



Key Exchange with Tight (Full) Forward Secrecy via Key Confirmation

Jiaxin Pan¹, Doreen Riepel², and Runzhi Zeng³

¹ University of Kassel, Kassel, Germany
jiaxin.pan@uni-kassel.de

² University of California San Diego, La Jolla, USA
driepel@ucsd.edu

³ Norwegian University of Science and Technology, Trondheim, Norway
runzhi.zeng@ntnu.no

Abstract. Weak forward secrecy (wFS) of authenticated key exchange (AKE) protocols is a passive variant of (full) forward secrecy (FS). A natural mechanism to upgrade from wFS to FS is the use of key confirmation messages which compute a message authentication code (MAC) over the transcript. Unfortunately, Gellert, Gjøsteen, Jacobson and Jager (GGJJ, CRYPTO 2023) show that this mechanism inherently incurs a loss proportional to the number of users, leading to an overall non-tight reduction, even if wFS was established using a tight reduction.

Inspired by GGJJ, we propose a new notion, called one-way verifiable weak forward secrecy (OW-VwFS), and prove that OW-VwFS can be transformed *tightly* to FS using key confirmation in the random oracle model (ROM). To implement our generic transformation, we show that several tightly wFS AKE protocols additionally satisfy our OW-VwFS notion tightly. We highlight that using the recent lattice-based protocol from Pan, Wagner, and Zeng (CRYPTO 2023) can give us the first lattice-based tightly FS AKE via key confirmation in the classical random oracle model. Besides this, we also obtain a Decisional-Diffie-Hellman-based protocol that is considerably more efficient than the previous ones.

Finally, we lift our study on FS via key confirmation to the quantum random oracle model (QROM). While our security reduction is overall non-tight, it matches the best existing bound for wFS in the QROM (Pan, Wagner, and Zeng, ASIACRYPT 2023), namely, it is square-root-and-session-tight. Our analysis is in the multi-challenge setting, and it is more realistic than the single-challenge setting as in Pan et al.

Keywords: Authenticated key exchange · forward secrecy · key confirmation · tight security · (quantum) random oracles

1 Introduction

Forward secrecy (FS) is an essential security requirement for authenticated key exchange (AKE) protocols. It states that even if an active adversary corrupts a

user’s long-term secret key, all session keys agreed before should remain secret to the adversary. A weaker form of FS is called weak FS (wFS), where an adversary is not allowed to perform active attacks, namely, it does not actively interfere with the protocol transcripts of the session that it attacks.

Key confirmation is simple and arguably the most efficient way in achieving FS and has been used in many works, e.g., [10, 15, 24]. Essentially, it generically transforms an AKE protocol with wFS to FS. More precisely, two parties firstly run a wFS AKE protocol to agree on a session key k , and then they exchange key confirmation messages derived from k . These messages are usually message authentication codes (MAC) on the protocol transcripts using k as the MAC key. Apart from key confirmation, one can use a digital signature scheme to sign a passively secure key exchange protocol as in the signed Diffie-Hellman protocol [19, 27] to provide FS. Considering that using a MAC or hash function is much more efficient than digital signatures, the signature-based approach is often inefficient and less desirable.

SECURITY MODELS FOR AKE. Defining the security for AKE protocols is a complex task, and there are many different security models for AKE (e.g., [3, 7, 25]). In this paper, we consider active adversaries that can modify, drop, or inject some messages. Moreover, they may adaptively corrupt users’ long-term secret keys via CORR oracle and reveal session keys via REVEAL oracle. Some of the models even allow adversaries to learn ephemeral states (which are usually randomness in generating protocol messages) via REV-STATE oracle. We formalize key secrecy via TEST, where an adversary \mathcal{A} chooses a *fresh* session, receives either a real or random key for it, and shall distinguish between the two. We consider the single-bit guessing, multi-challenge security, namely, \mathcal{A} can query TEST multiple times and each time TEST responds using the same bit in deciding real or random. Composability for this notion was initially proven for password-based key exchange [1], and we refer to [21] for further discussion on why this is the realistic and meaningful notion. For forward secrecy, keys of these TEST-sessions must be computed before CORR is queried to either parties of a TEST-session. Depending on the type of forward secrecy, freshness is defined differently. If it is wFS, then \mathcal{A} must perform only passive attacks on this fresh session. Otherwise, \mathcal{A} can perform active attacks, for instance, modify or inject some messages.

SECURITY LOSS FOR FS VIA KEY CONFIRMATION. The complexity of AKE models makes it challenging to prove security of an AKE protocol, in particular, giving tight security proofs for AKE. The security of modern cryptographic protocols is often proven by reductions. A reduction \mathcal{R} uses an adversary \mathcal{A} against protocol Π to break the security of the underlying primitive P . By doing so, we can conclude the concrete security bound, $\varepsilon_{\mathcal{A}} \leq \ell \cdot \varepsilon_{\mathcal{R}}$, where $\varepsilon_{\mathcal{A}}$ and $\varepsilon_{\mathcal{R}}$ are the success probability of \mathcal{A} and \mathcal{R} , respectively. ℓ is called the security loss. Assuming \mathcal{A} and \mathcal{R} have roughly the same running time, if ℓ is a small constant, we say protocol Π has tight security, and non-tight security, otherwise. A tight security reduction is highly desirable, since it allows protocols to be instantiated with optimal parameters without compensation for the security loss.

A natural question to ask is whether the key confirmation approach preserves the tightness of the underlying wFS AKE. Due to its high efficiency, it would be ideal to have an affirmative answer to this question, since it means that we do not need to increase the security parameter of the wFS AKE to compensate any security loss.

Intuitively, there should not be a tightness loss when going from wFS to FS, which was even falsely claimed by the work of Cohn-Gordon et al. [10] previously. At CRYPTO 2023, Gellert, Gjøsteen, Jacobsen, and Jager (GGJJ) [18] identified a flaw in [10] and proposed a fix by using a selective variant of wFS (called selective key secrecy in [18]). The selective wFS is essentially the same as wFS, except that an adversary \mathcal{A} has to select a user of which \mathcal{A} will not corrupt the long-term secret key. Unfortunately, when we construct a reduction \mathcal{R} to prove FS based on this selective wFS, \mathcal{R} has to guess the non-corrupted user, which leads to a security loss of $O(\mu)$ where μ is the maximal number of users. This security loss is proven to be inherent (and thus optimal) in [18] when starting from a wFS AKE with key indistinguishability.

However, a linear loss in the number of users is undesirable, since in the real world the number of users can be massive. According to the impossibility result in [10], it seems inherent to have this security loss. Hence, it motivates us to propose a different modularization that potentially requires strong security for the underlying wFS AKE in achieving tight FS.

1.1 Our Contribution I: Tight Forward Secrecy via Key Confirmation

We revise the security proof for the wFS-to-FS transformation.

TIGHT FS FROM VERIFIABLE WFS. We propose a new variant of wFS, called One-Wayness against key Verification attacks and weak Forward Secrecy (OW-VwFS). In the OW-VwFS security game, an adversary has the same capability as in the usual wFS game, but additionally it can verify whether a session key is the valid one of a particular session. Hence, the adversary capability of OW-VwFS is stronger than that of wFS and it is the main reason why we bypass the optimality result from Gellert et al. [18]. In terms of security goals, OW-VwFS is weaker than wFS, namely, OW-VwFS only requires an adversary cannot compute the session key of a fresh session, while wFS requires a session key to be indistinguishable from a random key.

Using key confirmation, we prove that OW-VwFS *tightly* implies FS in the random oracle model. Our transformation is the same as the standard wFS-to-FS transformation, but ours preserves the tightness of the underlying OW-VwFS protocol, and it enables tight FS in contrast to the selective notion in [18]. An important consequence of our work is that the future AKE design can aim at OW-VwFS, since its transformation to FS is the same as the standard wFS-to-FS one, but tightness-preserving. Moreover, our analysis considers security against (ephemeral) state reveals. Such a strong form of attacks was not considered in the work of Gellert et al. [18], which is why we bypass their impossibility.

CONSTRUCTING (TIGHTLY) VERIFIABLE wFS. Furthermore, we show that several tightly wFS protocols satisfy our new OW-VwFS notion tightly, in particular, the lattice-based protocol of Pan, Wagner, and Zeng [30]¹. Subsequently, this yields the *first* AKE protocol with tight FS from lattices.

Essentially, we show that a One-Way Checkable against Chosen-Ciphertext Attacks (OW-ChCCA) [30] secure key encapsulation mechanism (KEM) tightly implies a OW-VwFS AKE protocol. Once again, our analysis allows adversaries to reveal ephemeral state in the AKE protocol. Roughly speaking, the OW-ChCCA game is a multi-user, multi-challenge variant of the standard IND-CCA game: Besides the oracles provided by the standard IND-CCA security, it allows adversaries to corrupt some of the user’s decryption keys and decrypt some of the challenge ciphertexts, and, most importantly, it allows an adversary to check if a key is valid with respect to a ciphertext. The adversary goal is to invert a fresh challenge ciphertext. As shown in [30], we can construct OW-ChCCA KEM tightly from the Decisional Diffie-Hellman (DDH) and Learning-With-Errors (LWE) assumptions, respectively. As a technical note, our proof requires only a slightly weaker version of OW-ChCCA, where adversaries are not allowed to ask for a decryption, but to verify whether a ciphertext can be decapsulated.

EFFICIENCY COMPARISON AMONG DDH-BASED PROTOCOLS. Besides having the first lattice-based AKE with tight FS, we also obtain the most efficient DDH-based protocol against state reveal attacks. In Table 1, we compare efficiency among well-known DDH-based AKE with tight or “optimal” tight FS (namely, with security loss $O(\mu)$) to show the practicality of our work. Our estimation focuses on communication and computation complexity for both parties to agree on a session key. For computation complexity, we only count the number of exponentiations, since they are the most costly operations. For concrete efficiency, we instantiate the protocols at 128-bit security and assume that the number of users $\mu \approx 2^{30}$. This is about the number of monthly active users in a social media app². We instantiate the fully tight protocols with a NIST P256 curve and “optimal” tight ones with a NIST P384 (since they require a 158-bit hard DLog assumption). Our benchmarks for an exponentiation in a P256 and P384 are 0.5 ms and 1 ms, using Apple M1 Max, 32GB of RAM and macOS Ventura 13.3.1 (a).

We observe that the DDH-based non-committing KEM in [21] is tightly OW-ChCCA secure (cf. [30, Footnote 1]). Our analysis shows that the wFS JKRS in [21] is tightly OW-VwFS, and after adding key confirmation to the wFS JKRS in [21] we get JKRS_{KC} that is tightly FS. According to Table 1, our tight security proofs allow one to implement JKRS_{KC} with about 30% shorter transcripts and 50% faster speed than the one with the “optimal”, non-tight proofs in [18] at 128-bit security. Considering security against State Reveals, JKRS_{KC} is the

¹ Their lattice-based protocol is almost tight (similar to [9]), since it needs to lose a factor of $O(\lambda)$ to the LWE assumption, where λ is the security parameter. We call it tight as well, but specify the concrete loss in our theorems and proofs.

² Cf. <https://www.statista.com/statistics/264810/number-of-monthly-active-facebook-users-worldwide/>.

Table 1. Comparison of Diffie-Hellman-based AKE protocols with (tight or “optimal” tight) FS. Concrete efficiency is estimated for 128-bit security. “JKRS_{KC} [18]” is transforming the implicitly authenticated JKRS [21] via key confirmation. We estimate its efficiency, according to the “optimal”, non-tight security bound by Gellert et al. [18]. The last row is the same construction as the second last one, but with our tight security proof (cf. Theorem 4). In the upper arrows, schemes are using signatures, and we estimate the concrete bytes with the most efficient signature scheme in [13]. **Comm.** counts values exchanged during the protocol execution. \mathbb{G} counts the number of group elements, H the number of hashes or MACs, ‘Sign’ the number of signatures, and ‘other’ the additional data in bits. **Bytes** counts total data in bytes by instantiating \mathbb{G} with NIST P256 or P384 (for the non-tight JKRS_{KC}). **Exp.** counts the total numbers of exponentiation (which is the most costly computation in an AKE protocol) from both parties in agreeing a session key, and **Time** is the estimated time of computing those exponentiation in milliseconds.

Protocol	Comm. (\mathbb{G} , H , Sign, other)	Bytes	Exp.	Time (ms)	#Msg.	State Reveal	Security loss
TLS 1.3 [11, 14]	(2, 2, 2, 512)	384	32	16	3	no	$O(1)$
GJ [19]	(2, 1, 2, 0)	288	32	16	3	no	$O(1)$
LLGW [26]	(3, 0, 2, 0)	288	35	17.5	2	no	$O(1)$
JKRS [21]	(5, 1, 1, 0)	288	29	14.5	2	yes	$O(1)$
PQR [27]	(2, 0, 2, 0)	256	32	16	2	no	$O(1)$
CCGJJ _{KC} [10]	(2,2,0,0)	160	8	8	3	no	$O(\mu)$
JKRS _{KC} [18]	(5,2,0,0)	304	15	15	3	yes	$O(\mu)$
JKRS _{KC} (Ours)	(5,2,0,0)	224	15	7.5	3	yes	$O(1)$

most efficient DDH-based protocol, due to our tight security proofs. It is worth mentioning that the CCGJJ_{KC} has shorter protocol transcripts, but it is insecure under State Reveals.

Interestingly, although the signature-based JKRS uses relatively inefficient primitives as signatures, its tight security proof allows an instantiation that is slightly more efficient than the non-tight, signature-less JKRS (namely, JKRS_{KC} with proofs in [18]).

RELATION TO THE WORK OF GELLERT ET AL. [18]. We circumvent the impossibility result of Gellert et al. [18] by using a different wFS notion, OW-VwFS, and random oracles. As discussed earlier, the key checking oracle makes our notion stronger. The security definition in [18] does not have such an oracle and thus their impossibility result does not apply to our proof. At the same time, we opted for the weakest definition which allows a tight reduction (i.e., one-wayness and also no Reveal oracle), which makes our definition and that of [18] incomparable (neither implies the other). Moreover, their impossibility is in the standard model, while ours is in the random oracle model. For these reasons, our results do not contradict the impossibility result in [18], but rather provides an alternative way to prove security while enabling full tightness.

