XHMQV: Better Efficiency and Stronger Security for Signal's Initial Handshake based on HMQV^{*}

Rune Fiedler¹, Felix Günther², Jiaxin Pan³, and Runzhi Zeng³

 ¹ Technische Universität Darmstadt, Darmstadt, Germany rune.fiedler@cryptoplexity.de
 ² IBM Research Europe - Zurich, Rüschlikon, Switzerland mail@felixguenther.info
 ³ University of Kassel, Kassel, Germany {jiaxin.pan,runzhi.zeng}@uni-kassel.de

Abstract. The Signal protocol is the most widely deployed end-to-endencrypted messaging protocol. Its initial handshake protocol X3DH allows parties to asynchronously derive a shared session key without the need to be online simultaneously, while providing implicit authentication, forward secrecy, and a form of offline deniability. The X3DH protocol has been extensively studied in the cryptographic literature and is acclaimed for its strong "maximum-exposure" security guarantees, hedging against compromises of users' long-term keys and medium-term keys but also the ephemeral randomness used in the handshake. This maximum-exposure security is achieved by deriving keys from the concatenation of 3–4 Diffie– Hellman (DH) secrets, each combining two long-term, medium-term, or ephemeral DH shares.

Remarkably, X3DH's approach of concatenating plain DH combinations is sub-optimal, both in terms of maximum-exposure security and performance. Indeed, Krawczyk's well-known HMQV protocol (Crypto '05) is a high-performance, DH-based key exchange that provides strong security against long-term and ephemeral key compromise. One might hence wonder: why not base Signal's initial handshake on HMQV?

In this work, we study this question and show that a carefully adapted variant of HMQV, which we call XHMQV, indeed enables stronger security and efficiency while matching the constraints of Signal's initial handshake. Most notably, HMQV does not work as a drop-in replacement for X3DH, as the latter's asynchronicity requires the protocol to handle cases where one party runs out of ephemeral keys (pre-uploaded to the Signal server). Our XHMQV design hence augments HMQV with medium-term keys analogous to those used in X3DH. We prove that XHMQV provides security in all 3–4 compromise scenarios where X3DH does and *additionally* in 1–2 further scenarios, strengthening the handshake's maximum-exposure guarantees while using more efficient group operations. We further confirm that our XHMQV design achieves deniability guarantees comparable to X3DH. Our security model is the first

^{*} A preliminary version of this paper appears in the proceedings of the 45th Annual International Cryptology Conference (CRYPTO 2025), © IACR 2025. This is the full version.



Fig. 1. Signal's X3DH handshake (simplified), involving long-term keys a and b of users Alice resp. Bob, semi-static key s, and ephemeral keys x and y.

to capture Signal's long-term key reuse between DH key exchange and signatures, which may be of independent interest.

1 Introduction

The Signal protocol [49] for end-to-end-encrypted messaging is used by billions of people daily, underlying major messaging apps such as Facebook Messenger, Google Messages, Skype, Signal, or WhatsApp. Its initial handshake protocol X3DH ("extended triple Diffie–Hellman") is used to establish a session key between two users that want to start a conversation. It is designed in a way that the two users do not have to be online simultaneously: users upload a "prekey bundle" to the Signal server which initiators of a conversation can fetch to asynchronously run their part of the X3DH protocol and start sending messages, while responders can catch up later by completing their part of the handshake. X3DH further provides implicit authentication of the users, forward secrecy in case of future compromise of long-term keys, and a form of deniability allowing users to dispute having been involved in a conversation.

Sub-optimal solution for "maximum-exposure". Perhaps most importantly, X3DH aims at strong security guarantees, sometimes referred to as "maximum-exposure" security: its design hedges against compromises of users' long-term keys and medium-term keys as well as exposure of the ephemeral randomness used in the handshake. This maximum-exposure security is achieved by deriving the session key from the concatenation of 3–4 Diffie–Hellman (DH) secrets g^{as} , g^{xb} , g^{xs} , and (optionally) g^{xy} , each combining two long-term (g^a , g^b for users Alice and Bob), medium-term (g^s for user Bob only), or ephemeral DH shares (g^x , g^y). See Figure 1 for a (simplified) illustration of the X3DH handshake. Signal's handshake and related designs have been studied extensively in the cryptographic literature [11,12,35,53,27,28,19,8,24,6,13,23,29] and Signal is widely recognized as the "gold standard" for secure messaging.

Notably, X3DH uses only a partial coverage of the potential DH combinations, even when considering only the long-term and ephemeral keys. For instance, it does not include the long-/long-term combination g^{ab} , which is useful to hedge against compromise of ephemeral randomness and medium-term secrets. This is a security property that the current Signal handshake cannot achieve. A

Protocol	Group	o exp.	Key Indistinguishability					Deniał	oility	
	Alice	Bob	LT-SS	E-LT	E-SS	E-E	LT-LT	LT-E	semi-hon.	malic.
X3DH	4 (3)	4 (3)	1	1	1	(⁄)	X	×	✓	+
XHMQV	3(2)	2	✓	✓	1	(<	✓	(\checkmark)	\checkmark	×

Table 1. Performance, security, and deniability properties of X3DH and XHMQV. Key indistinguishability conditions AA-BB demand security as long as Alice's AA and Bob's BB keys are uncompromised, where LT, SS, E denote long-term, medium-term/semi-static, and ephemeral keys, respectively. Parenthesized entries are for reduced mode (without Bob's ephemeral key present). We discuss the intricacies of malicious deniability below.

straightforward modification is to incorporate this additional long-term/longterm combination, yet simply concatenating plain DH values is *sub-optimal in terms of performance*. On the other hand, the well-known HMQV protocol by Krawczyk [36], building on the MQV protocol [43,39], is a high-performance, DH-based key exchange that provides strong security against long-term and ephemeral key compromise. This naturally raises the question:

> Can Signal's initial handshake be (more) efficiently instantiated based on HMQV to provide (even) stronger security?

The short answer is: Yes, but it requires careful adaptation and rigorous proofs.

1.1 Contributions

In this work, we study the above question and introduce XHMQV, a carefully adapted variant of HMQV that enables stronger security and improved efficiency over X3DH, while adhering to the real-world constraints imposed on Signal's initial handshake. We emphasize that HMQV does not work as a drop-in replacement for X3DH, as the latter's asynchronicity requires the protocol to handle cases where one party runs out of ephemeral keys (pre-uploaded to the Signal server). Our XHMQV design hence augments HMQV with semi-static keys analogous to those used in X3DH.

We prove that XHMQV's session keys are secure under all 3–4 compromise scenarios targeted by X3DH and *additional* 1–2 scenarios (see Table 1 for details), and also that it matches Signal's deniability targets (we discuss technical subtleties below). Notably, our security model and proof is the first to capture the re-use of long-term keys in Signal for key exchange and signatures, which may be of independent interest.

In more detail, our contributions are as follows.

XHMQV: *HMQV*-based Signal handshake. We propose XHMQV, a new Signal handshake protocol based on Krawczyk's HMQV protocol [36], illustrated in

4



Fig. 2. Our XHMQV handshake protocol (simplified), involving long-term keys a and b of users Alice resp. Bob, semi-static key s, and ephemeral keys x and y, as well as hash functions h_0 , h_1 , h_2 . If Bob's ephemeral key $Y := g^y$ is unavailable, the resulting reduced handshake mode omits elements highlighted in light blue.

Figure 2. The core idea is to augment HMQV with a semi-static key on Bob's side $S = g^s$, similar to the semi-static keys used in Signal's X3DH for delayed forward secrecy in case the Signal server runs out of pre-key bundles for Bob including a fresh ephemeral share $Y = g^y$. (Indeed, removing $S = g^s$ and all related computations from Figure 2 essentially yields HMQV, modulo interaction and hashing context.)

In terms of performance, XHMQV involves 3 or 2 group exponentiations (for Alice and Bob, respectively; in reduced mode only 2 for Alice) whereas X3DH involves 4 group exponentiations per party (in full mode, 3 in reduced mode). While XHMQV additionally involves group multiplications, these are generally much more efficient than exponentiations, resulting in an overall performance advantage.⁴

Stronger "maximal-exposure" security. Compared to X3DH, our XHMQV design ensures security under two additional compromise scenarios, namely compromise of everything except (1) both initiator and responder long-term keys, as well as (2) the initiator long-term key and the responder ephemeral randomness. In both cases, an adversary can trivially compute all DH combinations in the key derivation of X3DH to obtain the session key, whereas our XHMQV protocol inherits protection against such compromise from the original HMQV protocol.

We argue that these strengthened security guarantees are not only theoretical: The former scenario (long-term–long-term uncompromised) hedges against bad randomness used when generating ephemeral and semi-static keys (while long-term keys generated only once may benefit from better randomness sources).

⁴ Note that we use multiplicative notation here, in line with prior Signal analyses. Signal's X3DH handshake is using elliptic-curve cryptography, specifically X25519 based on Curve25519 [5,38] for key exchange and XEdDSA signatures [46] over Ed25519 [38]. Notably, X3DH is re-using long-term Diffie-Hellman keys for both DH key exchange and signing a user's pre-key bundle, involving a (cheap) conversion between Curve25519 and Ed25519. When implementing XHMQV with X25519, one would likewise convert between Curve25519 and Ed25519 to benefit from the latter's fast, constant-time point addition.

The latter scenario (long-term–ephemeral uncompromised) mirrors the protection sought after by X3DH through combining initiator ephemeral and responder long-term keys. We stress that both additional security guarantees essentially come "for free" through employing HMQV's clever key derivation.

Security proof capturing reuse of long-term keys. We prove security for XHMQV in a computational, game-based security model capturing Bellare–Rogawaystyle [4] key indistinguishability. To this end, we adopt the security model of Brendel et al. [8,7] and Fiedler and Günther [23], which in turn builds on the initial computational analysis of Signal by Cohn-Gordon et al. [11,12]. In this model, the adversary actively interacts with multiple users and sessions, is allowed to adaptively compromise long-term, semi-static, and ephemeral secrets, and is tasked with distinguishing from random the session key in a chosen *test session* which is not trivially compromised and "*clean*". Cleanliness of a session captures the "maximum-exposure" security aimed at by the XHMQV protocol, asking for security as long as certain combination of keys are uncompromised for the test session. The stronger security achieved by XHMQV compared to X3DH corresponds to our model demanding security in (two) more "clean" cases.

As a notable technical difference to prior models, our security model is the first to capture Signal's key reuse of long-term user keys for both in the X25519 Diffie-Hellman key exchange and as signing key for XEdDSA [46] signatures. Prior work either modeled these signatures using separate signing keys [35,53,24,6,23,29] or not at all [11,12,8].

As a result of modeling the key reuse, we require a more fine-grained version of signature security. The technical challenge is that, since the long-term key in the DH key exchange is the same as the public key of the signature scheme, we cannot compute the session key when we construct a reduction to signature security, since we do not know the corresponding long-term secret key. To resolve this, we introduce a new notion of security, namely, one-per message unforgeability with a DDH oracle, where DDH($X = g^x, Y = g^y, Z$) outputs whether $Z = g^{xy}$. This is a notion that can be satisfied by signatures using the CDH or DLog assumptions. In Section 2.2 we argue that both the Schnorr and (EC)DSA signature schemes satisfy this notion. As discussed in [22], the (weaker) one-per message unforgeability is sufficient for AKE protocols.

We give more technical details about why this additional DDH oracle can help us to simulate the session key. In our XHMQV protocol, our reduction to the signature security needs to compute $\mathsf{K} := \mathsf{KDF}(DH, A, B, S, X)$ for adversary \mathcal{A} who is attacking XHMQV, where $DH = (XA^d)^{e_1b+s}$, imagining the simplest reduced handshake mode. Since the long-term keys A and B are coming from the signature challenger, the reduction knows neither a nor b. Our solution is to always answer $\mathsf{KDF}(DH, A, B, S, X)$ with a random value k, although the reduction cannot compute DH. To prevent the adversary from noticing this, KDF is modeled as a random oracle and if at some point \mathcal{A} queries KDF with the correct DH secret $DH = (XA^d)^{e_1b+s}$, then the reduction will program $\mathsf{KDF}(DH, A, B, S, X)$ to k. To check whether $DH = (XA^d)^{e_1b+s}$ holds, we need the DDH oracle. 6

Deniability. Signal aims for deniability of the initial handshake protocol [41, Section 4.4] in the following sense: the transcript of a protocol execution should not prove either party's involvement in the execution, even if the distinguisher gets access to both party's secret keys. In particular, this excludes the case of one party cooperating with the distinguisher during the protocol execution (online deniability), which Unger and Goldberg [51, Section 6.6] deem unachievable for asynchronous key exchange with forward secrecy. Hence, we focus on offline deniability. Arguably, Signal's requirement is restricted to semi-honest adversaries, i.e., adversaries that follow the protocol.

Fiedler and Janson [24] show that offline deniability against semi-honest adversaries holds for X3DH. Additionally, they show deniability for Alice against malicious adversaries limited to honestly generated long-term and semi-static keys. Furthermore, Vatandas, Gennaro, Ithurburn, and Krawczyk [53] show that X3DH is deniable against malicious adversaries based on their novel Extended Knowledge of Diffie–Hellman (EKDH) assumption. However, Fiedler and Langrehr [25] show that the EKDH assumption does not hold and that X3DH is not deniable as soon as simple auxiliary input is allowed.

Our protocol XHMQV matches the deniability guarantees of X3DH against semi-honest adversaries. Against malicious adversaries limited to honestly generated long-term and semi-static keys, our XHMQV does not provide deniability for Alice (unlike X3DH). Specifically, if Bob provably does not know his own ephemeral secret key, then the long-term–ephemeral contribution proves Alice's involvement. Arguably, the same deniability issue arises already in X3DH when Bob's semi-static keys are (also) allowed to be maliciously generated, which we deem realistic albeit not modeled in [24].⁵ In that regard, XHMQV fares comparably to X3DH in terms of deniability against realistic malicious adversaries. Against malicious adversaries in general, we expect that the findings of Fiedler and Langrehr [25] discussed above likewise apply to XHMQV as to X3DH.

Post-quantum readiness. Our XHMQV design can be readily extended to protect against harvest-now-decrypt-later quantum attacks, following a similar approach as Signal's recent post-quantum revision PQXDH [37] of the initial handshake. Recall that PQXDH essentially adds a (signed) KEM flow to the X3DH design, mixing the derived KEM shared secret into the key derivation. Such a KEM flow and the derived KEM shared secret can be added in an identical fashion to our XHMQV design, lifting the resulting security to a quantum-safe version as for PQXDH. We refer to [6,24,23] for detailed analyses of the design and security of PQXDH.

Deployment considerations. Since our XHMQV protocol only requires a change in how session keys are derived, but does not change the wire format of the protocol,

⁵ Even when assuming the Signal server does not collude in a deniability attack (and asks users to prove possession of keys they upload), we expect this is done only for long-term keys, but not semi-static keys. The honest key registration oracle of [24] however enforces that long-term *and semi-static* keys are honestly generated.

we deem XHMQV readily deployable through a next breaking protocol update of Signal. Observe that, in contrast to X3DH, the key derivation in XHMQV also includes as context the involved DH public keys; this is good practice and supports our proof technique. Similar suggestions have been made in the security analyses of Signal's PQXDH protocol [6,23] for a more robust design safeguarding against KEM re-encapsulation attacks [14].

1.2 Related Work and Further Discussion

Analyses of the HMQV protocol. The HMQV protocol was proven secure in [36] based on the Gap Diffie–Hellman assumption using the Forking Lemma [47]. In [34] Kiltz et al. gave a different security proof of HMQV with a modular approach. They first proposed a multi-user variant of the Computational DH problem, the <u>Challenge-Response Gap DH</u> with <u>Corr</u>uption (CorrCRGapDH) assumption, and they proved that the CorrCRGapDH assumption tightly implies the security of HMQV. Then they showed that the CorrCRGapDH assumption is non-tightly implied by the Gap DH assumption [1] using the Forking Lemma. Although their proof has the same loss as the original one in [36], its modular approach allows us to focus on the core algebraic assumption for the security of HMQV. In this paper, we follow the same modular approach to simplify our proof.

Tightness of the security proof. Similar to prior work analyzing HMQV [34], X3DH [12], and PQXDH [23], our security bound is not tight, as illustrated in Table 2. However, none of them models the reuse of long-term keys for DH and signing operations. Our analysis of XHMQV is the first one to account for the key reuse. While the result of [34] for HMQV is the only one to allow for multiple TEST queries, the protocol does not involve semi-static keys, nor does it take signatures on public keys or key reuse into account. The analysis of X3DH takes the semi-static keys into account, and the analysis of PQXDH additionally considers the signatures on the semi-static keys.

Our result relies on a challenge-response version CRGapDH of the Gap Diffie-Hellman assumption introduced by Kiltz et al. [34] (under the name (2,1)-CorrCRGapDH assumption). Kiltz et al.'s use of the Forking Lemma to reduce CorrCRGapDH to GapDH means that, to base our results on the GapDH assumption, we require rewinding that comes with a square root security loss. However, the security proof for DLog-based signature schemes such as XEdDSA (used in X3DH) and Schnorr signatures is in the ROM and uses rewinding, which leads to a square root loss as well. Therefore, an analysis of X3DH (or any other protocol for Signal's initial handshake) that takes the signatures into account already incurs a square root loss. Hence, the use of the CRGapDH assumption does not incur much penalty for rewinding. We discuss the relation between the GapDH assumption and its challenge-response variant in more detail in Section 2.1.

Analyses of the Signal protocol. The Signal protocol (both the initial handshake X3DH and the Double Ratchet protocol) were first analyzed with a game-based

8

Protocol	#Test queries	Semi- static keys	Key reuse	Model	Security bound
HMQV (+context) [34]	Т	×	×	ROM GGM	$\frac{\frac{(q_{\rm G}+n_u+n_s+1)^2}{p} + \frac{2q_{\rm RO}}{p}}{+\frac{3(n_u+n_s+1)^2(2q_{\rm RO}+1)}{p}}$
X3DH [12]	1	1	×	ROM	$\frac{\frac{(n_u+n_u\cdot n_{ss}+n_s)^2}{p}}{+n_u^2\cdot n_s\cdot \left(4n_{ss}+2+n_s\right)\left(\frac{1}{p}+\epsilon_{\rm GDH}\right)}$
PQXDH [23]	1	J J	X	ROM	$ \frac{(n_u + n_u n_{ss} + n_s)^2}{p} + \gamma_{\text{coll}}(n_u n_{ss} + n_s) $ $+ n_s \cdot \delta_{\text{corr}} + \epsilon_{LEAK^{+r}} $ $+ (n_u^2 n_{ss} + n_s n_u + n_u n_{ss} n_s) \cdot \epsilon_{\text{GDH}} $ $+ 4n_u \cdot \epsilon_{\text{SIG}} $ $+ (n_u n_{ss} n_s + n_s^2) \cdot \min(\epsilon_{\text{GDH}}, \epsilon_{\text{CCA}}) $ $+ n_u n_s^2 \cdot q_{\text{RO}} \cdot \epsilon_{\text{CCA}} $
XHMQV (this work)	1	11	1	ROM	$ \begin{array}{l} \frac{(n_u + n_u n_{ss} + n_s)^2}{p} + n_u \cdot \epsilon_{\text{EUF-opCMA-DDH}} \\ + \left(2n_u^2 + 2n_u n_s + n_u^2 n_{ss}\right) \cdot \delta_{sim} \\ + \left(n_u^2 + n_u n_{ss} + n_u^2 n_{ss} + n_u n_s \\ + n_s n_u n_{ss} + n_{ss}^2\right) \cdot \left(\epsilon_{\text{CRGapDH}} + \frac{Q_{\text{RO}}}{p}\right) \end{array} $

Table 2. Comparing the tightness of the bounds for HMQV (with additional context as input to the KDF), X3DH, PQXDH, and XHMQV (against classical adversaries). Column 2 denotes the number of allowed TEST queries; column 3 if semi-static keys are modeled (\checkmark not at all, \checkmark yes but without being signed, $\checkmark \checkmark$ yes and signed); column 4 indicates if the key reuse of the long-term keys for both DH and signing operations is modeled; column 5 states the model; and the final column gives the bound. For HMQV, $q_{\rm G}$ denotes the number of queries to the group oracle (of the generic group model).

approach by Cohn-Gordon, Cremers, Dowling, Garratt, and Stebila [11,12] and with a tool-based approach by Kobeissi, Bhargavan, and Blanchet [35] in the symbolic and computational model. Signal's initial handshake has since been replaced with PQXDH [37], which extends X3DH with a KEM in order to keep the classical guarantees of X3DH. Bhargavan, Jacomme, Kiefer, and Schmidt [6] have analyzed PQXDH with a tool-based approach in the symbolic and computational model and Fiedler and Günther [23] with a game-based approach. Several post-quantum replacements for X3DH have been proposed: Hashimoto, Katsumata, Kwiatkowski, and Prest [27,28] proposed SC-AKE, SC-DAKE, and SC-DAKE' based on KEMs, ring signatures and NIZKs achieving varying levels of deniability. Brendel, Fiedler, Günther, Janson, and Stebila [8] proposed SPQR, which is based on KEMs and designated verifier signatures. Dobson and Galbraith [19] proposed SI-X3DH based on supersingular isogenies, which was broken by the SIDH attack [10,40,48]. Collins, Huguenin-Dumittan, Nguyen, Rolin, and Vaudenay [13] proposed K-Waay based on KEMs and splitKEMs. Hashimoto, Katsumata, and Wiggers [29] proposed RingXKEM, which batches computation of pre-key bundles for better efficiency.

Deniability of X3DH was first analyzed by Vatandas, Gennaro, Ithurburn, and Krawczyk [53]. Fiedler and Janson [24] studied the deniability of PQXDH.

Key reuse. Haber and Pinkas [26] and An, Dodis, and Rabin [2] considered reuse of keys between signature schemes and public key encryption schemes and offer definitions for combined security. Degabriele, Lehmann, Paterson, Smart, and Strefler [17] and Paterson, Schuldt, Stam, and Thomson [44] looked at the same setting and define security for one scheme in the presence of an oracle for the other scheme, i.e., IND-CCA in the presence of a signing oracle and EUF-CMA in the presence of a decryption oracle. Patton and Shrimpton [45] generically compose two schemes while taking the context into account. Thormaker [50] analyzed the joint security of an X25519 based KEM and the Ed25519 signature scheme.

On the multi-test security from the multi-instance CorrCRGapDH. Kiltz et al. [34] proved that the security of HMQV can be tightly reduced to the (m, n)-CorrCRGapDH assumption, a multi-instance variant of CRGapDH. Moreover, their AKE security model considers a multi-test setting, namely, the adversary is allowed to query the TEST oracle multiple times. This naturally raises the question of whether our XHMQV protocol can be proven secure with a tight reduction under the (m, n)-CorrCRGapDH assumption in the multi-test setting.

We leave it for future study for the following reasons: First, we aim to remain consistent with prior analyses of Signal's handshake protocols [12,23]. Second, we would need to carefully modify our proof such that the underlying (m, n)-**CorrCRGapDH** assumption does not use an n or m depending on T, the number of TEST queries. Otherwise, T would appear in the security loss to the Gap DH assumption, since the reduction from (m, n)-CorrCRGapDH to Gap DH loses m^2n (cf. Theorem 1 and Figure 5 in [34]).

2 Preliminaries

We assume that all our algorithms are efficiently computable and probabilistic unless we state it.

2.1 Diffie–Hellman Assumptions

We use a weaker variant of the <u>Challenge-Response Gap Diffie-Hellman</u> with <u>Corr</u>uption (CorrCRGapDH) assumption of Kiltz et al. [34] to prove our XHMQV protocol secure. In [34] Kiltz et al. gave a direct security proof of HMQV from CorrCRGapDH, and by using the Forking Lemma [47] they proved that the Gap DH (GapDH) assumption [1] implies the CorrCRGapDH assumption. Hence, by using the CorrCRGapDH assumption, we do not require the direct use of the

$\mathcal{G}^{CRGapDH}_{(\mathbb{G},p,g)}(\mathcal{A},Q_{ ext{Ddh}},Q_{ ext{Ch}})$:	$\underline{\text{DDH}(X,Y,Z)}:$
$\begin{array}{l} 1 a,b \leftarrow \mathbb{S} \mathbb{Z}_p, \ A := g^a, B := g^b \\ 2 C \leftarrow \mathbb{S} \mathcal{A}^{\mathrm{DDH,CH}}((\mathbb{G},p,g),A,B) \end{array}$	// at most Q_{DDH} queries 4 return $\llbracket Z = X^{DL_g(Y)} \rrbracket$
3 return $\llbracket C \in \mathcal{L}_{CRGapDH} \rrbracket$	$CH(R \in \mathbb{G})$:
$ \begin{array}{l} \mathcal{G}^{GapDH}_{(\mathbb{G},p,g)}(\mathcal{A},Q_{\mathrm{DbH}}):\\ 1 a,b \leftarrow \mathbb{Z}_p, \ A:=g^a, B:=g^b\\ 2 C \leftarrow \mathbb{Z}_p, \ A:=g^a, B:=g^b\\ 3 \mathbf{return} \ \llbracket C = A^b \rrbracket $	// at most Q_{CH} queries 5 $h \leftarrow \mathbb{S}\mathbb{Z}_p$ 6 $\mathcal{L}_{CRGapDH} := \mathcal{L}_{CRGapDH} \cup \{(R \cdot A^h)^b, (R \cdot B^h)^a\}$ 7 return h

Fig. 3. Security game for the CRGapDH problem defined in Definition 1.

Forking Lemma in our security proof for XHMQV, simplifying the presentation. By [34, Theorem 1], the security of XHMQV is based on the GapDH assumption. The original CorrCRGapDH assumption is in the multi-user, multi-challenge setting, and ours is in the single-pair-of-user, single-challenge setting, and thus our assumption does not have any corruption query and is denoted by CRGapDH, which is (2, 1)-CorrCRGapDH in the notation of [34].

