_p$ for $i \in [m]$, as well as the coin tosses of \mathcal{A} .

Lemma 1 (Random self-reducibility of DDH [17]). For any adversary \mathcal{C} against the m -fold DDH problem, there exists an adversary \mathcal{B} against the DDH problem with roughly the same running time such that

$$\text{Adv}_{\text{GGen}}^{m\text{-DDH}}(\mathcal{C}) \leq \text{Adv}_{\text{GGen}}^{\text{DDH}}(\mathcal{B}) + \frac{1}{p-1}.$$

The following theorem establishes that the construction given in Fig. 15 is an N -receiver non-committing encapsulation mechanism under the DDH assumption.

Theorem 4. Under the DDH assumption and in the random oracle model, KEM_{DDH} is an N -receiver non-committing key encapsulation mechanism. In particular, for any N -NCKE-CCA adversary \mathcal{A} against KEM_{DDH} and Sim_{DDH} that issues at most q_E queries per user to ENCAPS , q_D queries to DECAPS and at most q_H queries to each random oracle H_n , $n \in [N]$, there exists an adversary \mathcal{B} against DDH with roughly the same running time such that

$$\text{Adv}_{\text{KEM}_{\text{DDH}}, \text{Sim}_{\text{DDH}}}^{N\text{-NCKE-CCA}}(\mathcal{A}) \leq \text{Adv}_{\text{GGen}}^{\text{DDH}}(\mathcal{B}) + \frac{N \cdot q_E \cdot (q_H + q_D + 1)}{p} + \frac{1}{p-1},$$

where Sim_{DDH} is the simulator defined relative to KEM_{DDH} .

Proof. We apply Theorem 1 and analyze the entropy of the underlying HPS. The key space \mathcal{Z} is \mathbb{Z}_p . For $\text{sk} = (x_1, x_2) \xleftarrow{\$} \mathbb{Z}_p^2$, $\text{pk} = \mu(\text{sk}) = g_1^{x_1} g_2^{x_2}$ and $Z = \text{Priv}(\text{sk}, (c_1, c_2)) = c_1^{x_1} c_2^{x_2}$, where $(c_1, c_2) = (g_1^r, g_2^{r'})$ and $(r, r') \xleftarrow{\$} \mathbb{Z}_p^2$, we have

$$\begin{pmatrix} \log_{g_1} \text{pk} \\ \log_{g_1} Z \end{pmatrix} = M \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}, \text{ where } M = \begin{pmatrix} 1 & w \\ r & wr' \end{pmatrix}.$$

If $r \neq r'$, then $\det M = w(r' - r) \neq 0$, which implies that \mathbf{pk} and Z are random and independent group elements as long as x_1, x_2 are unknown. Thus, for all $Z' \in \mathcal{Z}$, holds that $\Pr[Z = Z'] = 1/p$. In Definition 3, all values r_i and r'_i are drawn uniformly at random from \mathbb{Z}_p . The probability that $r_i = r'_i$ for any $i \in [N \cdot q_E]$ is upper bounded by $N \cdot q_E/p$. Furthermore, the probability that a specific challenge ciphertext is issued to DECAPS before it is output by ENCAPS is at most q_D/p . It follows that

$$\text{Adv}_{\text{KEM,Sim}}^{\text{N-NCKE-CCA}}(\mathcal{A}) \leq \text{Adv}_{\text{GGen}}^{m\text{-DDH}}(\mathcal{B}) + \frac{N \cdot q_E}{p} + \frac{N \cdot q_E \cdot q_H}{p} + \frac{N \cdot q_E \cdot q_D}{p} .$$

Now Theorem 4 follows directly from Lemma 1. \square

7.2 Concrete Instantiation of AKE Protocols

We instantiate protocols AKE_{wFS} (Sect. 5) and AKE_{FS} (Sect. 6.2) with KEM_{DDH} (Sect. 7.1) for both KEM_{CPA} and KEM_{CCA} . We will not give a concrete instantiation of the signature scheme used in AKE_{FS} at this point. The resulting protocols $\text{AKE}_{\text{wFS,DDH}}$ and $\text{AKE}_{\text{FS,DDH}}$ are shown in Fig. 1 in the introduction.