1.2 Our Contribution II: Forward Secrecy via Key Confirmation in the QROM

Our second contribution is proposing the first security proof for FS via key confirmation in the quantum random oracle model (QROM) [6], where a quantum adversary can have quantum access to the hash function. Our analysis considers the KEM-based AKE protocol (via key confirmation) and assumes a Multi-User, Multi-Challenge Chosen-Ciphertext Attacks (MUC-CCA) KEM and a Multi-Challenge CCA (MC-CCA) KEM. The main reason of doing so is that we do not know how to tightly prove OW-VwFS implies FS in the QROM, since it will trigger the Oneway-to-Hiding Lemma [32] and lead to a square-root-loss such as $\sqrt{\varepsilon}$, where ε is the advantage of breaking the underlying KEM. We still think that our tight lattice-based protocol in the classical ROM is interesting, since it is the first protocol with tight FS from post-quantum assumptions. Of course, one may alternatively rephrase our analysis in the classical ROM with the suitable KEMs, but it may lower the readability. More importantly, our OW-VwFS notion is more generic and gives more freedom to designers to construct their OW-VwFS protocols that will lead to FS in a tightness-preserving manner.

Our security bound in the QROM is unfortunately non-tight. More precisely, ignoring the statistically negligible terms, **our security bound for FS** in the QROM is

$$\varepsilon_{\text{FS}}^{\text{our}} \leq O(\mu) \cdot \varepsilon_{\text{MC-CCA}} + O(1) \cdot \varepsilon_{\text{MUC-CCA}}. \quad (1)$$

where μ is the number of users, $\varepsilon_{\text{MC-CCA}}$ is the advantage of MC-CCA, and $\varepsilon_{\text{MUC-CCA}}$ is that of MUC-CCA. It matches the best known bound for wFS in the QROM proposed by Pan, Wagner, and Zeng (PWZ) [31]. In this sense our FS bound preserves the tightness of the KEM-based AKE protocol with wFS. It is worth mentioning that, as shown in [31], we can tightly instantiate MC-CCA and MUC-CCA from the LWE assumption in the QROM.

We also improve PWZ’s analysis in the sense that their analysis considers only one single TEST query, but our security bound (as stated in Eq. (1)) is established in the context of multiple challenges, where an adversary is allowed to query TEST multiple times. The multi-challenge setting is more realistic and well-established for public-key primitives [2, 10, 16, 17], since in the real world, an adversary usually wants to attack multiple instances of a primitives. Although the security bound of PWZ can be extended to the multi-TEST setting with a multiplicative factor t (which is the number of TEST-queries), ours does not need to lose such a factor. In practice, t can be up to the total number of established sessions, which can be much larger than the number of users. We stress that the analysis of PWZ is only for wFS, and transforming it to FS, it may lose another multiplicative factor μ by applying the analysis of GGJJ [18]. Hence, combining the analysis of PWZ and GGJJ leads to a bound for FS in the QROM as

$$\varepsilon_{\text{FS}}^{\text{PWZ \& GGJJ}} \leq O(\mu^2 t) \cdot \varepsilon_{\text{MC-CCA}} + O(\mu t) \cdot \varepsilon_{\text{MUC-CCA}}. \quad (2)$$

Strictly speaking, the bound above is only an estimation and not theoretically sound, since the analysis of GGJJ is in the classical setting. Their bound may

change, if an adversary can query the hash function or key derivation function with a quantum state.

MORE RELATED WORK IN THE QROM. Another work on the KEM-based AKE protocol in the QROM is due to Hövelmanns, Kiltz, Schäge, and Unruh [20], and it has a square-root-loss, namely, its security bound is

$$O(S^2 + S \cdot \mu) \cdot \left(\varepsilon_{\text{CPA}} + \sqrt{Q \cdot \varepsilon_{\text{CPA}}} \right), \quad (3)$$

where S , μ , and Q are the numbers of total sessions, users, and random oracle queries, respectively, and ε_{CPA} is the advantage of breaking the underlying CPA secure PKE. Similar to the work of PWZ, Eq. (3) is only for wFS and in the single-TEST setting. Upgrading to FS in the multi-TEST setting requires an additional multiplicative loss in μt . It is usually less desirable to have the square-root-loss as in Eq. (3), since it reduces the security guarantee of the underlying PKE in half.

2 Preliminaries

For $n \in \mathbb{N}$, let $[n] = \{1, \dots, n\}$. For a finite set \mathcal{S} , we denote the sampling of a uniform random element x by $x \xleftarrow{\$} \mathcal{S}$. By $\llbracket B \rrbracket$ we denote the bit that is 1 if the evaluation of the Boolean statement B is **true** and 0 otherwise.

ALGORITHMS. For an algorithm \mathcal{A} which takes x as input, we denote its computation by $y := \mathcal{A}(x)$ if \mathcal{A} is deterministic, and $y \leftarrow \mathcal{A}(x)$ if \mathcal{A} is probabilistic. We assume all the algorithms (including adversaries) in this paper to be probabilistic unless stated differently. We denote an algorithm \mathcal{A} with access to an oracle \mathcal{O} by $\mathcal{A}^{\mathcal{O}}$. In terms of running time, if a reduction’s running time t' is dominated by that of an adversary t (more precisely, $t' = t + s$ where $s \ll t$), we write $t' \approx t$.

GAMES. We use code-based games [4] to present our definitions and proofs. We implicitly assume all Boolean flags to be initialized to 0 (**false**), numerical variables to 0, sets to \emptyset and strings to \perp . We make the convention that a procedure terminates once it has returned an output. $G^{\mathcal{A}} \Rightarrow b$ denotes the final (Boolean) output b of game G running adversary \mathcal{A} , and if $b = 1$ we say \mathcal{A} wins G . The randomness in $\Pr[G^{\mathcal{A}} \Rightarrow 1]$ is over all random coins in game G . More generically, we write $\Pr[\text{Event} : G]$ to denote the probability that **Event** happens in the game G . If the context is clear, we simply write it as $\Pr[\text{Event}]$. Within a procedure, “**abort**” means that we terminate the run of an adversary \mathcal{A} .

3 Three-Message Authenticated Key Exchange

We recall the AKE security model from [21] and adapt it to three-message protocols. A three-message key exchange protocol $\text{AKE} := (\text{Setup}, \text{Gen}_{\text{AKE}}, \text{Init}_{\text{I}}, \text{Init}_{\text{R}}, \text{Der}_{\text{I}}, \text{Der}_{\text{R}})$ consists of five algorithms which are executed interactively by two parties as shown in Fig. 1.

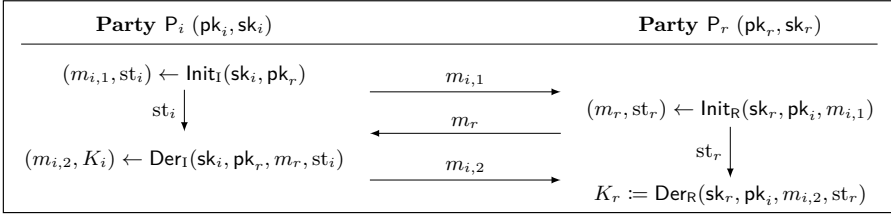


Fig. 1. Running a three-message AKE protocol between two parties.

Setup is the setup algorithm for the system parameters. 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 initiator’s 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 $m_{i,1}$ and the initiator’s state st_i . The responder’s initialization algorithm Init_R inputs the responder’s long-term secret key sk_r and the initiator’s long-term public key pk_i , and outputs a message m_r and the responder’s state st_r . The initiator’s derivation algorithm Der_I takes as input sk_i, pk_r , a message m_r and the state st_i . It computes the final message $m_{i,2}$ and a session key K . The responder’s derivation algorithm Der_R takes as input the sk_r, pk_i , a message $m_{i,2}$ and the state st_r . It computes the session key K . Here K can be \perp meaning that the session is rejected during the execution. Correctness of an AKE protocol states that an honest execution between two parties should yield the same session key.

Definition 1 (Correctness of three-message AKE). Let $\text{AKE} := (\text{Setup}, \text{Gen}_{\text{AKE}}, \text{Init}_I, \text{Init}_R, \text{Der}_I, \text{Der}_R)$ be a three-message AKE protocol. We say AKE is ρ -correct if

$$\Pr \left[\begin{array}{l} \text{par} \leftarrow \text{Setup}(1^\lambda), \\ (pk_i, sk_i) \leftarrow \text{Gen}_{\text{AKE}}(\text{par}), (pk_r, sk_r) \leftarrow \text{Gen}_{\text{AKE}}(\text{par}), \\ (m_{i,1}, st_i) \leftarrow \text{Init}_I(sk_i, pk_r), \\ (m_r, st_r) \leftarrow \text{Init}_R(sk_r, pk_i, m_{i,1}), \\ (m_{i,2}, K_i) \leftarrow \text{Der}_I(sk_i, pk_r, m_r, st_i), \\ K_r := \text{Der}_R(sk_r, pk_i, m_{i,2}, st_r) \end{array} \right] \geq \rho,$$

where the probability is taken over the randomness of $\text{Setup}, \text{Gen}_{\text{AKE}}, \text{Init}_I, \text{Init}_R$, and Der_I .

We give a security game written in pseudocode focusing on (full) forward secrecy, rather than implicit or explicit authentication. We refer readers to [12] for more details on different types of authentication for key exchange protocols, and their connections to forward secrecy in [18].

EXECUTION ENVIRONMENT. We consider μ parties P_1, \dots, P_μ with long-term key pairs $(pk_n, sk_n), n \in [\mu]$. When two parties A and B want to communicate, the initiator, say, A first creates a session. To identify this session, we increase the

<p>GAME IND-FS, $\boxed{\text{IND-FS-St}}$</p> <pre> 00 for $n \in [\mu]$ 01 $(pk_n, sk_n) \leftarrow \text{Gen}_{\text{AKE}}$ 02 $b \xleftarrow{\\$} \{0, 1\}$ 03 $b' \leftarrow \mathcal{A}^O(pk_1, \dots, pk_\mu)$ 04 for $sID^* \in \mathcal{S}_{\text{test}}$ 05 if $\text{FRESH}(sID^*) = \text{false}$ //session not fresh 06 or $\text{VALID}(sID^*) = \text{false}$ //no valid attack 07 return b 08 return $\llbracket b = b' \rrbracket$ </pre> <p>$\text{SESSION}_R((i, r) \in [\mu]^2, m_{i,1})$</p> <pre> 09 cnt$_S$ ++ 10 sID := cnt$_S$ 11 (Init[sID], Resp[sID]) := (i, r) 12 Type[sID] := "Re" 13 (m_r, st_r) $\leftarrow \text{Init}_R(sk_r, pk_i, m_{i,1})$ 14 (Msg$_{i,1}$[sID], Msg$_R$[sID]) := ($m_{i,1}, m_r$) 15 st[sID] := st$_r$ 16 return (sID, m_r) </pre> <p>$\text{DER}_R(sID \in [\text{cnt}_S], m_{i,2})$</p> <pre> 17 if $\text{SK}[sID] \neq \perp$ or $\text{Type}[sID] \neq \text{"Re"}$ 18 return \perp //no re-use 19 (i, r) := (Init[sID], Resp[sID]) 20 st$_r$:= ST[sID] 21 peerPreCor[sID] := cor[i] 22 $K := \text{Der}_R(sk_r, pk_i, m_{i,2}, st_r)$ 23 if $K \neq \perp$ 24 SK[sID] := K 25 else 26 SK[sID] := "reject" 27 Msg$_{i,2}$[sID] := $m_{i,2}$ 28 return ε </pre> <p>$\text{REV-STATE}(sID)$</p> <pre> 29 revST[sID] := true 30 return ST[sID] </pre>	<p>$\text{SESSION}_I((i, r) \in [\mu]^2)$</p> <pre> 31 cnt$_S$ ++ 32 sID := cnt$_S$ 33 (Init[sID], Resp[sID]) := (i, r) 34 Type[sID] := "In" 35 ($m_{i,1}, st_i$) $\leftarrow \text{Init}_I(sk_i, pk_r)$ 36 (Msg$_{i,1}$[sID], ST[sID]) := ($m_{i,1}, st_i$) 37 return (sID, $m_{i,1}$) </pre> <p>$\text{DER}_I(sID \in [\text{cnt}_S], m_r)$</p> <pre> 38 if $\text{SK}[sID] \neq \perp$ or $\text{Type}[sID] \neq \text{"In"}$ 39 return \perp //no re-use 40 (i, r) := (Init[sID], Resp[sID]) 41 st$_i$:= ST[sID] 42 peerPreCor[sID] := cor[r] 43 ($m_{i,2}, K$) $\leftarrow \text{Der}_I(sk_i, pk_r, m_r, st_i)$ 44 (Msg$_R$[sID], Msg$_{i,2}$[sID]) := ($m_r, m_{i,2}$) 45 if $K \neq \perp$ 46 SK[sID] := K 47 else 48 SK[sID] := "reject" 49 return $m_{i,2}$ </pre> <p>$\text{REVEAL}(sID)$</p> <pre> 50 revSK[sID] := true 51 return SK[sID] </pre> <p>$\text{CORR}(n \in [\mu])$</p> <pre> 52 cor[n] := true 53 return sk$_n$ </pre> <p>$\text{TEST}(sID)$</p> <pre> 54 if sID $\in \mathcal{S}_{\text{test}}$ return \perp //already tested 55 if SK[sID] $\in \{\perp, \text{"reject"}\}$ return \perp 56 $\mathcal{S}_{\text{test}} := \mathcal{S}_{\text{test}} \cup \{sID\}$ 57 $K_0^* := \text{SK}[sID]$ 58 $K_1^* \xleftarrow{\\$} \mathcal{K}$ 59 return K_b^* </pre>
--	--

Fig. 2. Games IND-FS and IND-FS-St for AKE. REV-STATE is only available in IND-FS-St. In IND-FS, \mathcal{A} has access to oracles $\mathcal{O} := \{\text{SESSION}_I, \text{SESSION}_R, \text{DER}_I, \text{DER}_R, \text{REVEAL}, \text{CORR}, \text{TEST}\}$. In IND-FS-St, \mathcal{A} has access to the oracles in IND-FS and the REV-STATE oracle. Helper procedures FRESH and VALID are defined in Fig. 3. If there exists any test session which is neither fresh nor valid, the game will return b .

global identification number sID and assign the current state of sID to identify this session owned by A. The state of sID will increase after every assignment. Moreover, a message will be sent to the responder. The responder then similarly creates a corresponding session which is assigned the current state of sID. Hence each conversation includes two sessions. We then define variables in relation to the identifier sID:

- $\text{Init}[sID] \in [\mu]$ denotes the initiator of the session.
- $\text{Resp}[sID] \in [\mu]$ denotes the responder of the session.
- $\text{Type}[sID] \in \{\text{"In"}, \text{"Re"}\}$ denotes the session's view, i. e. whether the initiator or the responder computes the session key.