Let \mathbb{G} be a cyclic group with prime order p and generator g. We write (\mathbb{G}, p, g) as the group description of \mathbb{G} .

Definition 1 (Challenge-Response Gap Diffie-Hellman Problem). We say the $(t, \epsilon, Q_{\text{DDH}}, Q_{\text{CH}})$ -CRGapDH assumption holds in (\mathbb{G}, p, g) if for any adversary \mathcal{A} with running time at most t making at most Q_{DDH} queries to its DDH oracle and Q_{CH} queries to its CH oracle, we have that

$$\mathsf{Adv}^{\mathsf{CRGapDH}}_{(\mathbb{G},p,g)}(\mathcal{A},Q_{\mathrm{DDH}},Q_{\mathrm{CH}}) := \Pr\left[\mathcal{G}^{\mathsf{CRGapDH}}_{(\mathbb{G},p,g)}(\mathcal{A},Q_{\mathrm{DDH}},Q_{\mathrm{CH}})\right] \leq \epsilon,$$

where $\mathcal{G}_{(\mathbb{G},p,g)}^{\mathsf{CRGapDH}}(\mathcal{A}, Q_{\mathrm{DDH}}, Q_{\mathrm{CH}})$ is defined in Fig. 3. For simplicity, when the numbers of oracle queries are clear, we may ignore t, Q_{DDH} , and Q_{CH} so that we can focus on security bounds. We may then write the game as $\mathcal{G}_{(\mathbb{G},p,g)}^{\mathsf{CRGapDH}}(\mathcal{A})$ and the advantage function as $\mathsf{Adv}_{(\mathbb{G},p,g)}^{\mathsf{CRGapDH}}(\mathcal{A})$.

Lemma 1 (GapDH $\xrightarrow{\text{rewind}}$ CRGapDH [34, Lemma 6]). For any adversary \mathcal{A} that breaks the $(t, \epsilon, Q_{\text{DDH}}, Q_{\text{CH}})$ -CRGapDH assumption, there exists an adversary \mathcal{B} that breaks the $(t', \epsilon', Q'_{\text{DDH}})$ -GapDH assumption, where $t' \approx 2t, \epsilon' \approx Q_{\text{CH}}\epsilon^2$, and $Q'_{\text{DDH}} \approx 2Q_{\text{DDH}}$.

2.2 Signature

In Signal's X3DH and PQXDH protocols, user's semi-static keys are signed under their long-term keys, and the signatures included in the pre-key bundles. These signatures ensure that no one can upload a malicious semi-static key unless the user's long-term key is compromised.

We recall that a signature scheme Sig := (KGen, Sign, Ver) consists of the following algorithms:

$\mathcal{G}^{EUF-opCMA-DDH}_{Sig[(\mathbb{G},p,g)]}(\mathcal{A},Q_{\mathrm{SIGN}},Q_{\mathrm{DDH}})$:	$\underline{\operatorname{SIGN}(m)}$:	$\overline{\text{DDH}(X,Y,Z)}$:
1 $(vk = g^{sk} \in \mathbb{G}, sk \in \mathbb{Z}_p) \leftarrow s KGen$ 2 $(m^*, \sigma^*) \leftarrow s \mathcal{A}^{SIGN, DDH}((\mathbb{G}, p, g), vk)$ 3 $\mathbf{if} (m^*, \sigma^*) \notin \mathcal{L}_{\mathtt{sign}}$ $\wedge Ver(vk, \sigma^*, m^*) = 1$ 4 $\mathbf{return} \ 1$ 5 $\mathbf{else \ return} \ 0$	$ \begin{array}{l} \text{ // at most } \mathcal{Q}_{\text{SiGN}} \text{ queries} \\ \textbf{6 if } \exists \sigma \text{ s.t. } (m, \sigma) \in \mathcal{L}_{\texttt{sign}} \\ \textbf{7 return } \sigma \\ \textbf{8 } \sigma \leftarrow \texttt{s} \text{ Sign}(\texttt{sk}, m) \\ \textbf{9 return } \sigma \\ \textbf{10 } \mathcal{L}_{\texttt{sign}} := \mathcal{L}_{\texttt{sign}} \cup \{(m, \sigma)\} \\ \textbf{11 return } \sigma \end{array} $	// at most Q_{DDH} queries 12 return $[\![Z = X^{\text{DL}_g(Y)}]\!]$

Fig. 4. Game EUF-opCMA-DDH for a signature scheme Sig.

- KGen generates a public and secret key pair (vk, sk).
- Sign(sk, m) generates a signature σ of message m under secret key sk.
- $\operatorname{Ver}(\mathsf{vk}, m, \sigma)$ deterministically outputs 1 if σ is a valid signature of m under vk or 0 otherwise.

We require a signature scheme to be correct, namely, for all $(vk, sk) \leftarrow s KGen$ and all messages m, we have Ver(vk, m, Sign(sk, m)) = 1.

Existential UnForgeability against Chosen-Message Attacks (EUF-CMA) is a standard security notion for signature schemes and is also used in proving security of Signal's PQXDH handshake when modeling signatures in the previous computational analysis of [23]. Here, we require only a weaker form of EUF-CMA, namely, one-per message unforgeability [22] where an adversary is restricted to ask at most one signature per message. This is because, in proving XHMQV (similar to the computational proofs for X3DH [11] and PQXDH [23]), we rule out collisions in honestly generated keys early on, meaning that each semi-static public key will only be singed once by a user. Since Signal's initial handshake protocol uses the same long-term key both for the DH key exchange and as secret key in its XEdDSA signature scheme, we consider the following customized oneper message security notion, which augments the regular notion with a decisional Diffie-Hellman oracle DDH.

Definition 2 (One-per message unforgeability with a DDH **oracle).** Let Sig be a signature scheme based on a group description (\mathbb{G} , p, g). We say Sig is $(t, \epsilon, Q_{\text{SIGN}}, Q_{\text{DDH}})$ -EUF-opCMA-DDH-secure if for any adversary \mathcal{F} with running time at most t making at most Q_{SIGN} queries to its SIGN oracle and Q_{DDH} queries to its DDH oracle, we have

$$\mathsf{Adv}^{\mathsf{EUF-opCMA-DDH}}_{\mathsf{Sig}[(\mathbb{G},p,g)]}(\mathcal{A},Q_{\mathrm{SIGN}},Q_{\mathrm{DDH}}) := \Pr\left[\mathcal{G}^{\mathsf{EUF-opCMA-DDH}}_{\mathsf{Sig}[(\mathbb{G},p,g)]}(\mathcal{A},Q_{\mathrm{SIGN}},Q_{\mathrm{DDH}})\right] \leq \epsilon,$$

where $\mathcal{G}_{\mathsf{Sig}[(\mathbb{G},p,g)]}^{\mathsf{EUF}-\mathsf{opCMA-DDH}}(\mathcal{A}, Q_{\mathrm{SigN}}, Q_{\mathrm{DDH}})$ is defined in Figure 4. For simplicity, when the group and the numbers of oracle queries are clear, we may write the game as $\mathcal{G}_{\mathsf{Sig}}^{\mathsf{EUF}-\mathsf{opCMA-DDH}}(\mathcal{A})$ and the advantage function as $\mathsf{Adv}_{\mathsf{Sig}}^{\mathsf{EUF}-\mathsf{opCMA-DDH}}(\mathcal{A})$.

The Schnorr signature scheme can be proven EUF-opCMA-DDH secure by using a variant of the Discrete Logarithm assumption, where an attacker has to solve a discrete logarithm when given a DDH oracle. This additional DDH oracle

12 Rune Fiedler, Felix Günther, Jiaxin Pan, and Runzhi Zeng

$\mathcal{G}^{Real}_{Sig,Q}(\mathcal{A})$:	$\mathcal{G}^{Sim}_{Sig,Q}(\mathcal{A})$:	$RO(\mathit{qry})$ // In \mathcal{G}^{Sim}_{Sig} :
1 (vk, sk) \leftarrow Sig.KGen, $\mathcal{L}_{RO} := \emptyset$ 2 (m ₁ ,, m _Q) \leftarrow $^{\$}\mathcal{A}()$	1 $(vk,sk) \leftarrow sSig.KGen, \mathcal{L}_{RO} := \emptyset$ 2 $(m_1,, m_Q) \leftarrow s\mathcal{A}()$	 s if ∃h s.t. (qry, h) ∈ L_{RO} s return h
3 for $i \in \{1,, Q\}$ 4 $\sigma_i \leftarrow \text{s Sign}^{\text{RO}}(\text{sk}, m_i)$	3 for $i \in \{1,, Q\}$ 4 $(\sigma_i, qry_i, h_i) \leftarrow \text{sSIMsign}(vk, m_i)$	10 $h \leftarrow $ s RO.image 11 $\mathcal{L}_{RO} := \mathcal{L}_{RO} \cup \{(ary, h)\}$
5 $b_{guess} \leftarrow \mathcal{A}^{RO}(vk,sk,\{\sigma_i\}_{1 \leq i \leq Q})$ 6 return b_{guess}	5 $\mathcal{L}_{RO} := \mathcal{L}_{RO} \cup \{(qry_i, h_i)\}$ 6 $b_{guess} \leftarrow s \mathcal{A}^{RO}(vk, sk, \{\sigma_i\}_{1 \le i \le Q})$ 7 return b_{guess}	12 return h

Fig. 5. Security game for the δ -simulatability in Definition 3. Here, \mathcal{A} is stateful.

$DSA.Sign^{RO}(sk=x,m)$:	SIMsign(vk = X, m):
1 $r \leftarrow \mathbb{S} \mathbb{Z}_p$ 2 $R := g^r$ 3 $t := f(R)$ 4 $h := RO(m)$ 5 $s := r^{-1} \cdot (h + xt)$ 6 $\mathbf{return} (s, t)$	1 $(a, b) \leftarrow \mathbb{Z}_p^2$ 2 $R := X^a \cdot g^b$ 3 $t := f(R)$ 4 $h := b \cdot t \cdot a^{-1}$ 5 $s := t \cdot a^{-1}$ 6 return $((s, t), m, h)$ // Here gry is m .
Schnorr.Sign ^{RO} (sk = x, m):	SIMsign(vk = X, m):
1 $r \leftarrow \mathbb{Z}_p, R := g^r$ 2 $h := RO(R, m)$ 3 $s := (r + x \cdot h) \in \mathbb{Z}_p$ 4 $\mathbf{return} (s, R)$	$ \frac{1}{h} \stackrel{h \leftrightarrow \mathbb{Z}_p}{\underset{s \leftarrow \mathbb{Z}_p}{2 \ s \leftarrow \mathbb{Z}_p}} \\ \frac{2}{3} R := g^s \cdot (X^{-1})^h \\ 4 \text{ return } ((s, R), (R, m), h) // \text{Here } gry \text{ is } (R, m). $

Fig. 6. Signing (Sign) and simulated signing (SIMsign) algorithms for DSA and Schnorr signatures. For both schemes, the secret key is a random $x \leftarrow \mathbb{Z}_p$ and public key is $X := g^x$.

is to simulate the DDH oracle in Fig. 4 and the remainder of the proof reduces forging a signature to computing a discrete logarithm as in [47,33]. This variant of the Discrete Log assumption is implied by the standard Gap DH assumption. We can show that the (EC)DSA scheme is also EUF-opCMA-DDH secure by following the proof in [22] and introducing a DDH oracle in the underlying assumption.

To stay as close as possible to the original Signal X3DH handshake, which uses long-term keys for both DH key exchange and signing, our XHMQV protocol does the same. We thus require the following δ -simulatability notion, adapted from [22, Definition 5]. Essentially, δ -simulatability states that there is an algorithm that can efficiently generate signatures which are statistically close to real signatures, when allowed programming the random oracle.

Definition 3 (δ -Simulatability, adapted from [22, Definition 5]). Let Sig be a signature scheme whose signing algorithm uses a function RO modeled as a random oracle. Consider the two games defined in Figure 5. We say Sig is δ -simulatable if there exists an algorithm SIMsign such that for any (unbounded) adversary \mathcal{A} and $Q \in \mathbb{N}$, we have

$$\left|\Pr\left[\mathcal{G}_{\mathsf{Sig},Q}^{Real}(\mathcal{A})=1\right]-\Pr\left[\mathcal{G}_{\mathsf{Sig},Q}^{Sim}(\mathcal{A})=1\right]\right| \leq \delta.$$

Both the DSA and the schnorr signature schemes satisfy the δ -simulatability definiton with negligible δ . Fig. 6 presents the SIMsign algorithms for the DSA and Schnorr signature schemes. The one for DSA is from [22, Figure 5]. The idea here is that the construction of h in the SIMsign of DSA ensures that h is distributed uniformly at random unless a is not invertible (which happens with probability $\frac{1}{p}$). By a union bound, we have $\delta \leq \frac{Q}{p}$. For the full proof of the simulatability of DSA scheme, we refer to [22, Lemma 1]. For the Schnorr scheme, we have $\delta = 0$.

3 Security Model

We analyze the security of our XHMQV protocol with a security model for authenticated key exchange, following the seminal Bellare–Rogaway (BR)–style key exchange model [4], adapted to the setting of Signal's initial handshake in a series of prior work [11,12,8,23]. Our analysis is game-based and yields a computational guarantee.

BR-style key exchange models allow an adversary to interact with multiple honest users with multiple sessions each. It can fully control the network and deliver, modify, inject and suppress messages (through a SEND oracle). Furthermore, the adversary can (partially) compromise users: It can corrupt long-term keys (via the CORRUPTLTKEY oracle), corrupt semi-static keys (via the CORRUPTSSKEY oracle), reveal the ephemeral randomness of a session (via the REVEALRAND oracle), and reveal the session key of a session (via the REVEALSESSKEY oracle). This fine-grained control over user compromise allows us to model the maximum exposure guarantees that Signal strives for. Lastly, the adversary picks a *test session* (via a TEST oracle) for which it needs to distinguish between the real session key and a randomly sampled session key. Our model precludes the adversary from being able to trivially compute the session key in the test session. Formally, we do this via a "freshness" condition, which includes checking for protocol specific compromise scenarios with the so-called clean predicate. We give more details on this below.

Our model is directly based on the revision by [8,23], adapted mainly to capture the additional exposure scenarios that our XHMQV design protects against. We hence closely follow the exposition of [23]. Based on this line of work, we consider Bob, who first uploads his pre-key bundle, as responder, and Alice, who thereafter sends the first peer-specific message, as initiator.

3.1 Syntax and Notation

We follow the syntax of [23], who explicitly model the signatures on semi-static keys.⁶

⁶ As [23] analyzed PQXDH, their model includes two types of semi-static keys (DH and KEM keys), while ours focuses on DH keys only.

- 14 Rune Fiedler, Felix Günther, Jiaxin Pan, and Runzhi Zeng
- $\mathsf{KGenLT}() \cong (ltpk, ltsk)$: The long-term key generation algorithm that outputs a party's public-key/secret-key pair.
- KGenSS(ltsk) $s \rightarrow (sspk, sssk)$: The semi-static key generation algorithm that takes as input a long-term secret key ltsk and outputs a public-key/secret-key pair.
- Run(*ltsk*, *sssk*, *ltpk*, π , *m*) $s \rightarrow (\pi', m')$: The session execution algorithm that takes as input a party's long-term secret key *ltsk*, that party's semi-static secret keys *sssk*, all parties' long-term public keys *ltpk*, a session state π , and an incoming message *m*, and outputs an updated session state π' and a (possibly empty) outgoing message *m'*. The session sending the first message is set up by calling Run with a distinguished message m = (create, (ssid, type)), where ssid indicates the semi-static key to be used and type whether a full (type = full) or reduced (type = reduced) handshake should be performed.

Parties and sessions. In the model, each party $P \in [n_p]$ holds a long-term public-key/secret-key pair generated by KGenLT and may run multiple instances of the protocol (simultaneously or sequentially); we denote the *i*th such session of party P by π_P^i . The security game maintains the following information for each session:

- oid $\in [n_p]$: The identity of the session owner.
- pid ∈ $[n_p] \cup \{\star\}$: The identity of the intended peer, which may initially be unknown (indicated by \star).
- role \in {initiator, responder}: The role of the party.
- $st_{exec} \in \{\bot, running, accepted, rejected\}$: The status of this session's execution.
- sid ∈ $\{0,1\}^* \cup \{\bot\}$: A session identifier defining partnering.
- cid ∈ $\{0,1\}^* \cup \{\bot\}$: A contributive identifier, defining a preliminary form of partnering (often as a substring or prefix of the session identifier) for the case the session is not yet bound to an authenticated peer [21].
- $\mathsf{K} \in \mathcal{K}_{\mathsf{KE}} \cup \{\bot\}$: The session key established in this session, initialized to \bot .
- type \in {full, reduced}: Indicator whether an ephemeral pre-key was used (type = full) for key establishment, or not (type = reduced).
- coins ∈ \mathcal{R}_{KE} : The random coins from the randomness space \mathcal{R}_{KE} used in the execution of Run; set by the game and read-only thereafter.
- $-\operatorname{sspk} \in \{0,1\}^* \cup \{\bot\}$: The semi-static public key used in this session.
- state $\in \{0,1\}^*$: Potential session state of the protocol.

For bookkeeping in the security game we additionally introduce the following flags, which are not accessible by the protocol sessions:

- revrand \in {true, false} indicates whether the random coins π .coins have been revealed via a REVEALRAND query. The default value is false.
- $pcorr \in \{true, false\}$ indicates whether the peer's long-term key was corrupted at the point in time when this session accepted. The default value is false.

Session partnering. We say that two sessions π_U^i, π_V^j are partnered if they agree on the session identifier, i.e., if π_U^i .sid = π_V^j .sid $\neq \perp$. Contributive identifiers indicate that two sessions may eventually derive the same session key but are not fully partnered (yet).

15

3.2 Security Game

We formalize security as key indistinguishability through game $\mathcal{G}_{\mathsf{KE}}^{\mathsf{KI}}(\mathcal{A})$ in Figure 7. The game first fixes a bit $b_{\mathsf{test}} \leftarrow \{0, 1\}$. Next, it samples long-term and semi-static keys for all users (one long-term key and n_{ss} for each of the n_p users). Then, the adversary \mathcal{A} gets all public keys ltpk, sspk and can access several oracles to control interaction with and between honest users.

- SEND(U, i, m): Sends message m to session π_U^i , which corresponds to executing Run $(ltsk_U, sssk_U, ltpk, \pi_U^i, m)$, saving the updated session state π' as π_U^i , and returning the outgoing message m' to the adversary.
- CORRUPTLTKEY(U): Returns party U's long-term secret key $ltsk_U$ to the adversary; recorded through the flag corrltk_U.
- CORRUPTSSKEY(U, ssid): Returns party U's semi-static secret key $sssk_U^{ssid}$ to the adversary; recorded through the flag corrssk $_{sspk_{T}^{ssid}}$.⁷
- REVEALRAND(U, i): Returns the randomness π_U^i .coins of session π_U^i to the adversary, recorded in the session through the flag revrand.
- REVEALSESSKEY(U, i): Returns the session key π_U^i .K of session π_U^i to the adversary, recorded in the session through the flag revealed.
- TEST(U, i): If a TEST query has been made before or session π_U^i has not accepted, then return \perp . Otherwise; if $b_{\text{test}} = 0$, return π_U^i .K, otherwise return a randomly sampled session key from the protocol's key space \mathcal{K}_{KE} . Record the test session as $\pi^* \leftarrow \pi_U^i$.

Finally, the adversary outputs a bit b_{guess} . If the adversary was able to violate soundness of the protocol, the game outputs 1. If the adversary can trivially compute the session key of the test session, we say that it violates *freshness* and the game outputs a random bit. Otherwise, the game returns whether the adversary's guess b_{guess} correctly decided the test bit b_{test} , i.e., if the adversary can distinguish between the real session key and a randomly sampled session key. The advantage of the adversary is its chance of winning minus its chance of randomly guessing correctly $(\frac{1}{2})$.

Definition 4 (Key Indistinguishability [23]). Let KE be a key exchange protocol and \mathcal{A} an adversary against the key indistinguishability (KI) game $\mathcal{G}_{\mathsf{KE}}^{\mathsf{KI}}(\mathcal{A})$ in Figure 7. We say that KE achieves $(t, \epsilon, (q_{\mathrm{Snd}}, q_{\mathrm{CorrLT}}, q_{\mathrm{CorrSS}}, q_{\mathrm{RevR}}, q_{\mathrm{RevSK}}))$ key indistinguishability, if for any adversary \mathcal{A} against $\mathcal{G}_{\mathsf{KE}}^{\mathsf{KI}}(\mathcal{A})$ with running time at most t and making at most $q_{\mathrm{Snd}}, q_{\mathrm{CorrLT}}, q_{\mathrm{CorrSS}}, q_{\mathrm{RevR}}, resp. q_{\mathrm{RevSK}}$ queries to its SEND, CORRUPTLTKEY, CORRUPTSSKEY, REVEALRAND, resp. REVEALSESSKEY oracles, we have that $\mathsf{Adv}_{\mathsf{KE}}^{\mathsf{KI}}(\mathcal{A}) = \Pr\left[\mathcal{G}_{\mathsf{KE}}^{\mathsf{KI}}(\mathcal{A})\right] - \frac{1}{2} \leq \epsilon$. Note that the model restricts the adversary to a single query to the TEST oracle.

⁷ Following [23], we base the **corrssk** flag on the semi-static public key and not on the semi-static id **ssid**, since initiator sessions do not learn the **ssid**.

	$\begin{array}{l} \underbrace{fresh(\pi^*):}_{13} & \text{if } \pi^*.revealed = true: return false \\ \text{// test session is revealed} \\ 14 & \text{if } \exists \pi_V^j \neq \pi^*: (\pi_V^j.sid = \pi^*.sid \land \pi_V^j.revealed = true): \\ & \textbf{return false} \\ \text{// test session's partner is revealed} \\ 15 & \textbf{return clean}_{KE}^{KE}(\pi^*) \\ & \text{// test session is clean wrt. to its handshake type (full resp. reduced)} \\ & \textbf{sound}(): \end{array}$
 9 bguess ←s A((tpk, sspk) // run adversary 10 if sound() = false: return true // adversary wins if it breaks soundness 11 if fresh(π*) = false: bguess ← 0 // attack invalid if test session is not fresh 12 return [[bguess = btest]] // determine win or loss 	16 return \forall distinct $\pi, \pi', \pi'' : ($ 17 $(\pi.\text{sid} = \pi'.\text{sid} \neq \bot \implies \pi.\text{K} = \pi'.\text{K} \land \pi.\text{type} = \pi'.\text{type} \land \pi.\text{cid} = \pi'.\text{cid})$ 18 and $(\pi.\text{sid} = \pi'.\text{sid} \neq \bot \land \pi.\text{role} = \text{initiator} \implies \pi'.\text{role} = \text{responder})$ 19 and $(\pi.\text{sid} = \pi'.\text{sid} = \pi''.\text{sid} \neq \bot \implies \pi.\text{type} = \text{reduced}))$ 19 and $(\pi.\text{sid} = \pi'.\text{sid} = \pi''.\text{sid} \neq \bot \implies \pi.\text{type} = \text{reduced}))$
$ \frac{\text{SEND}(U, i, m):}{20 \text{ if } \pi_U^i = \bot:} \\ \text{// initiate session: for responders, we have m = (\text{create, (sid, t)}) \\ 21 \pi_U^i \text{.oid} \leftarrow U \text{ // set owner identity} \\ 22 \text{if } m = (\text{create,}): \pi_U^i \text{.role} \leftarrow \text{responder} \\ \text{// set responder role} \\ 23 \text{else } \pi_U^i \text{.role} \leftarrow \text{initiator} \\ \text{// set initiator role} (m \text{ is first message}) \\ 24 \pi_U^i \text{.coins} \leftarrow \Re_{\text{KE}} \text{ // sample session randomness} \\ 25 \pi_U^i \text{.stexec} \leftarrow \text{running} \\ 26 (\pi_U^i, m') \leftarrow \text{Run}(ltsk_U, sssk_U, ltpk, \pi_U^i, m) \\ \text{// run session, random coin in } \pi_U^i \\ 27 \text{if } \pi_U^i \text{.stexec} = \text{accepted}: \\ \text{// flag if peer was corrupted upon acceptance} \\ 28 \pi_U^i \text{.pcorr} \leftarrow \text{corrlik}_{\pi_U^i, \text{pid}} \\ 29 \text{return } (m', \pi_U^i \text{.stexec}) \text{ // return message and session} \\ \hline \frac{\text{TEST}(U, i):}{30 \text{if } \pi_U^i = \bot \text{ or } \pi_U^i \text{.stexec} \neq \text{accepted or } \pi^* \\ \text{return } \bot \\ \text{// session does not exist, has not accepted yet, or test alread} \\ 31 \pi^* \leftarrow \pi_U^i \text{ // record test session} \\ 32 \text{K}_0 \leftarrow \pi_U^i \text{.K} \\ 33 \text{K}_1 \leftarrow \text{s} \text{K}_{\text{KE}} \end{cases}$	$\frac{\text{CORRUPTLTKEY}(U):}{35 \text{ corrlt} k_U \leftarrow \text{true } // \text{mark long-term key corrupted}}{36 \text{ return } llsk_U // return long-term secret key}$ $\frac{\text{CORRUPTSSKEY}(U, ssid):}{37 \text{ corrssk}_{sspksod} \leftarrow \text{true } // \text{mark semi-static key corrupted}}{38 \text{ return } sssk_U[\text{ssid}] // return semi-static secret key}}$ $\frac{\text{REVEALRAND}(U, i):}{9 \text{ if } \pi_U^i = \bot : \text{return } \bot // \text{session does not exist}}{41 \text{ return } \pi_U^i. \text{coins } // \text{ return session's random coins}}$ $\frac{\text{REVEALSESSKEY}(U, i):}{\sqrt{2} \text{ session does not exist}} + 2 \text{ if } \pi_U^i = \bot \text{ or } \pi_U^i. \text{ session does not exist}}{\sqrt{2} \text{ session does not exist}} + 42 \text{ if } \pi_U^i = \bot \text{ or } \pi_U^i. \text{ state}} + 2 \text{ if } \pi_U^i = \bot \text{ or } \pi_U^i. \text{ states on the exist of has not yet derived session key}}{44 \text{ return } \pi_U^i. K // \text{ return session key}} + 2 \text{ is}$

Fig. 7. Key indistinguishability (KI) game for key exchange protocol KE (top), in which adversary \mathcal{A} has access to oracles SEND, TEST, CORRUPTLTKEY, CORRUPTSSKEY, REVEALRAND, and REVEALSESSKEY (bottom), and wrt. to (protocol-specific) clean predicates clean^{KE} for KE. The clean predicates satisfied by our XHMQV protocol are given in Figure 8.

