

Security Analysis of Signal’s PQXDH Handshake

Rune Fiedler¹

Felix Günther²

¹ Cryptoplexity, Technische Universität Darmstadt, Germany

² IBM Research Europe – Zurich, Switzerland

rune.fiedler@cryptoplexity.de, mail@felixguenther.info

Abstract. Signal recently deployed a new handshake protocol named PQXDH to protect against “harvest now, decrypt later” attacks of a future quantum computer. To this end, PQXDH adds a post-quantum KEM to the Diffie–Hellman combinations of the prior X3DH handshake.

In this work, we give a reductionist security analysis of Signal’s PQXDH handshake in a game-based security model that captures the targeted “maximum-exposure” security, allowing fine-grained compromise of user’s long-term, semi-static, and ephemeral key material. We augment prior such models to capture not only the added KEM component but also the signing of public keys, which prior analyses did not capture but which adds an additional flavor of post-quantum security in PQXDH. We then establish a fully parameterized, concrete security bound for the session key security of PQXDH, in particular shedding light on a KEM binding property we require for PQXDH’s security, and how to avoid it.

Our discussion of KEM binding complements the tool-based analysis of PQXDH by Bhargavan, Jacomme, Kiefer, and Schmidt, which pointed out a potential re-encapsulation attack if the KEM shared secret does not bind the public key. We show that both Kyber (used in PQXDH) and its current NIST draft standard ML-KEM (foreseen to replace Kyber once standardized) satisfy a novel binding notion we introduce and rely on for our PQXDH analysis, which may be of independent interest.

1 Introduction

Billions of people today use messaging apps such as Facebook Messenger, Google Messages, Skype, Signal, or WhatsApp, in which the Signal end-to-end encryption protocol [Sig] secures the communication. The Signal protocol consists of two main components: The initial handshake protocol, which allows two parties to derive a shared session key while authenticating each other, and the Double Ratchet protocol, which allows renewing the session key in an ongoing session to achieve forward secrecy and post-compromise security. Until 2023, Signal used X3DH [MP16] as initial handshake protocol. In 2023, Signal released PQXDH [KS23, KS24], which modifies X3DH to add protection against quantum adversaries; in particular against “harvest now, decrypt later” attacks where an adversary records communication today to break its confidentiality when a cryptographically-relevant quantum computer becomes available in the future.

In X3DH, each user has Diffie–Hellman (DH) keys of different lifetimes (long-term, semi-static, and ephemeral). Combining several DH keys of two users Alice and Bob in the initial key agreement ensures mutual authentication and, as long as one of the combinations of keys remains uncompromised, that the derived session key is secure. To protect against “harvest now, decrypt later” quantum adversaries, PQXDH [KS24] adds a quantum-safe key encapsulation mechanism (KEM) to X3DH while keeping the existing, well-understood [CCD⁺17] handshake structure unmodified. Furthermore, X3DH signs the involved semi-static DH key and PQXDH additionally the KEM key under the the user’s long-term key.

In a nutshell, the protocol message flow is as follows (where the KEM appears only in PQXDH): Bob sends his long-term, semi-static, and ephemeral DH shares, as well as an ephemeral KEM public key for a full handshake; a reduced handshake omits the ephemeral keys and has a semi-static KEM public key instead. Alice generates an ephemeral DH key pair, encapsulates against Bob’s KEM key, and sends her ephemeral DH key and KEM ciphertext to Bob. Both parties derive the session key from DH shared secrets from three or four DH key combinations (long-term/semi-static, ephemeral/long-term, ephemeral/semi-static, and, if present, ephemeral/ephemeral) and the KEM shared secret.

Previous Security Analyses of X3DH and PQXDH

The development of PQXDH was accompanied by formal, tool-based verification conducted by Bhargavan, Jacomme, Kiefer, and Schmidt (BJKS) [BJKS23, BJK23]. Their analysis provided improvements and security assurance for PQXDH (leading to revisions of the protocol description [KS23, KS24]), using both the symbolic-analysis tool ProVerif [BC] and the computational-analysis tool CryptoVerif [Bla]. Security of Signal’s original X3DH handshake (as well as its ratcheting protocol) was established by Cohn-Gordon, Cremers, Dowling, Garratt, and Stebila (CCDGS) [CCD+17] through a reductionist security analysis in what we refer to as a “maximum-exposure”¹, game-based security model. Vatandas, Gennaro, Ithurburn, and Krawczyk [VGIK20] analyzed the deniability of X3DH, and Fiedler and Janson [FJ24] the deniability of PQXDH.

The core difference between these prior analyses of PQXDH and X3DH is that while the tool-based analyses of BJKS provide machine-checked assurance, the ProVerif results remain on the symbolic, protocol-logic level and the CryptoVerif model does not account for all “maximum-exposure” attack vectors that Signal aims to protect against. Concretely, for proof complexity reasons, the CryptoVerif model of BJKS [BJKS23, BJK23] only captures the compromise of long-term user keys (but not of semi-static and ephemeral keys) and separately considers attacks against the classic DH and the post-quantum KEM security. Such compromise, however, is a main reason for the many key combination Signal’s X3DH and PQXDH handshake, and combined classic/post-quantum hybrid guarantees a main reason behind the PQXDH design specifically.

Our Contributions

This work completes the analysis picture for PQXDH by providing a reductionist security analysis of PQXDH (Revision 3 [KS24], throughout this paper) in a “maximum-exposure”, game-based security model following that for X3DH by CCDGS [CCD+17].

Augmented game-based security model. To this end, we first augment the coded game-based security model of Brendel, Fiedler, Günther, Janson, and Stebila (BFGJS) [BFG+22, BFG+21] to capture the added KEM component in PQXDH. In particular, we model that the KEM key and the semi-static DH keys in PQXDH are signed (instead of assuming them to be authentically distributed like in the prior game-based models [CCD+17, BFG+22, BFG+21]). This requires careful adaptation of related corruption modeling, which adds both model and proof complexity but leads to a security model which more closely² captures the real-world deployment.

¹The adversary may compromise nearly all secrets as long as trivial wins are excluded.

²Security modeling still always requires trade-offs in abstraction. For example, like all prior analyses [CCD+17, BFG+22, BFG+21, BJKS23, BJK23], we assume dedicated signing keys whereas Signal’s implementation re-uses the long-term DH user keys for signing.

Concrete security bound for PQXDH. We then analyze the PQXDH handshake, establishing its security based on the Gap Diffie–Hellman (GapDH) assumption holding in the DH group, the KEM providing one-way security under chosen-ciphertext attacks (OW-CCA), a certain form of binding property (which we expand on below), and low key-collision probability and correctness errors of the KEM, the signature scheme being existentially unforgeable, and the key derivation function behaving like a random oracle. In particular, we confirm that combining DH and KEM keys yields a combined, *hybrid* security bound, compared to the separate analysis in BJKS [BJKS23]. Our resulting security theorem for PQXDH is fully parameterized, giving concrete security bounds for each component.

Design discussion and KEM binding. We conclude with a discussion of the PQXDH design, with insights that complement those by BKJS [BJKS23]. In particular, we provide a point of view on the key derivation approach taken by Signal’s X3DH and PQXDH, why it leads to requiring a certain *binding* property from the KEM in our analysis, and how this requirement could be easily avoided.