FRESH(sID*)	VALID(sID*)
00 $\mathfrak{M}(sID^*) := \text{MATCH}(sID^*)$	06 $\mathfrak{M}(sID^*) := \text{MATCH}(sID^*)$
01 if $\text{revSK}[sID^*]$ or $(\exists sID \in \mathfrak{M}(sID^*) : \text{revSK}[sID] = \text{true})$	07 $\mathfrak{P}(sID^*) := \text{PARTIALMATCH}(sID^*)$
02 return false	08 if $ \mathfrak{M}(sID^*) > 1$ or $ \mathfrak{P}(sID^*) \geq 1$ return true
03 if $\exists sID \in \mathfrak{M}(sID^*)$ s. t. $sID \in \mathcal{S}_{\text{test}}$	09 for $\text{attack} \in \text{Table 2}$ and $\text{attack} \in [28, \text{Table 4}]$
04 return false	10 if $\text{attack} = \text{true}$ return true
05 return true	11 return false
<i>//matching sessions</i>	
12 $\mathfrak{M}(sID^*) := \{sID \mid (\text{Init}[sID], \text{Resp}[sID]) = (\text{Init}[sID^*], \text{Resp}[sID^*]) \wedge (\text{Msg}_{I,1}[sID], \text{Msg}_R[sID],$ $\text{Msg}_{I,2}[sID]) = (\text{Msg}_{I,1}[sID^*], \text{Msg}_R[sID^*], \text{Msg}_{I,2}[sID^*]) \wedge \text{Type}[sID] \neq \text{Type}[sID^*]\}$	
13 return $\mathfrak{M}(sID^*)$	
<i>//partially matching sessions</i>	
14 $\mathfrak{P}(sID^*) := \{sID \mid (\text{Init}[sID], \text{Resp}[sID]) = (\text{Init}[sID^*], \text{Resp}[sID^*]) \wedge (\text{Msg}_{I,1}[sID], \text{Msg}_R[sID]) =$ $(\text{Msg}_{I,1}[sID^*], \text{Msg}_R[sID^*]) \wedge \text{Type}[sID] \neq \text{Type}[sID^*] \wedge \text{Type}[sID] = \text{"Re"}\}$	
15 return $\mathfrak{P}(sID^*)$	

Fig. 3. Helper procedures FRESH and VALID for games IND-FS and IND-FS-St defined in Fig. 2. Procedure FRESH checks if the adversary performed some trivial attack. In procedure VALID, each attack is evaluated by the set of variables in attack tables and checks if an allowed attack was performed. The attack table for IND-FS is shown in Table 2 and the table for game IND-FS-St is given in [28, Table 4], where the latter includes session-state reveal attacks. If the values of the variables are set as in the corresponding row, the attack was performed, i. e. $\text{attack} = \text{true}$, and thus the session is valid.

- $\text{Msg}_{I,1}[sID]$ denotes the first message that was computed by the initiator.
- $\text{Msg}_R[sID]$ denotes the message that was computed by the responder.
- $\text{Msg}_{I,2}[sID]$ denotes the final message that was computed by the initiator.
- $\text{ST}[sID]$ denotes the (secret) state information, i. e. ephemeral secret keys.
- $\text{SK}[sID]$ denotes the session key. If the session terminates without a valid session key, we set this variable to the special string “reject”.

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 the second one of type “Re”. In order to complete a session, oracles DER_I and DER_R have to be queried. The adversary has also access to oracles CORR , REVEAL and REV-STATE to obtain secret information. (The latter is only available if state reveal attacks are considered.) We use the following boolean values to keep track of which queries the adversary made:

- $\text{cor}[n]$ denotes whether the long-term secret key of party P_n was given to the adversary.
- $\text{peerPreCor}[sID]$ denotes whether the peer of the session was corrupted and its long-term key was given to the adversary *before* the owner’s session key was computed, which is important for forward security.
- $\text{revST}[sID]$ denotes whether the session state was given to the adversary.
- $\text{revSK}[sID]$ denotes whether the session key was given to the adversary.

The adversary can forward messages between sessions or modify them. By that, we can define the relationship between two sessions:

- **Matching Session:** Two sessions sID and sID' *match* if the same parties are involved, the messages sent and received are the same they are of different types (cf. line 12 in Fig. 3).
- **Partially Matching Session:** A session sID has a *partially matching* session sID' if the same parties are involved, the messages sent and received are the same without considering the last message and they are of different types, where sID' is of type “Re” (cf. line 14 in Fig. 3).

Finally, the adversary is given access to oracle TEST which can be queried multiple times and which will return either the session key of the specified session or a uniformly random key. We use one bit b for all queries, and store test sessions in a set $\mathcal{S}_{\text{test}}$. For each test session, we require that the adversary does not issue queries such that the session key can be trivially computed. In Fig. 3 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. For game IND-FS-St, we capture this with attacks listed in our full paper [28, Table 4]. For game IND-FS, we use Table 2 to capture valid attacks.

If the protocol does not use appropriate randomness, it should not be considered secure. In this case, there can be multiple matching sessions to a test session, which an adversary can take advantage of. We capture this as part of the validity property (cf. line 08). For an honest run of the protocol, the underlying min-entropy ensures that this attack will only happen with negligible probability.

We define validity of different attack strategies in Table 2, using variables to indicate which queries the adversary may (not) make. The purpose is to make our proofs precise by listing all the possible and non-trivial attacks. Attacks covered in the IND-FS model capture *forward secrecy* (FS) and *key compromise impersonation* (KCI) attacks. We provide a more detailed description of Table 2 and the full table for IND-FS-St in our full paper [28, Appendix B]. For all test sessions, at least one attack has to evaluate to true. Then, the adversary wins if he distinguishes the session keys from uniformly random keys which he obtains through queries to the TEST oracle.

Definition 2 (Key Indistinguishability of AKE). *We define games IND-FS and IND-FS-St as in Figs. 2 and 3. We say AKE is $(t, \varepsilon, \mu, S, T, Q_{\text{COR}})$ -IND-FS-secure resp. $(t', \varepsilon', \mu, S, T, Q_{\text{COR}}, Q_{\text{ST}})$ -IND-FS-St-secure if for all adversaries A attacking the protocol in time t resp. t' with μ users, S sessions, T test queries, Q_{COR} corruptions, and Q_{ST} state reveals, we have*

$$\left| \Pr[\text{IND-FS}_{\text{AKE}}^A \Rightarrow 1] - \frac{1}{2} \right| \leq \varepsilon \quad \text{resp.} \quad \left| \Pr[\text{IND-FS-St}_{\text{AKE}}^A \Rightarrow 1] - \frac{1}{2} \right| \leq \varepsilon' .$$

Table 2. Table of attacks for adversaries against three-message protocols with FS. 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 and **F** means “false”. This table is obtained from [28, Table 3] by excluding all trivial attacks.

	\mathcal{A} gets (Initiator, Responder)	$\text{cor}[i^*]$	$\text{cor}[r^*]$	$\text{peerPreCor}[sID^*]$	$\text{Type}[sID^*]$	$ \mathfrak{M}(sID^*) $	$ \mathfrak{P}(sID^*) $
1	(long-term, long-term)	–	–	–	“In”	–	1
2	(long-term, long-term)	–	–	–	“Re”	1	–
5	(long-term, long-term)	–	–	F	“In”	–	0
6	(long-term, long-term)	–	–	F	“Re”	0	–

Note that if there exists a session which is neither fresh nor valid, the game outputs the bit b , which implies that $\Pr[\text{IND-FS}_{\text{AKE}}^{\mathcal{A}} \Rightarrow 1] = \frac{1}{2}$ or $\Pr[\text{IND-FS-St}_{\text{AKE}}^{\mathcal{A}} \Rightarrow 1] = \frac{1}{2}$, giving the adversary an advantage equal to 0. This captures that an adversary will not gain any advantage by performing a trivial attack.

4 Verifiable Authenticated Key Exchange

To build a tightly secure three-message AKE protocol with key confirmation from a two-message AKE protocol, we define two security notions of the two-message protocol: The first one is One-Way against key Verification attacks and weak Forward Secrecy, or OW-VwFS for short, and the second one is OW-VwFS with state-reveal attacks, or OW-VwFS-St for short.

We define the syntax of a two-message key exchange protocol in a similar fashion as the three-message AKE. Let $\text{AKE}' := (\text{Setup}', \text{Gen}', \text{Init}'_I, \text{Init}'_R, \text{Der}'_I)$, where Setup' , Gen' and Init'_I are defined exactly as in the three-message protocol. Instead of a state, the responder’s algorithm Init'_R computes a session key K . The initiator’s algorithm Der'_I does not output a second message, but only the session key. Correctness is defined similarly to the three-message case.

Definition 3 (Correctness of two-message AKE). Let $\text{AKE}' := (\text{Setup}', \text{Gen}', \text{Init}'_I, \text{Init}'_R, \text{Der}'_I)$ be an AKE protocol. We say AKE' is ρ -correct if

$$\Pr \left[\begin{array}{l} \text{par}' \leftarrow \text{Setup}'(1^\lambda), \\ (\text{pk}_i, \text{sk}_i) \leftarrow \text{Gen}'(\text{par}), (\text{pk}_r, \text{sk}_r) \leftarrow \text{Gen}'(\text{par}), \\ (m_i, \text{st}_i) \leftarrow \text{Init}'_I(\text{sk}_i, \text{pk}_r), \\ (m_r, K_r) \leftarrow \text{Init}'_R(\text{sk}_r, \text{pk}_i, m_i), \\ K_i := \text{Der}'_I(\text{sk}_i, \text{pk}_r, m_r, \text{st}_i) \end{array} \right] \geq \rho,$$

where the probability is taken over the randomness of Setup' , Gen' , Init'_I , and Init'_R .

GAME OW-VwFS, [OW-VwFS-St]	SESSION' _i ((i, r) ∈ [μ] ²)
00 for n ∈ [μ]	20 cnt _S ++
01 (pk _n , sk _n) ← Gen'	21 sID := cnt _S
02 (sID*, k*) ← A ^O (pk ₁ , ..., pk _μ)	22 (Init[sID], Resp[sID]) := (i, r)
03 if sID* > cnt _S or VALID(sID*) = false	23 Type[sID] := "In"
04 return 0	24 (m _i , st _i) ← Init' _i (sk _i , pk _r)
05 return KVER(sID*, k*)	25 (Msg[sID], ST[sID]) := (m _i , st _i)
DER' _R ((i, r) ∈ [μ] ² , m _i)	26 return (sID, m _i)
06 cnt _S ++	DER' _i (sID ∈ [cnt _S], m _r)
07 sID := cnt _S	27 if SK[sID] ≠ ⊥ or Type[sID] ≠ "In"
08 (Init[sID], Resp[sID]) := (i, r)	28 return ⊥ // no re-use
09 Type[sID] := "Re"	29 (i, r) := (Init[sID], Resp[sID])
10 (m _r , k) ← Init' _R (sk _r , pk _i , m _i)	30 st _i := ST[sID]
11 (Msg[sID], Msg _R [sID]) := (m _i , m _r)	31 k := Der' _i (sk _i , pk _r , m _r , st _i)
12 SK[sID] := k	32 (Msg _R [sID], SK[sID]) := (m _r , k)
13 return (sID, m _r)	33 return ε
CORR'(n ∈ [μ])	VALID'(sID*) // Helper procedure
14 cor[n] := true	34 (i, r) := (Init[sID], Resp[sID])
15 return sk _n	35 if Type[sID*] = "In"
KVER(sID, k)	and revST[sID*] = false
16 if k = ⊥ return ⊥	36 if cor[r] = false or ∃(sID*) ≠ ∅
17 return [SK[sID] = k]	37 return true
REV-STATE'(sID)	38 if Type[sID*] = "Re"
18 revST[sID] := true	39 if cor[i] = false or ∃(sID*) ≠ ∅
19 return ST[sID]	40 return true
	41 return false

Fig. 4. Games OW-VwFS (without dashed boxes) and OW-VwFS-St (including dashed boxes) for AKE'. \mathcal{A} has access to oracles $\mathcal{O} := \{\text{SESSION}'_i, \text{DER}'_R, \text{DER}'_i, \text{CORR}', \text{KVER}\}$. In OW-VwFS-St, \mathcal{A} also has access to $\text{REV-STATE}'$. In two-message AKE, responder sessions do not have state. So, $\text{REV-STATE}'(\text{sID})$ will return \perp if sID is a responder session. Further, partially matching session is defined as $\mathfrak{P}(\text{sID}^*) := \{\text{sID} \mid \text{Type}[\text{sID}] = \text{"In"} \wedge (\text{Init}[\text{sID}], \text{Resp}[\text{sID}]) = (\text{Init}[\text{sID}^*], \text{Resp}[\text{sID}^*]) \wedge \text{Msg}_i[\text{sID}] = \text{Msg}_i[\text{sID}^*]\}$.

OW-VwFS is similar to the standard weak forward secrecy, but an adversary is additionally allowed to check if a key corresponds to some generated transcripts. The security notion OW-VwFS-St, based on OW-VwFS, allows the adversary to reveal session states. Moreover, these two security notions do not have REVEAL and TEST oracles. Our notion is motivated by the one-wayness against honest and key verification attacks in [27], but it is stronger in the sense that it allows active attacks. These are formally defined by Definition 4 with security games OW-VwFS and OW-VwFS-St as in Fig. 4.

Definition 4 (OW-VwFS and OW-VwFS-St security). A two-message authenticated key exchange protocol AKE' is $(t, \varepsilon, \mu, S, Q_{\text{COR}}, Q_{\text{VER}})$ -OW-VwFS secure resp. $(t', \varepsilon', \mu, S, Q_{\text{COR}}, Q_{\text{VER}}, Q_{\text{ST}})$ -OW-VwFS-St secure, where μ is the number of users, S is the number of sessions, Q_{VER} is the number of calls to KVER and Q_{ST} is the number of calls to REV-STATE', if for all adversaries \mathcal{A} attacking the protocol in time at most t resp. t' , we have

$$\Pr[\text{OW-VwFS}_{\text{AKE}'}^{\mathcal{A}} \Rightarrow 1] \leq \varepsilon \quad \text{resp.} \quad \Pr[\text{OW-VwFS-St}_{\text{AKE}'}^{\mathcal{A}} \Rightarrow 1] \leq \varepsilon'.$$

Valid attacks are defined via VALID' . For the session sid^* for which the adversary aims to compute the session key, we basically allow two types of attacks: If there is a (partially) matching session, then both parties may be corrupted. Otherwise, the adversary must not corrupt the peer of the session. Additionally for the model with state reveal attacks, the state for sid^* must not be revealed in any case.