Soundness. The model captures soundness, i.e., that session identifiers correctly reflect protocol executions, via the **sound** predicate. Specifically, it excludes the following three cases:



Fig. 8. Clean predicates for the XHMQV protocol. Predicates satisfied by XHMQV but not X3DH are highlighted.

- (i) Two sessions accept with the same session identifier, but derive different session keys, indicate different handshake types (full vs. reduced), or do not agree on their contributive identifiers (Fig. 7, line 17).
- (ii) Two initiator sessions accept with the same session identifier (Fig. 7, line 18).
- (iii) Three sessions accept with the same session identifier in full handshake type (Fig. 7, line 19).

Freshness. The security model excludes cases which allow the adversary to win trivially: If the adversary has revealed the test session (line 13), if it has revealed a session partnered to the test session (line 14), or if it can compute the session key with the keys it has compromised via queries to the CORRUPTLTKEY, CORRUPTSSKEY, REVEALRAND oracles (line 15). The clean predicates are protocol-specific since different protocols require different (combinations of) keys to compute the session key.

3.3 Clean Predicates for XHMQV

In our XHMQV protocol, key derivation combines both of Alice's keys (long-term and ephemeral) with each of Bob's keys (long-term, semi-static, and ephemeral), yielding six combinations in a full handshake. If Bob's ephemeral key is unavailable (because the Signal server ran out of pre-key bundles for Bob), the number of combinations accordingly reduce to four. As long as one of these six (or four) combinations is not known to the adversary, it should not be able to compute the session key.

Formally, we model this with the clean^{XHMQV} predicates given in Figure 8, more specifically clean^{XHMQV} and clean^{XHMQV}_{reduced} capturing the expected security of a full resp. reduced handshake. These predicates ensure that at least one of the combinations long-term–semi-static (clean_{LTSS}), ephemeral–long-term (clean_{ELT}), ephemeral–semi-static (clean_{ESS}), or long-term–long-term (clean_{LTLT}) is uncompromised. For a full handshake, the predicate is also satisfied if the ephemeral– ephemeral (clean_{EE}) or long-term–ephemeral (clean_{LTE}) combination is uncompromised.

To check that an ephemeral contribution is not compromised, we use the $\mathsf{clean}_{\mathsf{peerE}}$ sub-predicate: Here, we check that for an initiator session the randomness of a contributively partnered (i.e., via cid) responder session is not revealed or for a responder session the randomness of a partnered (via cid) initiator session is nor revealed.

For cases involving semi-static keys, i.e., $clean_{LTSS}$ and $clean_{ESS}$, simply requiring the semi-static key to be uncompromised is not sufficient: We additionally check that the responder's long-term key is uncompromised at the time when the initiator accepts via the pcorr flag or that the semi-static key was honestly generated. The first additional check was introduced by [23] to model that the responder's semi-static key is authenticated via a signature under the responder's long-term key. The second additional check is a novel strengthening to capture that even if the responder's long-term key is compromised, we expect security as long as the semi-static key is not compromised.

Compared to X3DH, the clean predicates of our XHMQV protocol give a strictly stronger guarantee: For all cases in which X3DH is clean, our XHMQV protocol is clean as well. Furthermore, if only the long-term–long-term or only the long-term–ephemeral combination is clean, then our clean predicate for XHMQV demands security, but X3DH does not satisfy this. This is the stronger maximum-exposure guarantee of our XHMQV protocol.

4 Our XHMQV Protocol

We now introduce our XHMQV protocol, using a group (\mathbb{G}, p, g) and a signature scheme Sig (like in Signal, Sig may build on \mathbb{G}). Further, denoting by \mathcal{K} the session key space, we use two key derivation functions $\mathsf{KDF}_f, \mathsf{KDF}_r \colon \{0, 1\}^* \to \mathcal{K}$ for deriving session keys, and four hash functions $h_0, h_1, h_2, h_3 \colon \{0, 1\}^* \to \mathbb{Z}_p$. (The indices indicate that in practice, we expect those to be instantiated using the same algorithm using domain separation.) We fully define XHMQV in Fig. 9, along with a variant XHMQV⁺. In XHMQV, long-term keys are used for both the DH key exchange key and for signing semistatic keys (as is done in X3DH). XHMQV⁺ instead use independent signing keys, ensuring proper key separation. We will focus on and prove security for XHMQV, i.e., with key reuse, given this is closest to Signal' deployment. We emphasize that our analysis capturing this key reuse is in contrast to prior work, which either modeled the signatures using separate signing keys [35,53,24,6,23,29] or not at all [11,12,8]. Our results naturally apply to XHMQV⁺, under weaker assumptions.

Adapting HMQV. To understand the design of XHMQV, we start by introducing how we adapt the idea from HMQV to the Signal setting. Let (A, a) and (X, x)be the long-term (identity) key pair and ephemeral key pair of Alice, and let (B, b), (S, s), and (Y, y) be the long-term, semi-static, and ephemeral key pairs of Bob, respectively. Following X3DH, a handshake session is either in the full mode or in the reduced mode.

Recall that in HMQV only involves two key pairs (long-term, ephemeral) on each side. In a full handshake, we hence need to integrate the additional semi-static key of the responder (S, s) into the combination of DH secrets. Our approach is to view (S, s) as another "long-term secret" of Bob (in the HMQV sense), and integrate it into DH combinations in the following way (from Alice's view):

$$(YB^{e_1})^{x+da}//\operatorname{HMQV} \longrightarrow (YB^{e_1}S^{e_2})^{x+da}//\operatorname{XHMQV}$$

where d, e_1 , e_2 are exponents derived using hash functions, following HMQV's design.

In a reduced handshake, (Y, y) is unavailable and the semi-static key pair (S, s), which is regularly rotated, serves to provide (delayed) forward secrecy. Here, we can essentially fall back to the HMQV equation, dropping (Y, y) and e_2 , and using the semi-static key in place of the responder's ephemeral key (again, from Alice's view):

$$(YB^{e_1})^{x+da} // \operatorname{HMQV} \longrightarrow (B^{e_1}S)^{x+da} // \operatorname{XHMQV},$$

Let us discuss how to derive d, e_1 , and e_2 . These three exponents are chosen to prevent a linear relationship between Y, B, and S. We keep $d := h_0(A, X)$ as in HMQV and let $e_1 := h_1(B, Y, S)$ and $e_2 := h_2(B, Y, S)$. That is, both e_1 and e_2 include the long-term, semi-static, and ephemeral information of Bob. This choice is not only for excluding trivial attacks (e.g., if S is not involved in e_2 , then the adversary can contruct a malicious Y to eliminate S), but also essential for our security proof, namely, allows us to construct reductions from CRGapDH assumption (cf. Definition 1).

Similar to X3DH, we sign semi-static keys in XHMQV using the long-term key, to protect against active attackers replacing them. While adding signature upgrades the protocol from weak forward secrecy security to full forward secrecy, the key reuse issue requires the signature scheme to have stronger security

than the standard unforgeability definition (cf. Definition 2). For completeness, we also provide a version that use independent signing key, $XHMQV^+$, see the highlighted differences in Fig. 9.

Security of XHMQV. The security theorem for XHMQV is given in Theorem 1 below. One can use it to derive a similar result for $XHMQV^+$ (except that we would require Sig to achieve EUF-opCMA instead of EUF-opCMA-DDH).

Theorem 1. The XHMQV protocol given in Fig. 9 achieves $(t, \epsilon, (q_{\text{Snd}}, q_{\text{CorrLT}}, q_{\text{CorrSS}}, q_{\text{RevR}}, q_{\text{RevSK}}))$ -key indistinguishability defined in Definition 4.

Specifically, let n_u , n_{ss} , and n_s be the numbers of users, semi-static keys per user, and sessions in the key exchange game \mathcal{G}_{KE}^{KI} , respectively. Let (\mathbb{G}, p, g) be a group and both XHMQV and the signature scheme Sig of XHMQV be built on this group. If the $(t_0, \epsilon_{CRGapDH}, Q_{DDH}, Q_{CH})$ -CRGapDH assumption holds in \mathbb{G} and Sig has $(t_1, \epsilon_{EUF-opCMA-DDH}, Q_{SIGN}, Q'_{DDH})$ -EUF-opCMA-DDH security and δ_{sim} simulatability, then for any adversary \mathcal{A} with running time t, we have

$$\begin{split} \epsilon &\leq \frac{(n_u + n_u n_{ss} + n_s)^2}{p} + n_u \cdot \epsilon_{\text{EUF-opCMA-DDH}} \\ &+ \left(2n_u^2 + 2n_u n_s + n_u^2 n_{ss}\right) \cdot \delta_{sim} \\ &+ \left(n_u^2 + n_u n_{ss} + n_u^2 n_{ss} \\ &+ n_u n_s + n_s n_u n_{ss} + n_{ss}^2\right) \cdot \left(\epsilon_{\text{CRGapDH}} + \frac{Q_{RO}}{p}\right) \end{split}$$

where Q_{RO} is the number of random oracle queries issued by \mathcal{A} across h_0 , h_1 , h_2 , h_3 , KDF_f , and KDF_r , and $t \approx t_0 \approx t_1$, $Q_{CH} \approx Q_{RO}$, $Q_{DDH} \approx Q_{RO}$, $Q'_{DDH} \approx Q_{RO}$, and $Q_{SIGN} \approx n_{ss}$. Without loss of generality, we assume that \mathcal{A} does not issue repeated queries since they return the same response, so we also have $q_{Snd} \leq 2n_s$, $q_{CorrLT} \leq n_u$, $q_{CorrSS} \leq n_u \cdot n_{ss}$, $q_{RevR} \leq n_s$, and $q_{RevSK} \leq n_s$.

Remark 1 (Use the same KDF or hash function, with domain separation). For the key derivation functions KDF_f and KDF_r , as well as the hash functions h_0 , h_1 , h_2 , and h_3 , the indices indicate that in practice, we expect those to be instantiated using the same algorithm using domain separation [3]. This can be achieved, e.g., via prefixing labels, similar to domain separation already done in Signal [46, Section 2.5].

Remark 2 (On the "Another Look at HMQV".). Menezes [42] identified several vulnerabilities in HQMV, including small-subgroup and unknown key-share attacks, and proposed corresponding countermeasures. These considerations are also reflected in our XHMQV design. For instance, XHMQV includes the session context in key derivation process to prevent unknown key-share attacks. In real-world implementations, XHMQV must additionally validate all group elements, ensuring the correct order or group membership of long-term public keys, semi-static public keys, and ephemeral pre-keys.



Fig. 9. Our XHMQV protocol and its variant $XHMQV^+$. Elements in gray boxes are only present in XHMQV, those in dashed boxes are only present in $XHMQV^+$.

5 Security Proof of XHMQV (Theorem 1)

We start with \mathbf{G}_0 which is the original key exchange security game $\mathcal{G}_{\mathsf{KE}}^{\mathsf{KI}}(\mathcal{A})$ with two modifications to simplify our proof: (1) For each session π_{U}^i , we add a variable $\pi_{\mathsf{U}}^i.ctxt$ into the session to store the context value ctxt used in computing the session key. By using this variable, we can easily verify if a hash query to KDF_f or KDF_r matches the public keys (i.e., long-term, semi-static, and ephemeral keys) used in a session. (2) We use random oracles to simulate all hash functions. Since the proof is in the ROM, these changes are conceptual, so we have

$$\operatorname{Adv}_{\operatorname{KE}}^{\operatorname{KI}}(\mathcal{A}) = \operatorname{Adv}_{\operatorname{XHMQV}}^{\operatorname{G}_0}(\mathcal{A}).$$

<u>GAME</u> **G**₁: This game excludes collisions of DH keys. Specifically, we abort the game (i.e., overwrite \mathcal{A} 's output with 0) if any two honestly generated DH key pairs collide. In the game, there are at most n_u long-term key pairs, n_{ss} semi-static key pairs for each user, and n_s ephemeral DH key pairs. By the birthday bound (over \mathbb{G} with $|\mathbb{G}| = p$), we have

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathbf{G}_0}(\mathcal{A}) \leq \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathbf{G}_1}(\mathcal{A}) + \frac{(n_u + n_u \cdot n_{ss} + n_s)^2}{p}$$

<u>GAME G2</u>: We introduce a new abort event Forge in G2: If \mathcal{A} forges a signature on a semi-static key with respect to an uncorrupted long-term key, then we abort the game. We denote this event as Forge. If Forge does not happen, then the distribution of G2 is the same as the one of G1. In the following, we construct an adversary \mathcal{F} against the EUF-opCMA-DDH security of Sig to bound Pr [Forge]. In the description of \mathcal{F} , we <u>underline</u> how we embed the challenge into the game and how we get the forgery if Forge happens.

Reduction $\mathcal{F}^{\text{DDH},\text{SIGN}}(\mathsf{vk}^* = B^*)$:

– Initialize \mathbf{G}_2 :

- 1. Initialize two patch key lists $\mathcal{L}_{patch-f}$ (for full sessions) and $\mathcal{L}_{patch-r}$ (for reduced sessions) to store all keys that are needed to be patched to the random oracle. For example, if there is a tuple (A, B, S, Y, X, K) in $\mathcal{L}_{patch-f}$, then it means that there is a session π such that $\pi.K = K$ and $\pi.ctxt = (A, B, S, Y, X)$.
- 2. Initialize a candidate key list (for full sessions) $\mathcal{L}_{\mathsf{skey-f}}$ to deal with the case that the adversary first computes the key itself (and queries the RO) and finishes the session later. Similarly, initialize and a key candidate list $\mathcal{L}_{\mathsf{skey-r}}$ for reduced sessions. For example, if there is a tuple (A, B, S, Y, X, K) in $\mathcal{L}_{\mathsf{skey-f}}$, then it means that \mathcal{A} has queried KDF_f on (A, B, S, Y, X, DH)(where DH is the correct XHMQV secret). Here (A, B, S, Y, X) may correspond to the context of a session (that is not yet completed). If \mathcal{A} completes such a session, then we can set its session key as K.
- 3. Guess a party B^* and set its long-term public key to B^* . When generating semi-static keys S of B^* , \mathcal{F} queries SIGN(S) to get the signatures. Note that we exclude collisions of all DH key pairs, so all S are unique and we only need to query SIGN(S) once for each S.

- 4. For all other parties, run KGenLT and KGenSS faithfully. After this step, \mathcal{F} has ltpk, sspk, and ltsk (except for the long-term secret key of B^{*}).
- 5. Run $b_{guess} \leftarrow \mathcal{A}(ltpk, sspk)$ and proceed the game as in \mathbf{G}_1 . The simulation of random oracles and game oracles are simulated in the way described below.
- 6. If the adversary never forges a valid signature on one of B^* 's semi-static keys before corrupting B^* , then abort the simulation and return a random forgery.
- Simulation of oracles defined in $\mathcal{G}_{\mathsf{KF}}^{\mathsf{KI}}$:
 - All oracles except SEND and CORRUPTLTKEY are simulated as defined in \mathcal{G}_{KE}^{Kl} . For CORRUPTLTKEY, \mathcal{F} aborts if \mathcal{A} queried for the long-term secret key of B^{*} (in which case Forge cannot happen afterwards) and answers with the corresponding secret key for all other users. For SEND queries, if the query does not involve B^{*}, then we also simulate the oracles in the same way as in \mathcal{G}_{KE}^{Kl} .
 - SEND(B^* , i, m): We only describe the changes to the Run algorithm and the setup of session keys.
 - * $\operatorname{Run}(ltsk_{B^*}, sssk_{B^*}, ltpk, \pi_{B^*}^i, (create, ssid, type))$ responder creation: Simulated as specified.
 - * $\operatorname{Run}(\operatorname{ltsk}_{B^*}, \operatorname{sssk}_{B^*}, \operatorname{ltpk}, \pi^i_{B^*}, m = (U, S, Y, \sigma))$ initiator: We describe handling a full handshake $(Y \neq \bot)$; handling the reduced mode works analogously.
 - 1. If σ is a valid signature of B^* and $S \notin sspk_{B^*}$, then we abort the simulation and output (S, σ) .
 - 2. Generate the ephemeral key pair (X, x).
 - 3. Let U be the long-term public key of U. If there exists an entry (B, U, S, Y, X, K) in $\mathcal{L}_{\mathsf{skey-f}}$, then retrieve this K. Otherwise, sample a random K and record (B, U, S, Y, X, K) in $\mathcal{L}_{\mathsf{patch-f}}$.
 - 4. Set $\pi^i_{\mathsf{B}^*}.\mathsf{K} := K$.
 - * $\operatorname{Run}(\operatorname{ltsk}_{B^*}, \operatorname{sssk}_{B^*}, \operatorname{ltpk}, \pi^i_{B^*}, m' = (U, epk_U = X))$ responder completion: We describe handling a full handshake; handling the reduced mode works analogously.
 - 1. Get the semi-static secret s and the ephemeral secret y (and their public keys) of this session. Let U be U's long-term public key.
 - 2. If there exists an entry (U, B, S, Y, X, K) in $\mathcal{L}_{\mathsf{skey-f}}$, then retrieve this K. Otherwise, sample a random K and record (U, B, S, Y, X, K) in $\mathcal{L}_{\mathsf{patch-f}}$.
 - 3. Set $\pi_{B^*}^i.K := K$.
 - SEND($\mathbf{U}, i, m = (\mathbf{B}^*, S, \sigma)$): If σ is a valid signature and $S \notin sspk_{\mathbf{B}^*}$, then we abort the simulation and output (S, σ) .
- Simulation of random oracles h_0 , h_1 , h_2 , h_3 , KDF_f , and KDF_r : We only discuss new queries. For repeated queries, \mathcal{F} returns the previously recorded values.
 - $\mathsf{KDF}_f(DH, A, B, S, Y, X)$:
 - 1. If $DDH(XA^d, YB^{e_1}S^{e_2}, DH)$ outputs 1:
 - * If there exists (A, B, S, Y, X, K) in $\mathcal{L}_{patch-f}$, retrieve this K.

- * If not, then sample a random key K and record (A, B, S, Y, X, K) in \mathcal{L}_{skey-f} .
- * Set $KDF_f(DH, A, B, S, Y, X) := K$
- * Return K.
- 2. Otherwise, simulate via lazy sampling.
- $\mathsf{KDF}_r(DH, A, B, S, X)$:
 - 1. If $DDH(XA^d, SB^{e_1}, DH)$ outputs 1:
 - * If there exists an entry (A, B, S, X, K) in $\mathcal{L}_{patch-r}$, retrieve this K.
 - * If not, then sample a random key K and record (A, B, S, X, K) in $\mathcal{L}_{\mathsf{skev}-\mathsf{r}}$.
 - * $\mathsf{KDF}_r(DH, A, B, S, X) := K.$
 - * Return K.
 - 2. Otherwise, simulate via lazy sampling.

If Forge occurs, meaning that \mathcal{A} forges a signature for a party's semi-static key before corrupting the party, then \mathcal{F} can use this forgery to win the unforgeability game if B^* is the party whose signature \mathcal{A} has forged. Since the guess of \mathcal{B}^* is independent of \mathcal{A} 's view, we have

$$\frac{1}{n_u} \cdot \Pr\left[\texttt{Forge}\right] \leq \Pr\left[\mathsf{Adv}_{\mathsf{Sig}}^{\mathsf{EUF-opCMA-DDH}}(\mathcal{F})\right]$$

Therefore, we have

$$\begin{split} \mathsf{Adv}^{\mathbf{G}_1}_{\mathsf{XHMQV}}(\mathcal{A}) &\leq \mathsf{Adv}^{\mathbf{G}_2}_{\mathsf{XHMQV}}(\mathcal{A}) + \Pr\left[\texttt{Forge}\right] \\ &\leq \mathsf{Adv}^{\mathbf{G}_2}_{\mathsf{XHMQV}}(\mathcal{A}) + n_u \cdot \mathsf{Adv}^{\mathsf{EUF}\text{-opCMA-DDH}}_{\mathsf{Sig}}(\mathcal{F}). \end{split}$$

Splitting G₂ **into clean cases:** We now divide the proof into six sub-cases according to which clean predicate the test session π^* satisfies. Let $\mathbf{G}_2 \wedge \operatorname{clean}_X$ denote \mathbf{G}_2 conditioned on clean predicate clean_X being satisfied. We can bound the advantage of \mathcal{A} in \mathbf{G}_2 by the sum of \mathcal{A} 's advantages in $\mathbf{G}_2 \wedge \operatorname{clean}_X$, for $\operatorname{clean}_X \in \{\operatorname{clean}_{\mathsf{LTLT}}, \operatorname{clean}_{\mathsf{LTE}}, \operatorname{clean}_{\mathsf{LTSS}}, \operatorname{clean}_{\mathsf{ELT}}, \operatorname{clean}_{\mathsf{ELT}}\}$ being one of the clauses in $\operatorname{clean}^{\mathsf{XHMQV}}$:

$$\mathsf{Adv}^{\mathbf{G}_2}_{\mathsf{XHMQV}}(\mathcal{A}) \leq \sum_{\substack{\mathsf{clean}_X \in \{\mathsf{clean}_{\mathsf{LTLT}}, \mathsf{clean}_{\mathsf{LTE}}, \mathsf{clean}_{\mathsf{LTSS}}, \\ \mathsf{clean}_{\mathsf{ELT}}, \mathsf{clean}_{\mathsf{ESS}}, \mathsf{clean}_{\mathsf{EE}}\}} \mathsf{Adv}^{\mathbf{G}_2 \wedge \mathsf{clean}_X}_{\mathsf{XHMQV}}(\mathcal{A}).$$

Note that predicates $\mathsf{clean}_{\mathsf{LTE}}$ and $\mathsf{clean}_{\mathsf{EE}}$ only apply to a test session in full handshake mode (π^* .type = full).

We bound the probabilities of each predicate in Appendix A. For space reasons, we summarize the resulting bound in the following lemma and only give a proof sketch illustrating the common proof strategy across the clean cases.

Lemma 2. With the notations and assumptions from Theorem 1, let $\epsilon_{CRGapDH}$ be the advantage bound for the CRGapDH assumption in \mathbb{G} and δ_{sim} be the bound

for simulatability of Sig, we have

$$\begin{split} \mathsf{Adv}^{\mathsf{G}_2}_{\mathsf{XHMQV}}(\mathcal{A}) &\leq n_u n_u & \cdot (2\delta_{sim} + \epsilon_{\mathsf{CRGapDH}} + Q_{RO}/p) & //\operatorname{clean_{LTLT}} \\ &+ n_u n_s & \cdot (\delta_{sim} + \epsilon_{\mathsf{CRGapDH}} + Q_{RO}/p) & //\operatorname{clean_{LTE}} \\ &+ n_u^2 n_{ss} & \cdot (\delta_{sim} + \epsilon_{\mathsf{CRGapDH}} + Q_{RO}/p) & //\operatorname{clean_{LTSS}} \\ &+ n_u n_s & \cdot (\delta_{sim} + \epsilon_{\mathsf{CRGapDH}} + Q_{RO}/p) & //\operatorname{clean_{LTSS}} \\ &+ n_s n_u n_{ss} & \cdot (\epsilon_{\mathsf{CRGapDH}} + Q_{RO}/p) & //\operatorname{clean_{ELT}} \\ &+ n_s^2 & \cdot (\epsilon_{\mathsf{CRGapDH}} + Q_{RO}/p) & //\operatorname{clean_{ESS}} \\ &+ n_s^2 & \cdot (\epsilon_{\mathsf{CRGapDH}} + Q_{RO}/p) & . & //\operatorname{clean_{EESS}} \end{split}$$

Let us illustrate the proof strategy at the example of the $\mathsf{clean}_{\mathsf{LTLT}}$ case, i.e., bounding $\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathbf{G}_2 \wedge \mathsf{clean}_{\mathsf{LTLT}}}(\mathcal{A})$. We proceed via the following steps:

- 1. We first guess the initiator and responder party of the test session π^* , picking two user identities A^* and B^* uniformly at random and aborting if that guess was incorrect. This leads to the multiplicative security loss of n_u^2 .
- 2. We then use the δ -simulatability of Sig to replace signatures of A^* and B^* . More specifically, we change the game such that all signatures of these two users are generated by running SIMsign and programming random oracles. This step allows us to simulate signatures only using public keys, which is essential for later reductions because we do not have the secret keys when constructing the CRGapDH reduction.
- 3. We then use the CRGapDH assumption over \mathbb{G} to exclude the case that the adversary queries the random oracle on the hash input of π^* 's session key. To this end, we will embed the CRGapDH challenge into the public keys of the two users \mathbb{A}^* and \mathbb{B}^* . Roughly, if \mathcal{A} queries KDF_f or KDF_r on (DH, ctxt), where $ctxt = \pi^*.ctxt$ and DH is the correct XHMQV's DH shared secret, then we can extract a CRGapDH solution from such a query.