For background, BKJS in their analysis [BJKS23, BJK23] discovered the following potential “KEM re-encapsulation attack” on PQXDH: An adversary can learn a KEM shared secret by compromising the involved KEM secret key, and then re-encapsulating this shared secret under another, uncompromised KEM public key. In consequence, decapsulations under two distinct public keys yield the same shared secret, and two sessions with different KEM public keys compute the same session key, violating the protocol’s security goals.

The BKJS analysis excludes this attack by modeling that the KEM public key is included in the associated data when AEAD-encrypting messages, a suggestion that the PQXDH description however only follows optionally: it argues that Kyber, used in the Signal implementation, already prevents the re-encapsulation attack [KS24, Sections 3.3 and 4.12]. Our analysis covers Signal’s approach, providing fine-grained insight into the properties required of the KEM to ensure security of the initial key agreement, *without* relying on the AEAD. Concretely, we require that the derived shared secret *binds* the involved public key and ciphertext, in a setting where the adversary knows the involved keys *and the randomness* to generate them. We prove that both Kyber [SAB⁺] (currently deployed in PQXDH) and the current NIST draft of ML-KEM [NIS23] (provisioned to replace Kyber in PQXDH once standardized) satisfy this binding property.

We view these KEM binding results as being of independent interest: They complement a range of binding notions for KEMs defined in concurrent work by Cremers, Dax, and Medinger (CDM) [CDM23], yet their notion that would fit the PQXDH setting is *not* satisfied by ML-KEM [Sch24]. We discuss the framework of CDM in more detail in Section 3 and relate the binding notion we introduce and require in the PQXDH analysis to theirs. Binding properties also arise in other key-combiner settings, e.g., in multi-recipient KEMs [AHK⁺23] and X-Wing [BCD⁺24], where the shared secret shall bind the ciphertext.

Related work

Several fully post-quantum replacements for X3DH have been considered in the literature: Hashimoto, Katsumata, Kwiatkowski, and Prest [HKKP22] propose SC-DAKE, a generic construction based on KEMs and a ring signature. BFGJS [BFG⁺22] propose a similar construction called SPQR based on KEMs and a designated verifier signature. Dobson and Galbraith [DG22] lift X3DH to supersingular isogenies with their construction called SI-X3DH, which is however broken by the SIDH attack [CD23, MMP⁺23, Rob23]. Collins, Huguenin-Dumittan, Nguyen, Rolin, and Vaudenay [CHDN⁺24] propose K-Waay, based on adapting split KEMs [BFG⁺20]. In this work, we focus on Signal’s deployed PQXDH protocol.

$\mathcal{G}_{\text{SIG}}^{\text{EUF-CMA}}(\mathcal{A}):$ 1 $(pk, sk) \leftarrow_{\$} \text{SIG.KGen}()$ 2 $Q \leftarrow \emptyset$ 3 $(m^*, \sigma^*) \leftarrow_{\$} \mathcal{A}^{\text{SIGN}}(pk)$ 4 return $\llbracket \text{Vf}(pk, m^*, \sigma^* \wedge m^* \notin Q) \rrbracket$	$\text{SIGN}(m):$ 5 $Q \leftarrow Q \cup \{m\}$ 6 return $\text{SIG.Sig}(sk, m)$
--	---

Figure 1: EUF-CMA security for a signature scheme $\text{SIG} = (\text{KGen}, \text{Sig}, \text{Vf})$.

2 Preliminaries

2.1 Notation

We refer to the first element of a list \vec{l} with $\vec{l}[1]$. For $n \in \mathbb{N}$, $[n]$ denotes $\{1, \dots, n\}$. We write a deterministic or randomized algorithm A with input x and output y as $y \leftarrow A(x)$ resp. $y \leftarrow_{\$} A(x)$ (or also $A(x) \rightarrow y$ resp. $A(x) \text{ s} \rightarrow y$), for the latter with explicit randomness r we write $y \leftarrow B(x; r)$. We sample an element e uniformly at random from a set S with $e \leftarrow_{\$} S$. We write $\Pr[\mathcal{G}(\mathcal{A})]$ for the probability that a game \mathcal{G} involving an adversary \mathcal{A} outputs true.

2.2 Signatures

Definition 2.1 (Signature scheme). *A signature scheme $\text{SIG} = (\text{KGen}, \text{Sig}, \text{Vf})$ with associated message space \mathcal{M} is a triple of PPT algorithms:*

- $\text{KGen}() \text{ s} \rightarrow (pk, sk)$: *The probabilistic key generation outputs a public-key/secret-key pair (pk, sk) .*
- $\text{Sig}(sk, m) \text{ s} \rightarrow \sigma$: *The probabilistic signing takes as input a secret key sk and a message m and outputs a signature σ .*
- $\text{Vf}(pk, m, \sigma) \rightarrow d$: *The deterministic verification takes as input a public key pk , a message m , and a signature σ and outputs a decision $d \in \{0, 1\}$. If $d = 1$ we say the signature is valid.*

A signature scheme is correct, if for all $(pk, sk) \leftarrow_{\$} \text{KGen}()$ and for all $m \in \mathcal{M}$ it holds that $\Pr[\text{Vf}(pk, m, \text{Sig}(sk, m)) = 1]$.

Definition 2.2 (EUF-CMA Security of Signatures). *Let $\text{SIG} = (\text{KGen}, \text{Sig}, \text{Vf})$ be a signature scheme. We say that SIG is $(t, \epsilon, q_{\text{Sig}})$ -EUF-CMA-secure, if for any adversary \mathcal{A} against Game $\mathcal{G}_{\text{SIG}}^{\text{EUF-CMA}}(\mathcal{A})$ defined in Figure 1 with running time at most t and making at most q_{Sig} queries to the SIGN oracle, we have that*

$$\text{Adv}_{\text{SIG}}^{\text{EUF-CMA}}(\mathcal{A}) := \Pr[\mathcal{G}_{\text{SIG}}^{\text{EUF-CMA}}(\mathcal{A})] \leq \epsilon.$$

2.3 Hash functions

Definition 2.3 (Preimage resistance of hash functions). *Let $H: \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a (hash) function. We say that H is $(t, \epsilon_d, \mathcal{M})$ -preimage resistant for randomly sampled preimages (PR-d-secure) and (t, ϵ_r) -preimage resistant for randomly sampled images (PR-r-secure), respectively, if for any adversary \mathcal{A} with running time at most t we have that*

$$\begin{aligned} \text{Adv}_{H, \mathcal{M}}^{\text{PR-d}}(\mathcal{A}) &:= \Pr[d = H(m') \mid m \leftarrow_{\$} \mathcal{M}, d \leftarrow H(m), m' \leftarrow_{\$} \mathcal{A}(d)] \leq \epsilon_d, \quad \text{resp.} \\ \text{Adv}_H^{\text{PR-r}}(\mathcal{A}) &:= \Pr[d = H(m') \mid d \leftarrow_{\$} \{0, 1\}^n, m' \leftarrow_{\$} \mathcal{A}(d)] \leq \epsilon_r. \end{aligned}$$

$\mathcal{G}_{(\mathbb{G}, g, q)}^{\text{GDH}}(\mathcal{A}):$ 1 $a, b \leftarrow_{\$} \mathbb{Z}_q$ 2 $g^z \leftarrow_{\$} \mathcal{A}^{\text{DDH}}((\mathbb{G}, g, q), a, b)$ 3 return $\llbracket g^z = g^{ab} \rrbracket$	$\text{DDH}(g^x, g^y, g^z):$ 4 return $\llbracket \text{CDH}(g^x, g^y) = g^z \rrbracket$	$\mathcal{G}_{\text{KEM}}^{\text{OW-CCA}}(\mathcal{A}):$ 1 $(pk, sk) \leftarrow_{\$} \text{KGen}()$ 2 $(ct^*, ss^*) \leftarrow_{\$} \text{Enc}(pk)$ 3 $ss' \leftarrow_{\$} \mathcal{A}^{\text{DECAPS}}(pk, ct^*)$ 4 return $\llbracket ss' = ss^* \rrbracket$	$\text{DECAPS}(ct):$ 5 if $ct = ct^*$ 6 return \perp 7 else 8 return $\text{Dec}(sk, ct)$
--	--	--	---

Figure 2: Security game for the GapDH problem defined in Definition 2.6 (left) and for OW-CCA security of KEM = (KGen, Enc, Dec) defined in Definition 3.3 (right).