MIN-ENTROPY. We require that public keys have γ bits of min-entropy, i. e., for all $(\text{pk}_0, \text{sk}_0) \leftarrow \text{Gen}'$, $(\text{pk}_1, \text{sk}_1) \leftarrow \text{Gen}'$, we have $\Pr[\text{pk}_0 = \text{pk}_1] \leq 2^{-\gamma}$. Similarly, we require that messages have α bits of min-entropy, i. e., for all messages m' we have $\Pr[m = m'] \leq 2^{-\alpha}$, where m is output by either Init'_I or Init'_R .

5 AKE with Key Confirmation

We now build a three-message AKE protocol AKE_{KC} with key confirmation from a two-message AKE protocol AKE' and three hash functions $\text{G}_I, \text{G}_R, \text{H}$. An overview is given in Fig. 5. Hash functions G_I, G_R and H are defined as follows: $\text{G}_I, \text{G}_R : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ and $\text{H} : \{0, 1\}^* \rightarrow \mathcal{K}$, where λ is the length of key confirmation tags and \mathcal{K} is the key space of AKE_{KC} .³

Let $\text{AKE}' = (\text{Setup}', \text{Gen}', \text{Init}'_I, \text{Init}'_R, \text{Der}'_I)$. We define AKE_{KC} as follows: $\text{Setup}, \text{Gen}_{\text{AKE}}, \text{Init}_I$ will be the same as $\text{Setup}', \text{Gen}'$ and Init'_I , respectively. Init_R first runs Init'_R to obtain the responder's message m_r and the key k of AKE' , where the latter is used to derive the final session key and key confirmation messages. In particular, the responder first computes the key confirmation tag $\pi_r := \text{G}_R(k, \text{ctxt})$, where ctxt is defined as the two parties' public keys and the initial messages (cf. Fig. 5). It then also computes the expected key confirmation tag π'_i and session key K' using G_I and H on the same input. It sends (m_r, π_r) to the initiator and keeps (π'_i, K') as state. The initiator runs Der_I which is defined as follows: First, it runs Der'_I to get k and then performs the same computations as the responder to compute key confirmation tags π_i, π'_r and the final session key K . It accepts K if $\pi'_r = \pi_r$ and sends π_i as the final message. The responder's derivation algorithm Der_R checks whether the key confirmation tag is valid, i. e., $\pi_i = \pi'_i$, and if this is the case it sets the session key to K' .

Whenever an equality check fails or the underlying algorithms of AKE' return \perp , the parties terminate the session, i. e., they *reject*, and return \perp .

CORRECTNESS. The correctness of AKE_{KC} follows directly from the correctness of AKE' . In particular, if AKE' is $(1 - \delta)$ -correct, then so is AKE_{KC} .

SECURITY. We prove IND-FS security of AKE_{KC} based on OW-VwFS security of AKE' and modeling G_I, G_R and H as random oracles.

Theorem 1. *Let AKE' be $(1 - \delta)$ -correct and have public keys with γ bits of entropy and messages with α bits of entropy. Let AKE_{KC} be as defined in Fig. 5, where $\text{G}_I, \text{G}_R : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ and $\text{H} : \{0, 1\}^* \rightarrow \mathcal{K}$*

³ We define three different hash functions here which allows us to model them as independent random oracles. When instantiating the hash functions with the same function, one would need to use appropriate domain separation.

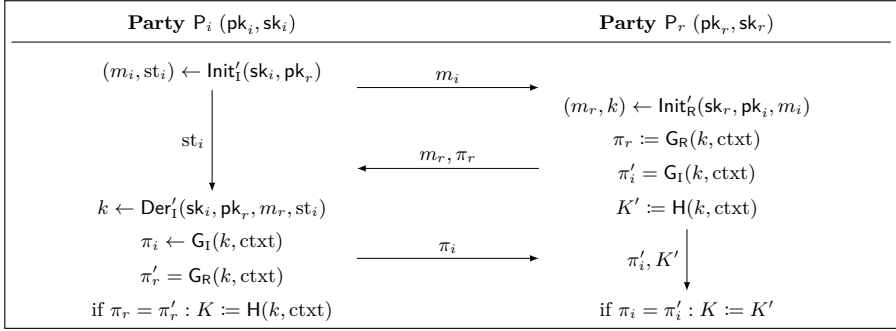


Fig. 5. AKE protocol AKE_{KC} from AKE' and key confirmation. The context is defined as $\text{ctxt} := (pk_i, pk_r, m_i, m_r)$. G_I , G_R and H are independent random oracles.

are modelled as random oracles. For every adversary \mathcal{A} that breaks the $(t, \varepsilon, \mu, S, T, Q_{\text{COR}})$ -IND-FS-security of AKE_{KC} , there exists an adversary \mathcal{B} that breaks the $(t', \varepsilon', \mu, S, Q_{\text{COR}}, Q_{\text{Ver}})$ -OW-VwFS security of AKE' with $t' \approx t$ and

$$\varepsilon \leq \varepsilon' + 2S \cdot \delta + (S + S^2) \cdot 2^{-\lambda} + \mu^2 \cdot 2^{-\gamma} + S(S + Q_{G_I} + Q_{G_R} + Q_H) \cdot 2^{-\alpha},$$

where Q_{G_I} , Q_{G_R} and Q_H are the number of queries to random oracles G_I , G_R and H and $Q_{\text{Ver}} \leq S + Q_{G_I} + Q_{G_R} + Q_H$.

The idea of the proof is that we can simulate the key confirmation tags and session keys without knowing the key k of the underlying two-message protocol as long as it has not been queried to (one of) the random oracles. For this we have to keep track of whether the adversary trivially knows k because the session is not fresh anymore. We can handle this case and still simulate correctly using KVER oracle. The only way to win the game is to compute k for a fresh and valid session, thus breaking one-wayness of the underlying protocol. We now prove the theorem formally.

Proof. Let \mathcal{A} be an adversary against IND-FS security of AKE_{KC} . We consider the sequence of games G_0 - G_3 in Figs. 6 and 7.

GAME G_0 . The first game G_0 is the original IND-FS security game, however we exclude that public keys or messages collide (which means that if such events happen, then the game will abort and return a random bit). This also includes the key confirmation tags. Thus we get

$$\Pr[G_0^{\mathcal{A}} \Rightarrow 1] \leq \Pr[\text{IND-FS}_{\text{AKE}_{\text{KC}}}^{\mathcal{A}} \Rightarrow 1] + \mu^2 \cdot 2^{-\gamma} + S^2 \cdot 2^{-\alpha} + S^2 \cdot 2^{-\lambda}.$$

Note that this means there can be at most one (partially) matching session for each session.

GAME G_1 . In game G_1 , we want to ensure that π_r , π_i and K have not been queried to the respective random oracle before they are determined. Note that

GAMES G_0 - G_3		$\text{SESSION}_I((i, r) \in [\mu]^2)$
00 for $n \in [\mu]$		38 cnts ++
01 $(pk_n, sk_n) \leftarrow \text{Gen}'$		39 $\text{sID} := \text{cnts}$
02 $b \stackrel{\$}{\leftarrow} \{0, 1\}$		40 $(\text{Init}[\text{sID}], \text{Resp}[\text{sID}]) := (i, r)$
03 $b' \leftarrow \mathcal{A}^O(pk_1, \dots, pk_\mu)$		41 $\text{Type}[\text{sID}] := \text{"In"}$
04 for $\text{sID}^* \in \mathcal{S}_{\text{test}}$		42 $(m_i, st_i) \leftarrow \text{Init}'_I(sk_i, pk_r)$
05 if $\text{FRESH}(\text{sID}^*) = \text{false}$		43 $(\text{Msg}_{G_1}[\text{sID}], \text{ST}[\text{sID}]) := (m_i, st_i)$
or $\text{VALID}(\text{sID}^*) = \text{false}$		44 $\text{ctxt}[\text{sID}] := (pk_i, pk_r, m_i, \perp)$ // G_2 - G_3
return b		45 return (sID, m_i)
06 return $\llbracket b = b' \rrbracket$		
$\text{SESSION}_R((i, r) \in [\mu]^2, m_i)$		$\text{DER}_I(\text{sID} \in [\text{cnts}], (m_r, \pi_r))$
08 cnts ++		46 if $\text{SK}[\text{sID}] \neq \perp$ or $\text{Type}[\text{sID}] \neq \text{"In"}$ return \perp
09 $\text{sID} := \text{cnts}$		47 $(i, r) := (\text{Init}[\text{sID}], \text{Resp}[\text{sID}])$
10 $(\text{Init}[\text{sID}], \text{Resp}[\text{sID}]) := (i, r)$		48 $st_i := \text{ST}[\text{sID}]$
11 $\text{Type}[\text{sID}] := \text{"Re"}$		49 $\text{peerPreCor}[\text{sID}] := \text{cor}[r]$
12 $(m_r, k) \leftarrow \text{Init}'_R(sk_r, pk_i, m_i)$		50 $k := \text{Der}'_I(sk_i, pk_r, m_r, st_i)$
13 if $k = \perp$		51 if $k = \perp$
14 $\text{SK}[\text{sID}] := \text{"reject"}$		52 $\text{SK}[\text{sID}] := \text{"reject"}$
15 return \perp		53 return \perp
16 $\pi_r := G_R(k, pk_i, pk_r, m_i, m_r)$ // G_0		54 if $\pi_r \neq G_R(k, pk_i, pk_r, m_i, m_r)$ // G_0 - G_1
17 $\pi_i := G_I(k, pk_i, pk_r, m_i, m_r)$ // G_0		55 $\text{SK}[\text{sID}] := \text{"reject"}$ // G_0 - G_1
18 $K := H(k, pk_i, pk_r, m_i, m_r)$ // G_0		56 return \perp // G_0 - G_1
19 if $\exists k' \text{ s.t. } G_R[k', pk_i, pk_r, m_i, m_r] \neq \perp$ // G_1 - G_3		57 $\pi_i := G_I(k, pk_i, pk_r, m_i, m_r)$ // G_0 - G_1
or $G_I[k', pk_i, pk_r, m_i, m_r] \neq \perp$ // G_1 - G_3		58 $K := H(k, pk_i, pk_r, m_i, m_r)$ // G_0 - G_1
or $H[k', pk_i, pk_r, m_i, m_r] \neq \perp$ // G_1 - G_3		59 $k[\text{sID}] := k$ // G_2 - G_3
20 $\text{BadEntropy} := \text{true}; \text{abort}$ // G_1 - G_3		60 Replace \perp in $\text{ctxt}[\text{sID}]$ with m_r // G_2 - G_3
21 $\pi_r \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda, \pi_i \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda, K \stackrel{\$}{\leftarrow} \mathcal{K}$ // G_1 - G_3		61 if $\exists \text{sID}' \text{ s.t. } \text{ctxt}[\text{sID}'] = \text{ctxt}[\text{sID}]$ // G_2 - G_3
22 $G_R[k, pk_i, pk_r, m_i, m_r] := \pi_r$ // G_1		62 if $\pi_r \neq G_R[\diamond, pk_i, pk_r, m_i, m_r]$ // G_2 - G_3
23 $G_I[k, pk_i, pk_r, m_i, m_r] := \pi_i$ // G_1		63 $\text{SK}[\text{sID}] := \text{"reject"}$ // G_2 - G_3
24 $H[k, pk_i, pk_r, m_i, m_r] := K$ // G_1		64 return \perp // G_2 - G_3
25 $\text{ctxt}[\text{sID}] := (pk_i, pk_r, m_i, m_r)$ // G_2 - G_3		65 $\pi_i := G_I[\diamond, pk_i, pk_r, m_i, m_r]$ // G_2 - G_3
26 $k[\text{sID}] := k$ // G_2 - G_3		66 $K := H[\diamond, pk_i, pk_r, m_i, m_r]$ // G_2 - G_3
27 if $\exists \text{sID}' \text{ s.t. } \text{ctxt}[\text{sID}'] = (pk_i, pk_r, m_i, \perp)$ // G_2 - G_3		67 else // G_2 - G_3
or $\text{cor}[i] = \text{false}$ // G_2 - G_3		68 if $G_R[k, pk_i, pk_r, m_i, m_r] = \pi_r$ // G_2 - G_3
28 $G_R[\diamond, pk_i, pk_r, m_i, m_r] := \pi_r$ // G_2 - G_3		69 if $\text{cor}[r] = \text{false}$ // G_3
29 $G_I[\diamond, pk_i, pk_r, m_i, m_r] := \pi_i$ // G_2 - G_3		70 $\text{QueryRO} := \text{true}; \text{abort}$ // G_3
30 $H[\diamond, pk_i, pk_r, m_i, m_r] := K$ // G_2 - G_3		71 $\pi_i := G_I(k, pk_i, pk_r, m_i, m_r)$ // G_2 - G_3
31 else // G_2 - G_3		72 $K := H(k, pk_i, pk_r, m_i, m_r)$ // G_2 - G_3
32 $G_R[\oplus, pk_i, pk_r, m_i, m_r] := \pi_r$ // G_2 - G_3		73 else // G_2 - G_3
33 $G_I[\oplus, pk_i, pk_r, m_i, m_r] := \pi_i$ // G_2 - G_3		74 $G_R[\diamond, pk_i, pk_r, m_i, m_r] \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$ // G_2 - G_3
34 $H[\oplus, pk_i, pk_r, m_i, m_r] := K$ // G_2 - G_3		75 if $\pi_r = G_R[\diamond, pk_i, pk_r, m_i, m_r]$ // G_2 - G_3
35 $(\text{Msg}_{G_1}[\text{sID}], \text{Msg}_R[\text{sID}]) := (m_i, (m_r, \pi_r))$ // G_2 - G_3		76 $\text{RandKC} := \text{true}; \text{abort}$ // G_2 - G_3
36 $\text{ST}[\text{sID}] := (\pi_i, K)$ // G_2 - G_3		77 $\text{SK}[\text{sID}] := \text{"reject"}$ // G_2 - G_3
37 return $(\text{sID}, (m_r, \pi_r))$		78 return \perp // G_2 - G_3
		79 $(\text{Msg}_R[\text{sID}], \text{Msg}_{G_1,2}[\text{sID}]) := (m_r, \pi_i)$
		80 $\text{SK}[\text{sID}] := K$
		81 return π_i