The other proof cases have a similar structure: First, we guess where to embed the CRGapDH challenge, leading to different guessing losses depending on the clean case. Second, for all clean cases involving a long-term key, i.e., the first four cases, we use the simulatability of Sig to replace the long-term key, resulting in a δ_{sim} term. Thirdly, we reduce to CRGapDH, obtaining the term $\epsilon_{\text{CRGapDH}} + Q_{\text{RO}}/p$. By combining all probability bounds, we obtain

$$\begin{split} \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathsf{KI}}(\mathcal{A}) &\leq \frac{(n_u + n_u n_{ss} + n_s)^2}{p} + n_u \cdot \epsilon_{\mathsf{EUF-opCMA-DDH}} \\ &+ \left(2n_u^2 + n_u n_{ss} + n_u^2 n_{ss} + n_u n_s\right) \cdot \delta_{sim} \\ &+ \left(n_u^2 + n_u n_{ss} + n_u^2 n_{ss} \\ &+ n_u n_s + n_s n_u n_{ss} + n_{ss}^2\right) \cdot \left(\epsilon_{\mathsf{CRGapDH}} + \frac{Q_{\mathrm{RO}}}{p}\right), \end{split}$$

This concludes our proof of Theorem 1.

6 Deniability of **XHMQV**

Signal strives for deniability of its initial handshake protocol, as discussed in their descriptions of X3DH and PQXDH [41,37, both Section 4.4]. The deniability

property that Signal aims for considers transcripts of protocol executions. Given a transcript, no third party should be able to tell if the transcript is legitimate or if it was created using solely one party's (long-term and semi-static) secret keys; even if the third party has access to both parties' secret keys. Precisely this notion was formalized by Brendel, Fiedler, Günther, Janson, and Stebila [8] in a game-based setting.

We consider types of *offline deniability*, i.e., the distinguisher has to decide after the fact. In contrast, *online deniability* [20,51,52] allows the distinguisher to interact during the protocol execution. However, Unger and Goldberg's "Iron Triangle" [51, Section 6.6] presumes it is impossible to combine asynchronicity, forward secrecy, and online deniability in a key exchange protocol.

Deniable key exchange. Prior work has produced several deniability definitions for key exchange [18,32,15,16,54,30,53,31,28,8,24,13]. Most definitions follow the same blueprint⁸: An adversary queries a challenge oracle, which either acts as a real user (by executing the Run algorithm and using that user's secret key) or *pretends* to be that user (by executing the Fake algorithm or the simulator, depending on the formalism, but without that user's secret key) to answer the adversary. Afterwards, the distinguisher gets the transcript of the adversary's interaction with the challenge oracle and has to guess if the challenge oracle used Run or Fake. Intuitively, this guarantees that a transcript does not prove involvement of a party since the transcript could have been produced by the Fake algorithm.

Fiedler and Janson [24] provide a modular definition, which eases comparing deniability notions, with the following parameters:

- 1. adversarial power (is the adversary *malicious*, i.e., it may deviate from the protocol flow, or *semi-honest*?);
- 2. power of the Fake algorithm (can anybody fake a transcript (1-out-of- ∞ deniability) or only the peer (1-out-of-2 deniability)? When producing a message for Bob, does the Fake algorithm need access to another session of Bob?);
- power of the distinguisher (does the distinguisher get all honestly generated secret keys (*big brother model*)?);
- 4. auxiliary inputs (do we assume the adversary, the Fake algorithm, and potentially the distinguisher get common extra input, e.g., transcripts of honest protocol executions?).

We recount the full definition in Appendix B.

Deniability of Signal's initial handshake. Vatandas, Gennaro, Ithurburn, and Krawczyk [53] show deniability of HMQV and X3DH against malicious adversaries by relying on the novel Extended Knowledge of Diffie–Hellman assumption

⁸ We describe the blueprint for the game-based setting, since our proof uses the game-based setting as well. For definitions in the simulation-based setting the formalism differs but the idea remains the same.

(EKDHA). For X3DH, they additionally rely on the auxiliary input to obtain a signature on Bob's semi-static key. Unfortunately, Fiedler and Langrehr [25] disprove the EKDHA and deniability of X3DH against malicious adversaries by relying on simple auxiliary input. Fiedler and Janson [24, Table 2] compare the deniability guarantees for X3DH and PQXDH by considering Alice and Bob separately: To achieve deniability for Bob, they model two ways for the Fake algorithm to obtain a signature on Bob's semi-static key and for PQXDH additionally on an ephemeral key that the distinguisher never sees.⁹ Against malicious adversaries, they require the EKDHA for both protocols and for PQXDH additionally plaintext awareness of the KEM.¹⁰ The two following results hold even in the big brother model: Against semi-honest adversaries, Fiedler and Janson show that either party can produce the transcript (1-out-of-2 deniability) by relying on the commutativity of DH operations (i.e., $(q^x)^y = (q^y)^x$). Similarly, against malicious adversaries restricted to honestly generated long-term and semi-static keys, they show deniability for Alice without assumptions: There, the Fake algorithm uses Bob's semi-static secret key to compute the long-term-semi-static contribution and a freshly sampled ephemeral key for Alice to compute the other three contributions. With this approach the Fake algorithm cannot compute a long-term–ephemeral contribution (which is not present in X3DH and PQXDH). but, looking ahead, prevents us from proving deniability for XHMQV in this setting. Though, it is not clear to us which scenario we model by forcing the adversary to sample semi-static keys honestly but not the ephemeral keys.

Deniability of post-quantum replacements for Signal's initial handshake. Hashimoto, Katsumata, Kwiatkowski, and Prest [28] propose SC-DAKE' based on KEMs, ring signatures, and NIZKs, which they prove deniable under the model of [18]. Dobson and Galbraith [19] propose SI-X3DH based on supersingular isogenies, broken by the SIDH attack [10,40,48], and sketch a deniability argument based on X3DH. Brendel et al. [8] propose SPQR based on KEMs and designated verifier signatures (DVS) and show its deniability against semi-honest adversaries in the big brother model for Alice with their novel deniability definition matching the requirements of X3DH [41, Section 4.4]. Collins, Huguenin-Dumittan, Nguyen, Rolin, and Vaudenay [9] propose K-Waay, which combines KEMs with a splitKEM and show its deniability under a novel definition close to that of [8].

Deniability of XHMQV. In the following, we show that our protocol XHMQV achieves the same deniability guarantees against semi-honest adversaries in the big brother model as X3DH and PQXDH.¹¹ For malicious adversaries limited to honestly generated (long-term and semi-static) keys our XHMQV protocol does not provide the deniability guarantees of X3DH. We cannot avoid this since the long-term–ephemeral contribution proves Alice's involvement if Bob provably

⁹ The Fake learns the signature via the auxiliary input or with the oracle to another session of Bob.

¹⁰ Note that Signal's implementation uses Kyber as KEM, which does not achieve plaintext awareness.

 $^{^{11}}$ SPQR is proven to achieve the same guarantee only for Alice, but not for Bob.

does not know his ephemeral secret key. For malicious adversaries (which are not artificially limited to honestly generated long-term and semi-static keys) the currently only known proof strategy for X3DH involves the broken EKDHA assumption. Hence, XHMQV is the same as X3DH and PQXDH for not having a (viable) proof of deniability against malicious adversaries.

We use the model of [24], formally given in Definition 5, which considers Bob as initiator and Alice as responder. We split the theorems by role, beginning with the case of Bob, for whom the Fake algorithm needs to obtain signatures on the semi-static keys. Following [24], we do this via the auxiliary input and note that it can similarly be done with the USER oracle that allows Fake to access another session with Bob. When answering for Alice, the Fake algorithm does not need auxiliary input or the USER oracle, since Alice does not sign anything. The proof strategy follows [24, Theorem 4.1]: The Fake algorithm uses its SK oracle to learn the peer's long-term and semi-static secret keys and the extracts the ephemeral secret key from the adversary's randomness. Using these secret keys it computes the session key. We give the proofs in Appendix B.

Theorem 2. The XHMQV protocol as shown in Figure 9 is $(n_p, n_{ss}, q_{O_A}, q_{O_F}, q_{O_D}, t_A, t_{Dist}, \epsilon_{Dist})$ -deniable with respect to ({REGHON, CHALLHONINIT}, {SK}, {SKS, AUX})-oracles and auxiliary info sampled with AuxPrep yielding a valid pre-key bundle per ssid and user, and aux known to the distinguisher, where n_p is the number of parties and n_{ss} the number of semi-static keys per party, and $q_{O_A}, q_{O_F}, q_{O_D}, t_A, t_{Dist}$ are arbitrary and $\epsilon_{Dist} = 0$.

Theorem 3. The XHMQV protocol as shown in Figure 9 is $(n_p, n_{ss}, q_{O_A}, q_{O_F}, q_{O_D}, t_A, t_{Dist}, \epsilon_{Dist})$ -deniable with respect to ({REGHON, CHALLHONRESP}, {SK}, {SKS, AUX})-oracles and auxiliary info sampled with arbitrary AuxPrep, and aux known to the distinguisher, where n_p is the number of parties and n_{ss} the number of semi-static keys per party, and $q_{O_A}, q_{O_F}, q_{O_D}, t_A, t_{Dist}$ are arbitrary and $\epsilon_{Dist} = 0$.

Acknowledgments

We thank the anonymous reviewers for their insightful comments. R.F. was supported by the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE.

29

References

- Abdalla, M., Bellare, M., Rogaway, P.: The oracle Diffie-Hellman assumptions and an analysis of DHIES. In: Naccache, D. (ed.) CT-RSA 2001. LNCS, vol. 2020, pp. 143–158. Springer, Berlin, Heidelberg (Apr 2001). https://doi.org/10.1007/ 3-540-45353-9_12
- An, J.H., Dodis, Y., Rabin, T.: On the security of joint signature and encryption. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 83–107. Springer, Berlin, Heidelberg (Apr / May 2002). https://doi.org/10.1007/ 3-540-46035-7_6
- Bellare, M., Davis, H., Günther, F.: Separate your domains: NIST PQC KEMs, oracle cloning and read-only indifferentiability. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part II. LNCS, vol. 12106, pp. 3–32. Springer, Cham (May 2020). https://doi.org/10.1007/978-3-030-45724-2_1
- Bellare, M., Rogaway, P.: Entity authentication and key distribution. In: Stinson, D.R. (ed.) CRYPTO'93. LNCS, vol. 773, pp. 232–249. Springer, Berlin, Heidelberg (Aug 1994). https://doi.org/10.1007/3-540-48329-2_21
- Bernstein, D.J.: Curve25519: New Diffie-Hellman speed records. In: Yung, M., Dodis, Y., Kiayias, A., Malkin, T. (eds.) PKC 2006. LNCS, vol. 3958, pp. 207–228. Springer, Berlin, Heidelberg (Apr 2006). https://doi.org/10.1007/11745853_14
- 6. Bhargavan, K., Jacomme, C., Kiefer, F., Schmidt, R.: Formal verification of the PQXDH post-quantum key agreement protocol for end-to-end secure messaging. In: Balzarotti, D., Xu, W. (eds.) USENIX Security 2024. USENIX Association (Aug 2024), https://www.usenix.org/conference/usenixsecurity24/presentation/ bhargavan
- Brendel, J., Fiedler, R., Günther, F., Janson, C., Stebila, D.: Post-quantum asynchronous deniable key exchange and the Signal handshake. Cryptology ePrint Archive, Report 2021/769 (2021), https://eprint.iacr.org/2021/769
- Brendel, J., Fiedler, R., Günther, F., Janson, C., Stebila, D.: Post-quantum asynchronous deniable key exchange and the Signal handshake. In: Hanaoka, G., Shikata, J., Watanabe, Y. (eds.) PKC 2022, Part II. LNCS, vol. 13178, pp. 3–34. Springer, Cham (Mar 2022). https://doi.org/10.1007/978-3-030-97131-1_1
- Carlini, N., Hayes, J., Nasr, M., Jagielski, M., Sehwag, V., Tramèr, F., Balle, B., Ippolito, D., Wallace, E.: Extracting training data from diffusion models. In: Calandrino, J.A., Troncoso, C. (eds.) USENIX Security 2023. pp. 5253– 5270. USENIX Association (Aug 2023), https://www.usenix.org/conference/ usenixsecurity23/presentation/carlini
- Castryck, W., Decru, T.: An efficient key recovery attack on SIDH. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023, Part V. LNCS, vol. 14008, pp. 423–447. Springer, Cham (Apr 2023). https://doi.org/10.1007/978-3-031-30589-4_15
- Cohn-Gordon, K., Cremers, C., Dowling, B., Garratt, L., Stebila, D.: A formal security analysis of the Signal messaging protocol. In: 2017 IEEE European Symposium on Security and Privacy. pp. 451–466. IEEE Computer Society Press (Apr 2017). https://doi.org/10.1109/EuroSP.2017.27
- Cohn-Gordon, K., Cremers, C., Dowling, B., Garratt, L., Stebila, D.: A formal security analysis of the Signal messaging protocol. Journal of Cryptology 33(4), 1914–1983 (Oct 2020). https://doi.org/10.1007/s00145-020-09360-1
- Collins, D., Huguenin-Dumittan, L., Nguyen, N.K., Rolin, N., Vaudenay, S.: K-waay: Fast and deniable post-quantum X3DH without ring signatures. In: Balzarotti, D., Xu, W. (eds.) USENIX Security 2024. USENIX Association (Aug

2024), https://www.usenix.org/conference/usenixsecurity24/presentation/ collins

- Cremers, C., Dax, A., Medinger, N.: Keeping up with the KEMs: Stronger security notions for KEMs and automated analysis of KEM-based protocols. In: Luo, B., Liao, X., Xu, J., Kirda, E., Lie, D. (eds.) ACM CCS 2024. pp. 1046–1060. ACM Press (Oct 2024). https://doi.org/10.1145/3658644.3670283
- Cremers, C., Feltz, M.: One-round strongly secure key exchange with perfect forward secrecy and deniability. Cryptology ePrint Archive, Report 2011/300 (2011), https://eprint.iacr.org/2011/300
- Dagdelen, Ö., Fischlin, M., Gagliardoni, T., Marson, G.A., Mittelbach, A., Onete, C.: A cryptographic analysis of OPACITY - (extended abstract). In: Crampton, J., Jajodia, S., Mayes, K. (eds.) ESORICS 2013. LNCS, vol. 8134, pp. 345–362. Springer, Berlin, Heidelberg (Sep 2013). https://doi.org/10.1007/ 978-3-642-40203-6_20
- Degabriele, J.P., Lehmann, A., Paterson, K.G., Smart, N.P., Strefler, M.: On the joint security of encryption and signature in EMV. In: Dunkelman, O. (ed.) CT-RSA 2012. LNCS, vol. 7178, pp. 116–135. Springer, Berlin, Heidelberg (Feb / Mar 2012). https://doi.org/10.1007/978-3-642-27954-6_8
- Di Raimondo, M., Gennaro, R., Krawczyk, H.: Deniable authentication and key exchange. In: Juels, A., Wright, R.N., De Capitani di Vimercati, S. (eds.) ACM CCS 2006. pp. 400–409. ACM Press (Oct / Nov 2006). https://doi.org/10.1145/ 1180405.1180454
- Dobson, S., Galbraith, S.D.: Post-quantum Signal key agreement from SIDH. In: Cheon, J.H., Johansson, T. (eds.) Post-Quantum Cryptography - 13th International Workshop, PQCrypto 2022. pp. 422–450. Springer, Cham (Sep 2022). https://doi.org/10.1007/978-3-031-17234-2_20
- Dodis, Y., Katz, J., Smith, A., Walfish, S.: Composability and on-line deniability of authentication. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 146–162. Springer, Berlin, Heidelberg (Mar 2009). https://doi.org/10.1007/ 978-3-642-00457-5_10
- Dowling, B., Fischlin, M., Günther, F., Stebila, D.: A cryptographic analysis of the TLS 1.3 handshake protocol candidates. In: Ray, I., Li, N., Kruegel, C. (eds.) ACM CCS 2015. pp. 1197–1210. ACM Press (Oct 2015). https://doi.org/10. 1145/2810103.2813653
- Fersch, M., Kiltz, E., Poettering, B.: On the one-per-message unforgeability of (EC)DSA and its variants. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017, Part II. LNCS, vol. 10678, pp. 519–534. Springer, Cham (Nov 2017). https://doi.org/ 10.1007/978-3-319-70503-3_17
- Fiedler, R., Günther, F.: Security analysis of Signal's PQXDH handshake. In: Jager, T., Pan, J. (eds.) PKC 2025, Part II. LNCS, vol. 15675, pp. 137–169. Springer, Cham (May 2025). https://doi.org/10.1007/978-3-031-91823-0_5
- 24. Fiedler, R., Janson, C.: A deniability analysis of Signal's initial handshake PQXDH. PoPETs 2024(4), 907-928 (Oct 2024). https://doi.org/10.56553/ popets-2024-0148
- Fiedler, R., Langrehr, R.: On deniable authentication against malicious verifiers. Cryptology ePrint Archive, Paper 2025/470 (2025), https://eprint.iacr.org/ 2025/470
- 26. Haber, S., Pinkas, B.: Securely combining public-key cryptosystems. In: Reiter, M.K., Samarati, P. (eds.) ACM CCS 2001. pp. 215–224. ACM Press (Nov 2001). https://doi.org/10.1145/501983.502013

- Hashimoto, K., Katsumata, S., Kwiatkowski, K., Prest, T.: An efficient and generic construction for Signal's handshake (X3DH): Post-quantum, state leakage secure, and deniable. In: Garay, J. (ed.) PKC 2021, Part II. LNCS, vol. 12711, pp. 410–440. Springer, Cham (May 2021). https://doi.org/10.1007/978-3-030-75248-4_15
- Hashimoto, K., Katsumata, S., Kwiatkowski, K., Prest, T.: An efficient and generic construction for Signal's handshake (X3DH): Post-quantum, state leakage secure, and deniable. Journal of Cryptology 35(3), 17 (Jul 2022). https://doi.org/10. 1007/s00145-022-09427-1
- 29. Hashimoto, K., Katsumata, S., Wiggers, T.: Bundled authenticated key exchange: A concrete treatment of (post-quantum) Signal's handshake protocol. In: 34th USENIX Security Symposium, USENIX Security 2025. USENIX Association (Aug 2025), to appear. Available as Cryptology ePrint Archive Report 2025/040, https: //eprint.iacr.org/2025/040.
- 30. Jiang, S.: Timed encryption with application to deniable key exchange. Theor. Comput. Sci. 560, 172–189 (2014). https://doi.org/10.1016/J.TCS.2014.02. 005, https://doi.org/10.1016/j.tcs.2014.02.005
- Jiang, S., Chee, Y.M., Ling, S., Wang, H., Xing, C.: A new framework for deniable secure key exchange. Inf. Comput. 285(Part), 104866 (2022). https://doi.org/ 10.1016/J.IC.2022.104866, https://doi.org/10.1016/j.ic.2022.104866
- 32. Jiang, S., Safavi-Naini, R.: An efficient deniable key exchange protocol (extended abstract). In: Tsudik, G. (ed.) FC 2008. LNCS, vol. 5143, pp. 47–52. Springer, Berlin, Heidelberg (Jan 2008). https://doi.org/10.1007/978-3-540-85230-8_4
- Kiltz, E., Masny, D., Pan, J.: Optimal security proofs for signatures from identification schemes. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part II. LNCS, vol. 9815, pp. 33–61. Springer, Berlin, Heidelberg (Aug 2016). https: //doi.org/10.1007/978-3-662-53008-5_2
- Kiltz, E., Pan, J., Riepel, D., Ringerud, M.: Multi-user CDH problems and the concrete security of NAXOS and HMQV. In: Rosulek, M. (ed.) CT-RSA 2023. LNCS, vol. 13871, pp. 645–671. Springer, Cham (Apr 2023). https://doi.org/ 10.1007/978-3-031-30872-7_25
- Kobeissi, N., Bhargavan, K., Blanchet, B.: Automated verification for secure messaging protocols and their implementations: A symbolic and computational approach. In: 2017 IEEE European Symposium on Security and Privacy. pp. 435–450. IEEE Computer Society Press (Apr 2017). https://doi.org/10.1109/EuroSP. 2017.38
- Krawczyk, H.: HMQV: A high-performance secure Diffie-Hellman protocol. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 546–566. Springer, Berlin, Heidelberg (Aug 2005). https://doi.org/10.1007/11535218_33
- 37. Kret, E., Schmidt, R.: The PQXDH key agreement protocol (Jan 2024), https: //signal.org/docs/specifications/pqxdh/
- 38. Langley, A., Hamburg, M., Turner, S.: Elliptic Curves for Security. RFC 7748 (Informational) (Jan 2016). https://doi.org/10.17487/RFC7748, https://www. rfc-editor.org/rfc/rfc7748.txt
- Law, L., Menezes, A., Qu, M., Solinas, J.A., Vanstone, S.A.: An efficient protocol for authenticated key agreement. DCC 28(2), 119–134 (2003). https://doi.org/ 10.1023/A:1022595222606
- Maino, L., Martindale, C., Panny, L., Pope, G., Wesolowski, B.: A direct key recovery attack on SIDH. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023, Part V. LNCS, vol. 14008, pp. 448–471. Springer, Cham (Apr 2023). https:// doi.org/10.1007/978-3-031-30589-4_16

- 32 Rune Fiedler, Felix Günther, Jiaxin Pan, and Runzhi Zeng
- 41. Marlinspike, M., Perrin, T.: The X3DH key agreement protocol (Nov 2016), https: //signal.org/docs/specifications/x3dh/
- Menezes, A.: Another look at HMQV. Cryptology ePrint Archive, Report 2005/205 (2005), https://eprint.iacr.org/2005/205
- Menezes, A., Qu, M., , Vanstone, S.A.: Some new key agreement protocols providing mutual implicit authentication. 2nd Workshop on Selected Areas in Cryptography (SAC 1995) (1995)
- Paterson, K.G., Schuldt, J.C.N., Stam, M., Thomson, S.: On the joint security of encryption and signature, revisited. In: Lee, D.H., Wang, X. (eds.) ASI-ACRYPT 2011. LNCS, vol. 7073, pp. 161–178. Springer, Berlin, Heidelberg (Dec 2011). https://doi.org/10.1007/978-3-642-25385-0_9
- Patton, C., Shrimpton, T.: Security in the presence of key reuse: Contextseparable interfaces and their applications. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part I. LNCS, vol. 11692, pp. 738–768. Springer, Cham (Aug 2019). https://doi.org/10.1007/978-3-030-26948-7_26
- 46. Perrin, T.: The XEdDSA and VXEdDSA signature schemes (Oct 2016), https: //signal.org/docs/specifications/xeddsa/
- Pointcheval, D., Stern, J.: Security arguments for digital signatures and blind signatures. Journal of Cryptology 13(3), 361–396 (Jun 2000). https://doi.org/10.1007/s001450010003
- Robert, D.: Breaking SIDH in polynomial time. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023, Part V. LNCS, vol. 14008, pp. 472–503. Springer, Cham (Apr 2023). https://doi.org/10.1007/978-3-031-30589-4_17
- 49. Signal: Technical information, https://signal.org/docs/
- 50. Thormarker, E.: On using the same key pair for Ed25519 and an X25519 based KEM. Cryptology ePrint Archive, Report 2021/509 (2021), https://eprint. iacr.org/2021/509
- 51. Unger, N., Goldberg, I.: Deniable key exchanges for secure messaging. In: Ray, I., Li, N., Kruegel, C. (eds.) ACM CCS 2015. pp. 1211–1223. ACM Press (Oct 2015). https://doi.org/10.1145/2810103.2813616
- Unger, N., Goldberg, I.: Improved strongly deniable authenticated key exchanges for secure messaging. PoPETs 2018(1), 21-66 (Jan 2018). https://doi.org/10. 1515/popets-2018-0003
- Vatandas, N., Gennaro, R., Ithurburn, B., Krawczyk, H.: On the cryptographic deniability of the Signal protocol. In: Conti, M., Zhou, J., Casalicchio, E., Spognardi, A. (eds.) ACNS 2020, Part II. LNCS, vol. 12147, pp. 188–209. Springer, Cham (Oct 2020). https://doi.org/10.1007/978-3-030-57878-7_10
- 54. Yao, A.C.C., Zhao, Y.: OAKE: a new family of implicitly authenticated Diffie-Hellman protocols. In: Sadeghi, A.R., Gligor, V.D., Yung, M. (eds.) ACM CCS 2013. pp. 1113–1128. ACM Press (Nov 2013). https://doi.org/10.1145/ 2508859.2516695

A Full Proof of Bounding G₂ in Theorem 1

In this section, we formally bound the probability from Theorem 1:

$$\mathsf{Adv}^{\mathbf{G}_2}_{\mathsf{XHMQV}}(\mathcal{A}) \leq \sum_{\substack{\mathsf{clean}_{X} \in \{\mathsf{clean}_{\mathsf{LTLT}}, \mathsf{clean}_{\mathsf{LTE}}, \mathsf{clean}_{\mathsf{LTSS}}, \\ \mathsf{clean}_{\mathsf{ELT}}, \mathsf{clean}_{\mathsf{ESS}}, \mathsf{clean}_{\mathsf{EE}}\}} \mathsf{Adv}^{\mathbf{G}_2 \wedge \mathsf{clean}_X}_{\mathsf{XHMQV}}(\mathcal{A}).$$

We prove each case in different subsections below; concrete security bounds for each case can be found in these subsections. Here we only give the final security bound of $\operatorname{Adv}_{\mathsf{XHMQV}}^{\mathbf{G}_2}(\mathcal{A})$, i.e., the bound in Lemma 2 where $\epsilon_{\mathsf{CRGapDH}}$ denotes the advantage bound for the CRGapDH assumption in \mathbb{G} and δ_{sim} the bound for simulatability of Sig:

$$\begin{split} \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathsf{G}_2}(\mathcal{A}) &\leq n_u n_u & \cdot (2\delta_{sim} + \epsilon_{\mathsf{CRGapDH}} + Q_{\mathsf{RO}}/p) & //\operatorname{clean_{\mathsf{LTLT}}} \\ &+ n_u n_s & \cdot (\delta_{sim} + \epsilon_{\mathsf{CRGapDH}} + Q_{\mathsf{RO}}/p) & //\operatorname{clean_{\mathsf{LTE}}} \\ &+ n_u^2 n_{ss} & \cdot (\delta_{sim} + \epsilon_{\mathsf{CRGapDH}} + Q_{\mathsf{RO}}/p) & //\operatorname{clean_{\mathsf{LTSS}}} \\ &+ n_u n_s & \cdot (\delta_{sim} + \epsilon_{\mathsf{CRGapDH}} + Q_{\mathsf{RO}}/p) & //\operatorname{clean_{\mathsf{LTSS}}} \\ &+ n_s n_u n_{ss} & \cdot (\epsilon_{\mathsf{CRGapDH}} + Q_{\mathsf{RO}}/p) & //\operatorname{clean_{\mathsf{ELT}}} \\ &+ n_s^2 & \cdot (\epsilon_{\mathsf{CRGapDH}} + Q_{\mathsf{RO}}/p) & //\operatorname{clean_{\mathsf{ELSS}}} \\ &+ n_s^2 & \cdot (\epsilon_{\mathsf{CRGapDH}} + Q_{\mathsf{RO}}/p) & . & //\operatorname{clean_{\mathsf{ELSS}}} \end{split}$$

A.1 Case clean_{LTLT}

By the definition of $\mathsf{clean}_{\mathsf{LTLT}}$, the long-term keys of both $\pi^*.\mathsf{oid}$ and $\pi^*.\mathsf{pid}$ are uncompromised. We begin with $\mathcal{G}_0^{\mathsf{ltlt}} := \mathbf{G}_2 \wedge \mathsf{clean}_{\mathsf{LTLT}}(\pi^*)$, so we have

$$\mathsf{Adv}^{\mathbf{G}_2 \wedge \mathsf{clean}_{\mathsf{LTLT}}(\pi^*)}_{\mathsf{XHMQV}}(\mathcal{A}) = \mathsf{Adv}^{\mathcal{G}^{\mathsf{ltt}}_0}_{\mathsf{XHMQV}}(\mathcal{A})$$

We use game sequence $\mathcal{G}_1^{\mathsf{ltlt}}$ - $\mathcal{G}_4^{\mathsf{ltlt}}$ to bound the winning advantage of \mathcal{A} when π^* satisfies $\mathsf{clean}_{\mathsf{LTLT}}$. The transitions $\mathcal{G}_1^{\mathsf{ltlt}}$ - $\mathcal{G}_3^{\mathsf{ltl}}$ enable the game to simulate signatures without using long-term keys. We use the simulatability (cf. Definition 3) of the signature scheme to bound the probability difference. For $\mathcal{G}_3^{\mathsf{ltl}}$ - $\mathcal{G}_4^{\mathsf{ltl}}$, we use the $\mathsf{CRGapDH}$ assumption to prove that if the adversary cannot query KDF_f (or KDF_r) on the correct hash input of π^* , then it cannot distinguish the real session key of π^* from a random key.