Definition 2.4 (Collision resistance of hash functions). *Let $H: \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a (hash) function. We say that H is (t, ϵ) -collision-resistant (CR-secure), if for any adversary \mathcal{A} with running time at most t we have that*

$$\text{Adv}_{\text{H}}^{\text{CR}}(\mathcal{A}) := \Pr[H(m) = H(m') \wedge m \neq m' \mid (m, m') \leftarrow_{\$} \mathcal{A}] \leq \epsilon_{\text{H}}.$$

Our analysis further uses that finding two input values that collide in the truncated (256-bit) output is hard. Prior work discusses the related (weak) near collision resistance [MvV97, PS14, JKN21], partial collision resistance [WY05, LT11, YCW14], and truncated collision resistance [JK18]. We model this as collision resistance of a hash function H truncated to the first 256 bits, which we denote by H_{256} .

SHA3. Bertoni, Daemen, Peters, and Van Assche [BDPV08] show that SHA3 is indistinguishable from a random oracle when the underlying permutation is modeled as a random permutation. Thus, finding a collision or preimage in SHA3 or SHA3₂₅₆ is as hard as finding it in a random oracle yielding digests of the same length.

2.4 Diffie–Hellman Key Exchange

Definition 2.5 (Diffie–Hellman Key Exchange). *A Diffie–Hellman Key Exchange (DH) scheme is a tuple of algorithms (KGen, DH), defined as follows:*

- $\text{KGen}() \xrightarrow{\$} (pk, sk)$: *This probabilistic algorithm returns a key pair (pk, sk)*
- $\text{DH}(pk_A, sk_B) \rightarrow \text{DH}_{AB}$: *On input a public key pk_A and a secret key sk_B , this deterministic algorithm returns the shared DH secret DH_{AB} .*

We say that a DH key exchange DH is correct if, for every $(pk_A, sk_A), (pk_B, sk_B) \leftarrow_{\$} \text{KGen}()$, it holds that

$$\Pr[\text{DH}(pk_A, sk_B) = \text{DH}(pk_B, sk_A)] = 1.$$

For notational convenience we allow the arguments to be in arbitrary order, i.e., $\text{DH}(pk_A, sk_B) = \text{DH}(sk_A, pk_B)$. We write $\text{CDH}(pk_A, pk_B)$ to refer to the DH shared secret DH_{AB} when we do not care which secret key is used.

We rely on the GapDH problem [OP01, CCD⁺17], i.e., that the computational DH problem is hard given a decisional DH oracle.

Definition 2.6 (GapDH problem). *Let (\mathbb{G}, g, q) be a group of order q with generator g . We say that the $(t, \epsilon_{\text{GDH}}, q_{\text{DDH}})$ -GapDH problem holds in (\mathbb{G}, g, q) if for any adversary \mathcal{A} with running time at most t making at most q_{DDH} queries to its DDH oracle, we have that*

$$\text{Adv}_{(\mathbb{G}, g, q)}^{\text{GDH}}(\mathcal{A}) := \Pr[\mathcal{G}_{(\mathbb{G}, g, q)}^{\text{GDH}}(\mathcal{A})] \leq \epsilon_{\text{GDH}},$$

where $\mathcal{G}_{(\mathbb{G}, g, q)}^{\text{GDH}}(\mathcal{A})$ is defined in Figure 2.

3 Key Encapsulation Mechanisms and Binding Properties

We first recap the basic syntax, correctness, and security of key encapsulation mechanisms (KEMs), before discussing the novel KEM binding properties.

Definition 3.1 (Key Encapsulation Mechanisms). *A key encapsulation mechanism $\text{KEM} = (\text{KGen}, \text{Enc}, \text{Dec})$ consists of the following three algorithms:*

- $\text{KGen}() \xrightarrow{s} (pk, sk)$: *The probabilistic key generation with randomness space $\mathcal{R}_{\text{KEM.KGen}}$ outputs a public-key/secret-key pair with $(pk, sk) \in \mathcal{PK} \times \mathcal{SK}$.*
- $\text{Enc}(pk) \xrightarrow{s} (ct, ss)$: *The probabilistic encapsulation algorithm with randomness space $\mathcal{R}_{\text{KEM.Enc}}$ takes as input a public key $pk \in \mathcal{PK}$ and outputs a ciphertext $ct \in \mathcal{C}$ and the therein encapsulated shared secret $ss \in \mathcal{SS}$.*
- $\text{Dec}(sk, ct) \rightarrow ss'$: *The deterministic decapsulation algorithm takes as input a ciphertext $ct \in \mathcal{C}$ and secret key sk and outputs $ss' \in \mathcal{SS} \cup \{\perp\}$, where \perp indicates an error.*

We say that a KEM $\text{KEM} = (\text{KGen}, \text{Enc}, \text{Dec})$ is δ -correct if, for every key pair $(pk, sk) \leftarrow_s \text{KGen}()$, and every encapsulation $(ct, ss) \leftarrow_s \text{Enc}(pk)$, we have

$$\Pr[ss' \neq ss \mid ss' \leftarrow \text{Dec}(sk, ct)] \leq \delta.$$

In our analysis, we sometimes need to rule out that honestly generated KEM public keys collide. We capture that probability in the following. For Kyber [SAB⁺] and ML-KEM [NIS23], such public-key collisions boil down to random 256-bit seeds colliding under SHA3-512.

Definition 3.2 (KEM public-key collision probability). *We define the public-key collision probability of a KEM via a function $\gamma_{\text{coll}}: \mathbb{N} \rightarrow [0, 1]$, letting $\gamma_{\text{coll}}(n)$ denote the probability that two among n honestly generated public keys collide:*

$$\gamma_{\text{coll}}(n) := \Pr[pk_i = pk_j \wedge i \neq j \mid (pk_i, sk_i) \leftarrow_s \text{KGen}() \text{ for } i \in [1, n]].$$

Our analysis requires KEM *one-way security under chosen-ciphertext attacks* (OW-CCA), which is implied by standard IND-CCA security.