Fig. 6. Games G_0 - G_3 for the proof of Theorem 1. \mathcal{A} has access to oracles $O := \{\text{SESSION}_I, \text{SESSION}_R, \text{DER}_I, \text{DER}_R, \text{REVEAL}, \text{CORR}, \text{TEST}, G_I, G_R, H\}$. Helper procedures FRESH and VALID are defined in Fig. 3.

all three values will be determined in SESSION_R when m_r and k are computed. Thus, whenever SESSION_R is queried, we check whether there already exists a query $(k', pk_i, pk_r, m_i, m_r)$ to G_R, G_I or H for some k' (line 19). If this is the case, we raise flag BadEntropy and abort. If BadEntropy is not raised, we draw fresh values for π_r, π_i and K and explicitly assign them to the corresponding

DERR($sID \in [\text{cnts}], \pi_i$)		GR(k, pk_i, pk_r, m_i, m_r)	
00 if $SK[sID] \neq \perp$ or $Type[sID] \neq \text{"Re"}$		27 if $GR[\diamond, pk_i, pk_r, m_i, m_r] = \pi \neq \perp$	//G ₂ -G ₃
01 return \perp		28 $S := \{sID \mid \text{ctxt}[sID] = (pk_i, pk_r, m_i, m_r)\}$	//G ₂ -G ₃
02 $(i, r) := (\text{Init}[sID], \text{Resp}[sID])$		29 for $sID \in S$ // note $ S \leq 2$	//G ₂ -G ₃
03 $(\pi'_i, K') := ST[sID]$		30 if $k[sID] = k$	//G ₂ -G ₃
04 $\text{peerPreCor}[sID] := \text{cor}[i]$		31 QueryRO := true; abort	//G ₃
05 if $\pi_i \neq \pi'_i$	//G ₀ -G ₁	32 return π	//G ₂ -G ₃
06 $SK[sID] := \text{"reject"}$	//G ₀ -G ₁	33 elseif $GR[\oplus, pk_i, pk_r, m_i, m_r] = \pi \neq \perp$	//G ₂ -G ₃
07 return \perp	//G ₀ -G ₁	34 Find sID s.t. $\text{ctxt}[sID] = (pk_i, pk_r, m_i, m_r)$	//G ₂ -G ₃
08 $K := K'$	//G ₀ -G ₁	35 if $k[sID] = k$	//G ₂ -G ₃
09 if $\exists sID'$ s.t. $\text{ctxt}[sID'] = \text{ctxt}[sID]$	//G ₂ -G ₃	36 Replace \oplus with k	//G ₂ -G ₃
10 if $\pi_i \neq G_1[\diamond, pk_i, pk_r, m_i, m_r]$	//G ₂ -G ₃	37 return π	//G ₂ -G ₃
11 $SK[sID] := \text{"reject"}$	//G ₂ -G ₃	38 if $GR[k, pk_i, pk_r, m_i, m_r] = \pi \neq \perp$	
12 return \perp	//G ₂ -G ₃	39 return π	
13 $K := H[\diamond, pk_i, pk_r, m_i, m_r]$	//G ₂ -G ₃	40 $\pi \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$	
14 else	//G ₂ -G ₃	41 $GR[k, pk_i, pk_r, m_i, m_r] := \pi$	
15 if $G_1[k, pk_i, pk_r, m_i, m_r] = \pi_i$	//G ₂ -G ₃	42 return π	
16 if $\text{cor}[i] = \text{false}$	//G ₃	CORR($n \in [n]$)	
17 QueryRO := true; abort	//G ₃	43 $\text{cor}[n] := \text{true}$	
18 $K := H(k, pk_i, pk_r, m_i, m_r)$	//G ₂ -G ₃	44 return sk_n	
19 else	//G ₂ -G ₃	TEST(sID)	
20 $G_1[\diamond, pk_i, pk_r, m_i, m_r] \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$	//G ₂ -G ₃	45 if $sID \in \mathcal{S}_{\text{test}}$ return \perp	
21 if $\pi_i = G_1[\diamond, pk_i, pk_r, m_i, m_r]$	//G ₂ -G ₃	46 if $SK[sID] \in \{\perp, \text{"reject"}\}$ return \perp	
22 RandKC := true; abort	//G ₂ -G ₃	47 $\mathcal{S}_{\text{test}} := \mathcal{S}_{\text{test}} \cup \{sID\}$	
23 $SK[sID] := \text{"reject"}$	//G ₂ -G ₃	48 $K_0^* := SK[sID]$	
24 return \perp	//G ₂ -G ₃	49 $K_1^* \stackrel{\$}{\leftarrow} \mathcal{K}$	
25 $(\text{Msg}_{1,2}[sID], SK[sID]) := (\pi_i, K)$		50 return K_b^*	
26 return ε			

Fig. 7. Oracles for games G_0 - G_3 for the proof of Theorem 1. G_1 and H are defined analogously to G_R .

entry of the respective random oracle (lines 21–24). We make this explicit here to prepare for the next step where we have to do a case distinction. Note that G_0 and G_1 are the same, except if BadEntropy is raised. Thus,

$$|\Pr[K_0^A \Rightarrow 1] - \Pr[K_1^A \Rightarrow 1]| \leq \Pr[\text{BadEntropy}] \leq S(Q_{G_R} + Q_{G_1} + Q_H) \cdot 2^{-\alpha},$$

where we bound the event by the entropy of AKE' . The message m_r is computed by the game directly before we check for this event. We then use the union bound over the maximum number of sessions S .

GAME G_2 . In game G_2 , we want to compute π_r , π_i and K without using k explicitly and prepare for the reduction to OW-VwFS . For this, we have to make a distinction between fresh and non-fresh sessions. We add two additional variables $\text{ctxt}[sID]$ and $k[sID]$ for each session which store the context and the session key of the underlying AKE' . When SESSION_R is queried, we no longer assign the random oracle entries $[k, pk_i, pk_r, m_i, m_r]$. Instead, we use a special placeholder symbol for the key k . In particular, if the session is still fresh and valid (i.e., there exists a session with a matching context up to this point, or the intended peer is not (yet) corrupted), we use the symbol \diamond (lines 28–30). Otherwise, in case the session is not valid, we use the symbol \oplus (lines 32–34). This distinction will be necessary to patch the random oracle correctly. Note that an adversary might be able to compute the correct key k for a non-valid session.

We describe how the random oracles are patched below. First, we explain how to change Der_I and Der_R accordingly. For each query to Der_I , we first update the context with the message m_r that was used to query the oracle. Then we check whether there exists a potential partnered session with the same context (line 61). In this case we know the corresponding values π_r, π_i and K which are stored with the symbol \diamond . We check whether the tag π_r is correct (if not, the session rejects) and assign the session key (lines 62–66). If there is no other session with the same context, we have to make another case distinction. In the first case, we use k explicitly to check whether there has been a query to the random oracle G_R such that the tag π_r matches (line 68). In this case, we proceed normally. Looking ahead, this will be a critical point in the next game modification. If there exists no such query to G_R , then the game makes the query and chooses a tag uniformly at random (line 74). If this tag is the same as the one provided by the adversary, we raise flag RandKC and abort (line 76). Otherwise, the session simply rejects. We modify Der_R in the exact same way, except that we are now looking at π_i (Fig. 7, lines 09–24).

Before bounding RandKC , we explain the simulation of random oracles as described in Fig. 7. We explain G_R in more detail. (G_I and H are modeled in exactly the same way). For each query $(k, \text{pk}_i, \text{pk}_r, m_i, m_r)$, we first check for entries with the special symbol. More specifically, if there exists an entry with the given context and the symbol \diamond , we look for the sID with this context. Note that there can be at most two sessions (one of type “In” and one of type “Re” which will be matching sessions), which we capture by computing a set \mathcal{S} containing the corresponding $\text{sID}(s)$ (line 29). If the key k of the random oracle query corresponds to the one stored in $\mathbf{k}[\text{sID}]^4$, then G_R simply outputs the stored value π (line 32). We do the same for special symbol \oplus (line 34, note that here sID is always unique), except that we also update the entry accordingly, i. e., replace \oplus with k if $k = \mathbf{k}[\text{sID}]$ (line 36). This way, the simulation is consistent with the other oracles.

Overall, the two games only differ when flag RandKC is raised. Note that we can bound the probability that a tag is valid without the random oracle being queried by the length of the tag. Union bound over S session gives us

$$|\Pr[\text{G}_1^A \Rightarrow 1] - \Pr[\text{G}_2^A \Rightarrow 1]| \leq \Pr[\text{RandKC}] \leq S \cdot 2^{-\lambda} + S \cdot \delta.$$

GAME G_3 . In the final game G_3 , we raise flag QueryRO if the adversary ever queries the random oracle on a key k of a fresh session. Depending on the order of queries, this event can occur for different oracles: Der_I (Fig. 6, line 70), Der_R (Fig. 7, line 17) or one of the random oracles (Fig. 7, line 31). First, we look at sessions that do not have a matching session with the same context. For queries to oracle Der_I we check the validity of π_r in line 68. If the peer r is not corrupted,

⁴ Since we cannot check correctness efficiently in the reduction which we will build in the next step, we explicitly perform the test here for all sessions in \mathcal{S} . However, if \mathcal{S} indeed contains two sessions, then by correctness, this key (and thus the outcome) will be the same.

then the session is still fresh and valid. Thus, we raise `QueryRO` if there has been a query to G_R on the correct key and context such that the output is indeed π_r . We proceed similarly for responder sessions when DER_R is queried, checking for queries to G_I . This means that all sessions where the peer is uncorrupted and no session with a matching context exist will reject (or abort).

We now look at sessions that have a session with matching context and whose relevant random oracle entries are marked with \diamond . Whenever one of the random oracles is queried, we check whether the key matches the one stored in $k[sID]$ (as described earlier) and if this is the case, we also raise `QueryRO` and abort.

We claim that

$$|\Pr[G_2^A \Rightarrow 1] - \Pr[G_3^A \Rightarrow 1]| \leq \Pr[\text{QueryRO}] \leq \varepsilon' + S \cdot \delta.$$

Before proving the claim, note that in G_3 we have $\Pr[G_3^A \Rightarrow 1] = 1/2$. For this, observe that all sessions must have a (partially) matching session and that random oracle H is never queried on k for any of those sessions. Thus, the session key is indistinguishable from a uniformly random key.

BOUNDING EVENT `QueryRO`. We now describe an adversary \mathcal{B} against OW-VwFS security of the underlying AKE' to bound event `QueryRO`. A pseudocode description is given in Fig. 8. The idea is that whenever \mathcal{A} queries one of the random oracles on the underlying key k of a fresh and valid session (either in order to forge a key confirmation tag or to distinguish the actual session key), we can use this to break OW-VwFS security of AKE' , where the verification oracle `KVER` is used to simulate the random oracles consistently.

We now describe \mathcal{B} in more detail. It gets as input μ public keys and forwards them to \mathcal{A} . \mathcal{B} simulates queries to oracle `SESSION_I` in a straightforward way by querying its own oracle `SESSION'_I` which returns (sID, m_i) . After assigning the corresponding variables, \mathcal{B} forwards the output to \mathcal{A} . Queries to `SESSION_R` are simulated as in game G_3 . \mathcal{B} first queries DER'_R to receive (sID, m_r) . Instead of checking whether $k = \perp$, \mathcal{B} checks whether $m_r = \perp$. If this is the case, it rejects and outputs \perp . Otherwise, it proceeds as described in G_3 , preparing random oracle assignments by assigning fresh values to π_i , π_r and K and returning $(sID, (m_r, \pi_r))$.

When \mathcal{A} queries DER_I , \mathcal{B} queries DER'_I . \mathcal{B} will not be able to explicitly check whether the session key was computed successfully, however, we will argue that the simulation is consistent by correctness of AKE' and the validity of the key confirmation tag. Thus, \mathcal{B} directly proceeds as described in G_3 . Whenever there exists a session with the same context, then the key confirmation tag must be the same as the one computed by that session, up to correctness of AKE' . Thus the simulation is perfect except with probability $S \cdot \delta$. Whenever there exists no (partially) matching session, \mathcal{B} needs to check whether G_R was already queried on the correct k . For this it checks all random oracle queries that have output π_r provided by \mathcal{A} . If such a query exists, it will be unique since we excluded collisions in the first game. \mathcal{B} checks whether the respective key of the query is the correct key using its oracle `KVER`. If this is the case, we further distinguish two cases, based on whether the session still qualifies for a valid test session