<u>GAME</u> $\mathcal{G}_1^{\text{ltlt}}$: We guess the identities \mathbb{A}^* and \mathbb{B}^* of the initiator and responder involved in the test session (i.e., $(\pi^*.\text{oid}, \pi^*.\text{pid}) = (\mathbb{A}^*, \mathbb{B}^*)$). If this guess is incorrect, then we overwrite \mathcal{A} 's bit guess with 0. Since there are at most n_u users involved in the game and the guess is uniformly at random in \mathcal{A} 's view, we have

$$\mathsf{Adv}^{\mathcal{G}_0^{\mathrm{int}}}_{\mathsf{XHMQV}}(\mathcal{A}) = n_u^2 \cdot \mathsf{Adv}^{\mathcal{G}_1^{\mathrm{int}}}_{\mathsf{XHMQV}}(\mathcal{A})$$

For the sake of simplicity, let (a^*, A^*) and (b^*, B^*) be the long-term key pairs of A^* and B^* , respectively. In the game transitions $\mathcal{G}_1^{\text{ltl}}$ - $\mathcal{G}_3^{\text{ltl}}$, we change the simulation of h_{sig} (the random oracle used by the signature scheme) and the

33

34 Rune Fiedler, Felix Günther, Jiaxin Pan, and Runzhi Zeng

$KGenSS(\mathit{ltsk}_{B^*} = b^*) \text{ for } \mathcal{G}_1^{ltlt}$	$KGenSS(\mathit{ltsk}_{B^*} = b^*) \text{ for } \mathcal{G}_2^{ltlt}:$
1 $s \leftarrow \mathbb{Z}_p, S := g^s$	1 $s \leftarrow \mathbb{Z}_p, S := g^s$
2 $\sigma \leftarrow Sign^{h_{sig}}(b^*, S)$	$(\sigma, qry, h) \leftarrow SIMsign(B^*, S)$
3 $sssk_U := s, sspk_U := (S, \sigma)$	3 $\mathcal{L}_{\texttt{sim}} := \mathcal{L}_{\texttt{sim}} \cup \{(qry,h)\}$
4 return $(sssk_U, sspk_U)$	$4 \overline{\sigma \leftarrow Sign^{h_{sig}}(b^*, S)}$
$\underline{h_{sig}(qry)}$:	5 $sssk_U := s, sspk_U := (S, \sigma)$ 6 $\mathbf{return} (sssk_U, sspk_U)$
// Lazy sampling	
64 if $\exists h \text{ s.t. } (qry, h) \in \mathcal{L}_{h_{sig}}$	$h_{sig}(qry)$:
65 return h	69 if $\exists h$ s.t. $(ary, h) \in \mathcal{L}_{sim}$
66 $h \leftarrow h_{sig}$.image	70 return h //Program the BO st h : (gru) = h
67 $\mathcal{L}_{h_{sig}} := \mathcal{L}_{h_{sig}} \cup (qry, h)$	71 · · · // As the lary sampling in G^{lth}
68 return h	// As the lazy sampling in \mathcal{Y}_1

Fig. 10. Code of KGenSS and h_{sig} for games \mathcal{G}_1^{ltlt} (left) and \mathcal{G}_2^{ltlt} (right)

generation of A^* 's and B^* 's semi-static keys before running A.

<u>GAME</u> $\mathcal{G}_2^{\text{lth}}$: This game changes the random oracle h_{sig} (used by the signature scheme) and the key generation of semi-static keys for B^{*}. Since the code is presented in Fig. 10, here we only give the high-level idea:

- When running KGenSS($ltsk_{B^*}$), we use SIMsign to generate simulated signatures and some simulation information. The simulation information is recorded in \mathcal{L}_{sim} and will be used to program h_{sig} later.
- We use the outputs of $SIMsign(B^*, \cdot)$ to simulate h_{sig} 's responses to queries that may be related to the signatures of B^* 's semi-static keys.

In $\mathcal{G}_2^{\text{ltlt}}$, the differences compared to $\mathcal{G}_1^{\text{ltlt}}$ are (1) we use the signature from SIMsign, and (2) we use the simulated hash value from SIMsign to simulate the random oracle h_{sig} . We construct a straight-forward reduction from δ -simulatability of Sig to bound these differences: Let \mathcal{R} be an adversary against the simulatability of Sig. \mathcal{R} plays against either $\mathcal{G}_{\text{Sig}}^{Real}$ or $\mathcal{G}_{\text{Sig}}^{Sim}$ and simulates either $\mathcal{G}_1^{\text{ltl}}$ or $\mathcal{G}_2^{\text{ltl}}$ for \mathcal{A} , respectively. \mathcal{R} works as follows:

- \mathcal{R} first samples b_{test} and a party id B^* . It generates B^* 's semi-static keys (s_i, S_i) for $1 \leq i \leq n_{ss}$, and outputs $(S_i)_{1 \leq i \leq n_{ss}}$ to the game. It then gets $(\forall \mathsf{k} = B^*, \mathsf{sk} = b^*, (\sigma_i)_{1 \leq i \leq [n_{ss}]})$ and random oracle access to RO.
- Then, \mathcal{R} sets up all variables (long-term key pairs and semi-static key pairs of all other parties, session initialization, etc.) and oracles required in $\mathcal{G}_2^{\mathsf{ltl}}$ and runs \mathcal{A} . All h_{sig} queries from \mathcal{A} are forwarded to RO.
- When \mathcal{A} outputs b_{guess} , \mathcal{R} returns $\llbracket b_{test} = b_{guess} \rrbracket$.

It is straight-forward to see that if \mathcal{R} is playing against \mathcal{G}_{Sig}^{Real} , then it perfectly simulates \mathcal{G}_{1}^{ltlt} for \mathcal{A} . If \mathcal{R} is playing against \mathcal{G}_{Sig}^{Real} , then we have $h_{sig}(qry') = h' =$ $\mathsf{RO}(qry)$ if (σ', qry', h') is generated by SIMsign. Hence, in this case \mathcal{R} perfectly simulates \mathcal{G}_{2}^{ltlt} . Therefore, we have

$$\left|\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_1^{\mathsf{litt}}}(\mathcal{A}) - \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_2^{\mathsf{litt}}}(\mathcal{A})\right| = \left|\Pr\left[\mathcal{G}_{\mathsf{Sig}}^{Real} = 1\right] - \Pr\left[\mathcal{G}_{\mathsf{Sig}}^{Sim} = 1\right]\right| \le \delta$$

35

and

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_1^{\operatorname{litt}}}(\mathcal{A}) \leq \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_2^{\operatorname{litt}}}(\mathcal{A}) + \delta.$$

From \mathcal{G}_2^{lth} on, we no longer need the long-term secret key of B^* to generate signatures on B^* 's semi-static keys.

<u>GAME</u> $\mathcal{G}_3^{\text{ltlt}}$: We apply the same approach used in $\mathcal{G}_2^{\text{ltlt}}$ to \mathbb{A}^* . Namely, we use a modified KGenSS and change the simulation of h_{sig} so that we can simulate signatures on \mathbb{A}^* 's semi-static keys without using its long-term key. Similar to $\mathcal{G}_2^{\text{ltlt}}$, we have

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_2^{\mathsf{litt}}}(\mathcal{A}) \leq \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_3^{\mathsf{litt}}}(\mathcal{A}) + \delta,$$

and from $\mathcal{G}_3^{\text{lth}}$ on, we no longer need the long-term secret key of \mathbb{A}^* .

GAME $\mathcal{G}_{4}^{\text{htt}}$: We introduce a new abort event QueryDH_{ttt} in $\mathcal{G}_{4}^{\text{htt}}$: Specifically, if

- \mathcal{A} queries KDF_f on (DH, A^*, B^*, S, Y, X) , where $d = h_0(A^*, X), e_1 = h_1(B^*, Y, S), e_2 = h_2(B^*, Y, S)$, and $\mathsf{DDH}(X(A^*)^d, Y(B^*)^{e_1}S^{e_2}, DH) = 1$, and $(A^*, B^*, S, Y, X) = \pi^*.ctxt$, or
- \mathcal{A} queries KDF_r on (DH, A^*, B^*, S, X) , where $d = h_0(A^*, X)$, $e_1 = h_3(B^*, S)$, and $\mathsf{DDH}(X(A^*)^d, S(B^*)^{e_1}, DH) = 1$ and $(A^*, B^*, S, X) = \pi^*.ctxt$,

then the game aborts (i.e., return 0 as \mathcal{A} 's bit guess). In other words, we abort the game if \mathcal{A} queries KDF_f or KDF_r on any DH value that contains the information of $(B^*)^{a^*} (= (A^*)^{b^*})$.

To bound the probability of this event, we construct a CRGapDH adversary $\mathcal{B}_{\mathsf{ltlt}}$ such that if such an abort event happens, then \mathcal{B} can output a solution to its CRGapDH challenge. By Definition 1, the input to $\mathcal{B}_{\mathsf{ltlt}}$ is (A^*, B^*) and $\mathcal{B}_{\mathsf{ltlt}}$ has access to oracles DDH and CH. We <u>underline</u> how we embed the challenge into the game.

Reduction $\mathcal{B}_{\mathsf{ltlt}}^{\mathsf{DDH},\mathsf{CH}}(A^*,B^*)$:

- 1. Choose two users A^* and B^* uniformly at random and set the long-term public key of A^* and B^* as A^* and B^* , respectively.
- 2. Initialize lists $\mathcal{L}_{patch-f}$, \mathcal{L}_{skey-f} , $\mathcal{L}_{patch-r}$, and \mathcal{L}_{skey-r} for the simulation of session keys using DDH and random oracle patching (as in \mathbf{G}_2). We also initialize \mathcal{L}_{sim} to simulate the signatures on \mathbf{A}^* 's and \mathbf{B}^* 's semi-static keys (as in Fig. 10).
- 3. For parties A* and B*, run the modified KGenSS algorithm defined in Fig. 10. After this step, B_{ltlt} has *ltpk*, *sspk*, and *ltsk* (except for the longterm secret keys of A* and B*). For all other parties, follow the KGenLT and KGenSS defined in XHMQV faithfully.
- 4. Run $b_{guess} \leftarrow A(ltpk, sspk)$ and proceed the game as in \mathcal{G}_{KE}^{KI} . If the initiator and responder of the test session are not A^* and B^* , respectively, then overwrite \mathcal{A} 's bit with 0.

The random oracles and game oracles are simulated as described below.

⁻ Initialize $\mathcal{G}_4^{\mathsf{ltlt}}$:

- Simulation of oracles defined in $\mathcal{G}_{\mathsf{KF}}^{\mathsf{KI}}$:
 - All oracles except SEND and CORRUPTLTKEY are simulated in the same way defined in \mathcal{G}_{KE}^{KI} .
 - By clean_{LTLT}, the adversary does not query CORRUPTLTKEY on A* or B*. CORRUPTLTKEY queries on other parties are simulated as usual (since we have long-term keys of all other parties).
 - For SEND, if the query does not involve A^* and B^* , then we also simulate the oracle in the same way as in \mathcal{G}_{KF}^{K1} .
 - SEND(A^{*}, *i*, *m*) (We only present how to simulate the Run algorithms and set up the session keys):
 - * Case m = (create, ssid, type) (as the session responder): Simulate as usual.
 - * Case $m = (\mathbf{U}, X)$ (as the session responder and type = full):
 - 1. Get the semi-static secret s, the ephemeral secret y (and their public keys) of this session. Let U be the long-term public key of the initiator party U.
 - 2. If there exists an entry (U, A^*, S, Y, X, K) in $\mathcal{L}_{\mathsf{skey-f}}$, then retrieve this K.
 - 3. Otherwise, sample a random K and record (U, A^*, S, Y, X, K) in the list $\mathcal{L}_{patch-f}$.
 - 4. Set $\pi_{A^*}^i.K := K$.
 - * Case $m = (\mathbf{U}, X)$ (as the session responder and type = reduced). This case is similar to the case of type = full. We present the details for completeness:
 - 1. Get the semi-static key pair (s, S) of this session. Let U be the long-term public key of U.
 - 2. If there exists (U, A^*, S, X, K) in $\mathcal{L}_{\mathsf{skey-r}}$, then retrieve this K.
 - 3. Otherwise, sample a random K and record (U, A^*, S, X, K) in $\mathcal{L}_{patch-r}$.
 - 4. Set $\pi^i_{\mathbb{A}^*}.\mathsf{K} := K$.
 - * Case $m = (\mathbf{U}, S, Y, \sigma)$ (as the session initiator and type = full):
 - 1. Generate the ephemeral key pair (X, x) of this session. Let U be the long-term public key of the responder party U.
 - 2. If there exists (A^*, U, S, Y, X, K) in $\mathcal{L}_{\mathsf{skey-f}}$, then retrieve this K.
 - 3. Otherwise, sample a random K and record (A^*, U, S, Y, X, K) in $\mathcal{L}_{patch-f}$.
 - 4. Set $\pi^i_{\mathbf{A}^*}.\mathsf{K} := K$
 - * Case $m = (U, S, Y = \bot, \sigma)$ (as the session initiator) and type = reduced works analogously to the case of type = full.
 - SEND(B^*, j, m): The simulation is analogous to SEND(A^*, i, m).
- Simulation of Random Oracles:
 - For fresh RO queries to $\mathsf{KDF}_f(A, B, S, Y, X, DH)$ and $\mathsf{KDF}_r(A, B, S, X, DH)$, where $(A, B) \neq (A^*, B^*)$, we simulate in the same way as in \mathbf{G}_2 . Similarly, for fresh RO queries to h_0, h_1, h_2 , and h_3 that do not involve A^* and B^* , we simulate as usual.
 - $h_0(A^*, X): h \leftarrow CH(X), h_0(A^*, X) := h.$

- $h_1(B^*, Y, S): e_2 := h_2(B^*, Y, S), R := YS^{e_2}, h \leftarrow CH(R), h_1(B^*, Y, S) := h.$
- $h_2(B^*, Y, S)$: Simulate as usual.
- $h_3(B^*, S): h \leftarrow CH(S), h_0(B^*, S) := h.$
- $\mathsf{KDF}_f(DH, A^*, B^*, S, Y, X)$:
 - 1. If $(A^*, B^*, S, Y, X) = \pi^* . ctxt$ and DDH $(X(A^*)^d, Y(B^*)^{e_1}S^{e_2}, DH) = 1$, then extract the value

$$DH' := \begin{cases} \left(DH / (X(A^*)^d)^{y+se_2} \right)^{e_1^{-1}}, & \pi^* \text{ is a responder session} \\ \left(DH / (Y(B^*)^{e_1} S^{e_2})^x \right)^{d^{-1}}, & \pi^* \text{ is an initiator session} \end{cases}$$

where $d := h_0(A^*, X), e_1 := h_1(B, Y, S)$, and $e_2 := h_2(B, Y, S)$. If π^* is an initiator session, then we have its ephemeral key x. If π^* is a responder session, then we have the ephemeral key y and the semistatic key s of the responder B^* . Once such a DH' value is found, $\mathcal{B}_{\mathsf{ltlt}}$ aborts the simulation and returns DH' as the CRGapDH solution.

- 2. Otherwise, simulate by lazy sampling.
- $KDF_r(DH, A^*, B^*, S, X)$:
 - 1. If $(A^*, B^*, S, X) = \pi^* . ctxt$ and $DDH(X(A^*)^d, S(B^*)^{e_1}) = 1$, then extract the value

$$DH' := \begin{cases} \left(DH/(X(A^*)^d)^s \right)^{e_1^{-1}}, & \pi^* \text{ is a responder session} \\ \left(DH/\left((S(B^*)^{e_1})^x \right) \right)^{d^{-1}}, & \pi^* \text{ is an initiator session} \end{cases}$$

where $d := h_0(A^*, X)$ and $e_1 := h_3(B, S)$. If π^* is an initiator session, then we have its ephemeral key x. If π^* is a responder session, then we have the semi-static key s of the responder B^* . Once such a DH' value is found, \mathcal{B}_{lth} aborts the simulation and returns DH' as the CRGapDH solution.

2. Otherwise, simulate by lazy sampling.

To bound $\Pr[\texttt{QueryDH}_{\mathsf{ltit}}]$, we first prove that $\mathcal{B}_{\mathsf{ltit}}$ can extract a correct $\mathsf{CRGapDH}$ solution from KDF_f and KDF_r (specifically, when the DDH oracle outputs 1 and the query matches $\pi^*.ctxt$). In KDF_f , if π^* is an initiator session, then the extracted DH value DH' equals $(Y(B^*)^{e_1}S^{e_2})^{a^*}$. By the simulation of $h_1(B^*, Y, S)$, we have $e_2 = h_1(B^*, Y, S) = h$ where $h \leftarrow \mathrm{CH}(YS^{e_2})$. Therefore, we have

$$DH' = (Y(B^*)^{e_1} S^{e_2})^{a^*} = ((YS^{e_2}) \cdot (B^*)^h)^{a^*},$$

where $h = CH(YS^{e_2})$. So, DH' is a solution to CRGapDH. A similar argument applies to KDF_r and the case that π^* is a responder session. Therefore, as long as \mathcal{B}_{ltlt} aborts the game when simulating KDF_f or KDF_r , the extracted solution is a correct CRGapDH solution.

By definition, if QueryDH_{Itlt} happens, then \mathcal{A} must have queried $\mathsf{KDF}_f(DH, A^*, B^*, S, Y, X)$ or $\mathsf{KDF}_r(DH, A^*, B^*, S, X)$ which makes the DDH oracle output 1,

and \mathcal{B}_{ltlt} will extract a CRGapDH solution. Therefore, if QueryDH_{ltlt} happens, then \mathcal{B}_{ltlt} solves its CRGapDH challenge, and thus we have

$$\begin{split} \mathsf{Adv}^{\mathcal{G}^{\mathrm{lit}}_{\mathsf{X}\mathsf{H}\mathsf{M}\mathsf{Q}\mathsf{V}}}_{\mathsf{X}\mathsf{H}\mathsf{M}\mathsf{Q}\mathsf{V}}(\mathcal{A}) &\leq \mathsf{Adv}^{\mathcal{G}^{\mathrm{lit}}_{\mathsf{X}\mathsf{H}\mathsf{M}\mathsf{Q}\mathsf{V}}}_{\mathsf{X}\mathsf{H}\mathsf{M}\mathsf{Q}\mathsf{V}}(\mathcal{A}) + \mathrm{Pr}\left[\mathtt{Query}\mathtt{D}\mathtt{H}_{\mathsf{It}\mathsf{It}}\right] \\ &\leq \mathsf{Adv}^{\mathcal{G}^{\mathrm{lit}}_{\mathsf{X}\mathsf{H}\mathsf{M}\mathsf{Q}\mathsf{V}}}_{\mathsf{X}\mathsf{H}\mathsf{M}\mathsf{Q}\mathsf{V}}(\mathcal{A}) + \mathsf{Adv}^{\mathsf{CRGapDH}}_{(\mathbb{G},p,g)}(\mathcal{B}_{\mathsf{It}\mathsf{It}}) + Q_{\mathrm{RO}}/p, \end{split}$$

where the last term counts in the probability that there exists a non-invertible h value from the CH oracle.

Now we can bound the adversary's advantage in $\mathcal{G}_4^{\text{lth}}$. In this game, the session keys of all sessions are computed via lazy sampling. These keys are distributed uniformly at random in \mathcal{A} 's view until it queries KDF_f or KDF_r on the correct hash input. By the abort event $\mathsf{QueryDH}_{\text{lth}}$, the adversary cannot query the correct hash input of the test session, so the session key of π^* is distributed uniformly at random in \mathcal{A} 's view, independent of the game's challenge bit b_{test} . Therefore, the probability that $b_{\text{guess}} = b_{\text{test}}$ is $\frac{1}{2}$, and we have

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_4^{\mathsf{lift}}}(\mathcal{A}) = 0.$$

We assume that the ϵ -CRGapDH assumption holds over (\mathbb{G}, p, g). Combining all the probability bounds, we have the final bound for the case clean_{LTLT}(π^*):

$$\mathsf{Adv}^{\mathbf{G}_2 \wedge \mathsf{clean}_{\mathsf{LTLT}}}_{\mathsf{XHMQV}}(\mathcal{A}) \leq n_u^2 \cdot (2\delta + \epsilon + Q_{\mathrm{RO}}/p)$$

A.2 Case clean_{LTSS}

By the definition of $\mathsf{clean}_{\mathsf{LTSS}}$, we have (1) the initiator of π^* is uncompromised, (2) the semi-static key of π^* 's responder is uncompromised, and (3) if π^* is an initiator session, then either the responder's long-term key was uncompromised upon acceptance, or the semi-static key used in π^* is generated by the security game (rather than generated by the adversary).

By \mathbf{G}_2 , we always have $\pi^*.sspk \in sspk_{\pi^*,pid}$ as long as $\mathsf{clean}_{\mathsf{LTSS}}(\pi^*)$. This is because, on the one hand, if $\mathsf{clean}_{\mathsf{LTSS}}(\pi^*)$ and $\pi^*.pcorr = 0$, then we have $\pi^*.sspk \in sspk$ unless \mathcal{A} forged a valid signature on the semi-static key of an uncompromised party, which has been excluded in \mathbf{G}_2 . On the other hand, by definition, we have $\mathsf{clean}_{\mathsf{LTSS}}(\pi^*) \wedge \pi^*.pcorr \implies \pi^*.sspk \in sspk_{\pi^*.pid}$. Therefore, the third condition described above and the modification in \mathbf{G}_2 ensure that the semi-static key used in π^* has been generated by the game. This allows us to embed a CRGapDH challenge into the semi-static key.

We begin with $\mathcal{G}_0^{\mathsf{ltss}} := \mathbf{G}_2 \wedge \mathsf{clean}_{\mathsf{LTSS}}(\pi^*)$, so we have

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathbf{G}_2 \wedge \mathsf{clean}_{\mathsf{LTSS}}(\pi^*)}(\mathcal{A}) = \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_0^{\mathsf{ltss}}}(\mathcal{A})$$

We use the game sequence $\mathcal{G}_1^{\mathsf{ltss}}$ - $\mathcal{G}_4^{\mathsf{ltss}}$ to bound the winning advantage of \mathcal{A} when π^* satisfies $\mathsf{clean}_{\mathsf{LTSS}}$. Similar to the proof of the case $\mathsf{clean}_{\mathsf{LTLT}}(\pi^*)$, we use game transitions $\mathcal{G}_1^{\mathsf{ltss}}$ - $\mathcal{G}_3^{\mathsf{ltss}}$ and the simulatability (cf. Definition 3) of the

signature scheme to simulate signatures of the initiator of π^* without using its long-term key. Then, we use \mathcal{G}_3^{ltss} - \mathcal{G}_4^{ltss} and CRGapDH to prove that the adversary cannot query KDF_f (or KDF_r) on the correct hash input of π^* , and thus it cannot distinguish whether π^* 's key is real or random.