Note that for $\text{AKE}_{\text{wFS,DDH}}$ we can improve efficiency by sending only one ciphertext for both $\widetilde{\mathbf{pk}}$ and \mathbf{pki} in the second message, as KEM_{DDH} is a multi-recipient KEM. We establish Theorem 5 and give a proof sketch.

Theorem 5 (IND-wFS-St security of $\text{AKE}_{\text{wFS,DDH}}$). *Under the DDH assumption, $\text{AKE}_{\text{wFS,DDH}}$ is IND-wFS-St secure in the random oracle model. In particular, for any IND-wFS-St adversary \mathcal{A} against $\text{AKE}_{\text{wFS,DDH}}$ with N parties that establishes at most S sessions and issues at most T queries to the test oracle TEST, q_G queries to random oracle G , $q_{\widetilde{\mathbf{H}}}$, $q_{\mathbf{H}_n}$ queries to each random oracle $\widetilde{\mathbf{H}}_{\text{SID}}$ and \mathbf{H}_n and at most q_H queries to random oracle \mathbf{H} , there exists an adversary \mathcal{B} against DDH with roughly the same running time such that*

$$\begin{aligned} \text{Adv}_{\text{AKE}_{\text{wFS,DDH}}}^{\text{IND-wFS-St}}(\mathcal{A}) \leq & 2 \cdot \text{Adv}_{\text{GGen}}^{\text{DDH}}(\mathcal{B}) + T \cdot \frac{q_G + q_H}{2^\kappa} + (N + S)^2 \cdot \frac{1}{p} + N^2 \cdot \frac{1}{2^\kappa} \\ & + S^2 \cdot \left(\frac{2}{p} + \frac{1}{2^\kappa} \right) + 2S \cdot \left(\frac{q_G}{2^{2\kappa}} + \frac{q_{\widetilde{\mathbf{H}}} + q_{\mathbf{H}_n} + 1}{p} \right) + \frac{2}{p-1} , \end{aligned}$$

where κ is a security parameter.

Due to the improved construction, we cannot apply Theorem 2 directly, but we give a proof sketch from the DDH assumption and show that the same technique as in the proofs of Theorems 2 and 4 can be used.

Proof. We proceed similar and consider collisions first. We assume that all key pairs generated by Gen_{DDH} are different. Note that we also have to consider collisions between long-term and ephemeral public keys. It holds that

$$\Pr[x_1, x_2, x'_1, x'_2 \stackrel{\$}{\leftarrow} \mathbb{Z}_p : g_1^{x_1} g_2^{x_2} = g_1^{x'_1} g_2^{x'_2}] = 1/p .$$

Union bound yields $(N+S)^2/p$, as we have N long-term public keys and at most S ephemeral public keys. For ciphertexts $(c_1, c_2) \in \mathcal{C}$ output by the encapsulation algorithm $\text{Encaps}_{\text{DDH}}$, it holds that $\Pr[r \xleftarrow{\$} \mathbb{Z}_p : (c_1, c_2) = (g_1^r, g_2^r)] = 1/p$, which yields an upper bound for collisions of S^2/p , as there are at most S sessions with one ciphertext. We also assume that values IV are different in all sessions and keys k_n are different for all parties.

We use the secret keys to compute keys K_i, K_r and \tilde{K} . Next, we replace all ciphertexts by uniformly random group elements at the same time, reducing to the S -fold DDH assumption and use the random self-reducibility property. In addition to that, we ensure that all ciphertexts are indeed invalid by adding S/p which is the probability that exponents are the same for any ciphertext.