$\mathcal{B}^{\text{SESSION}'_I, \text{DER}'_R, \text{CORR}'_I, \text{KVER}}(\text{pk}_1, \dots, \text{pk}_\mu)$ 00 $b' \leftarrow \mathcal{A}^O(\text{pk}_1, \dots, \text{pk}_\mu)$ 01 return \perp $\text{SESSION}_R((i, r) \in [\mu]^2, m_i)$ 02 $(\text{sID}, m_r) \leftarrow \text{DER}'_R((i, r), m_i)$ 03 $(\text{Init}[\text{sID}], \text{Resp}[\text{sID}]) := (i, r)$ 04 $\text{Type}[\text{sID}] := \text{"Re"}$ 05 if $m_r = \perp$ 06 $\text{SK}[\text{sID}] := \text{"reject"}$ 07 return \perp 08 if $\exists k' \text{ s. t. } \text{G}_R[k', \text{pk}_i, \text{pk}_r, m_i, m_r] \neq \perp$ 09 or $\text{G}_I[k', \text{pk}_i, \text{pk}_r, m_i, m_r] \neq \perp$ 10 or $\text{H}[k', \text{pk}_i, \text{pk}_r, m_i, m_r] \neq \perp$ 11 abort 12 $\pi_r \xleftarrow{\$} \{0, 1\}^\lambda, \pi_i \xleftarrow{\$} \{0, 1\}^\lambda, K \xleftarrow{\$} \mathcal{K}$ 13 $\text{ctxt}[\text{sID}] := (\text{pk}_i, \text{pk}_r, m_i, m_r)$ 14 if $\exists \text{sID}' \text{ s. t. } \text{ctxt}[\text{sID}'] = (\text{pk}_i, \text{pk}_r, m_i, \perp)$ 15 or $\text{cor}[i] = \text{false}$ 16 $\text{G}_R[\diamond, \text{pk}_i, \text{pk}_r, m_i, m_r] := \pi_r$ 17 $\text{G}_I[\diamond, \text{pk}_i, \text{pk}_r, m_i, m_r] := \pi_i$ 18 $\text{H}[\diamond, \text{pk}_i, \text{pk}_r, m_i, m_r] := K$ 19 else 20 $\text{G}_R[\oplus, \text{pk}_i, \text{pk}_r, m_i, m_r] := \pi_r$ 21 $\text{G}_I[\oplus, \text{pk}_i, \text{pk}_r, m_i, m_r] := \pi_i$ 22 $\text{H}[\oplus, \text{pk}_i, \text{pk}_r, m_i, m_r] := K$ 23 $(\text{Msg}_I[\text{sID}], \text{Msg}_R[\text{sID}]) := (m_i, (m_r, \pi_r))$ 24 return $(\text{sID}, (m_r, \pi_r))$ $\text{REVEAL}(\text{sID})$ 25 $\text{revSK}[\text{sID}] := \text{true}$ 26 return $\text{SK}[\text{sID}]$ $\text{CORR}(n \in [\mu])$ 27 $\text{cor}[n] := \text{true}$ 28 $\text{sk}_n \leftarrow \text{CORR}'(n)$ 29 return sk_n	$\text{SESSION}_I((i, r) \in [\mu]^2)$ 27 $(\text{sID}, m_i) \leftarrow \text{SESSION}'_I(i, r)$ 28 $(\text{Init}[\text{sID}], \text{Resp}[\text{sID}]) := (i, r)$ 29 $\text{Type}[\text{sID}] := \text{"In"}$ 30 $\text{Msg}_{i,1}[\text{sID}] := m_i$ 31 $\text{ctxt}[\text{sID}] := (\text{pk}_i, \text{pk}_r, m_i, \perp)$ 32 return (sID, m_i) $\text{DER}_I(\text{sID}, (m_r, \pi_r))$ 33 if $\text{SK}[\text{sID}] \neq \perp$ or $\text{Type}[\text{sID}] \neq \text{"In"}$ 34 return \perp 35 $(i, r) := (\text{Init}[\text{sID}], \text{Resp}[\text{sID}])$ 36 $\text{peerPreCor}[\text{sID}] := \text{cor}[r]$ 37 $\text{DER}'_I(\text{sID}, m_r)$ 38 Replace \perp in $\text{ctxt}[\text{sID}]$ with m_r 39 if $\exists \text{sID}' \text{ s. t. } \text{ctxt}[\text{sID}'] = \text{ctxt}[\text{sID}]$ 40 if $\pi_r \neq \text{G}_R[\diamond, \text{pk}_i, \text{pk}_r, m_i, m_r]$ 41 $\text{SK}[\text{sID}] := \text{"reject"}$ 42 return \perp 43 $\pi_i := \text{G}_I[\diamond, \text{pk}_i, \text{pk}_r, m_i, m_r]$ 44 $K := \text{H}[\diamond, \text{pk}_i, \text{pk}_r, m_i, m_r]$ 45 else 46 if $\exists k \text{ s. t. } \text{G}_R[k, \text{pk}_i, \text{pk}_r, m_i, m_r] = \pi_r$ 47 and $\text{KVER}(k, \text{sID})$ 48 if $\text{cor}[r] = \text{false}$ 49 Stop with (sID, k) 50 $\pi_i := \text{G}_I(k, \text{pk}_i, \text{pk}_r, m_i, m_r)$ 51 $K := \text{H}(k, \text{pk}_i, \text{pk}_r, m_i, m_r)$ 52 else 53 $\text{G}_R[\diamond, \text{pk}_i, \text{pk}_r, m_i, m_r] \xleftarrow{\$} \{0, 1\}^\lambda$ 54 if $\pi_r = \text{G}_R[\diamond, \text{pk}_i, \text{pk}_r, m_i, m_r]$ abort 55 $\text{SK}[\text{sID}] := \text{"reject"}$ 56 return \perp 57 $(\text{Msg}_R[\text{sID}], \text{Msg}_{i,2}[\text{sID}]) := (m_r, \pi_i)$ 58 $\text{SK}[\text{sID}] := K$ 59 return π_i
--	---

Fig. 8. Adversary \mathcal{B} against OW-VwFS. \mathcal{A} has access to oracles $O := \{\text{SESSION}_I, \text{SESSION}_R, \text{DER}_I, \text{DER}_R, \text{REVEAL}, \text{CORR}, \text{TEST}, \text{G}_I, \text{G}_R, \text{H}\}$. Helper procedures FRESH and VALID are defined in Fig. 3. Oracles DER_R , TEST and G_R are defined in Fig. 9, and G_I , H are defined analogously.

or not. If the peer of the session has not been corrupted yet, then this is a valid session and \mathcal{B} outputs (sID, k) as solution in its own game. Otherwise, it proceeds. Oracle DER_R is simulated similarly, looking at G_I instead of G_R .

Oracle REVEAL and CORR can be simulated in a straightforward way. The latter requires \mathcal{B} to query its own oracle CORR' . Queries to TEST will always return the real session key. Note that this is a perfect simulation since session keys are perfectly hidden unless QueryRO happens in which case \mathcal{B} stops because it breaks OW-VwFS security.

It remains to describe the simulation of random oracles. In Fig. 9 we give a description of G_R . G_I and H are simulated in the same way. As in G_3 , \mathcal{B} first checks whether there exists an entry with the special symbol \diamond . If this is the case, it finds the corresponding sID and uses the KVER oracle to check whether the key k provided by \mathcal{A} belongs to this session. Since \diamond is used to mark sessions

<pre> DER_R(sID, π_i) 00 if SK[sID] ≠ ⊥ or Type[sID] ≠ “Re” return ⊥ 01 (i, r) := (Init[sID], Resp[sID]) 02 peerPreCor[sID] := cor[i] 03 if ∃sID' s. t. ctxt[sID'] = ctxt[sID] 04 if π_i ≠ G_I[◊, pk_i, pk_r, m_i, m_r] 05 SK[sID] := “reject” 06 return ⊥ 07 K := H[◊, pk_i, pk_r, m_i, m_r] 08 else 09 if ∃k s. t. G_I[k, pk_i, pk_r, m_i, m_r] = π_i 10 and KVER(k, sID) 11 if peerPreCor[sID] = false 12 Stop with (sID, k) 13 K := H(k, pk_i, pk_r, m_i, m_r) 14 else 15 G_I[◊, pk_i, pk_r, m_i, m_r] ←_{\$} {0, 1}^λ 16 if π_i = G_I[◊, pk_i, pk_r, m_i, m_r] abort 17 SK[sID] := “reject” 18 return ⊥ 19 (Msg_{1,2}[sID], SK[sID]) := (m_{i,2}, K) 20 return ε </pre>	<pre> G_R(k, pk_i, pk_r, m_i, m_r) 20 if G_R[◊, pk_i, pk_r, m_i, m_r] = π ≠ ⊥ 21 S := {sID ctxt[sID] = (pk_i, pk_r, m_i, m_r)} 22 for sID ∈ S 23 if KVER(sID, k) 24 Stop with (sID, k) 25 elseif G_R[⊕, pk_i, pk_r, m_i, m_r] = π ≠ ⊥ 26 Find sID s. t. ctxt[sID] = (pk_i, pk_r, m_i, m_r) 27 if KVER(sID, k) 28 Replace ⊕ with k 29 return π 30 if G_R[k, pk_i, pk_r, m_i, m_r] = π ≠ ⊥ 31 return π 32 π ←_{\$} {0, 1}^λ 33 G_R[k, pk_i, pk_r, m_i, m_r] := π 34 return π TEST(sID) 35 if sID ∈ S_{test} return ⊥ 36 if SK[sID] ∈ {⊥, “reject”} return ⊥ 37 S_{test} := S_{test} ∪ {sID} 38 return SK[sID] </pre>
--	---

Fig. 9. Oracles DER_R, TEST and G_R for adversary \mathcal{B} . G_I and H are defined analogously to G_R.

that have a (partially) matching session, \mathcal{B} can always use this key to win the OW-VwFS game. If there is no entry with \diamond but one with \oplus , \mathcal{B} again queries the KVER oracle, but this time it updates the corresponding entry with the correct key (if KVER returns **true**). This way, \mathcal{B} can perfectly simulate non-test sessions. If none of these cases happen or KVER has returned **false**, then \mathcal{B} proceeds as usual by lazy sampling.

This concludes the description of \mathcal{B} . Note that if QueryRO happens in game G₃, i.e., there exists a random oracle query for a fresh and valid session with correct key k , then \mathcal{B} wins game OW-VwFS. We get $\Pr[\text{QueryRO}] \leq \varepsilon' + S \cdot \delta$.

Further, note that \mathcal{B} issues at most $(S + Q_{G_I} + Q_{G_R} + Q_H)$ to KVER since we have excluded collisions of tags in the first game. The number of queries to all other oracles is preserved. This completes the proof of Theorem 1.

AKE WITH KEY CONFIRMATION AGAINST STATE REVEAL. Based on AKE_{KC}, we build a three-message AKE protocol AKE_{stKC} that is secure against state-reveal attacks (cf. Definition 2). Since AKE_{stKC} has a similar structure with AKE_{KC}, we follow the notations used in defining AKE_{KC} (cf. Fig. 5). An overview of AKE_{stKC} is given in Fig. 10.

AKE_{stKC} uses the state-encryption technique [21, 30] to protect session states. Concretely, let $G_{stI} : \{0, 1\}^\kappa \times \{0, 1\}^\kappa \rightarrow \{0, 1\}^{d_I}$ and $G_{stR} : \{0, 1\}^\kappa \times \{0, 1\}^\kappa \rightarrow \{0, 1\}^{d_R}$ be two hash functions. We assume that any initiator session state of the underlying two-message AKE protocol AKE' can be encoded as a d_I -bit string and $d_R = 2\lambda$ (the length of key confirmation tag plus the length of session key derived by AKE'). AKE_{stKC} proceeds the same as AKE_{KC} except that (1) the long-term secret key of user i in AKE_{stKC} also include a uniformly random key

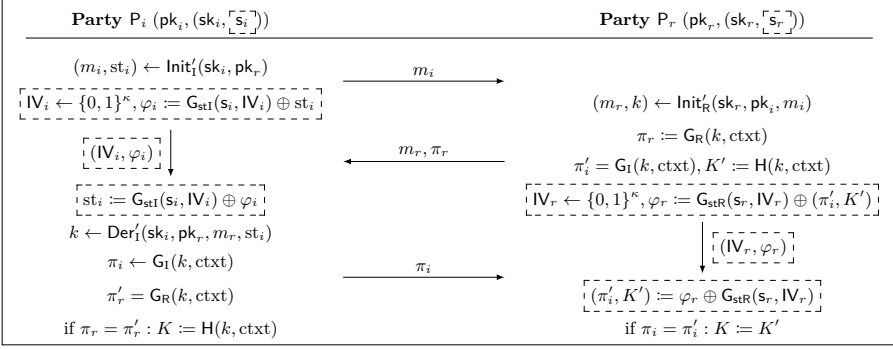


Fig. 10. AKE protocol AKE_{stKC} from AKE' , key confirmation, and state encryption. The context is defined as $\text{ctxt} := (\text{pk}_i, \text{pk}_r, m_i, m_r)$. G_I , G_R , and H are independent random oracles. Dashed parts show how we use the state encryption technique to protect the session states. G_{stI} and G_{stR} are independent random oracles used for state encryption.

$s_i \in \{0, 1\}^\kappa$, and (2) each session will sample a one-time key IV uniformly at random and encrypt the session state of AKE_{KC} via XORing with the one-time pad $\text{G}_{\text{stI}}(s_i, \text{IV})$. Now the session state (that the adversary can reveal in the state-reveal AKE model) is (IV, φ) . Dashed parts in Fig. 10 shows how this technique works.

CORRECTNESS. Similar to AKE_{KC} , the correctness of AKE_{stKC} follows directly from the correctness of AKE' . If AKE' is $(1 - \delta)$ -correct, then so is AKE_{stKC} .

SECURITY. In Theorem 2, we prove IND-FS-St security of AKE_{KC} based on OW-VwFS-St security of AKE' and modeling G_I , G_R , H , G_{stI} , and G_{stR} as random oracles. Here we sketch the proof idea. By using the state encryption technique, the adversary cannot learn the unencrypted states of the underlying two-message AKE , unless it reveals the encrypted session state and corrupts the owner of the session. But this makes the session invalid and thus, it cannot be tested. Therefore, for valid sessions, state-reveal queries do not give any advantage to the adversary, and thus we can use the proof idea of Theorem 1. The full proof of Theorem 2 is postponed to our full version [28, Appendix C].

Theorem 2. *Let AKE' be $(1 - \delta)$ -correct and have public keys with γ bits of entropy and messages with α bits of entropy. Let AKE_{stKC} be as defined in Fig. 5, where $\text{G}_I, \text{G}_R : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$, $\text{H} : \{0, 1\}^* \rightarrow \mathcal{K}$, $\text{G}_{\text{stI}} : \{0, 1\}^\kappa \times \{0, 1\}^\kappa \rightarrow \{0, 1\}^{d_I}$, and $\text{G}_{\text{stR}} : \{0, 1\}^\kappa \times \{0, 1\}^\kappa \rightarrow \{0, 1\}^{d_R}$ are modeled as random oracles. For every adversary \mathcal{A} that breaks the $(t, \varepsilon, \mu, S, T, Q_{\text{COR}}, Q_{\text{ST}})$ - IND-FS-St -security of AKE_{stKC} , there exists an adversary \mathcal{B} that breaks the $(t', \varepsilon', \mu, S, Q_{\text{COR}}, Q_{\text{Ver}}, S)$ - OW-VwFS-St security of AKE' with*

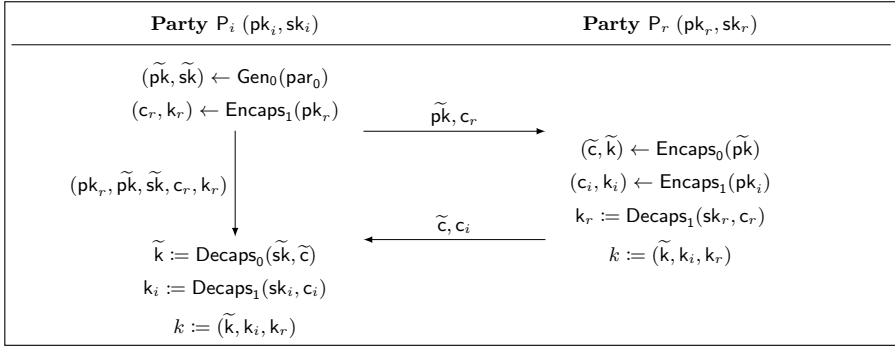


Fig. 11. AKE protocol AKE'_{kem} from KEM schemes KEM_1 , KEM_0 .