Definition 3.3 (OW-CCA Security of KEMs). *Let $\text{KEM} = (\text{KGen}, \text{Enc}, \text{Dec})$ be a KEM. We say that KEM is $(t, \epsilon, q_{\text{Dec}})$ -OW-CCA-secure, if for any adversary \mathcal{A} against Game $\mathcal{G}_{\text{KEM}}^{\text{IND-CCA}}(\mathcal{A})$ defined in Figure 2 with running time at most t and making at most q_{Dec} queries to the DECAPS oracle, we have that*

$$\text{Adv}_{\text{KEM}}^{\text{OW-CCA}}(\mathcal{A}) := \Pr[\mathcal{G}_{\text{KEM}}^{\text{OW-CCA}}(\mathcal{A})] \leq \epsilon.$$

3.1 Binding Properties

Due to the way PQXDH includes (only) the KEM shared secret in its key derivation (but not the KEM public key or ciphertext), our security analysis relies on a novel binding property of the KEM scheme. In a nutshell, we ask that it is hard to find two distinct ciphertexts or (honestly generated) public keys so that the corresponding decapsulations output the same shared secrets, even when given full control over the ciphertexts and knowing corresponding secret keys and the *randomness* used to generate these keys. The latter is a result of capturing the “maximum-exposure” security of PQXDH, allowing an adversary to reveal a session’s random coins (cf. Section 4). The PQXDH setting, in which an adversary interacts with many KEM keys, also motivates our notion being *multi-user*.

$$\mathcal{G}_{\text{KEM},n}^{\text{LEAK}^+\text{-BIND-SS-}\{CT,PK\}}(\mathcal{A}):$$

```

1 for  $k \in [n]$ 
2    $r_k \leftarrow \mathcal{R}_{\text{KEM.KGen}}$ 
3    $(pk_k, sk_k) \leftarrow \text{KEM.KGen}(\cdot; r_k)$ 
4    $(i, ct_i, j, ct_j) \leftarrow \mathcal{A}((pk_k, sk_k, r_k)_{k \in [n]})$ 
   //  $\mathcal{A}$  gets all key pairs and KGen randomness
5    $ss_i \leftarrow \text{KEM.Dec}(sk_i, ct_i)$ 
6    $ss_j \leftarrow \text{KEM.Dec}(sk_j, ct_j)$ 
7   if  $ss_i = \perp \vee ss_j = \perp$ : return false
8   return  $\llbracket ss_i = ss_j \wedge (ct_i \neq ct_j \vee pk_i \neq pk_j) \rrbracket$ 

```

$$\mathcal{G}_{\text{KEM}}^{X\text{-BIND-SS-}\{CT,PK\}}(\mathcal{A}):$$

```

11 if  $X = \text{LEAK}$ :
12    $(pk_0, sk_0) \leftarrow \text{KEM.KGen}()$ 
13    $(pk_1, sk_1) \leftarrow \text{KEM.KGen}()$ 
14    $(ct_0, ct_1) \leftarrow \mathcal{A}(pk_0, sk_0, pk_1, sk_1)$ 
15 if  $X = \text{MAL}$ :
16    $(pk_0, sk_0, pk_1, sk_1, ct_0, ct_1) \leftarrow \mathcal{A}()$ 
17    $ss_0 \leftarrow \text{KEM.Dec}(sk_0, ct_0)$ 
18    $ss_1 \leftarrow \text{KEM.Dec}(sk_1, ct_1)$ 
19   if  $ss_0 = \perp \vee ss_1 = \perp$ : return false
20   return  $\llbracket ss_0 = ss_1 \wedge (ct_0 \neq ct_1 \vee pk_0 \neq pk_1) \rrbracket$ 

```

Figure 3: Security games for our (multi-user) $\text{LEAK}^+\text{-BIND-SS-}\{CT,PK\}$ notion (left) and the $\text{LEAK-BIND-SS-}\{CT,PK\}$ and $\text{MAL-BIND-SS-}\{CT,PK\}$ notion from [CDM23] (right). In all games, the adversary’s goal is to produce colliding shared secrets ss under distinct ciphertexts ct or public keys pk .

Aligning with language of Cremers, Dax, and Medinger (CDM) [CDM23] from their concurrent work systematizing *binding* notions for KEMs, our notion asks that the KEM shared secret binds both public key and ciphertext used to produce it, under leakage of key generation randomness (and hence secret keys), which is a novel variant we denote as $\text{LEAK}^+\text{-BIND-SS-}\{CT,PK\}$. In general, their notion $X\text{-BIND-}P\text{-}Q$ captures that the components in set $P \in \{\{ss\}, \{ct\}, \{ss, ct\}\}$ bind the components in set $Q \in \{\{pk\}, \{ss\}, \{ct\}\}$, where the KEM keys are chosen by the adversary ($X = \text{MAL}$), honestly generated and leaked to the adversary ($X = \text{LEAK}$), or honestly generated without leakage of secrets ($X = \text{HON}$). We reproduce their LEAK and MAL notions for ss simultaneously binding pk and ct in our syntax³ on the right-hand side of Figure 3.

In the following, we formalize our new $\text{LEAK}^+\text{-BIND-SS-}\{CT,PK\}$ binding property and relate it to the notions in the CDM framework [CDM23]. In particular, we show in Section 3.3 below that both Kyber and the current NIST draft standard of ML-KEM satisfy the $\text{LEAK}^+\text{-BIND-SS-}\{CT,PK\}$ property, while the notion from [CDM23] that implies ours, $\text{MAL-BIND-SS-}\{CT,PK\}$, is not satisfied by ML-KEM [Sch24].⁴ While Signal’s current implementation of PQXDH uses Kyber [SAB⁺], it is already prepared⁵ to transition to the NIST standard ML-KEM [NIS23].

Definition 3.4 ($\text{LEAK}^+\text{-BIND-SS-}\{CT,PK\}$). *Let $\text{KEM} = (\text{KGen}, \text{Enc}, \text{Dec})$ be a KEM and $n \geq 2$. We say that KEM is $(t, \epsilon, n)\text{-LEAK}^+\text{-BIND-SS-}\{CT,PK\}$ -secure, if for any adversary \mathcal{A} against Game $\mathcal{G}_{\text{KEM},n}^{\text{LEAK}^+\text{-BIND-SS-}\{CT,PK\}}(\mathcal{A})$ defined in Figure 3 with running time at most t , we have that*

$$\text{Adv}_{\text{KEM},n}^{\text{LEAK}^+\text{-BIND-SS-}\{CT,PK\}}(\mathcal{A}) := \Pr \left[\mathcal{G}_{\text{KEM},n}^{\text{LEAK}^+\text{-BIND-SS-}\{CT,PK\}}(\mathcal{A}) \right] \leq \epsilon.$$

Figure 4 visualizes the relations between $\text{LEAK}^+\text{-BIND-SS-}\{CT,PK\}$ and the corresponding notions from [CDM23]: The MAL variant is strictly stronger than the LEAK^+ variant, which in turn is strictly stronger than the LEAK variant. Notably, ML-KEM separates the MAL and our LEAK^+ variants; we show the other implications and separations below, incl. relating our multi-user notion to the 2-user case.

³Note that in the syntax of [CDM23] the KEMs shared secret ss is simply called a “key”, denoted k . Accordingly, in their notation, $P \in \{\{k\}, \{ct\}, \{k, ct\}\}$ and $Q \in \{\{pk\}, \{k\}, \{ct\}\}$. We will stick to our syntax and hence, e.g., write $\text{MAL-BIND-SS-}CT$ corresponding to their notion $\text{MAL-BIND-}K\text{-}CT$.

⁴CDM [CDM23, Appendices B and E] conjecture that Kyber is $\text{MAL-BIND-SS-}\{CT,PK\}$ -secure as it hashes both pk and ct into the key derivation. Their generic argument [CDM23, Thm. E.1] requires including both via an injective function, a hash function as used in Kyber however is clearly not injective. We leave it as an open question whether Kyber achieves $\text{MAL-BIND-SS-}\{CT,PK\}$ security; for us, the LEAK^+ variant suffices.