Instead of the corresponding random oracles, we use internal hash functions \tilde{H}'_{sID} and H'_n for $\text{sID} \in [S]$ and $n \in [N]$ to compute keys K_i, K_r and \tilde{K} , but patch the random oracles if the secret key is known to the adversary. As there are at most S challenge keys computed with a long-term key pair and at most S challenge keys computed with an ephemeral key pair, the difference can be upper bounded by $S \cdot q_{H_n}/p + S \cdot q_{\tilde{H}}/p$ using a hybrid argument. Now we can replace K_i, K_r and \tilde{K} by uniformly random keys.

The rest of the proof is equal to the proof of Theorem 2. The size of the key space of KEM_{DDH} is 2^κ and the bound follows by collecting all probabilities. \square

For protocol $\text{AKE}_{\text{FS,DDH}}$, we apply Theorem 3 to show IND-FS-St security. The collision probabilities for KEM_{DDH} are already shown in the previous proof. Additionally, we need a strongly unforgeable signature scheme.

Theorem 6 (IND-FS-St security of $\text{AKE}_{\text{FS,DDH}}$). *For an N -SUF-CMA secure signature scheme SIG and under the DDH assumption, $\text{AKE}_{\text{FS,DDH}}$ is IND-FS-St secure in the random oracle model. In particular, for any IND-FS-St adversary \mathcal{A} against $\text{AKE}_{\text{FS,DDH}}$ with N parties that establishes at most S sessions and issues at most T queries to the test oracle TEST , q_G queries to random oracle G , q_F queries to random oracle F , $q_{\tilde{H}}, q_{H_n}$ queries to each random oracle \tilde{H}_{sID} and H_n and at most q_H queries to random oracle H , there exists an adversary \mathcal{B} against DDH and an adversary \mathcal{C} against N -SUF-CMA such that*

$$\begin{aligned} \text{Adv}_{\text{AKE}_{\text{FS,DDH}}}^{\text{IND-FS-St}}(\mathcal{A}) &\leq 4 \cdot \text{Adv}_{\text{GGen}}^{\text{DDH}}(\mathcal{B}) + 2 \cdot \text{Adv}_{\text{SIG}}^{N\text{-SUF-CMA}}(\mathcal{C}) + T \cdot \frac{q_F + q_G + q_H}{2^\kappa} \\ &\quad + N^2 \cdot \left(\frac{1}{2^{\mu_{\text{SIG}}}} + \frac{1}{p} + \frac{1}{2^\kappa} \right) + S^2 \cdot \left(\frac{2q_{\tilde{H}} + 6}{p} + \frac{1}{2^\kappa} \right) \\ &\quad + 2NS \cdot \frac{q_{H_n} + 2}{p} + 2S \cdot \frac{q_G}{2^{2\kappa}} + \frac{4}{p-1}, \end{aligned}$$

where μ_{SIG} is the collision probability of the key generation algorithm Gen_{SIG} and κ is a security parameter.

The signature scheme can be instantiated with the tight scheme based on the DDH and CDH assumption proposed by Gjøsteen and Jager in [21], which is also used in their authenticated key exchange protocol.

Acknowledgments. Tibor Jager was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme, grant agreement 802823. Eike Kiltz was supported by the BMBF iBlockchain project, the EU H2020 PROMETHEUS project 780701, DFG SPP 1736 Big Data, and the DFG Cluster of Excellence 2092 CASA. Doreen Riepel was supported by the Deutsche Forschungsgemeinschaft (DFG) Cluster of Excellence 2092 CASA. Sven Schäge was supported by the German Federal Ministry of Education and Research (BMBF), Project DigiSeal (16KIS0695) and Huawei Technologies Düsseldorf, Project vHSM.