$t' \approx t$ and

$$\begin{aligned} \varepsilon \leq \varepsilon' + 2S \cdot \delta + (\mu^2 + S^2 + \mu Q_{G_{\text{stl}}} + 2SQ_{G_{\text{stl}}}) \cdot 2^{-\kappa} \\ + \mu^2 \cdot 2^{-\gamma} + (S + S^2) \cdot 2^{-\lambda} + (Q_{G_R} + Q_{G_I} + Q_H + S) \cdot S \cdot 2^{-\alpha}, \end{aligned}$$

where Q_h is the number of queries to the respective random oracle h and $Q_{Ver} \leq S + Q_{G_I} + Q_{G_R} + Q_H$.

6 Applying Our Results to Existing Protocols

We first show how to construct verifiable AKE from KEMs which gives us tight AKE with key confirmation and perfect forward secrecy from lattices and DDH. The advantage is that we do not have to consider random oracles and the proofs are comparably simpler than those of full AKE security. We then show how we can recover the optimal tightness bound for the CCGJJ protocol [10] using our modular transformation rather than that of [18].

6.1 AKE from KEMs

We provide results for KEM-based AKE secure without and with state reveal, where the former allows for weaker assumptions. The protocol, denoted by AKE'_{kem} , to which we want to apply our compiler from the previous section is given in Fig. 11. Each party holds long-term keys of a KEM scheme KEM_1 and in each session, an ephemeral key using KEM_0 is exchanged. The session key then simply consists of three KEM keys. We also denote its variant with key confirmation by AKE_{kem} (i. e., combining Fig. 11 with Fig. 5) and the one resisting state reveals by $\text{AKE}_{\text{st,kem}}$ (i. e., combining Fig. 11 with Fig. 10).

ONE-WAY SECURITY OF KEM. Depending on whether the KEM is used for long-term keys or ephemeral keys and whether state reveals are allowed,

we need a different variant of one-way security: multi-user one-way security under plaintext checking and ciphertext validity attacks without (OW-PCVA) or with corruptions (OW-PCVA-C), and with corruptions and reveal queries (OW-PCVA-CR). These notions are weaker variants of OW-ChCCA security from Pan, Wagner and Zeng [30] and are tightly implied by OW-ChCCA. The formal security definitions are given in our full version [28, Appendix A].

ANALYSIS OF AKE'_{kem} AND AKE_{kem} . We prove that AKE'_{kem} protocol is a secure verifiable AKE protocol. When not considering state reveals, we can use weaker assumptions, namely OW-PCVA and OW-PCVA-C. We also prove security with state-reveals which uses the definition of OW-PCVA-CR security. We then apply Theorem 1 resp. Theorem 2 to obtain AKE_{kem} which has full forward secrecy.

CORRECTNESS AND ENTROPY. Let KEM_0 be $(1 - \delta_0)$ -correct and have public keys with γ_0 bits of entropy and messages with α_0 bits of entropy. Let KEM_1 be $(1 - \delta_1)$ -correct and have public keys with γ_1 bits of entropy and messages with α_1 bits of entropy. Then AKE'_{kem} is $(1 - \delta_0 - 2\delta_1)$ -correct. Further, AKE'_{kem} has public keys with γ_1 bits of entropy and messages with at least $\min(\gamma_0, \alpha_0, \alpha_1 - 1)$ bits of entropy.

We now establish OW-VwFS and OW-VwFS-St security of AKE'_{kem} and defer the proofs to our full paper [28, Appendix D.1].

Lemma 1. *For every adversary \mathcal{A} that breaks the $(t, \varepsilon, \mu, S, T, Q_{\text{COR}}, Q_{\text{VER}})$ -OW-VwFS security of AKE'_{kem} , there exist adversaries \mathcal{B}_1 and \mathcal{B}_2 that break $(t_1, \varepsilon_1, S, S, S, Q_{\text{VER}})$ -OW-PCVA security of KEM_0 and $(t_2, \varepsilon_2, \mu, S, S, 2Q_{\text{VER}}, Q_{\text{COR}})$ -OW-PCVA-C security of KEM_1 with $t_1 \approx t_2 \approx t$ and $\varepsilon \leq \varepsilon_1 + \varepsilon_2$.*

Lemma 2. *For every adversary \mathcal{A} that breaks the $(t, \varepsilon, \mu, S, T, Q_{\text{COR}}, Q_{\text{VER}}, Q_{\text{ST}})$ -OW-VwFS-St security of AKE'_{kem} , there exist adversaries \mathcal{B}_1 and \mathcal{B}_2 that break $(t_1, \varepsilon_1, S, S, S, Q_{\text{VER}}, Q_{\text{ST}})$ -OW-PCVA-C security of KEM_0 and $(t_2, \varepsilon_2, \mu, S, S, 2Q_{\text{VER}}, Q_{\text{COR}}, Q_{\text{ST}})$ -OW-PCVA-CR security of KEM_1 with $t_1 \approx t_2 \approx t$ and $\varepsilon \leq \varepsilon_1 + \varepsilon_2$.*

We now add key confirmation to AKE'_{kem} as described in Fig. 5 resp. Fig. 10. The following theorem then follows from combining Theorem 1 with Lemma 1 resp. Theorem 2 with Lemma 2.

Theorem 3. *Let KEM_0 be $(1 - \delta_0)$ -correct and have public keys with γ_0 bits of entropy and messages with α_0 bits of entropy. Let KEM_1 be $(1 - \delta_1)$ -correct and have public keys with γ_1 bits of entropy and messages with α_1 bits of entropy. Let AKE_{kem} resp. $\text{AKE}_{\text{st,kem}}$ be defined as described above by combining Fig. 11 with Fig. 5 resp. Fig. 10, where $\text{G}_I, \text{G}_R : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$, $\text{H} : \{0, 1\}^* \rightarrow \mathcal{K}$, $\text{G}_{\text{stI}} : \{0, 1\}^\kappa \times \{0, 1\}^\kappa \rightarrow \{0, 1\}^{d_I}$ and $\text{G}_{\text{stR}} : \{0, 1\}^\kappa \times \{0, 1\}^\kappa \rightarrow \{0, 1\}^{d_R}$ are modeled as random oracles. Let Q_h be the number of queries to the respective random oracle h .*

For any \mathcal{A} against the $(t, \varepsilon, \mu, S, T, Q_{\text{COR}})$ -IND-FS-security of AKE_{kem} , there exist adversaries \mathcal{B}_1 and \mathcal{B}_2 that break $(t_1, \varepsilon_1, S, S, S, Q_{\text{VER}})$ -OW-PCVA security

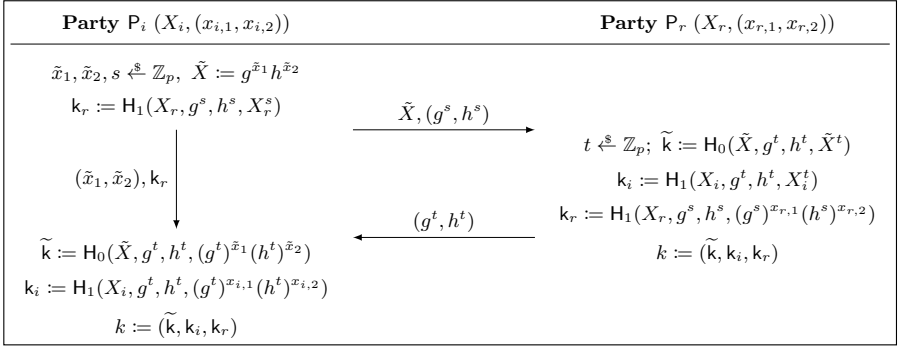


Fig. 12. AKE protocol JKRS. H_0, H_1 are independent random oracles. Protocol JKRS_{KC} is obtained by adding the transformation from Fig. 10.

of KEM_0 and $(t_2, \varepsilon_2, \mu, S, S, 2Q_{\text{Ver}}, Q_{\text{Cor}})$ -OW-PCVA-C security of KEM_1 , where $Q_{\text{Ver}} \leq S + Q_{\text{G}_I} + Q_{\text{G}_R} + Q_{\text{H}}$, with $t_1 \approx t_2 \approx t$ and

$$\varepsilon \leq \varepsilon_1 + \varepsilon_2 + 2S \cdot (\delta_0 + 2\delta_1) + (S + S^2) \cdot 2^{-\lambda} + \mu^2 \cdot 2^{-\gamma_1} + S(S + Q_{\text{G}_I} + Q_{\text{G}_R} + Q_{\text{H}}) \cdot (2^{-\gamma_0} + 2^{-\alpha_0} + 2^{-\alpha_1+1}).$$

Further, for every adversary \mathcal{A} that breaks the $(t, \varepsilon, \mu, S, T, Q_{\text{Cor}}, Q_{\text{ST}})$ -IND-FS-St-security of $\text{AKE}_{\text{st}, \text{kem}}$, there exist adversaries \mathcal{B}_1 and \mathcal{B}_2 that break $(t_1, \varepsilon_1, S, S, S, Q_{\text{Ver}}, Q_{\text{ST}})$ -OW-PCVA-C security of KEM_0 and $(t_2, \varepsilon_2, \mu, S, S, 2Q_{\text{Ver}}, Q_{\text{Cor}}, Q_{\text{ST}})$ -OW-PCVA-CR security of KEM_1 , where $Q_{\text{Ver}} \leq S + Q_{\text{G}_I} + Q_{\text{G}_R} + Q_{\text{H}}$, with $t_1 \approx t_2 \approx t$ and

$$\varepsilon \leq \varepsilon_1 + \varepsilon_2 + 2S \cdot (\delta_0 + 2\delta_1) + (\mu^2 + S^2 + \mu Q_{\text{G}_{\text{st}}} + 2SQ_{\text{G}_{\text{st}}}) \cdot 2^{-\kappa} + \mu^2 \cdot 2^{-\gamma_1} + (S + S^2) \cdot 2^{-\lambda} + S(Q_{\text{G}_R} + Q_{\text{G}_I} + Q_{\text{H}} + S) \cdot (2^{-\gamma_0} + 2^{-\alpha_0} + 2^{-\alpha_1+1}).$$

INSTANTIATION WITH NON-COMMITTING KEM. We can use a non-committing KEM as defined in [21] to instantiate a verifiable AKE protocol very efficiently, e. g., from DDH (cf. protocol JKRS in Fig. 12). We can easily show that a non-committing KEM implies OW-PCVA-CR security of that KEM. In our full version [28, Appendix A], we recall the formal definition of NC-CCA security for KEMs from [21] and show the implication. Adding key confirmation as described in Fig. 10 then yields protocol JKRS_{KC} .

SECURITY OF JKRS_{KC} . We now establish security of protocol JKRS_{KC} . Since the JKRS protocol is perfectly correct, so is JKRS_{KC} . Further, public keys and messages have $\log(p)$ bits entropy. Security is based on the DDH assumption which asks to distinguish between (g^x, g^y, g^{xy}) and (g^x, g^y, g^z) for $x, y, z \xleftarrow{\$} \mathbb{Z}_p$.

Theorem 4. Let JKRS_{KC} be defined as in Fig. 12, where $\text{G}_I, \text{G}_R : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$, $\text{H} : \{0, 1\}^* \rightarrow \mathcal{K}$, $\text{H}_0 : \{0, 1\}^* \rightarrow \text{KEM}_0.\mathcal{K}$, $\text{H}_1 : \{0, 1\}^* \rightarrow \text{KEM}_1.\mathcal{K}$,

$G_{\text{stI}} : \{0, 1\}^\kappa \times \{0, 1\}^\kappa \rightarrow \{0, 1\}^{d_1}$, and $G_{\text{stR}} : \{0, 1\}^\kappa \times \{0, 1\}^\kappa \rightarrow \{0, 1\}^{d_R}$ are modeled as random oracles.

For every adversary \mathcal{A} that breaks the $(t, \varepsilon, \mu, S, T, Q_{\text{COR}}, Q_{\text{ST}})$ -IND-FS-St-security of JKRS_{KC} , there exists an adversary \mathcal{B} that breaks (t', ε') -DDH with $t' \approx t$ and

$$\varepsilon \leq \varepsilon' + (\mu^2 + S^2 + \mu Q_{G_{\text{stI}}} + 2S Q_{G_{\text{stI}}}) \cdot 2^{-\kappa} + \mu^2 \cdot 2^{-\log(p)} \\ + (S + S^2) \cdot 2^{-\lambda} + S(Q_{G_R} + Q_{G_I} + Q_H + Q_{H_0} + Q_{H_1} + S + 1) \cdot 2^{-\log(p)},$$

where Q_h is the number of queries to the respective random oracle h .

The theorem follows from [28, Theorem 6], Lemma 2 and Theorem 2 in combination with [21, Theorem 5], where the latter deals with the optimization that only one ciphertext is sent in the second round.

INSTANTIATION FROM LATTICES. We can also instantiate the KEM-based verifiable AKE protocol using lattices assumptions. The scheme KEM_{LWE} described in [30, Section 3] satisfies OW-ChCCA security which implies OW-PCVA-CR security. This gives us an AKE protocol with key confirmation from LWE secure in the random oracle model.

6.2 The CCGJJ Protocol and Its Isogeny-Based Variant

It is easy to see that the core protocol from Cohn-Gordon et al. (CCGJJ) [10] is a verifiable AKE protocol, ignoring the session key hash. For completeness, we provide a formal treatment in our full paper [28, Appendix F].

ISOGENY-BASED AKE. The isogeny-based AKE protocol which was independently analyzed by de Kock, Gjøsteen and Veroni [23] and Kawashima et al. [22] follows the same blueprint as the CCGJJ protocol, relying on the group action structure of CSIDH [8] rather than prime-order groups. Thus, we also get an AKE protocol with key confirmation from isogenies, based on the same assumptions as the analysis in [22, 23]. This is particularly interesting because the only group action based and tightly-secure signature scheme supporting adaptive corruptions [29] is rather inefficient.

7 KEM-Based AKE with Key Confirmation in the QROM

We analyze FS via key confirmation in the quantum random oracle model (QROM). Following [31], we use IND-CCA-secure KEMs in the multi-user, multi challenge settings as building blocks. By the key confirmation technique, we lift the result of Pan, Wagner, and Zeng [31] to FS in the QROM. The work of Pan, Wagner, and Zeng only achieves weak FS in the QROM. Our result not only preserves the security loss of their protocol, but also achieves FS and allows multiple TEST queries with single challenge bit, while [31] allows at most one single TEST query.