<u>GAME</u> $\mathcal{G}_1^{\text{ltss}}$: We guess the identity \mathbb{A}^* of the initiator involved in the test session. If this guess is incorrect, then we overwrite \mathcal{A} 's bit guess with 0. We denote \mathbb{A}^* 's long-term key pair as (a^*, A^*) . We have

$$\mathsf{Adv}^{\mathcal{G}_0^{\mathsf{ltss}}}_{\mathsf{XHMQV}}(\mathcal{A}) \leq n_u \cdot \mathsf{Adv}^{\mathcal{G}_1^{\mathsf{ltss}}}_{\mathsf{XHMQV}}(\mathcal{A})$$

<u>GAME</u> $\mathcal{G}_2^{\text{ltss}}$: We guess the identity \mathbb{B}^* of the responder involved in π^* and the identifier ssid^{*} of the repsonder \mathbb{B}^* 's uncorrupted semi-static key. The semi-static key pair of π^* is denoted as (s^*, S^*) (i.e., $\pi^*.\text{ssid} = \text{ssid}^*$ and $(s^*, S^*) = (sssk_{\mathbb{B}^*}[ssid^*], sspk_{\mathbb{B}^*}[ssid^*])$). Since there are at most n_u users involved in this game and at most n_{ss} semi-static keys for each user, we have

$$\mathsf{Adv}^{\mathcal{G}^{\mathsf{ltss}}_{\mathsf{XHMQV}}}_{\mathsf{XHMQV}}(\mathcal{A}) \leq n_u n_{ss} \cdot \mathsf{Adv}^{\mathcal{G}^{\mathsf{ltss}}_{2s}}_{\mathsf{XHMQV}}(\mathcal{A})$$

<u>GAME</u> $\mathcal{G}_3^{\text{ltss}}$: In this game, we use the RO-simulatability of Sig to simulate A*'s signatures without using a^* . Similar to the case of clean_{LTLT}(π^*) = 1, we can use the δ -RO-simulatability of Sig to bound the probability difference, so we have

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_2^{\mathsf{Itss}}}(\mathcal{A}) \leq \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_3^{\mathsf{Itss}}}(\mathcal{A}) + \delta$$

GAME \mathcal{G}_4^{ltss} : We introduce an abort event $QueryDH_{ltss}$ in \mathcal{G}_4^{ltss} : Specifically, if

- \mathcal{A} queries KDF_f on (DH, A^*, B, S^*, Y, X) , where $d = h_0(A^*, X), e_1 = h_1(B, Y, S^*), e_2 = h_2(B, Y, S^*), \text{ DDH}(X(A^*)^d, YB^{e_1}(S^*)^{e_2}, DH) = 1$, and $(A^*, B, S^*, Y, X) = \pi^*.ctxt$,
- \mathcal{A} queries KDF_r on (DH, A^*, B, S^*, X) , where $d = h_0(A^*, X), e_1 = h_3(B, S^*)$, and $\mathsf{DDH}(X(A^*)^d, S^*B^{e_1}, DH) = 1$ and $(A^*, B, S^*, X) = \pi^*.ctxt$,

then the game aborts. In other words, we abort the game if \mathcal{A} queries KDF_f or KDF_r on any DH value that contains the information of $(A^*)^{s^*} (= (S^*)^{a^*})$. We construct a CRGapDH adversary $\mathcal{B}_{\mathsf{ltss}}$ to bound the probability of $\mathsf{QueryDH}_{\mathsf{ltss}}$.

Reduction $\mathcal{B}_{\mathsf{ltss}}^{\mathsf{Ddh},\mathsf{Ch}}(A^*,S^*)$:

- Initialize $\mathcal{G}_4^{\mathsf{ltss}}$:

- 1. Choose two users \mathbb{A}^* and \mathbb{B}^* uniformly at random. Pick ssid^{*} $\leftarrow s[n_{ss}]$. The long-term public key of \mathbb{A}^* is set as A^* .
- 2. Initialize lists $\mathcal{L}_{patch-f}$, \mathcal{L}_{skey-f} , $\mathcal{L}_{patch-r}$, and \mathcal{L}_{skey-r} for the simulation of session keys using DDH and random oracle patching (as in \mathbf{G}_2). We also initialize \mathcal{L}_{sim} to simulate the signatures of A^* 's and B^* 's semi-static keys (as we did in Fig. 10).

- 40 Rune Fiedler, Felix Günther, Jiaxin Pan, and Runzhi Zeng
 - 3. For party A^* , run the modified KGenSS algorithm defined in Fig. 10. After this step, \mathcal{B}_{ltss} has ltpk, sspk, and ltsk (except for the long-term secret key of A^*). \mathcal{B}_{ltss} further replaces $sspk_{B^*}[ssid^*]$ with S^* . For all other parties, follow the KGenLT and KGenSS defined in XHMQV faithfully.
 - 4. Run $b_{guess} \leftarrow \mathcal{A}(ltpk, sspk)$ and proceeds the game as in \mathcal{G}_{KE}^{Kl} . If the initiator and responder of the test session are not A^* and B^* respectively, then overwrite \mathcal{A} 's bit with 0.

The random oracles and game oracles are simulated as described below. – Simulation of oracles defined in $\mathcal{G}_{\mathsf{KE}}^{\mathsf{KI}}$:

- All oracles except SEND, CORRUPTLTKEY, and CORRUPTSSKEY are simulated in the same way defined in G^{KI}_{KF}.
- By the definition of $clean_{LTSS}$, \mathcal{A} does not query CORRUPTLTKEY on A^* , nor to CORRUPTSSKEY on $(B^*, ssid^*)$. Queries to CORRUPTLTKEY and CORRUPTSSKEY on other parties and sessions are simulated as usual.
- For a SEND query, if the query does not involve either A^* or B^* with ssid^{*}, then we also simulate the oracle in the same way as in \mathcal{G}_{KE}^{KI} . If the query involves A^* or B^* with ssid^{*}, then we simulate in the same ways as in \mathcal{B}_{Itlt} . Namely, we use $\mathcal{L}_{patch-f}$, \mathcal{L}_{skey-f} , $\mathcal{L}_{patch-r}$, and \mathcal{L}_{skey-r} to make sure that we can always create session keys that are consistent with the random oracles.
- Simulation of Random Oracles:
 - For fresh RO queries to KDF_f and KDF_r that do not match π^* , we simulate in the same way as in \mathbf{G}_2 . Similarly, for fresh RO queries to h_0, h_1, h_2 , and h_3 that do not involve either A^* or S^* , we simulate as usual.
 - $\underline{h_0(A^*, X): h \leftarrow CH(X), h_0(A^*, X):= h.}$
 - $\overline{h_2(B,Y,S^*)}: h \leftarrow CH(Y), h_0(B,Y,S^*):=h.$
 - $\overline{h_1}$ and h_3 : Simulate as usual.
 - $\mathsf{KDF}_f(DH, A^*, B^*, S, Y, X)$:
 - 1. If $(A^*, B^*, S^*, Y, X) = \pi^* . ctxt$ and DDH $(X(A^*)^d, Y(B^*)^{e_1}(S^*)^{e_2}, DH) = 1$, then extract the value

$$DH' := \begin{cases} \left(DH/(X(A^*)^d)^{y+b^*e_1} \right)^{e_2^{-1}}, & \pi^* \text{ is a responder session} \\ \left(DH/(Y(B^*)^{e_1}(S^*)^{e_2})^x \cdot A^{de_1b} \right)^{d^{-1}}, & \pi^* \text{ is an initiator session} \end{cases}$$

where $d := h_0(A^*, X), e_1 := h_1(B^*, Y, S^*)$, and $e_2 := h_2(B^*, Y, S^*)$. In this game, we have the long-term keys of all parties except A^* , so we have the long-term key b^* of the party B^* . If π^* is an initiator session, then we have its ephemeral key x. If π^* is a responder session, then we have the ephemeral key y. Once such a DH' is found, \mathcal{B}_{ltss} aborts the simulation and return DH' as the CRGapDH solution.

- 2. Otherwise, simulate by lazy sampling.
- $KDF_r(DH, A^*, B^*, S, X)$:
 - 1. If $(A^*, B^*, S^*, X) = \pi^* . ctxt$ and $DDH(X(A^*)^d, S^*(B^*)^{e_1}, DH) = 1$, then extract the value

$$DH' := \begin{cases} DH/(X(A^*)^d)^{be_1}, & \pi^* \text{ is a responder session} \\ \left(DH/\left(S^*(B^*)^{e_1}\right)^x \cdot A^{de_1b}\right)^{d^{-1}}, & \pi^* \text{ is an initiator session} \end{cases}$$

where $d := h_0(A^*, X)$ and $e_1 := h_3(B^*, Y, S^*)$. In this game, we have the long-term key b^* of B^* . If π^* is a responder session, then we have the ephemeral key y, and thus $\mathcal{B}_{\mathsf{ltss}}$ can calculate DH' and return DH'as the CRGapDH solution.

If π^* is an initiator session, then $DH' = \mathsf{DH}(A^*, S^*)$. However, DH'itself does not align with the format of a CRGapDH solution. To deal with it, $\mathcal{B}_{\mathsf{ltss}}$ first pick a random $r \leftarrow \mathbb{Z}_p$ and compute $R := g^r$, and queries $\mathrm{CH}(R)$ to obtain h. Then $\mathcal{B}_{\mathsf{ltss}}$ computes $DH'' := (S^*)^r \cdot (DH')^h$ and outputs DH'' as the CRGapDH solution. One can verify that $DH'' = (R(A^*)^h)^{s^*}$, where $h = \mathrm{CH}(R)$, which means that it is a valid CRGapDH solution.

2. Otherwise, simulate by lazy sampling.

Here we only prove that the simulation of KDF_f and KDF_r can extract the $\mathsf{CRGapDH}$ solution. In KDF_f , if π^* is an initiator session, then the extracted DH value DH' equals to $(Y(S^*)^{e_2})^{a^*}$, where $e_2 = \mathrm{CH}(B^*, Y, S^*) = h$ and h is from $\mathrm{CH}(Y)$, so DH' is a valid solution in this case. If π^* is a responder session, then $DH' = (X(A^*)^d)^{s^*}$ and d is from $\mathrm{CH}(X)$, and thus it is also a valid $\mathsf{CRGapDH}$ solution. A similar argument applies to KDF_r . Therefore, we have

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_{\mathsf{ltss}}^{\mathsf{ltss}}}(\mathcal{A}) \leq \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_{\mathsf{ltss}}}(\mathcal{A}) + \mathsf{Adv}_{(\mathbb{G},p,g)}^{\mathsf{CRGapDH}}(\mathcal{B}_{\mathsf{ltss}}) + Q_{\mathrm{RO}}/p$$

The remain argument to bound $\Pr[QueryDH_{ltss}]$ is the same as in the case of $clean_{LTLT}$, so we have

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_4^{\mathsf{ltss}}}(\mathcal{A}) = 0.$$

We assume that the ϵ -CRGapDH assumption holds over (\mathbb{G}, p, g). Combining all the probability bounds gives us the final bound of the case clean_{LTSS}:

$$\mathsf{Adv}^{\mathbf{G}_2 \wedge \mathsf{clean}_{\mathsf{LTSS}}}_{\mathsf{XHMOV}}(\mathcal{A}) \leq n_u^2 n_{ss} \cdot (\delta + \epsilon + Q_{\mathrm{RO}}/p)$$

A.3 Case clean_{LTE}(π^*)

In this case, π^* performs a full handshake as the ephemeral secret of the responder is involved. By definition of $\mathsf{clean}_{\mathsf{LTE}}$, the initiator's long-term key is uncompromised and the responder's session randomness is not revealed. We begin with $\mathcal{G}_0^{\mathsf{lte}} := \mathbf{G}_2 \wedge \mathsf{clean}_{\mathsf{LTE}}(\pi^*)$, so we have

$$\mathsf{Adv}^{\mathbf{G}_2 \wedge \mathsf{clean}_{\mathsf{LTE}}(\pi^*)}_{\mathsf{XHMQV}}(\mathcal{A}) = \mathsf{Adv}^{\mathcal{G}^{\mathsf{lte}}_0}_{\mathsf{XHMQV}}(\mathcal{A})$$

We use game sequence $\mathcal{G}_1^{\mathsf{lte}}$ - $\mathcal{G}_4^{\mathsf{lte}}$ to bound the winning advantage of \mathcal{A} when π^* satisfies $\mathsf{clean}_{\mathsf{LTE}}$. Similar to the proof of the case $\mathsf{clean}_{\mathsf{LTLT}}(\pi^*)$, we use game transitions $\mathcal{G}_1^{\mathsf{lte}}$ - $\mathcal{G}_3^{\mathsf{lte}}$ and the simulatability (cf. Definition 3) of the signature scheme to simulate signatures of the initiator of π^* without using its long-term key. Then, we use $\mathcal{G}_4^{\mathsf{lte}}$ - $\mathcal{G}_4^{\mathsf{lte}}$ and the CRGapDH assumption to prove that the adversary cannot query KDF_f (or KDF_r) on the correct hash input of π^* , and thus

cannot distinguish whether π^* 's key is real or random.

<u>GAME</u> $\mathcal{G}_1^{\text{lte}}$: We guess the responder session π_r^* involved in the test session π^* (i.e., either $\pi^* = \pi_r^*$ or $\pi^*.cid = \pi_r^*.cid$). If the guess is incorrect, then we overwrite \mathcal{A} 's bit guess with 0. We denote π_r^* 's ephemeral key pair as (y^*, Y^*) .

Specifically, we guess π_r^* in the following way (the code will be provided in the reduction described later): We first initialize a counter $\mathsf{cnt} := 0$ and choose $\mathsf{cnt}^* \leftarrow [n_s]$. When \mathcal{A} issues $\mathsf{SEND}(\mathsf{U}, i, m)$ queries with $m = (\mathsf{create}, *, *)$ (i.e., initializes a responder session), then cnt increments. If $\mathsf{cnt} = \mathsf{cnt}^*$, then we set $\pi_r^* := \pi_{\mathsf{U}}^*$.

In $\mathcal{G}_{1}^{\mathsf{lte}}$, there are at most n_s sessions during $\mathcal{G}_{\mathsf{KE}}^{\mathsf{KI}}$. Since cnt^* is independent of \mathcal{A} 's view, the probability that cnt^* -th responder session (regardless of the owner) is involved in π^* is $\frac{1}{n_c}$, and thus we have

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_0^{\mathsf{lte}}}(\mathcal{A}) = n_s \cdot \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_1^{\mathsf{lte}}}(\mathcal{A})$$

GAME $\mathcal{G}_2^{\text{lte}}$: We guess the identity \mathbf{A}^* of the initiator involved in the test session. If this guess is incorrect, then we overwrite \mathcal{A} 's bit guess with 0. We denote \mathbf{A}^* 's long-term key pair as (a^*, A^*) . Since there are at most n_u users involved in this game and the guess is uniformly at random in \mathcal{A} 's view, we have

$$\mathsf{Adv}^{\mathcal{G}_1^{\mathsf{lte}}}_{\mathsf{XHMQV}}(\mathcal{A}) \leq n_u \cdot \mathsf{Adv}^{\mathcal{G}_2^{\mathsf{lte}}}_{\mathsf{XHMQV}}(\mathcal{A})$$

<u>GAME</u> $\mathcal{G}_3^{\text{lte}}$: In this game, we use the RO-simulatability of Sig to simulate A*'s signatures only using A*'s long-term public key A*. Similar to $\mathcal{G}_2^{\text{ltl}}$ and $\mathcal{G}_3^{\text{ltl}}$ in the case of clean_{LTLT}(π^*), we can use the δ -RO-simulatability of Sig to bound the probability difference.

$$\mathsf{Adv}^{\mathcal{G}_2^{\mathsf{lte}}}_{\mathsf{XHMQV}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathcal{G}_3^{\mathsf{lte}}}_{\mathsf{XHMQV}}(\mathcal{A}) + \delta$$

GAME $\mathcal{G}_{4}^{\mathsf{lte}}$: We introduce an abort event $\mathsf{QueryDH}_{\mathsf{lte}}$ in $\mathcal{G}_{4}^{\mathsf{lte}}$: Specifically, if

- \mathcal{A} queries KDF_f on (DH, A^*, B, S, Y^*, X) , where $d = h_0(A^*, X), e_1 = h_1(B, Y^*, S), e_2 = h_2(B, Y^*, S)$, $\mathsf{DDH}(X(A^*)^d, Y^*B^{e_1}S^{e_2}, DH) = 1$, and $(A^*, B, S, Y^*, X) = \pi^*.ctxt$,

then the game aborts (i.e., return 0 as \mathcal{A} 's bit guess). In other words, we abort the game if \mathcal{A} queries KDF_f or KDF_r on any DH value that contains the information of $(\mathcal{A}^*)^{y^*} (= (Y^*)^{a^*})$. We construct a CRGapDH adversary $\mathcal{B}_{\mathsf{lte}}$ such that if such an abort event happens, then $\mathcal{B}_{\mathsf{lte}}$ can output a solution to its CRGapDH challenge.

Reduction $\mathcal{B}_{\mathsf{lte}}^{\mathsf{Ddh},\mathsf{Ch}}(A^*,Y^*)$:

– Initialize $\mathcal{G}_4^{\mathsf{lte}}$:

1. Choose a party A^* uniformly at random and set its long-term public key as A^* .

43

- 2. Initialize a counter $\mathsf{cnt} := 0$ and choose $\mathsf{cnt}^* \leftarrow [n_s]$. Initialize $\pi_r^* := \bot$ and set it up later.
- 3. Initialize lists $\mathcal{L}_{patch-f}$, \mathcal{L}_{skey-f} , $\mathcal{L}_{patch-r}$, and \mathcal{L}_{skey-r} , and \mathcal{L}_{sim} .
- 4. For party A^* , we use its long-term public key, the algorithm SIMsign, and \mathcal{L}_{sim} to simulate the signatures of its semi-static keys (as we did in Fig. 10). For all other parties, follow the KGenLT and KGenSS defined in XHMQV to generate *ltpk*, *sspk*, and *ltsk* faithfully.
- 5. Run $b_{guess} \leftarrow \mathcal{A}(ltpk, sspk)$ and proceed the game as in $\mathcal{G}_{KE}^{\mathsf{KI}}$. If the initiator of π^* is not \mathbb{A}^* , π_r^* is not π^* itself, or π_r^* is not the contributive partner session of π^* , then overwrite \mathcal{A} 's bit guess with 0.

The random oracles and game oracles are simulated as described below. – Simulation of oracles defined in $\mathcal{G}_{\mathsf{KE}}^{\mathsf{KI}}$:

- All oracles except SEND, CORRUPTLTKEY, and REVEALRAND are simulated in the same way defined in $\mathcal{G}_{\mathsf{KE}}^{\mathsf{KI}}.$
- If CORRUPTLTKEY(A^*) is queried, then abort the game and overwrite \mathcal{A} 's bit guess with 0. By definition of $\mathsf{clean}_{\mathsf{LTE}}$, if the guess is correct, then \mathcal{A} will not corrupt A^* . CORRUPTLTKEY queries on other parties are simulated as usual.
- If REVEALRAND(\mathbf{U}, i) is queried, where $\pi_i^{\mathbf{U}} = \pi_r^* \neq \bot$, then abort the game and overwrite \mathcal{A} 's bit with 0. By definition of $\mathsf{clean}_{\mathsf{LTE}}$, if the guess is correct, then \mathcal{A} will not reveal the session randomness of π_r^* . REVEALRAND queries on other sessions are simulated as usual.
- Send($U \neq A^*, i, m$):
 - * Case m = (create, ssid, type) (as the session responder): cnt := cnt + 1. If $\text{cnt} = \text{cnt}^*$ and type = full, then we use Y^* as the ephemeral public key of this session and set $\pi_r^* := \pi_i^{U}$. Otherwise, simulate as usual. * For other cases, simulate as usual.
- SEND (A^*, i, m) (We only present how to set up the session keys):
 - * Case m = (create, ssid, type) (as the session responder): Simulate as usual. If type = full, then cnt := cnt + 1.
 - * Case $m = (\mathbf{U}, X)$ (as the session responder): We apply the same approach as in case clean_{LTLT} for \mathcal{B}_{ltlt} . Namely, we use the lists $\mathcal{L}_{patch-f}$, $\mathcal{L}_{patch-r}$, \mathcal{L}_{skey-f} , and \mathcal{L}_{skey-r} and oracle patching for KDF_f and KDF_r to simulate the session keys and the RO outputs of KDF_f and KDF_r consistently.
 - * Case $m = (B, S, Y, \sigma)$ (as the session initiator): Simulate analogously to the case m = (U, X) (no matter whether $Y = Y^*$).
- Simulation of Random Oracles:
 - For fresh RO queries to $\mathsf{KDF}_r(DH, A, B, S, X)$ and $\mathsf{KDF}_f(DH, A, B, S, Y, X)$, where $(A, Y) \neq (A^*, Y^*)$, we simulate in the same way as in \mathbf{G}_2 . Similarly, for fresh RO queries to h_0, h_1, h_2 , and h_3 that do not involve \mathbf{A}^* and Y^* , we simulate as usual.
 - $h_0(A^*, X): h \leftarrow CH(X), h_0(A, X) := h.$
 - $h_1(B, Y^*, S)$: Simulate as usual.
 - $h_2(B, Y^*, S): h \leftarrow CH(S), h_1(B, Y^*, S) := h^{-1}$
 - $\operatorname{\mathsf{KDF}}_f(DH, A^*, B, S, Y^*, X)$:

1. If $DDH(X(A^*)^d, Y^*B^{e_1}S^{e_2}, DH)$ outputs 1, then extract the value

$$DH' := \begin{cases} DH/(X(A^*)^d)^{be_1 + se_2}, & \pi^* \text{ is a responder session} \\ \left(DH/\left(Y^*B^{e_1}S^{e_2}\right)^x \cdot (A^*)^{dbe_1} \right)^{d^{-1}e_2^{-1}}, & \pi^* \text{ is an initiator session} \end{cases}$$

where $d := h_0(A^*, X), e_1 := h_1(B, Y, S)$, and $e_2 := h_2(B, Y, S)$. If π^* is an initiator session, then we have its ephemeral key x. If π^* is a responder session, then we have the long-term key b of the responder B and the semi-static key s. $\mathcal{B}_{\mathsf{Ite}}$ aborts the simulation and return DH' as the CRGapDH solution.

2. Otherwise, simulate by lazy sampling.

We first prove that the simulation of KDF_f can successfully extract the $\mathsf{CRGapDH}$ solution. In KDF_f , if π^* is an initiator session, then the extracted DH value DH' equals to $(Y^*S^{e_2})^{a^*e_2^{-1}} = ((Y^*)^{e_2^{-1}}S)^{a^*}$. By the simulation of $h_2(B^*, Y, S)$, we have $e_2 = h_2(B, Y^*, S) = h^{-1}$ where $h \leftarrow \mathrm{CH}(S)$. Therefore, we have

$$DH' = ((Y^*)^{e_2^{-1}}S)^{a^*} = (S(Y^*)^h)^{a^*},$$

where h = CH(S). So, DH' is a solution to CRGapDH. Similarly, if π^* is a responder session, then we have

$$DH' = (X(A^*)^d)^{y^*},$$

where d = CH(X). This is also a valid solution to CRGapDH.

It is straight-forward to see that if $QueryDH_{lte}$ occurs, then \mathcal{B}_{lte} will extract a CRGapDH solution from KDF_f . So, we have

$$\begin{split} \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_4^{\mathsf{lte}}}(\mathcal{A}) &\leq \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_4^{\mathsf{lte}}}(\mathcal{A}) + \Pr\left[\mathtt{QueryDH}_{\mathsf{lte}}\right] \\ &\leq \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_4^{\mathsf{lte}}}(\mathcal{A}) + \mathsf{Adv}_{(\mathbb{G},p,g)}^{\mathsf{CRGapDH}}(\mathcal{B}_{\mathsf{lte}}) + Q_{\mathrm{RO}}/p, \end{split}$$

where the last term counts in the probability that there exists a non-invertible h value from the CH oracle. Finally, we use a similar argument as for bounding $\mathcal{G}_4^{\text{ltt}}$ to bound \mathcal{A} 's advantage in $\mathcal{G}_4^{\text{lte}}$. That is, if \mathcal{A} cannot trigger QueryDH_{lte}, then the session key of π^* is distributed uniformly at random in \mathcal{A} 's view. Therefore, we have

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_4^{\mathsf{ite}}}(\mathcal{A}) = 0.$$

We assume that the ϵ -CRGapDH assumption holds over (\mathbb{G}, p, g) . Combining all the probability bounds gives us the final bound of the case clean_{LTE} (π^*) :

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathbf{G}_2 \land \mathsf{clean}_{\mathsf{LTE}}}(\mathcal{A}) \le n_u n_s \cdot (\delta + \epsilon + Q_{\mathrm{RO}}/p)$$

A.4 Case clean_{ELT}

In this case, $\mathsf{clean}_{\mathsf{ELT}}$ ensures for π^* that (1) the initiator randomness involved in π^* is not revealed, (2) the long-term key of the responder of π^* is uncompromised, and (3) if π^* is a responder session, then there exists a partnered initiator session. Condition (3) ensures that the initiator ephemeral public key of π^* is always controlled by the game (so we can embed the CRGapDH challenge). Since $\mathsf{clean}_{\mathsf{LTE}}$ is similar to $\mathsf{clean}_{\mathsf{LTE}}$, we may reuse some arguments in proving $\mathsf{clean}_{\mathsf{LTE}}$ (cf. Appendix A.3). We begin with $\mathcal{G}_0^{\mathsf{elt}} := \mathbf{G}_2 \wedge \mathsf{clean}_{\mathsf{ELT}}$, so we have

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathbf{G}_2 \wedge \mathsf{clean}_{\mathsf{ELT}}}(\mathcal{A}) = \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_0^{\mathsf{elt}}}(\mathcal{A})$$

GAME $\mathcal{G}_{1}^{\text{elt}}$: We guess the initiator session π_{i}^{*} involved in the test session π^{*} (i.e., either $\pi^{*} = \pi_{i}^{*}$ or $\pi^{*}.\text{sid} = \pi_{i}^{*}.\text{sid}$). If the guess is incorrect, then we overwrite \mathcal{A} 's bit guess with 0. We denote π_{i}^{*} 's ephemeral key pair as (x^{*}, X^{*}) .