References

1. Bader, C., Hofheinz, D., Jager, T., Kiltz, E., Li, Y.: Tightly-secure authenticated key exchange. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part I. LNCS, vol. 9014, pp. 629–658. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46494-6_26
2. Bader, C., Jager, T., Li, Y., Schäge, S.: On the impossibility of tight cryptographic reductions. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 273–304. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_10
3. Bellare, M., Boldyreva, A., Micali, S.: Public-key encryption in a multi-user setting: security proofs and improvements. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 259–274. Springer, Heidelberg (2000). https://doi.org/10.1007/3-540-45539-6_18
4. Bellare, M., Rogaway, P.: Entity authentication and key distribution. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 232–249. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-48329-2_21
5. Bergsma, F., Jager, T., Schwenk, J.: One-round key exchange with strong security: an efficient and generic construction in the standard model. In: Katz, J. (ed.) PKC 2015. LNCS, vol. 9020, pp. 477–494. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46447-2_11
6. Boyd, C., Nieto, J.G.: On forward secrecy in one-round key exchange. In: Chen, L. (ed.) IMACC 2011. LNCS, vol. 7089, pp. 451–468. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25516-8_27
7. Brzuska, C., Fischlin, M., Warinschi, B., Williams, S.C.: Composability of Bellare-Rogaway key exchange protocols. In: Chen, Y., Danezis, G., Shmatikov, V. (eds.) ACM CCS 2011, pp. 51–62. ACM Press (2011). <https://doi.org/10.1145/2046707.2046716>
8. Canetti, R., Feige, U., Goldreich, O., Naor, M.: Adaptively secure multi-party computation. In: 28th ACM STOC, pp. 639–648. ACM Press (1996). <https://doi.org/10.1145/237814.238015>
9. Canetti, R., Krawczyk, H.: Analysis of key-exchange protocols and their use for building secure channels. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 453–474. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44987-6_28
10. Cohn-Gordon, K., Cremers, C., Gjøsteen, K., Jacobsen, H., Jager, T.: Highly efficient key exchange protocols with optimal tightness. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part III. LNCS, vol. 11694, pp. 767–797. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26954-8_25

11. Cramer, R., Shoup, V.: Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 45–64. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-46035-7_4
12. Cremers, C.: Examining indistinguishability-based security models for key exchange protocols: the case of CK, CK-HMQV, and eCK. In: Cheung, B.S.N., Hui, L.C.K., Sandhu, R.S., Wong, D.S. (eds.) ASIACCS 11, pp. 80–91. ACM Press (2011)
13. Cremers, C.J.F.: Session-state Reveal is stronger than Ephemeral Key Reveal: attacking the NAXOS authenticated key exchange protocol. In: Abdalla, M., Pointcheval, D., Fouque, P.-A., Vergnaud, D. (eds.) ACNS 2009. LNCS, vol. 5536, pp. 20–33. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-01957-9_2
14. Cremers, C., Feltz, M.: Beyond eCK: perfect forward secrecy under actor compromise and ephemeral-key reveal. In: Foresti, S., Yung, M., Martinelli, F. (eds.) ESORICS 2012. LNCS, vol. 7459, pp. 734–751. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33167-1_42
15. Davis, H., Günther, F.: Tighter proofs for the SIGMA and TLS 1.3 key exchange protocols. Cryptology ePrint Archive, Report 2020/1029 (2020). <https://eprint.iacr.org/2020/1029>
16. Diemert, D., Jager, T.: On the tight security of TLS 1.3: theoretically-sound cryptographic parameters for real-world deployments. Cryptology ePrint Archive, Report 2020/726 (2020). <https://eprint.iacr.org/2020/726>
17. Escala, A., Herold, G., Kiltz, E., Ràfols, C., Villar, J.: An algebraic framework for Diffie-Hellman assumptions. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 129–147. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40084-1_8
18. Freire, E.S.V., Hofheinz, D., Kiltz, E., Paterson, K.G.: Non-interactive key exchange. In: Kurosawa, K., Hanaoka, G. (eds.) PKC 2013. LNCS, vol. 7778, pp. 254–271. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36362-7_17
19. Fujioka, A., Suzuki, K., Xagawa, K., Yoneyama, K.: Strongly secure authenticated key exchange from factoring, codes, and lattices. In: Fischlin, M., Buchmann, J., Manulis, M. (eds.) PKC 2012. LNCS, vol. 7293, pp. 467–484. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-30057-8_28
20. Gay, R., Hofheinz, D., Kiltz, E., Wee, H.: Tightly CCA-secure encryption without pairings. In: Fischlin, M., Coron, J.-S. (eds.) EUROCRYPT 2016, Part I. LNCS, vol. 9665, pp. 1–27. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49890-3_1
21. Gjøsteen, K., Jager, T.: Practical and tightly-secure digital signatures and authenticated key exchange. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018. LNCS, vol. 10992, pp. 95–125. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96881-0_4
22. Hövelmanns, K., Kiltz, E., Schäge, S., Unruh, D.: Generic authenticated key exchange in the quantum random oracle model. In: Kiayias, A., Kohlweiss, M., Wallden, P., Zikas, V. (eds.) PKC 2020, Part II. LNCS, vol. 12111, pp. 389–422. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45388-6_14
23. Jager, T., Kiltz, E., Riepel, D., Schäge, S.: Tightly-secure authenticated key exchange, revisited. Cryptology ePrint Archive, Report 2020/1279 (2020). <https://eprint.iacr.org/2020/1279>