The MC-IND-CCA and MUC-IND-CCA security definitions of KEMs are given in our full paper [28, Appendix A]. We use notations introduced in Sect. 5 to present our protocol AKE_{kem} . Let KEM_1 and KEM_0 be two KEM schemes and $\text{G}_I, \text{G}_R : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$, and $\text{H} : \{0, 1\}^* \rightarrow \mathcal{K}$ be hash functions, where λ is the length of key confirmation tags and \mathcal{K} is the key space of AKE_{kem} . An overview of our AKE construction AKE_{kem} is given in Fig. 13. AKE_{kem} is essentially the KEM-based AKE protocol in [31] adding key confirmation, namely, it is obtained from combining Fig. 11 and Fig. 5.

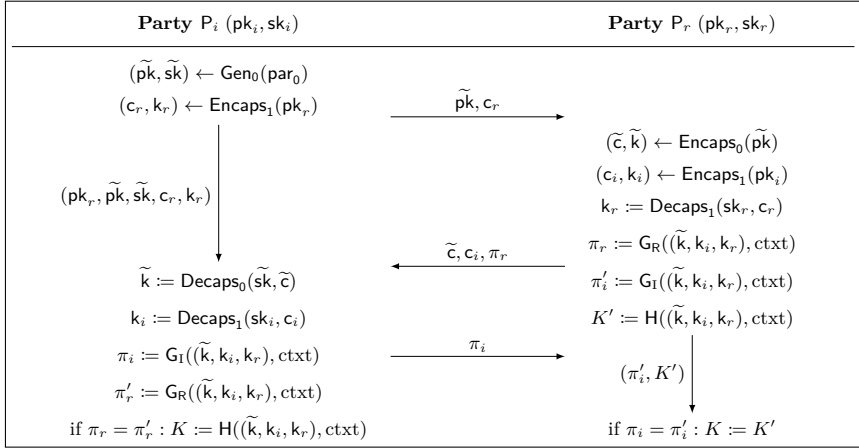


Fig. 13. AKE protocol AKE_{kem} from KEM schemes KEM_1 , KEM_0 , and key confirmation. The context is defined as $\text{ctxt} := (pk_i, pk_r, \tilde{pk}, \tilde{c}, c_i, c_r)$. G_I , G_R , and H are independent random oracles.

CORRECTNESS. The correctness of AKE_{kem} is due to KEM_1 and KEM_0 . Each session of AKE_{kem} includes two ciphertexts of KEM_1 and one ciphertext of KEM_0 . If KEM_1 is $(1 - \delta_1)$ -correct and KEM_0 is $(1 - \delta_0)$ -correct, then by the union bound, AKE_{kem} is $(1 - 2\delta_1 - \delta_0)$ -correct.

SECURITY. We prove IND-FS security of AKE_{kem} based on the MC-IND-CCA security of KEM_1 , the MUC-IND-CCA security of KEM_0 , and modeling G_I , G_R , and H as quantum-accessible random oracles, as stated in Theorem 5. The proof of Theorem 5 is postponed to our full version [28, Appendix E].

Theorem 5. *Let KEM_0 be $(1 - \delta_0)$ -correct and have public keys with γ_0 bits of entropy and messages with α_0 bits of entropy. Let KEM_1 be $(1 - \delta_1)$ -correct and have public keys with γ_1 bits of entropy and messages with α_1 bits of entropy. \mathcal{K}_0 and \mathcal{K}_1 are the KEM key spaces of KEM_0 and KEM_1 , respectively.*

Let AKE_{kem} be as defined in Fig. 13, where Let $\text{G}_I, \text{G}_R : \{0, 1\}^ \rightarrow \{0, 1\}^\lambda$ and $\text{H} : \{0, 1\}^* \rightarrow \mathcal{K}$. For every adversary \mathcal{A} that breaks the*

$(t, \varepsilon, \mu, S, T, Q_{\text{COR}})$ -IND-FS-security of AKE_{KC} , there exists an adversary \mathcal{B}_0 that breaks the $(t'_0, \varepsilon'_0, S, S)$ -MUC-IND-CCA security of KEM_0 and an adversary \mathcal{B}_1 that breaks the $(t'_1, \varepsilon'_1, S)$ -MC-IND-CCA security of KEM_1 with $t'_0 \approx t'_1 \approx t$ and

$$\varepsilon \leq 2\varepsilon'_0 + 2\mu\varepsilon'_1 + 2S(\delta_0 + \mu\delta_1) + \mu^2 2^{-\gamma_1} + \mu S 2^{-\lambda+1} \\ + S^2(2^{-\alpha_1} + 2^{-\gamma_0} + 2^{-\alpha_0}) + \frac{2\mu(Q_{\text{GR}} + Q_{\text{GI}})\sqrt{S}}{\sqrt{|\mathcal{K}_1|}} + \frac{2Q_{\text{H}}\sqrt{S}}{\sqrt{|\mathcal{K}_0|}},$$

where Q_{GI} , Q_{GR} and Q_{H} are the number of quantum-superposition queries to G_1 , G_R and H .

Remark 1 (Implicit Rejection). Following [31], when proving Theorem 5, we assume KEM_1 and KEM_0 have implicit rejection [5], namely, if the input ciphertext is invalid, then the decapsulation algorithm returns a pseudorandom KEM key. We use implicit-rejection KEM because it simplifies our AKE proof.

To adapt the proof of Theorem 5 to the one that uses explicit-rejection KEMs, we can add extra codes in the games sequence to deal with explicit rejections from KEM. Concretely, upon receiving an invalid KEM ciphertext, the session oracle (e.g., SESSION_R , DER_R , or DER_1) simply sets the session key as “reject” and returns \perp . This can be tightly simulated by MUC-IND-CCA and MC-IND-CCA secure KEMs with explicit rejection, and thus the security bound in Theorem 5 also applies to explicit-rejection KEMs.

Remark 2 (Instantiations with LWE). In [31], Pan et al. proposed lattice-based instantiations of MC-IND-CCA-secure KEM and MUC-IND-CCA-secure KEM that have (almost-)tight reduction from the well-known Learning With Errors (LWE) problem in the QROM. Here we only discuss the security loss of these KEM schemes and give the final security loss of our AKE protocol instantiated with these KEM schemes.

Let ε_{lwe} be the best computational advantage against LWE assumptions and λ be the security parameter (which decides the security level, the length of message, etc.). The two KEM schemes proposed in [31] have asymptotic bounds $\varepsilon'_1 \leq \Theta(\lambda) \cdot \varepsilon_{\text{lwe}}$ and $\varepsilon'_0 \leq \Theta(\lambda) \cdot \varepsilon_{\text{lwe}}$, where ε'_1 and ε'_0 are the computational advantages against MC-IND-CCA security and MUC-IND-CCA security of the KEM schemes in [31], respectively. By combining these bounds with the bounds given in Theorem 5, we have

$$\varepsilon \leq \Theta(\lambda) + \Theta(\mu) \cdot \Theta(\lambda) \cdot \varepsilon_{\text{lwe}} = \Theta(\mu) \cdot \Theta(\lambda) \cdot \varepsilon_{\text{lwe}},$$

where ε is the computational advantage against the resulting AKE protocol. This gives us a session-tight and square-root-tight (namely, does not suffer from the square-root security loss) LWE-based instantiation of AKE with full forward secrecy in the QROM.

Acknowledgements. We thank the anonymous reviewers for their valuable comments on better motivating our works and comparisons with the related work. Doreen Riepel was supported in part by Bellare’s KACST grant. Jiaxin Pan was supported in part by the Research Council of Norway (RCN) under Project No. 324235, and Runzhi Zeng were supported by the same project from RCN.

References

1. Abdalla, M., Fouque, P.A., Pointcheval, D.: Password-based authenticated key exchange in the three-party setting. In: Vaudenay, S. (ed.) PKC 2005. LNCS, vol. 3386, pp. 65–84. Springer, Heidelberg (2005). https://doi.org/10.1007/978-3-540-30580-4_6
2. 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
3. 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
4. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 409–426. Springer, Heidelberg (2006). https://doi.org/10.1007/11761679_25
5. Bernstein, D.J., Persichetti, E.: Towards KEM unification. Cryptology ePrint Archive, Report 2018/526 (2018). <https://eprint.iacr.org/2018/526>
6. Boneh, D., Dagdelen, Ö., Fischlin, M., Lehmann, A., Schaffner, C., Zhandry, M.: Random oracles in a quantum world. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 41–69. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25385-0_3
7. Canetti, R., Krawczyk, H.: Analysis of key-exchange protocols and their use for building secure channels. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 453–474. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44987-6_28
8. Castryck, W., Lange, T., Martindale, C., Panny, L., Renes, J.: CSIDH: an efficient post-quantum commutative group action. In: Peyrin, T., Galbraith, S. (eds.) ASIACRYPT 2018, Part III. LNCS, vol. 11274, pp. 395–427. Springer, Heidelberg (2018). https://doi.org/10.1007/978-3-030-03332-3_15
9. Chen, J., Wee, H.: Fully, (almost) tightly secure IBE and dual system groups. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 435–460. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40084-1_25
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, Heidelberg (2019). https://doi.org/10.1007/978-3-030-26954-8_25
11. Davis, H., Günther, F.: Tighter proofs for the SIGMA and TLS 1.3 key exchange protocols. In: Sako, K., Tippenhauer, N.O. (eds.) ACNS 2021, Part II. LNCS, vol. 12727, pp. 448–479. Springer, Heidelberg (2021). https://doi.org/10.1007/978-3-030-78375-4_18
12. Delpèch de Saint Guilhem, C., Fischlin, M., Warinschi, B.: Authentication in key-exchange: definitions, relations and composition. In: Jia, L., Küsters, R. (eds.) CSF 2020 Computer Security Foundations Symposium, pp. 288–303. IEEE Computer Society Press (2020). <https://doi.org/10.1109/CSF49147.2020.00028>
13. Diemert, D., Gellert, K., Jager, T., Lyu, L.: More efficient digital signatures with tight multi-user security. In: Garay, J. (ed.) PKC 2021, Part II. LNCS, vol. 12711, pp. 1–31. Springer, Heidelberg (2021). https://doi.org/10.1007/978-3-030-75248-4_1

14. Diemert, D., Jager, T.: On the tight security of TLS 1.3: theoretically sound cryptographic parameters for real-world deployments. *J. Cryptol.* **34**(3), 30 (2021). <https://doi.org/10.1007/s00145-021-09388-x>
15. Fischlin, M., Günther, F., Schmidt, B., Warinschi, B.: Key confirmation in key exchange: a formal treatment and implications for TLS 1.3. In: 2016 IEEE Symposium on Security and Privacy, pp. 452–469. IEEE Computer Society Press, May 2016. <https://doi.org/10.1109/SP.2016.34>
16. 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
17. 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
18. Gellert, K., Gjøsteen, K., Jacobsen, H., Jager, T.: On optimal tightness for key exchange with full forward secrecy via key confirmation. In: Handschuh, H., Lysyanskaya, A. (eds.) CRYPTO 2023. LNCS, Springer, Heidelberg (2023). https://doi.org/10.1007/978-3-031-38551-3_10
19. Gjøsteen, K., Jager, T.: Practical and tightly-secure digital signatures and authenticated key exchange. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part II. LNCS, vol. 10992, pp. 95–125. Springer, Heidelberg (2018). https://doi.org/10.1007/978-3-319-96881-0_4
20. 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, Heidelberg (2020). https://doi.org/10.1007/978-3-030-45388-6_14
21. Jager, T., Kiltz, E., Riepel, D., Schäge, S.: Tightly-secure authenticated key exchange, revisited. In: Canteaut, A., Standaert, F.X. (eds.) EUROCRYPT 2021, Part I. LNCS, vol. 12696, pp. 117–146. Springer, Heidelberg (2021). https://doi.org/10.1007/978-3-030-77870-5_5
22. Kawashima, T., Takashima, K., Aikawa, Y., Takagi, T.: An efficient authenticated key exchange from random self-reducibility on CSIDH. In: Hong, D. (ed.) ICISC 2020. LNCS, vol. 12593, pp. 58–84. Springer, Heidelberg (2020). https://doi.org/10.1007/978-3-030-68890-5_4
23. de Kock, B., Gjøsteen, K., Veroni, M.: Practical isogeny-based key-exchange with optimal tightness. In: Dunkelman, O., Jacobson Jr., M.J., O’Flynn, C. (eds.) SAC 2020. LNCS, vol. 12804, pp. 451–479. Springer, Heidelberg (2020). https://doi.org/10.1007/978-3-030-81652-0_18
24. 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
25. LaMacchia, B.A., 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 (Nov 2007)
26. 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, Heidelberg (2020). https://doi.org/10.1007/978-3-030-64834-3_27

27. Pan, J., Qian, C., Ringerud, M.: Signed (group) Diffie-Hellman key exchange with tight security. *J. Cryptol.* **35**(4), 26 (2022). <https://doi.org/10.1007/s00145-022-09438-y>
28. Pan, J., Riepel, D., Zeng, R.: Key exchange with tight (full) forward secrecy via key confirmation. In: *Cryptology ePrint Archive* (2024)
29. Pan, J., Wagner, B.: Lattice-based signatures with tight adaptive corruptions and more. In: Hanaoka, G., Shikata, J., Watanabe, Y. (eds.) *PKC 2022, Part II*. LNCS, vol. 13178, pp. 347–378. Springer, Heidelberg (2022). https://doi.org/10.1007/978-3-030-97131-1_12
30. Pan, J., Wagner, B., Zeng, R.: Lattice-based authenticated key exchange with tight security. In: Handschuh, H., Lysyanskaya, A. (eds.) *CRYPTO 2023*. LNCS, Springer, Heidelberg (2023). https://doi.org/10.1007/978-3-031-38554-4_20
31. Pan, J., Wagner, B., Zeng, R.: Tighter security for generic authenticated key exchange in the QROM. In: *ASIACRYPT 2023*. LNCS, Springer, Heidelberg (2023). https://doi.org/10.1007/978-981-99-8730-6_13, <https://eprint.iacr.org/2023/1380>
32. Unruh, D.: Revocable quantum timed-release encryption. In: Nguyen, P.Q., Oswald, E. (eds.) *EUROCRYPT 2014*. LNCS, vol. 8441, pp. 129–146. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-55220-5_8