We apply the same guessing approach used in proving $\mathcal{G}_1^{\text{lte}}$ to this proof case. Roughly, we use a counter cnt which increments each time \mathcal{A} queries SEND to initialize an initiator session. And we guess $\text{cnt}^* \leftarrow [n_s]$ and set $\pi_i^* := \pi_{U}^i$ if SEND(U, i, m) is the cnt*-th SEND query for initializing an initiator session. We have

$$\mathsf{Adv}^{\mathcal{G}_0^{\mathsf{elt}}}_{\mathsf{XHMQV}}(\mathcal{A}) = n_s \cdot \mathsf{Adv}^{\mathcal{G}_1^{\mathsf{elt}}}_{\mathsf{XHMQV}}(\mathcal{A})$$

GAME $\mathcal{G}_2^{\text{elt}}$: We guess the identity B^* of the responder involved in the test session. If this guess is incorrect, then we overwrite \mathcal{A} 's bit guess with 0. We denote B^* 's long-term key pair as (b^*, B^*) . Since there are at most n_u users involved in this game and the guess is uniformly at random in \mathcal{A} 's view, we have

$$\mathsf{Adv}^{\mathcal{G}_1^{\mathrm{elt}}}_{\mathsf{XHMQV}}(\mathcal{A}) \leq n_u \cdot \mathsf{Adv}^{\mathcal{G}_2^{\mathrm{elt}}}_{\mathsf{XHMQV}}(\mathcal{A})$$

<u>GAME</u> $\mathcal{G}_3^{\text{elt}}$: In this game, we use the RO-simulatability of Sig to simulate B^{*}'s signatures without using b^* . Similar to previous proof cases that involve uncompromised long-term secrets, we can use the δ -RO-simulatability of Sig to bound the probability difference.

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_2^{\mathsf{elt}}}(\mathcal{A}) \leq \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_3^{\mathsf{elt}}}(\mathcal{A}) + \delta$$

GAME $\mathcal{G}_4^{\text{elt}}$: We introduce an abort event $\text{QueryDH}_{\text{elt}}$ in $\mathcal{G}_4^{\text{elt}}$: Specifically, if

- \mathcal{A} queries KDF_f on (DH, A, B^*, S, Y, X^*) , where $d = h_0(A, X^*), e_1 = h_1(B, Y^*, S), e_2 = h_2(B^*, Y, S), \mathsf{DDH}(X^*A^d, Y(B^*)^{e_1}S^{e_2}, DH) = 1$, and $(A, B^*, S, Y, X^*) = \pi^*.ctxt$, or
- \mathcal{A} queries KDF_r on (DH, A, B^*, S, X^*) , where $d = h_0(A, X^*)$, $e_1 = h_3(B^*, S)$, and $\mathsf{DDH}(X^*(A)^d, S(B^*)^{e_1}, DH) = 1$ and $(A, B^*, S, X^*) = \pi^*.ctxt$,

then the game aborts (i.e., return 0 as \mathcal{A} 's bit guess). In other words, we abort the game if \mathcal{A} queries KDF_f or KDF_r on any DH value that contains the information of $(B^*)^{x^*} (= (X^*)^{b^*})$. We construct a CRGapDH adversary $\mathcal{B}_{\mathsf{elt}}$ such that if such an abort event happens, then \mathcal{B}_{elt} can output a solution to its CRGapDH challenge.

Reduction $\mathcal{B}_{\mathsf{elt}}^{\mathsf{DDH},\mathsf{CH}}(X^*,B^*)$:

- Initialize $\mathcal{G}_4^{\mathsf{elt}}$:
 - 1. Choose a party B^* uniformly at random and set the long-term public key of B^* as B^* .
 - 2. Initialize a counter cnt := 0 and choose $cnt^* \leftarrow [n_s]$. Initialize $\pi_i^* := \bot$ and set up it later.
 - 3. Initialize lists $\mathcal{L}_{patch-f}$, \mathcal{L}_{skey-f} , $\mathcal{L}_{patch-r}$, and \mathcal{L}_{skey-r} , and \mathcal{L}_{sim} .
 - 4. For party B^{*}, we use B^* , the algorithm SIMsign, and \mathcal{L}_{sim} to simulate the signatures of its semi-static keys (as we did in Fig. 10). For all other parties, follow the KGenLT and KGenSS defined in XHMQV to generate *ltpk*, *sspk*, and *ltsk* faithfully.
 - 5. Run $b_{guess} \leftarrow \mathcal{A}(ltpk, sspk)$ and proceeds the game as in \mathcal{G}_{KE}^{Kl} . If the responder of π^* is not B^* , π_i^* is not π^* , or π_i^* is not the matching session of π^* , then overwrite \mathcal{A} 's bit with 0.

The random oracles and game oracles are simulated as described below. – Simulation of oracles defined in $\mathcal{G}_{\mathsf{KF}}^{\mathsf{KI}}$:

- All oracles except SEND, CORRUPTLTKEY, and REVEALRAND are simulated in the same way defined in $\mathcal{G}_{\mathsf{KE}}^{\mathsf{KI}}$.
- If CORRUPTLTKEY(B^*) is queried, then abort the game and overwrite \mathcal{A} 's bit with 0. By definition of $\mathsf{clean}_{\mathsf{ELT}}$, if the guess is correct, then \mathcal{A} will not corrupt B^* . CORRUPTLTKEY queries on other parties are simulated as usual.
- If REVEALRAND(\mathbf{U}, j) is queried, where $\pi_j^{\mathbf{U}} = \pi_i^* \neq \bot$, then abort the game and overwrite \mathcal{A} 's bit with 0. By definition of $\mathsf{clean}_{\mathsf{ELT}}$, if the guess is correct, then \mathcal{A} will not reveal the session randomness of π_i^* . REVEALRAND queries on other sessions are simulated as usual.
- Send($U \neq B^*, j, m$):
 - * Case m = (B*, sspk, epk) (initiator session with intended partner B*): cnt := cnt + 1. If cnt = cnt*, then we use X* as the ephemeral public key of this session and set π_i* := π_j^U. Otherwise, simulate as usual.
 * For other cases, simulate as usual.
- SEND(B^{*}, j, m): We apply the same approach used in $\mathcal{B}_{\mathsf{ltlt}}$ to this case. Namely, we use the lists $\mathcal{L}_{\mathsf{patch-f}}$, $\mathcal{L}_{\mathsf{patch-r}}$, $\mathcal{L}_{\mathsf{skey-f}}$, and $\mathcal{L}_{\mathsf{skey-r}}$ and oralce patching of KDF_f and KDF_r to simulate the session keys and the RO outputs of KDF_f and KDF_r consistently.
- Simulation of Random Oracles:
 - For fresh RO queries to $\mathsf{KDF}_r(DH, A, B, S, X)$ and $\mathsf{KDF}_f(DH, A, B, S, Y, X)$, where $(B, X) \neq (B^*, X^*)$, we simulate in the same way as in \mathbf{G}_2 . Similarly, for fresh RO queries to h_0, h_1, h_2 , and h_3 that do not involve B^* and X^* , we simulate as usual.
 - $\frac{h_1(B^*, Y, S)$: Compute $e_2 := h_2(B^*, Y, S), R := YS^{e_2}$, and queries CH(R)and gets h. Set $h_1(B^*, Y, S) := h$.

47

- $h_3(B^*, S): h \leftarrow CH(S), h_3(B^*, S) := h.$
- $h_0(A, X^*)$ and $h_2(B, Y^*, S)$: Simulate as usual.
- $KDF_f(DH, A, B^*, S, Y, X^*)$:

1. If
$$(A, B^*, S, Y, X^*) = \pi^* . ctxt$$
 and DDH $(X^*A^d, Y(B^*)^{e_1}S^{e_2}, DH)$ outputs 1, then extract the value

$$DH' := \begin{cases} \left(DH/(X^*A^d)^{y+se_2} \cdot (B^*)^{e_1ad} \right)^{e_1^{-1}}, & \pi^* \text{ is a responder session} \\ DH/(Y(B^*)^{e_1}S^{e_2})^{ad}, & \pi^* \text{ is an initiator session} \end{cases}$$

where $d := h_0(A, X^*), e_1 := h_1(B^*, Y, S)$, and $e_2 := h_2(B^*, Y, S)$. We have a since we have long-term keys of all parties except B^* . If π^* is an initiator session, then $\mathcal{B}_{\mathsf{lte}}$ aborts the simulation and return DH'as the CRGapDH solution. If π^* is a responder session (so we have the ephemeral key y and the semi-static key s), then we have that $DH' = (X^*)^{b^*}$. Similar to $\mathcal{B}_{\mathsf{ltss}}$, in this case, $\mathcal{B}_{\mathsf{elt}}$ first pick a $r \leftarrow \mathbb{Z}_p$ and compute $R := g^r$, and queries CH(R) and gets h. Then $\mathcal{B}_{\mathsf{ltss}}$ computes $DH'' := (S^*)^r \cdot (DH')^h$ and outputs DH' as the CRGapDH solution.

- 2. Otherwise, simulate by lazy sampling.
- $KDF_r(DH, A, B^*, S, X^*)$:
 - 1. If $(A, B^*, S, X^*) = \pi^*.ctxt$ and $DDH(X^*A^d, S(B^*)^{e_1}, DH)$ outputs 1, then extract the value

$$DH' := \begin{cases} \left(DH/(X^*A^d)^s \cdot (B^*)^{e_1 a d} \right)^{e_1^{-1}}, & \pi^* \text{ is a responder session} \\ DH/(S(B^*)^{e_1})^{a d}, & \pi^* \text{ is an initiator session} \end{cases}$$

where $d := h_0(A, X^*)$ and $h_1(B^*, S)$. If π^* is an initiator session, then $\mathcal{B}_{\mathsf{lte}}$ aborts the simulation and return DH' as the CRGapDH solution. If π^* is a responder session, then we have $DH' = (X^*)^{b^*}$, and we can use the same approach in KDF_f to construct a valid CRGapDH solution. Otherwise, simulate by large sampling

2. Otherwise, simulate by lazy sampling.

Here we only prove that the simulation of KDF_f and KDF_r can extract the CRGapDH solution. In KDF_f , if π^* is an initiator session with full mode, then we have $DH' = (Y(B^*)^{e_1}S^{e_2})^{x^*} = ((YS^{e_2}) \cdot (B^*)^{e_1})^{x^*}$, where $e_1 = h_1(B^*, Y, S)$ is from $CH(YS^{e_2})$, so DH' is a valid solution in this case. A similar argument applies to KDF_r . Therefore, we have

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}^{\mathsf{elt}}_{\mathsf{I}}}(\mathcal{A}) \leq \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}^{\mathsf{elt}}_{\mathsf{I}}}(\mathcal{A}) + \mathsf{Adv}_{(\mathbb{G},p,g)}^{\mathsf{CRGapDH}}(\mathcal{B}_{\mathsf{ltss}}) + Q_{\mathrm{RO}}/p$$

The remain argument to bound $\Pr[\texttt{QueryDH}_{\mathsf{ltss}}]$ is the same as in previous cases. We have $\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_4^{\mathsf{et}}}(\mathcal{A}) = 0$. We assume that the ϵ -CRGapDH assumption holds over (\mathbb{G}, p, g) . Combining all the probability bounds gives us the final bound of the case clean_{LTSS}:

$$\operatorname{Adv}_{\mathsf{XHMQV}}^{\mathbf{G}_2 \wedge \operatorname{clean}_{\mathsf{ELT}}}(\mathcal{A}) \leq n_u n_s \cdot (\delta + \epsilon + Q_{\mathrm{RO}}/p)$$

A.5 Case clean_{ESS}

In this case, clean_{ESS} ensures for π^* that:

- The initiator randomness involved in π^* is not revealed.
- The responder's semi-static key is uncompromised.
- If π^* is an initiator session, then the responder's long-term key was uncompromised upon acceptance or $\pi^*.sspk \in sspk_{\pi^*.pid}$.
 - We can use the argument in Appendix A.2 to compress this condition into simply $\pi^*.sspk \in sspk_{\pi^*.pid}$ (without considering $\pi^*.pcorr$). Roughly, \mathbf{G}_2 ensures that \mathcal{A} cannot forge a signature on an uncompromised party, so we can argue that $clean_{ESS}(\pi^*) \implies \pi^*.sspk \in sspk_{\pi^*.pid}$. Therefore, the third condition described above and the modification in \mathbf{G}_2 ensure that the semistatic key used in π^* has been generated by the game. This enables us to embed a CRGapDH challenge into the semi-static key.
- If π^* is a responder session, then there exists a partner initiator session. This condition ensures that the initiator ephemeral key of π^* is controlled by the game so that we can embed CRGapDH challenge when doing the reduction.

In this proof case, we reuse some arguments from proving $\mathsf{clean}_{\mathsf{ELT}}$ and $\mathsf{clean}_{\mathsf{LTSS}}$ (cf. Appendices A.2 and A.4). The main difference is that the reductions always know the long-term secret keys of all parties, so we do not need to use the simulatability property of Sig. We begin with $\mathcal{G}_0^{\mathsf{ess}} := \mathbf{G}_2 \wedge \mathsf{clean}_{\mathsf{ESS}}$, so we have

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathbf{G}_2 \wedge \mathsf{clean}_{\mathsf{ESS}}}(\mathcal{A}) = \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_0^{\mathrm{css}}}(\mathcal{A})$$

<u>GAME</u> $\mathcal{G}_1^{\text{ess}}$: We guess the initiator session π_i^* involved in the test session π^* (i.e., either $\pi^* = \pi_i^*$ or $\pi^*.sid = \pi_i^*.sid$). If the guess is incorrect, then we overwrite \mathcal{A} 's bit guess with 0. We denote π_i^* 's ephemeral key pair as (x^*, X^*) . We apply the same guessing approach used in proving $\mathcal{G}_1^{\text{lte}}$ and $\mathcal{G}_1^{\text{elt}}$ (cf. Appendices A.3 and A.4) to this proof case. We have

$$\mathsf{Adv}^{\mathcal{G}_0^{\mathrm{ess}}}_{\mathsf{XHMQV}}(\mathcal{A}) = n_s \cdot \mathsf{Adv}^{\mathcal{G}_1^{\mathrm{ess}}}_{\mathsf{XHMQV}}(\mathcal{A})$$

<u>GAME</u> $\mathcal{G}_2^{\text{ess}}$: We guess the identity B^* of the responder involved in π^* and the identifier ssid^{*} $\in [n_{ss}]$ of the repsonder B^* 's uncompromised semi-static key. We denote the semi-static key pair as (s^*, S^*) . We have

$$\mathsf{Adv}^{\mathcal{G}_1^{\mathrm{ess}}}_{\mathsf{XHMQV}}(\mathcal{A}) \leq n_u n_{ss} \cdot \mathsf{Adv}^{\mathcal{G}_2^{\mathrm{ess}}}_{\mathsf{XHMQV}}(\mathcal{A})$$

GAME $\mathcal{G}_3^{\text{ess}}$: We introduce an abort event QueryDH_{ess} in $\mathcal{G}_3^{\text{ess}}$: Specifically, if

- \mathcal{A} queries KDF_f on $(DH, A, B^*, S^*, Y, X^*)$, where $d = h_0(A, X^*), e_1 = h_1(B^*, Y, S^*), e_2 = h_2(B^*, Y, S^*), \text{ DDH}(X^*A^d, YB^{e_1}(S^*)^{e_2}, DH) = 1$, and $(A, B^*, S^*, Y, X^*) = \pi^*.ctxt$, or
- \mathcal{A} queries KDF_r on (DH, A, B^*, S^*, X^*) , where $d = h_0(A, X^*)$, $e_1 = h_3(B^*, S^*)$, and $\mathsf{DDH}(X^*(A)^d, S^*B^{e_1}, DH) = 1$ and $(A, B^*, S^*, X^*) = \pi^*.ctxt$,

49

then the game aborts (i.e., return 0 as \mathcal{A} 's bit guess). In other words, we abort the game if \mathcal{A} queries KDF_f or KDF_r on DH value that contains the information of $(S^*)^{x^*} (= (X^*)^{s^*})$. We construct a CRGapDH adversary $\mathcal{B}_{\mathsf{ess}}$ such that if such an abort event happens, then $\mathcal{B}_{\mathsf{ess}}$ can output a solution to its CRGapDH challenge.

Reduction $\mathcal{B}_{ess}^{\text{DDH,CH}}(X^*, S^*)$:

- Initialize \mathcal{G}_3^{ess} :

- 1. Choose a party B^* uniformly at random and pick ssid^{*} $\leftarrow s[n_{ss}]$.
- 2. Initialize a counter cnt := 0 and choose $cnt^* \leftarrow [n_s]$. Initialize $\pi_i^* := \bot$ and set up it later.
- 3. Initialize lists $\mathcal{L}_{patch-f}$, \mathcal{L}_{skey-f} , $\mathcal{L}_{patch-r}$, and \mathcal{L}_{skey-r} .
- 4. For all parties, follow the KGenLT and KGenSS defined in XHMQV to generate ltpk, sspk, and ltsk. After this step, \mathcal{B}_{ess} has ltpk, sspk, and ltsk. \mathcal{B}_{ess} further replaces $sspk_{B^*}[ssid^*]$ with S^* .
- 5. Run $b_{guess} \leftarrow \mathcal{A}(ltpk, sspk)$ and proceeds the game as in \mathcal{G}_{KE}^{KI} . If the responder of π^* is not B^* , π_i^* is not π^* , or π_i^* is not the matching session of π^* , then overwrite \mathcal{A} 's bit with 0.
- The random oracles and game oracles are simulated as described below.
- Simulation of oracles defined in $\mathcal{G}_{\mathsf{KE}}^{\mathsf{KI}}$:
 - All oracles except SEND, CORRUPTSSKEY, and REVEALRAND are simulated in the same way defined in \mathcal{G}_{KE}^{KI} .
 - If CORRUPTSSKEY(B^* , ssid^{*}) is queried, then abort the game and overwrite \mathcal{A} 's bit with 0. CORRUPTSSKEY queries on other parties are simulated as usual.
 - If REVEALRAND(U, j) is queried, where $\pi_j^U = \pi_i^* \neq \bot$, then abort the game and overwrite \mathcal{A} 's bit with 0. REVEALRAND queries on other sessions are simulated as usual.
 - Send($U \neq B^*, j, m$):
 - * Case m = (B*, sspk, epk) (initiator session with intended partner B*): cnt := cnt + 1. If cnt = cnt*, then we use X* as the ephemeral pk of this session and set π_i^{*} := π_j^U. Otherwise, simulate as usual.
 * For other cases, simulate as usual.
 - * For other cases, simulate as usual.
 - For all other SEND queries, We apply the same approach used in \mathcal{B}_{ltlt} to simulate the session key. Namely, we use the lists $\mathcal{L}_{patch-f}$, $\mathcal{L}_{patch-r}$, \mathcal{L}_{skey-f} , and \mathcal{L}_{skey-r} and oralce patching of KDF_f and KDF_r (using DDH) to simulate the session keys and the RO outputs of KDF_f and KDF_r consistently.
- Simulation of Random Oracles:
 - For fresh RO queries to $\mathsf{KDF}_r(DH, A, B, S, X)$ and $\mathsf{KDF}_f(DH, A, B, S, Y, X)$ that do not involve (B^*, S^*, X^*) , we simulate in the same way as in \mathbf{G}_2 . Similarly, for fresh RO queries to h_0, h_1, h_2 , and h_3 that do not involve (B^*, S^*, X^*) , we simulate as usual.
 - $h_2(B^*, Y, S^*): h \leftarrow CH(Y), h_2(B^*, Y, S^*):=h$
 - h_0 , h_1 , and h_3 : Simulate as usual.
 - $KDF_f(DH, A, B^*, S^*, Y, X^*)$:

- 50 Rune Fiedler, Felix Günther, Jiaxin Pan, and Runzhi Zeng
 - 1. If $(A, B^*, S^*, Y, X^*) = \pi^* . ctxt$ and $DDH(X^*A^d, YB^{e_1}(S^*)^{e_2}, DH)$ outputs 1, then extract the value

$$DH' := \begin{cases} \left(DH / (YB^{e_1}(S^*)^{e_2})^{ad} \cdot (X^*)^{y+e_1b} \right)^{e_2^{-1}}, & \pi^* \text{ is a responder session} \\ DH / (YB^{e_1}(S^*)^{e_2})^{ad} \cdot (X^*)^{e_1b}, & \pi^* \text{ is an initiator session} \end{cases}$$

where $d := h_0(A, X^*), e_1 := h_1(B^*, Y, S^*)$, and $e_2 := h_2(B^*, Y, S^*)$. We have *a* since we have long-term keys of all parties. If π^* is an initiator session, then \mathcal{B}_{ess} aborts the simulation and return DH' as the CRGapDH solution. If π^* is a responder session (so we have the ephemeral key *y*), then we have that $DH' = (X^*)^{s^*}$, and \mathcal{B}_{ess} can use a similar approach in \mathcal{B}_{ess} to construct a valid CRGapDH solution.

- 2. Otherwise, simulate by lazy sampling.
- $KDF_r(DH, A, B^*, S^*, X^*)$:
 - 1. If $(A, B^*, S^*, X^*) = \pi^* . ctxt$ and $DDH(X^*A^d, (S^*)B^{e_1}, DH)$ outputs 1, then extract the value

$$DH' := DH/(S^*B^{e_1})^{ad} \cdot (X^*)^{e_1b}$$

where $d := h_0(A, X^*), e_1 := h_1(B^*, Y, S^*)$, and $e_2 := h_2(B^*, Y, S^*)$. We have *a* and *b* since we have long-term keys of all parties. Since we have $DH' = (X^*)^{s^*}$ in this case, \mathcal{B}_{ess} can use a similar approach in \mathcal{B}_{ess} to construct a valid CRGapDH solution.

2. Otherwise, simulate by lazy sampling.

Here we only prove that the simulation of KDF_f can extract the $\mathsf{CRGapDH}$ solution. If π^* is an initiator session with full mode, then we have $DH' = (Y(S*)^{e_2})^{x^*}$, where $e_2 = h_2(B^*, Y, S^*)$ is from CH(Y), so DH' is a valid solution in this case. Therefore, we have

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_3^{\mathrm{ess}}}(\mathcal{A}) \leq \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_4^{\mathrm{ess}}}(\mathcal{A}) + \mathsf{Adv}_{(\mathbb{G},p,g)}^{\mathsf{CRGapDH}}(\mathcal{B}_{\mathsf{ess}}) + Q_{\mathrm{RO}}/p$$

The remain argument to bound $\Pr[\texttt{QueryDH}_{ess}]$ is the same as in previous cases. We have $\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_3^{ess}}(\mathcal{A}) = 0$ and ϵ -CRGapDH assumption holds over (\mathbb{G}, p, g) . Combining all the probability bounds gives us the final bound of the case clean_{ESS}:

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathbf{G}_2 \land \mathsf{clean}_{\mathsf{ESS}}}(\mathcal{A}) \le n_s n_u n_{ss} \cdot (\epsilon + Q_{\mathrm{RO}}/p)$$

A.6 Case clean_{EE}

In this case, $\mathsf{clean}_{\mathsf{EE}}$ ensures for π^* that:

- The randomness of π^* is not revealed, and
- $-\pi^*$ has a partnered session whose randomness is also not revealed.

In this proof case, we reuse some arguments from proving $\mathsf{clean}_{\mathsf{ELT}}$ and $\mathsf{clean}_{\mathsf{LTE}}$ (cf. Appendices A.3 and A.4). Similar to the case of $\mathsf{clean}_{\mathsf{ESS}}(\pi^*)$, the reductions always know the long-term secret keys of all parties, so we do not need to use the simulatability property of Sig. We begin with $\mathcal{G}_0^{\mathsf{ee}} := \mathbf{G}_2 \wedge \mathsf{clean}_{\mathsf{EE}}$, so we have

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathbf{G}_2 \wedge \mathsf{clean}_{\mathsf{EE}}}(\mathcal{A}) = \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_0^{\mathsf{e}}}(\mathcal{A})$$

<u>GAME</u> $\mathcal{G}_{1}^{\text{ee}}$: We guess the initiator session π_{i}^{*} involved in the test session π^{*} (i.e., either $\pi^{*} = \pi_{i}^{*}$ or $\pi^{*}.sid = \pi_{i}^{*}.sid$). If the guess is incorrect, then we overwrite \mathcal{A} 's bit guess with 0. We denote π_{i}^{*} 's ephemeral key pair as (x^{*}, X^{*}) . We use the same guessing approach as for proving $\mathcal{G}_{1}^{\text{ee}}$ and $\mathcal{G}_{1}^{\text{ee}}$ (cf. Appendices A.3 and A.4). We have

$$\mathsf{Adv}^{\mathcal{G}^{\mathrm{ee}}_0}_{\mathsf{XHMQV}}(\mathcal{A}) = n_s \cdot \mathsf{Adv}^{\mathcal{G}^{\mathrm{ee}}_1}_{\mathsf{XHMQV}}(\mathcal{A})$$

GAME $\mathcal{G}_2^{\text{ee}}$: We guess the responder session π_r^* involved in the test session π^* (i.e., either $\pi^* = \pi_i^*$ or $\pi^*.cid = \pi_i^*.cid$). If the guess is incorrect, then we overwrite \mathcal{A} 's bit guess with 0. We denote π_r^* 's ephemeral key pair as (y^*, Y^*) . We use the same guessing approach in proving $\mathcal{G}_1^{\text{lte}}$ and $\mathcal{G}_2^{\text{elt}}$ (cf. Appendices A.3 and A.4). We have

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_1^{\mathrm{ee}}}(\mathcal{A}) = n_s \cdot \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_2^{\mathrm{ee}}}(\mathcal{A})$$

Now if the guess of π_i^* and π_r^* is correct, then either $\pi^* = \pi_i^*$ or $\pi^* = \pi_r^*$.