24. Jager, T., Kohlar, F., Schäge, S., Schwenk, J.: On the security of TLS-DHE in the standard model. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 273–293. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32009-5_17
25. Kiltz, E., Pietrzak, K., Stam, M., Yung, M.: A new randomness extraction paradigm for hybrid encryption. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 590–609. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-01001-9_34
26. Krawczyk, H.: HMQV: a high-performance secure Diffie-Hellman protocol. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 546–566. Springer, Heidelberg (2005). https://doi.org/10.1007/11535218_33
27. Kudla, C., Paterson, K.G.: Modular security proofs for key agreement protocols. In: Roy, B. (ed.) ASIACRYPT 2005. LNCS, vol. 3788, pp. 549–565. Springer, Heidelberg (2005). https://doi.org/10.1007/11593447_30
28. LaMacchia, B., Lauter, K., Mityagin, A.: Stronger security of authenticated key exchange. In: Susilo, W., Liu, J.K., Mu, Y. (eds.) ProvSec 2007. LNCS, vol. 4784, pp. 1–16. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-75670-5_1
29. Lauter, K., Mityagin, A.: Security analysis of KEA authenticated key exchange protocol. In: Yung, M., Dodis, Y., Kiayias, A., Malkin, T. (eds.) PKC 2006. LNCS, vol. 3958, pp. 378–394. Springer, Heidelberg (2006). https://doi.org/10.1007/11745853_25
30. Li, Y., Schäge, S.: No-match attacks and robust partnering definitions: defining trivial attacks for security protocols is not trivial. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017, pp. 1343–1360. ACM Press (2017). <https://doi.org/10.1145/3133956.3134006>
31. Li, Y., Schäge, S., Yang, Z., Bader, C., Schwenk, J.: New modular compilers for authenticated key exchange. In: Boureanu, I., Owesarski, P., Vaudenay, S. (eds.) ACNS 2014. LNCS, vol. 8479, pp. 1–18. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-07536-5_1
32. Liu, X., Liu, S., Gu, D., Weng, J.: Two-pass authenticated key exchange with explicit authentication and tight security. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020, Part II. LNCS, vol. 12492, pp. 785–814. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64834-3_27
33. Nielsen, J.B.: Separating random oracle proofs from complexity theoretic proofs: the non-committing encryption case. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 111–126. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45708-9_8
34. Schäge, S.: TOPAS: 2-pass key exchange with full perfect forward secrecy and optimal communication complexity. In: Ray, I., Li, N., Kruegel, C. (eds.) ACM CCS 2015, pp. 1224–1235. ACM Press (2015). <https://doi.org/10.1145/2810103.2813683>
35. Shoup, V.: Sequences of games: a tool for taming complexity in security proofs. Cryptology ePrint Archive, Report 2004/332 (2004). <http://eprint.iacr.org/2004/332>
36. Yoneyama, K.: One-round authenticated key exchange with strong forward secrecy in the standard model against constrained adversary. In: Hanaoka, G., Yamauchi, T. (eds.) IWSEC 2012. LNCS, vol. 7631, pp. 69–86. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34117-5_5