⁵<https://github.com/signalapp/libsignal/commit/0670f0d>

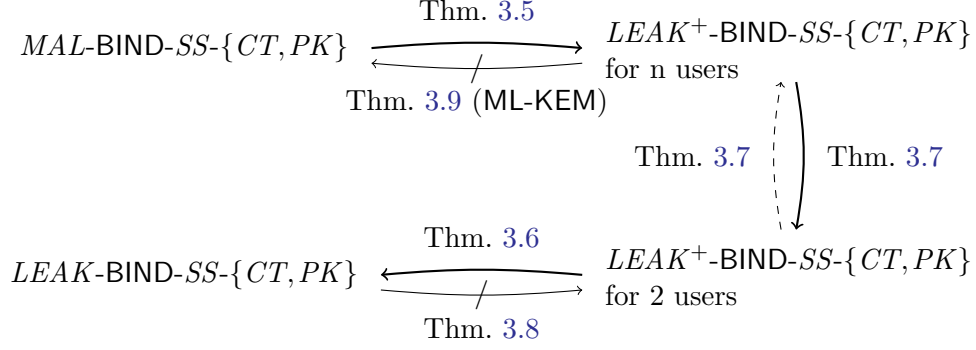


Figure 4: Relations between our $LEAK^+$ -BIND-SS- $\{CT, PK\}$ binding notion and the corresponding MAL , $LEAK$ notions from [CDM23] for KEMs. Solid arrows indicate implications, dashed ones loose implications, and crossed-out ones separations. Annotations indicate the respective theorems.

3.2 Relations of Binding Properties

First, we show the implication from MAL to $LEAK^+$ (Theorem 3.5) and from $LEAK^+$ to $LEAK$ (Theorem 3.6). Second, we show that our notion is equivalent for n and 2 users, up to a factor of n^2 (Theorem 3.7). Third, we give a separation between $LEAK$ and $LEAK^+$ (Theorem 3.8). Lastly, we show that ML-KEM and Kyber are $LEAK^+$ (Theorem 3.9 and Theorem 3.10), while ML-KEM is known to not be MAL [Sch24], giving us a separation for $LEAK^+$ and MAL .

Note that for the following theorems we assume KEMs with a public key space $|\mathcal{PK}| \geq 2$.

Theorem 3.5 ($MAL\text{-}BIND\text{-}SS\text{-}\{CT, PK\} \implies LEAK^+\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$). *A $MAL\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure KEM KEM is also $LEAK^+\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure. Concretely, for any $LEAK^+\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ adversary \mathcal{A} against KEM there exist an adversary \mathcal{B} such that*

$$\text{Adv}_{KEM, n}^{LEAK^+\text{-}BIND\text{-}SS\text{-}\{CT, PK\}}(\mathcal{A}) \leq \text{Adv}_{KEM}^{MAL\text{-}BIND\text{-}SS\text{-}\{CT, PK\}}(\mathcal{B}).$$

Proof. To win in the $LEAK^+\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ game, \mathcal{A} must produce ciphertexts ct_i, ct_j such that $ss_i = ss_j$ and one (or both) of $pk_i \neq pk_j$ and $ct_i \neq ct_j$ hold.

We let \mathcal{B} truthfully simulate $\mathcal{G}_{KEM, n}^{LEAK^+\text{-}BIND\text{-}SS\text{-}\{CT, PK\}}(\mathcal{A})$ for \mathcal{A} , i.e., honestly computing keys pk_k for $k \in [n]$, obtaining i, ct_i, j, ct_j from \mathcal{A} . Then, \mathcal{B} outputs $(pk_i, sk_i, pk_j, sk_j, ct_i, ct_j)$. If \mathcal{A} wins, we have that $ss_i = ss_j$ while $pk_i \neq pk_j$ or $ct_i \neq ct_j$, so \mathcal{B} also wins in the $MAL\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ game. \square

Theorem 3.6 ($LEAK^+\text{-}BIND\text{-}SS\text{-}\{CT, PK\} \implies LEAK\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$). *A $LEAK^+\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure KEM KEM is also $LEAK\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure. Concretely, for any $LEAK\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ adversary \mathcal{A} against KEM there exist an adversary \mathcal{B} such that*

$$\text{Adv}_{KEM}^{LEAK\text{-}BIND\text{-}SS\text{-}\{CT, PK\}}(\mathcal{A}) \leq \text{Adv}_{KEM, 2}^{LEAK^+\text{-}BIND\text{-}SS\text{-}\{CT, PK\}}(\mathcal{B}).$$

Proof. We let \mathcal{B} start \mathcal{A} on its own inputs except the key generation randomness, thereby truthfully simulating the $LEAK\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ game for \mathcal{A} . When \mathcal{A} terminates, \mathcal{B} returns the output of \mathcal{A} . If \mathcal{A} wins, then so does \mathcal{B} . \square

Theorem 3.7 ($LEAK^+\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$: n users \iff 2 users). *Let $n \geq 2$. A (t, ϵ, n) - $LEAK^+\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure KEM KEM is also $(t, \epsilon, 2)$ - $LEAK^+\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure and a $(t, \epsilon, 2)$ - $LEAK^+\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure KEM KEM is also $(t', n^2 \cdot \epsilon, n)$ - $LEAK^+\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure*

(for $t' \approx t$). Concretely, for any $LEAK^+$ -BIND-SS- $\{CT, PK\}$ adversary \mathcal{A} against KEM there exist an adversary \mathcal{B} such that

$$\begin{aligned} \text{Adv}_{\text{KEM},2}^{LEAK^+-\text{BIND-SS-}\{CT,PK\}}(\mathcal{A}) &\leq \text{Adv}_{\text{KEM},n}^{LEAK^+-\text{BIND-SS-}\{CT,PK\}}(\mathcal{A}) \quad \text{and} \\ \text{Adv}_{\text{KEM},n}^{LEAK^+-\text{BIND-SS-}\{CT,PK\}}(\mathcal{A}) &\leq n^2 \cdot \text{Adv}_{\text{KEM},2}^{LEAK^+-\text{BIND-SS-}\{CT,PK\}}(\mathcal{B}). \end{aligned}$$

Proof. The “ \implies ” direction straightforwardly holds by letting \mathcal{A} ignore the remaining $n - 2$ keys.