GAME $\mathcal{G}_3^{\text{ee}}$: We introduce an abort event $\operatorname{QueryDH}_{ee}$ in \mathcal{G}_3^{ee} : Specifically, if $(A, B, \overline{S, Y^*, X^*}) = \pi^*.ctxt$, and

-
$$\mathcal{A}$$
 queries $\mathsf{KDF}_f(DH, A, B, S, Y^*, X^*)$, where $d = h_0(A, X^*), e_1 = h_1(B, Y^*, S), e_2 = h_2(B, Y^*, S), \mathsf{DDH}(X^*A^d, Y^*B^{e_1}S^{e_2}, DH) = 1$,

then the game aborts (i.e., return 0 as \mathcal{A} 's bit guess). In other words, we abort the game if \mathcal{A} queries KDF_f on DH value that contains the information of $(Y^*)^{x^*} (= (X^*)^{y^*})$. We construct a CRGapDH adversary \mathcal{B}_{ee} such that if such an abort event happens, then \mathcal{B}_{ee} can output a solution to its CRGapDH challenge.

Reduction $\mathcal{B}_{ee}^{\text{DDH,CH}}(X^*, S^*)$:

- Initialize \mathcal{G}_3^{ee} :

- 1. Choose a B^{*} uniformly at random. Pick ssid^{*} \leftarrow s_{ss} .
- 2. Initialize counters $\operatorname{cnt}_1 := 0$ and $\operatorname{cnt}_2 := 0$ and choose $\operatorname{cnt}_1^*, \operatorname{cnt}_2^* \leftarrow [n_s]$. Initialize $\pi_i^* := \bot$ and Initialize $\pi_j^* := \bot$ and set up them later.
- 3. Initialize lists $\mathcal{L}_{patch-f}$, \mathcal{L}_{skey-f} , $\mathcal{L}_{patch-r}$, and \mathcal{L}_{skey-r} .
- 4. For all parties, follow the KGenLT and KGenSS defined in XHMQV to generate *ltpk*, *sspk*, and *ltsk*.
- 5. Run $b_{guess} \leftarrow \mathcal{A}(ltpk, sspk)$ and proceeds the game as in \mathcal{G}_{KE}^{KI} . If π^* and the partner session of π^* are not π_i^* and π_j^* , then overwrite b_{guess} with 0. The random oracles and game oracles are simulated as described below.

[–] Simulation of oracles defined in $\mathcal{G}_{\mathsf{KE}}^{\mathsf{KI}}$:

- 52 Rune Fiedler, Felix Günther, Jiaxin Pan, and Runzhi Zeng
 - All oracles except SEND and REVEALRAND are simulated in the same way defined in $\mathcal{G}_{\mathsf{KE}}^{\mathsf{KI}}.$
 - If REVEALRAND(U, j) is queried, where $\pi_j^U = \pi_i^* \neq \bot$ or $\pi_j^U = \pi_r^* \neq \bot$, then abort the game and overwrite \mathcal{A} 's bit with 0. REVEALRAND queries on other sessions are simulated as usual.
 - SEND (\mathbf{U}, j, m) :
 - * Case m = (V, sspk, epk): $cnt_1 := cnt_1 + 1$. If $cnt_1 = cnt_1^*$, then \mathcal{B}_{ee} uses X^* as the ephemeral pk of this session and set $\pi_i^* := \pi_j^{U}$. Otherwise, simulate as usual.
 - * Case m = (create, ssid, type): $\operatorname{cnt}_2 := \operatorname{cnt}_2 + 1$. If $\operatorname{cnt}_2 = \operatorname{cnt}_2^*$, then \mathcal{B}_{ee} uses Y^* as the ephemeral pk of this session and set $\pi_r^* := \pi_j^{U}$. \mathcal{B}_{ee} further sets $h_2(U, Y^*, S) := h^{-1}$ where U is the long-term public key of $U, h \leftarrow \operatorname{CH}(S), S$ is the semi-static pk of π_j^{U} . Otherwise, simulate as usual.
 - * For other cases, simulate as usual.
 - For all other SEND queries, we apply the same approach used in \mathcal{B}_{ltlt} to simulate the session key. Namely, we use the lists $\mathcal{L}_{patch-f}$, $\mathcal{L}_{patch-r}$, \mathcal{L}_{skey-f} , and \mathcal{L}_{skey-r} and oralce patching of KDF_f and KDF_r (using DDH) to simulate the session keys and the RO outputs of KDF_f and KDF_r consistently.
- Simulation of Random Oracles:
 - For fresh RO queries to $\mathsf{KDF}_f(DH, A, B, S, Y, X)$ that do not involve (X^*, Y^*) , we simulate in the same way as in \mathbf{G}_2 .
 - h_0 , h_1 , h_2 , h_3 , and KDF_r : Simulate as usual.
 - $KDF_f(DH, A, B, S, Y^*, X^*)$:
 - 1. If $(A, B, S, Y^*, X^*) = \pi^* . ctxt$ and DDH $(X^*A^d, Y^*B^{e_1}S^{e_2}, DH)$ outputs 1, then extract the value

$$DH' := \begin{cases} DH/(Y^*B^{e_1}S^{e_2})^{ad} \cdot (X^*)^{e_1b+e_2s}, & \pi^* \text{ is a responder session} \\ \left(DH/(Y^*B^{e_1}S^{e_2})^{ad} \cdot (X^*)^{e_1b}\right)^{e_2^{-1}}, & \pi^* \text{ is an initiator session} \end{cases}$$

where $d := h_0(A, X^*), e_1 := h_1(B, Y^*, S)$, and $e_2 := h_2(B, Y^*, S)$. We have a and b since we have long-term keys of all parties. If π^* is an initiator session, then \mathcal{B}_{ee} aborts the simulation and return DH' as the CRGapDH solution. If π^* is a responder session (and we have the semi-static key s), then we have that $DH' = (X^*)^{y^*}$, and \mathcal{B}_{ee} can use a similar approach in \mathcal{B}_{ee} to construct a valid CRGapDH solution.

2. Otherwise, simulate by lazy sampling.

If π^* is an initiator session with full mode, then we have $DH' = (Y^*(S)^{e_2})^{x^*e_2^{-1}}$, where $e_2 = h_2(B, Y^*, S) = h^{-1}$ and h is from CH(S). So, we also have

$$DH' = (Y^*(S)^{e_2})^{x^*e_2^{-1}} = (Y^*(S)^{h^{-1}})^{x^*h} = (S(Y^*)^h)^{x^*}$$

Therefore, DH' is a valid solution in this case, and we have

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}^{\mathsf{ee}}_3}(\mathcal{A}) \leq \mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}^{\mathsf{ee}}_4}(\mathcal{A}) + \mathsf{Adv}_{(\mathbb{G},p,g)}^{\mathsf{CRGapDH}}(\mathcal{B}_{\mathsf{ess}}) + Q_{\mathrm{RO}}/p$$

The remain argument to bound $\Pr[\mathsf{QueryDH}_{ee}]$ is the same as in previous cases. We have $\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathcal{G}_3^{ee}}(\mathcal{A}) = 0$. By assuming ϵ -CRGapDH holds over (\mathbb{G}, p, g) and combining all the probability bounds, we have the final bound of the case clean_{ESS}:

$$\mathsf{Adv}_{\mathsf{XHMQV}}^{\mathbf{G}_2 \wedge \mathsf{clean}_{\mathsf{EE}}}(\mathcal{A}) \leq n_s^2 \cdot (\epsilon + Q_{\mathrm{RO}}/p)$$

B Deferred Material on Deniability

We recall the formal definition for deniability of key exchange from Fiedler and Janson [24]. The model includes user messages, i.e., "Hey Bob, what's up?", as μ . We leave the user messages μ empty for clarity and since the subsequent AEAD encryption of user messages is not affected by the changes proposed in our protocol XHMQV. We refer the reader to [24, Section 3] for a full explanation of the model.

Definition 5 (Deniability for Key Exchange [24]). We say that a key exchange protocol $\mathsf{KE} = (\mathsf{KGenLT}, \mathsf{KGenSS}, \mathsf{Run})$ is $(n_p, n_{ss}, q_{\mathsf{O}_{\mathsf{A}}}, q_{\mathsf{O}_{\mathsf{P}}}, q_{\mathsf{O}_{\mathsf{D}}}, t_{\mathcal{A}}, t_{\mathsf{Dist}}, \epsilon_{\mathsf{Dist}})$ -deniable wrt. $(\mathsf{O}_A, \mathsf{O}_F, \mathsf{O}_D)$ -oracles, where $\mathsf{O}_A \subseteq \{\mathsf{REG}, \mathsf{INIT}, \mathsf{CHALLINIT}, \mathsf{CHALLRESP}, \mathsf{REGHON}, \mathsf{CHALLHONINIT}, \mathsf{CHALLHONRESP}\}, \mathsf{O}_F \subseteq \{\mathsf{SK}, \mathsf{USER}_{x,y}\}, \mathsf{O}_D \subseteq \{\mathsf{SKS}, \mathsf{AUX}\}, and auxiliary inputs sampled with AuxPrep, if for any adversary <math>\mathcal{A}$ running in time $t_{\mathcal{A}}$, making queries to O_A limited by q_{O_A} , there exists an algorithm Fake making queries to O_F limited by q_{O_F} such that for any distinguisher Dist making queries to O_D limited by q_{O_D} and running in time t_{Dist} , it holds that

$$\Pr[\mathcal{G}_{\mathsf{KE},\mathsf{AuxPrep},n_p,n_{ss}}^{\mathcal{O}_A,\mathcal{O}_F,\mathcal{O}_D}\text{-}den}(\mathcal{A},\mathsf{Dist})=1] \leq \frac{1}{2} + \epsilon_{\mathsf{Dist}},$$

where $\mathcal{G}_{\mathsf{KE},\mathsf{AuxPrep},n_p,n_{ss}}^{\mathcal{O}_A,\mathcal{O}_F,\mathcal{O}_D}$ -den (A, Dist) is defined in Figures 11 and 12.

B.1 Deferred proofs of deniability for XHMQV

Here, we give the deferred proofs for deniability, beginning with initiator (Bob) deniability. Note that the model of [24] considers the party who sends the first message as initiator, in contrast to the setting of Signal, where we usually consider Alice as initiator.

Proof (of Theorem 2). We give the Fake algorithm in Figure 13.

The Fake algorithm simulates Bob's pre-key bundle by sampling a fresh ephemeral key (if needed) and learns a signature on the semi-static key from the auxiliary info *aux*. Furthermore, to process Alice's message, the Fake algorithm

 $\mathcal{G}^{\mathcal{O}_{A},\mathcal{O}_{F},\mathcal{O}_{D}\text{-}den}_{\mathsf{KE},\mathsf{AuxPrep},n_{p},n_{ss}}(\mathcal{A},\mathsf{Dist})\text{:}$ $1 \ Q[\cdot] := []; \quad K[\cdot] := \bot; \quad C := \emptyset$ 2 (**pk**, **sk**) \leftarrow \$ KeyPrep (n_p, n_{ss}) $\exists aux \leftarrow AuxPrep(\mathbf{pk}, \mathbf{sk})$ 4 $b \leftarrow \{0, 1\}$ 5 $r \leftarrow \$ $\mathcal{R}_{\mathcal{A}}$ // randomness for the adversary 6 $\mathcal{A}^{\mathcal{O}_A}(\mathbf{pk}, aux; r)$ 7 $b' \leftarrow \text{sDist}^{\mathcal{O}_D}(\mathbf{pk}, r, Q, K)$ 8 return $\llbracket b' = b \rrbracket$ $KeyPrep(n_p, n_{ss})$: 9 for $U \in [n_p]$ // number of keys to prepare 10 $(ltpk_{U}, ltsk_{U}) \leftarrow \mathsf{KGenLT}()$ for ssid $\in [n_{ss}]$ 11 $(sspk_{U}^{ssid}, sssk_{U}^{ssid}) \leftarrow \mathsf{KGenSS}()$ 12 13 $ltpk := \{ltpk_U\}_{U \in [n_p]}$ 14 $sspk := \{sspk_U^{ssid}\}_{U \in [n_p]}^{ssid \in [n_{ss}]}$ 15 $ltsk := \{ltsk_U\}_{U \in [n_p]}$ 16 $sssk := \{sssk_U^{ssid}\}_{U \in [n_p]}^{ssid \in [n_{ss}]}$ 17 return ((ltpk, sspk), (ltsk, sssk))INIT (U, s, role, V) : 18 if $\pi_U^s = \bot \land U \notin C$ π^s_U .role := role 19 $\pi^s_U.{
m oid}:=U$ // set owner identity 20 $\pi^s_U.\mathsf{pid} := V$ 21 // set partner identity (* if post-specified peers) return success 22 23 return \perp USER_{x,y} (π, u, m, μ) : 24 $U := \pi.oid$ 25 $H := H^u_{\pi}$ // easier notation for u^{th} honest dummy for session π 26 if $\pi_U^H = \bot$ $(\mathsf{pk}_H,\mathsf{sk}_H) \gets \mathsf{KeyPrep}(1,1)$ 27 $\mathbf{pk} := \mathbf{pk} \cup \{\mathbf{pk}_H\}$ 28 if π .role = initiator 29 π_U^H .role := x30 else 31 π_U^H .role := y32 π_U^H .oid := U 33 π_U^H .pid := H 34 35 $(\pi_U^H, m', \mu') \leftarrow \text{\mathbb{s}} \operatorname{Run}(\operatorname{sk}_U, \operatorname{pk}, \pi_U^H, m, \mu)$ 36 return m'

Fig. 11. The deniability game for key exchange as defined in Definition 5. The challenge oracles are in Figure 12. The oracles are split into the adversary's capabilities $O_A \subseteq \{\text{Reg}, \text{INIT}, \text{CHALLINIT}, \text{CHALLRESP}, \text{RegHon}, \text{CHALLHONINIT}, \text{CHALLHONRESP}\}$, the capabilities of the **Fake** algorithm $O_F \subseteq \{\text{SK}, \text{USER}_{x,y}\}$, and the capabilities of the distinguisher Dist $O_D \subseteq \{\text{SKs}, \text{AUX}\}$.

 $\operatorname{RegHon}(U)$: 37 if $U \in [n_p] \cup C$ // abort if another key is already known for Ureturn \perp 38 39 $(\mathbf{pk}_U, \mathbf{sk}_U) \leftarrow \mathsf{KeyPrep}(1, n_{ss})$ 40 pk $\xleftarrow{+}$ pk_U; sk $\xleftarrow{+}$ sk_U 41 $C := C \cup \{U\}$ 42 return $(\mathbf{pk}_U, \mathbf{sk}_U)$ // simulates $\mathcal A$ honestly generating keys $\operatorname{Reg}(U, \mathbf{pk}_U)$: 43 if $U \in [n_p] \cup C$ // abort if another key is already known for \boldsymbol{U} ${f return} \perp$ 44 45 $C := C \cup \{U\}$ 46 $\mathbf{pk} \xleftarrow{+} \mathbf{pk}_U$ // includes semi-static keys 47 return success $SK(\pi)$: 48 return $sk_{\pi pid}$ SKs():49 return sk AUX():50 return aux

CHALLINIT (U, s, m, μ) CHALLRESP (U, s, m, μ) :

51 if $\pi^s_U = \perp$ //initialized session only $\vee \pi^s_U$.role = responder initiator return \perp 52 53 **if** b = 0 $(\pi_U^s, m') \leftarrow \text{sRun}(\mathsf{sk}_U, \mathsf{pk}, \pi_U^s, m, \mu)$ 54 55 **else** $(\pi_U^s, m') \leftarrow \mathsf{Fake}^{\mathcal{O}_F}(\mathsf{pk}, \pi_U^s, m, \mu, aux, r)$ 56 57 $Q[\pi_U^s] \xleftarrow{+} (m, m')$ 58 $K[\pi_U^s] \leftarrow \pi_U^s.\mathsf{K}$ 59 return m'CHALLHONINIT $(U, V, info_{create}, \mu, r_C)$: 60 if $U \in C$ // do not allow challenging a party under \mathcal{A} 's control 61 return \perp 62 π_U .oid $\leftarrow U$; π_{U} .pid $\leftarrow V$ 63 π_V .oid $\leftarrow V$; π_V .pid $\leftarrow U$ 64 π_U .role \leftarrow initiator; π_V .role \leftarrow responder 65 $m_1 \leftarrow (\text{create}, \text{info}_{\text{create}})$ 66 $(\mu_1,\ldots,\mu_{n_m}) \leftarrow \mu$ 67 for $i \in [1, 3, \ldots, n_m - 1]$ // until n_m if n_m is odd **if** b = 068 $(\pi_U, m_{i+1}) \leftarrow \text{sRun}(\text{sk}_U, \text{pk}, \pi_U, m_i, \mu_i)$ 69 else 70 $(\pi_{II}, m_{i+1}) \leftarrow \text{sFake}^{O_F}(\mathbf{pk}, \pi_{II}, m_i, \mu_i, aux, r_C)$ 71 $(\pi_V, m_{i+2}) \leftarrow \mathsf{Run}(\mathsf{sk}_V, \mathsf{pk}, \pi_V, m_{i+1}, \mu_{i+1}; r_C)$ 72 73 $Q[\pi_U] \leftarrow (m_i)_{i=1}^{n_m};$ $K[\pi_U] \leftarrow \pi_U.K$ 74 return $(Q[\pi_U], K[\pi_U])$ CHALLHONRESP $(U, V, info_{create}, \mu, r_C)$: 75 if $V \in C$ // do not allow challenging a party under \mathcal{A} 's control return \perp 76 77 π_U .oid $\leftarrow U$; π_U .pid $\leftarrow V$ $\pi_V.oid \leftarrow V; \quad \pi_V.pid \leftarrow U$ 78 79 π_U .role \leftarrow initiator; π_V .role \leftarrow responder 80 $m_1 \leftarrow (\text{create}, \text{info}_{\text{create}})$ 81 $(\mu_1,\ldots,\mu_{n_m}) \leftarrow \mu$

Solution (μ_1, \dots, μ_{m_m}) $\leftarrow \mu$ Solution (μ_1, \dots, μ_{m_m}) $\leftarrow \mu$ Solution (π_U, m_{i+1}) $\leftarrow \operatorname{Run}(\operatorname{sk}_U, \operatorname{pk}, \pi_U, m_i, \mu_i; r_C)$ Solution (π_V, m_{i+1}) $\leftarrow \operatorname{Run}(\operatorname{sk}_V, \operatorname{pk}, \pi_V, m_{i+1}, \mu_{i+1})$ Solution (π_V, m_{i+2}) $\leftarrow \operatorname{s} \operatorname{Run}(\operatorname{sk}_V, \operatorname{pk}, \pi_V, m_{i+1}, \mu_{i+1})$ Solution (π_V, m_{i+2}) $\leftarrow \operatorname{s} \operatorname{Fake}^{O_F}(\operatorname{pk}, \pi_V, m_{i+1}, \mu_{i+1}, aux, r_C)$ Solution (π_V, m_{i+2}) $\leftarrow \operatorname{s} \operatorname{Fake}^{O_F}(\operatorname{pk}, \pi_V, m_{i+1}, \mu_{i+1}, aux, r_C)$ Solution (π_V, m_{i+2}) $\leftarrow \operatorname{s} \operatorname{Fake}^{O_F}(\operatorname{pk}, \pi_V, m_{i+1}, \mu_{i+1}, aux, r_C)$ Solution (π_V, m_{i+2}) $\leftarrow \operatorname{Solution}(\pi_V, m_V) \leftarrow \pi_V.K$

89 return $(Q[\pi_V], K[\pi_V])$

Fig. 12. The challenge oracles for the deniability game in Definition 5, separate from Figure 11 for space reasons. The suffix of the oracle name indicates as which party the oracle acts. The infix HON indicates that this oracle enforces semi-honest behavior of the adversary.

	$Fake(\mathbf{pk}, \pi, m = (\mathbf{A}, X), \mu, aux, r_C):$
	11 π .pid := A, $A := ltpk[A]$
	12 $\left(a, \mathbf{s}, $
	13 $x := \mathrm{from}\; r_C$ // Alice's ephemeral secret key
$Fake(\mathbf{pk} = (ltpk, sspk), \pi, m, \mu, aux, r_C):$	14 $X := g^x$
(create, (ssid, type)) := m	15 $(B, vk_{B}) := ltpk[\pi.oid], S := \pi.sspk$
2 B := π .oid	16 $Y := \pi.cid[4]$ // Bob's ephemeral public key
$\pi.(pid,sspk,type) := (\star, sspk,type)$	17 $d := h_0(A, X)$
4 $sspk := sspk_{B}[ssid]$	18 if $Y \neq \bot$ // full handshake
5 $\sigma := \text{from } aux \text{ for } sspk$	19 $e_1 := h_1(B, Y, S)$
6 if type = full	20 $e_2 := h_2(B, Y, S)$
7 $y \leftarrow \mathbb{Z}_p, \ epk_{\mathtt{B}} := g^y$	21 $DH := (YB^{e_1}S^{e_2})^{x+da}$
8 else $epk_{\rm B} := \bot$	22 ctxt := (A, B, S, Y, X)
9 $m := (B, sspk, epk_{B}, [\sigma])$	23 $\pi.K := KDF_f(ctxt, DH)$
10 return (π, m, ε)	24 $else$ // reduced handshake
	25 $e_1 := h_3(B, S)$
	$26 \qquad DH := (B^{e_1}S)^{x+da}$
	27 $ctxt := (A, B, S, X)$
	28 $\pi.K := KDF_r(ctxt, DH)$
	29 return $(\pi, \varepsilon, \varepsilon)$

Fig. 13. The Fake algorithms simulating Bob's messages (initiator deniability), used in Theorem 2.

learns Alice's long-term secret key from the SK oracle and Alice's ephemeral secret key X from the randomness that was previously used to create Alice's message. Finally, the Fake algorithm computes the session key with Alice's ephemeral and long-term secret keys.

Since Fake produces a transcript and session key in the same way as Run, the distinguisher cannot have an advantage in winning its game, even if the distinguisher has access to all long-term and semi-static secret keys.

Next is responder (Alice) deniability.

Proof (of Theorem 3). We give the Fake algorithm in Figure 14.

The Fake algorithm verifies Bob's signature on the semi-static key. Next, it learns Bob's long-term and semi-static secret keys from the SK oracle and Bob's ephemeral secret key Y from the randomness that was previously used to create Bob's message. It samples a fresh ephemeral key X for Alice. Finally, the Fake algorithm computes the session key with Bob's ephemeral, semi-static, and long-term secret keys.

Since Fake produces a transcript and session key in the same way as Run, the distinguisher cannot have an advantage in winning its game, even if the distinguisher has access to all long-term and semi-static secret keys.

 $\mathsf{Fake}(\mathbf{pk} = (\mathbf{ltpk}, \mathbf{sspk}), \pi, m, \mu, aux, r_C):$

1 $(B, sspk, epk_B) := m, \quad \pi.pid := B,$ 2 $(B, vk_{B}) := ltpk[B]$ S := sspk, // prepare Bob's keys3 $\begin{pmatrix} b, \mathbf{s} \mathbf{k}_{\mathsf{B}}, \mathbf{s} \end{pmatrix}$:= extract from SK (π) 4 $y \leftarrow$ from r_C , $Y := epk_{\mathsf{B}}$ 5 **if** $Vf(B, \sigma, sspk) \neq 1$: **return** $(\pi, \varepsilon, \varepsilon)$ // verify signature 6 $\left[\mathbf{if} \ \nabla \mathbf{f}(\mathbf{vk}_{\mathsf{B}}, \sigma, sspk) \neq 1 : \mathbf{return} \ (\pi, \varepsilon, \varepsilon) \right]$ 7 $A := ltpk[\bar{\mathtt{A}}]$ // prepare Alice's keys $s \ x \leftarrow s \mathbb{Z}_p, \quad X := g^x$ 9 $d := h_0(A, X)$ 10 if $Y \neq \perp$ // full handshake $e_1 := h_1(B, Y, S)$ 11 $e_2 := h_2(B, Y, S)$ 12 $DH := (XA^d)^{y+e_1b+e_2s}$ 13 ctxt := (A, B, S, Y, X)14 $\pi.\mathsf{K} := \mathsf{KDF}_f(ctxt, DH)$ 15 16 else // reduced handshake $e_1 := h_3(\mathsf{B}, S)$ 17 $DH := (XA^d)^{e_1b+s}$ 18 ctxt := (A, B, S, X)19 $\pi.\mathsf{K} := \mathsf{KDF}_r(ctxt, DH)$ 20 21 return $(\pi, \varepsilon, \varepsilon)$

Fig. 14. The Fake algorithms simulating Alice's message (responder deniability), used in Theorem 3.