The “ \impliedby ” direction holds via a guessing argument: Let \mathcal{B} guess two distinct key indices $i, j \in [n]$, embed the two keys obtained in its game in these positions, and sample the remaining $n - 2$ keys itself. If \mathcal{A} uses i (and possibly j) for its attack, \mathcal{B} is successful if \mathcal{A} is. The chance of \mathcal{B} guessing correctly is $\frac{1}{n^2}$, establishing the claim. \square

Theorem 3.8 ($LEAK$ -BIND-SS- $\{CT, PK\} \not\Rightarrow LEAK^+$ -BIND-SS- $\{CT, PK\}$). *Assuming a $LEAK$ -BIND-SS- $\{CT, PK\}$ -secure KEM KEM' and a preimage-resistant (PR-d) hash function H (when sampling from $\{0, 1\}^{256}$), there exists a KEM KEM which is $LEAK$ -BIND-SS- $\{CT, PK\}$ -secure but not $LEAK^+$ -BIND-SS- $\{CT, PK\}$ -secure. Concretely, for any $LEAK$ -BIND-SS- $\{CT, PK\}$ adversary \mathcal{A} against KEM' there exist adversaries $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ such that*

$$\begin{aligned} \text{Adv}_{\text{KEM},2}^{LEAK^+-\text{BIND-SS-}\{CT,PK\}}(\mathcal{B}_1) &= 1 \quad \text{and} \\ \text{Adv}_{\text{KEM}}^{LEAK\text{-BIND-SS-}\{CT,PK\}}(\mathcal{A}) &\leq \text{Adv}_H^{\text{PR-d}}(\mathcal{B}_2) + \text{Adv}_{\text{KEM}'}^{LEAK\text{-BIND-SS-}\{CT,PK\}}(\mathcal{B}_3). \end{aligned}$$

Proof. Let $KEM' = (KGen', Enc', Dec')$ with shared secret space \mathcal{SS}' and randomness space $\mathcal{R}_{KGen'}$ for key generation be $LEAK$ -BIND-SS- $\{CT, PK\}$ -secure, and consider the following modification KEM with shared secret space $\mathcal{SS}' \cup \{ss^*\}$ where ss^* denotes a distinguished shared secret with $ss^* \notin \mathcal{SS}'$ and randomness space $\{0, 1\}^{256} \times \mathcal{R}_{KGen'}$ for key generation.

$KGen(; r)$:	$Enc(pk)$:	$Dec(sk, pk, ct)$:
<ol style="list-style-type: none"> 1 $(x, r_{KG}) \leftarrow r$ 2 $y \leftarrow H(x)$ 3 $(pk', sk') \leftarrow KGen'(; r_{KG})$ 4 return $(pk', (sk', y))$ 	<ol style="list-style-type: none"> 5 return $(Enc'(pk), \perp)$ 	<ol style="list-style-type: none"> 6 $(sk', y) \leftarrow sk$ 7 $(ct', z) \leftarrow ct$ 8 if $y = H(z)$ 9 return ss^* 10 return $Dec'(sk', pk, ct')$

KEM is not $LEAK^+$ -BIND-SS- $\{CT, PK\}$ -secure: An adversary \mathcal{B}_1 gets $(pk_1, sk_1, r_1, pk_2, sk_2, r_2)$ as input, parses the randomness as $(x_1, r_{KG,1}) \leftarrow r_1$ and $(x_2, r_{KG,2}) \leftarrow r_2$, and outputs the two ciphertexts $(0, x_1), (1, x_2)$ (omitting the indices for the two users). Both decapsulations trigger the special case. Hence, the shared secrets collide, while the ciphertexts differ (at least) in their first component and \mathcal{B}_1 always wins.

KEM is $LEAK$ -BIND-SS- $\{CT, PK\}$ -secure: If \mathcal{A} triggers the special condition in the decapsulation, then it has found a preimage under H (for randomly sampled 256-bit preimages), breaking preimage resistance. Otherwise, the adversary cannot take advantage of the extra components in the secret key and ciphertext. Hence, breaking $LEAK$ -BIND-SS- $\{CT, PK\}$ of KEM implies breaking $LEAK$ -BIND-SS- $\{CT, PK\}$ of KEM' , which we have excluded by assumption. \square

KGen():	Enc(pk):	Dec(sk, ct):
1 $z \leftarrow_{\$} \{0, 1\}^{256}$ 2 $(pk, sk') \leftarrow_{\$} \text{PKE.KGen}()$ 3 $sk \leftarrow (sk', pk, H(pk), z)$ 4 return (pk, sk)	5 $m \leftarrow_{\$} \{0, 1\}^{256}$ 6 $m \leftarrow H(m)$ 7 $(ss', r) \leftarrow G(m H(pk))$ 8 $ct \leftarrow \text{PKE.Enc}(pk, m; r)$ 9 $ss \leftarrow \text{KDF}(ss' H(ct))$ 10 return (ct, ss)	11 $(sk', pk, h, z) \leftarrow sk$ 12 $m' \leftarrow \text{PKE.Dec}(sk', ct)$ 13 $(ss', r') \leftarrow G(m' h)$ 14 $ct \leftarrow \text{PKE.Enc}(pk, m'; r')$ 15 if $ct = ct'$: 16 $ss \leftarrow \text{KDF}(ss' H(ct))$ 17 else 18 $ss \leftarrow \text{KDF}(z H(ct))$ 19 return ss

Figure 5: Algorithmic description of ML-KEM and Kyber using functions H, G, and J resp. KDF and a public key encryption scheme PKE. Solid boxes are exclusive to Kyber, dashed boxes are exclusive to ML-KEM.

3.3 Kyber and ML-KEM

Figure 5 gives an algorithmic description of Kyber [SAB⁺] and its corresponding current draft NIST standard ML-KEM [NIS23]. On a high level, the scheme uses a public key encryption scheme PKE in an FO transform [FO99], as well as three functions H, G, and J (ML-KEM) resp. KDF (Kyber), instantiated as SHA3-256, SHA3-512, and SHAKE256 with 256 bits output, respectively. Our interest being in the binding properties, we focus here on the internals of decapsulation. The decapsulation algorithm first decrypts the PKE ciphertext ct to m' . Then, it hashes m' and h (the hashed public key, stored in the secret key) under G onto ss', r' . It uses r' as randomness for re-encrypting m' . If the re-encrypted ciphertext matches ct , Kyber outputs $\text{KDF}(ss' || H(ct))$ as shared secret while ML-KEM outputs ss' directly. Otherwise, perform an implicit rejection, i.e., compute the shared secret as $\text{KDF}(z || H(ct))$ for Kyber resp. $J(z || ct)$ for ML-KEM, where z is a secret random 256-bit string, which is part of the secret key.

In the following, we show that ML-KEM and Kyber satisfy the $LEAK^+$ -BIND-SS- $\{CT, PK\}$ binding property, assuming collision resistance of all involved hash functions, (only for ML-KEM) random-oracle properties of H and J, and (only for Kyber) preimage resistance of G_{256} for randomly sampled images.

Theorem 3.9 (ML-KEM is $LEAK^+$ -BIND-SS- $\{CT, PK\}$ -secure). *Assuming H is (t, ϵ_H) -collision-resistant, G_{256} is (t, ϵ_G) -collision-resistant, J is (t, ϵ_J) -collision-resistant, and modeling H, J as independent random oracles, ML-KEM is (t', ϵ, n) - $LEAK^+$ -BIND-SS- $\{CT, PK\}$ -secure for $t' \approx t$ and*

$$\epsilon = \frac{n^2}{2^{256}} + 2\epsilon_J + \frac{q_{RO}^2}{2^{256}} + \epsilon_H + 2\epsilon_G,$$

where q_{RO} is the number of random oracle queries made by the adversary.

Proof. We analyze the probability of \mathcal{A} winning over a series of game hops.

Game 0. We start with the original binding game $\mathcal{G}_{\text{ML-KEM}, n}^{LEAK^+-\text{BIND-SS-}\{CT, PK\}}(\mathcal{A})$:

$$\text{Adv}_{\text{ML-KEM}, n}^{LEAK^+-\text{BIND-SS-}\{CT, PK\}}(\mathcal{A}) = \text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_0}(\mathcal{A}).$$

We distinguish several cases for the adversary to win, depending on how the adversary fulfills the winning condition and if the decapsulation accepts (line 15 in Figure 5 evaluates to true) or implicitly rejects (i.e.,

decapsulation branches into line 17). We denote line 15 being taken with A_x and line 17 with R_x , where $x \in \{i, j\}$ are the key indices in the binding game. We index any intermediate values from decapsulation of ct_i and ct_j with i and j , respectively. We assume a successful adversary, i.e., $ss_i = ss_j$. Note that R_x and A_x are mutually exclusive. The cases are:

- A. Both decapsulations reject and the ciphertexts are for two distinct public keys, i.e., $R_i \wedge R_j \wedge pk_i \neq pk_j$.
- B. Both decapsulations reject and the ciphertexts are distinct, i.e., $R_i \wedge R_j \wedge ct_i \neq ct_j$.
- C. One decapsulation rejects, i.e., wlog. $R_i \wedge A_j$.
- D. Both decapsulations accept and the ciphertexts are for two distinct public keys, i.e., $A_i \wedge A_j \wedge pk_i \neq pk_j$.
- E. Both decapsulations accept and the ciphertexts are distinct, i.e., $A_i \wedge A_j \wedge pk_i = pk_j \wedge ct_i \neq ct_j$.

We treat these cases as events in \mathcal{G}_0 and indicate the occurrence of event X by $\mathcal{G}_0[X]$. By the union bound we get:

$$\text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_0}(\mathcal{A}) \leq \sum_{X \in \{A, B, C, D, E\}} \text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_0[X]}(\mathcal{A})$$

Case A (Both decapsulations reject, distinct public keys: $R_i \wedge R_j \wedge pk_i \neq pk_j$).

Here, we bound the probability of the adversary producing ciphertexts that both get implicitly rejected during decapsulation and decapsulate to the same shared secret under distinct public keys.

Game A.0. This is the game conditioned on $R_i \wedge R_j \wedge pk_i \neq pk_j$ being satisfied.

$$\text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_{A.0}}(\mathcal{A}) = \text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_0[R_i \wedge R_j \wedge pk_i \neq pk_j]}(\mathcal{A}).$$

Game A.1 (Colliding z values). We let $\mathcal{G}_{A.0}$ return false if there is a collision among the z values in the n secret keys. Since the z values are 256-bit strings sampled uniformly at random, by the birthday bound we get:

$$\text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_{A.0}}(\mathcal{A}) \leq \text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_{A.1}}(\mathcal{A}) + \frac{n^2}{2^{256}}.$$

Game A.2 (Collision in J). Now, we have established $z_i \| ct_i \neq z_j \| ct_j$ (since $z_i \neq z_j$), yet $J(z_i \| ct_i) = J(z_j \| ct_j)$. We build a reduction \mathcal{B}_1 that outputs $z_i \| ct_i, z_j \| ct_j$, breaking (t, ϵ_J) -collision resistance of J:

$$\text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_{A.1}}(\mathcal{A}) \leq \epsilon_J.$$

Case B (Both decapsulations reject, distinct ciphertexts: $R_i \wedge R_j \wedge ct_i \neq ct_j$).

Here, we bound the probability of the adversary producing two distinct ciphertexts that both get implicitly rejected during decapsulation and decapsulate to the same shared secret.

Game B.0. This is the game conditioned on $R_i \wedge R_j \wedge ct_i \neq ct_j$ being satisfied.

$$\text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_{B.0}}(\mathcal{A}) = \text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_0[R_i \wedge R_j \wedge ct_i \neq ct_j]}(\mathcal{A}).$$

Since $ct_i \neq ct_j$, we also have $z_i \| ct_i \neq z_j \| ct_j$, yet $J(z_i \| ct_i) = J(z_j \| ct_j)$. We build a reduction \mathcal{B}_2 that outputs $z_i \| ct_i, z_j \| ct_j$, breaking (t, ϵ_J) -collision resistance of J:

$$\text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_{B.0}}(\mathcal{A}) \leq \epsilon_J.$$

Case C (One decapsulation rejects: $R_i \wedge A_j$).

Here, we bound the probability of the adversary producing exactly one ciphertext that gets implicitly rejected during decapsulation while both ciphertexts decapsulate to the same shared secret.

Game C.0. This is the game conditioned on $R_i \wedge A_j$ being satisfied.

$$\text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_{\text{C.0}}}(\mathcal{A}) = \text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_0[R_i \wedge A_j]}(\mathcal{A}).$$

For \mathcal{A} to win, it needs to create a collision $J(z_i \| ct_i) = ss_i = ss_j = G_{256}(m'_j \| h_j)$ between J and G . Assuming both J and G behave like (independent) random oracles, the probability of finding such a collision with q_{RO} many random oracle queries is upper bounded by:

$$\text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_{\text{C.0}}}(\mathcal{A}) \leq \frac{q_{\text{RO}}^2}{2^{256}}.$$

Case D (Both decapsulations accept, distinct public keys: $A_i \wedge A_j \wedge pk_i \neq pk_j$).

Here, we bound the probability of the adversary producing ciphertexts that both get accepted during decapsulation and decapsulate to the same shared secret under distinct public keys.

Game D.0. This is the game conditioned on $A_i \wedge A_j \wedge pk_i \neq pk_j$ being satisfied.

$$\text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_{\text{D.0}}}(\mathcal{A}) = \text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_0[A_i \wedge A_j \wedge pk_i \neq pk_j]}(\mathcal{A}).$$

Since we know that $pk_i \neq pk_j$ and $G_{256}(m_i \| H(pk_i)) = ss_i = ss_j = G_{256}(m_j \| H(pk_j))$, it must be that \mathcal{A} provides us with a collision either in H or in G_{256} . Once more, we can distinguish two cases:

1. The first case is $m_i \| H(pk_i) = m_j \| H(pk_j)$. Then we can build a reduction \mathcal{B}_3 that outputs pk_i, pk_j as collision in H , breaking (t, ϵ_H) -collision resistance of H .
2. Otherwise, $m_i \| H(pk_i) \neq m_j \| H(pk_j)$, and we can build reduction \mathcal{B}_4 that outputs $m_i \| H(pk_i), m_j \| H(pk_j)$ as collision in G_{256} . Hence, \mathcal{B}_2 breaks the (t, ϵ_G) -collision resistance of G_{256} .

Jointly,

$$\text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_{\text{D.0}}}(\mathcal{A}) \leq \epsilon_H + \epsilon_G.$$

Case E (Both decapsulations accept, distinct ciphertexts: $A_i \wedge A_j \wedge pk_i = pk_j \wedge ct_i \neq ct_j$).

Here, we bound the probability of the adversary producing ciphertexts that both get accepted during decapsulation and decapsulate to the same shared secret under the same public key.

Game E.0. This is the game conditioned on $A_i \wedge A_j \wedge pk_i = pk_j \wedge ct_i \neq ct_j$ being satisfied.

$$\text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_{\text{E.0}}}(\mathcal{A}) = \text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_0[A_i \wedge A_j \wedge pk_i = pk_j \wedge ct_i \neq ct_j]}(\mathcal{A}).$$

We know that $pk_i = pk_j$ (and hence $h_i = h_j$) for $G_{256}(m_i \| h_i) = ss_i = ss_j = G_{256}(m_j \| h_j)$. Note that we must have $m_i \neq m_j$: Otherwise, the re-encryption step in both decapsulations would result in identical ciphertexts $ct_i = ct_j$, which contradicts the condition of event E .

Now, we can build a reduction \mathcal{B}_5 that outputs $m_i \| H(pk_i), m_j \| H(pk_j)$ as collision in G_{256} . Since G collides in its first output, which is 256 bits, \mathcal{B}_3 breaks the (t, ϵ_G) -collision resistance of G_{256} .

$$\text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_{\text{E.0}}}(\mathcal{A}) \leq \epsilon_G.$$

Collecting the bounds yields the claim. □

For the Kyber proof, we do not require any collisions *across* hash functions, allowing us to avoid relying on the random oracle. Instead, we additionally require G_{256} to be (t, ϵ'_G) -preimage-resistant for randomly sampled images.

Theorem 3.10 (Kyber is $LEAK^+$ -BIND-SS- $\{CT, PK\}$ -secure). *Assuming H is (t, ϵ_H) -collision-resistant, G_{256} is (t, ϵ_G) -collision-resistant and (t, ϵ'_G) -preimage-resistant (PR-r, i.e., when sampling from the range), and KDF is (t, ϵ_{KDF}) -collision-resistant, Kyber is (t', ϵ, n) - $LEAK^+$ -BIND-SS- $\{CT, PK\}$ -secure for $t' \approx t$ and*

$$\begin{aligned} \epsilon &= \epsilon_H + \epsilon_{KDF} + \frac{n^2}{2^{256}} + \epsilon_{KDF} + \epsilon_{KDF} + n\epsilon'_G + \epsilon_{KDF} + \epsilon_G + \epsilon_H \\ &= 2\epsilon_H + \epsilon_G + n\epsilon'_G + 4\epsilon_{KDF} + \frac{n^2}{2^{256}}. \end{aligned}$$

Proof. We analyze the probability of \mathcal{A} winning over a series of game hops.

Game 0. We start with the original binding game $\mathcal{G}_{\text{Kyber}, n}^{LEAK^+ \text{-BIND-SS-}\{CT, PK\}}(\mathcal{A})$:

$$\text{Adv}_{\text{Kyber}, n}^{LEAK^+ \text{-BIND-SS-}\{CT, PK\}}(\mathcal{A}) = \text{Adv}_{\text{Kyber}}^{\mathcal{G}_0}(\mathcal{A}).$$

We distinguish several cases for the adversary to win, depending on how the adversary fulfills the winning condition and if the decapsulation accepts (line 15 in Figure 5 evaluates to true) or implicitly rejects (i.e., decapsulation branches into line 17). We denote line 15 being taken with A_x and line 17 with R_x , where $x \in \{i, j\}$ are the key indices in the binding game. We index any intermediate values from decapsulation of ct_i and ct_j with i and j , respectively. We assume a successful adversary, i.e., $ss_i = ss_j$. Note that R_x and A_x are mutually exclusive. The cases are:

- A. Both ciphertexts are distinct, i.e., $ct_i \neq ct_j$.
- B. Both decapsulations reject and the ciphertexts are for two distinct public keys, i.e., $R_i \wedge R_j \wedge pk_i \neq pk_j$.
- C. One decapsulation rejects and the ciphertexts are for two distinct public keys, i.e., wlog. $R_i \wedge A_j \wedge pk_i \neq pk_j$.
- D. Both decapsulations accept and the ciphertexts are for two distinct public keys, i.e., $A_i \wedge A_j \wedge pk_i \neq pk_j$.

We treat these cases as events in \mathcal{G}_0 and indicate the occurrence of event X by $\mathcal{G}_0[X]$. By the union bound we get:

$$\text{Adv}_{\text{Kyber}}^{\mathcal{G}_0}(\mathcal{A}) \leq \sum_{X \in \{A, B, C, D\}} \text{Adv}_{\text{Kyber}}^{\mathcal{G}_0[X]}(\mathcal{A})$$

Case A (distinct ciphertexts: $ct_i \neq ct_j$).

Here, we bound the probability of the adversary producing two distinct ciphertexts that decapsulate to the same shared secret.

Game A.0. This is the game conditioned on $ct_i \neq ct_j$ being satisfied.

$$\text{Adv}_{\text{Kyber}}^{\mathcal{G}_{A,0}}(\mathcal{A}) = \text{Adv}_{\text{Kyber}}^{\mathcal{G}_0[ct_i \neq ct_j]}(\mathcal{A}).$$

We know that $\text{KDF}(x_i \| H(ct_i)) = ss_i = ss_j = \text{KDF}(x_j \| H(ct_j))$ for x_i being either ss'_i (in case of an accepting decapsulation) or z_i (in case of an implicitly rejecting decapsulation), and x_j likewise. Since we

know that $ct_i \neq ct_j$, we can now either build a reduction \mathcal{B}_1 that outputs ct_i, ct_j , breaking (t, ϵ_H) -collision-resistance of H , or a reduction \mathcal{B}_2 that outputs $x_i \| H(ct_i), x_j \| H(ct_j)$, breaking $(t, \epsilon_{\text{KDF}})$ -collision-resistance of KDF:

$$\text{Adv}_{\text{Kyber}}^{\mathcal{G}_{\text{A.0}}}(\mathcal{A}) \leq \epsilon_H + \epsilon_{\text{KDF}}.$$

Case B (Both decapsulations reject, distinct public keys: $R_i \wedge R_j \wedge pk_i \neq pk_j$).

Here, we bound the probability of the adversary producing ciphertexts that both get implicitly rejected during decapsulation and decapsulate to the same shared secret under distinct public keys.

Game B.0. This is the game conditioned on $R_i \wedge R_j \wedge pk_i \neq pk_j$ being satisfied.

$$\text{Adv}_{\text{Kyber}}^{\mathcal{G}_{\text{B.0}}}(\mathcal{A}) = \text{Adv}_{\text{Kyber}}^{\mathcal{G}_0[R_i \wedge R_j \wedge pk_i \neq pk_j]}(\mathcal{A}).$$

Game B.1 (Colliding z values). We let $\mathcal{G}_{\text{B.0}}$ return false if there is a collision among the z values in the n secret keys. Since the z values are 256-bit strings sampled uniformly at random, by the birthday bound we get:

$$\text{Adv}_{\text{Kyber}}^{\mathcal{G}_{\text{B.0}}}(\mathcal{A}) \leq \text{Adv}_{\text{Kyber}}^{\mathcal{G}_{\text{B.1}}}(\mathcal{A}) + \frac{n^2}{2^{256}}.$$

Game B.2 (Collision in KDF). Now, we have established $z_i \| H(ct_i) \neq z_j \| H(ct_j)$ (since $z_i \neq z_j$), yet $\text{KDF}(z_i \| H(ct_i)) = \text{KDF}(z_j \| H(ct_j))$. We build a reduction \mathcal{B}_3 that outputs $z_i \| H(ct_i), z_j \| H(ct_j)$, breaking $(t, \epsilon_{\text{KDF}})$ -collision resistance of KDF:

$$\text{Adv}_{\text{Kyber}}^{\mathcal{G}_{\text{B.1}}}(\mathcal{A}) \leq \epsilon_{\text{KDF}}.$$

Case C (One decapsulation rejects, distinct public keys: $R_i \wedge A_j \wedge pk_i \neq pk_j$).

Here, we bound the probability of the adversary producing exactly one ciphertext that gets implicitly rejected during decapsulation, while both ciphertexts decapsulate to the same shared secret under distinct public keys.

Game C.0. This is the game conditioned on $R_i \wedge A_j \wedge pk_i \neq pk_j$ being satisfied.

$$\text{Adv}_{\text{Kyber}}^{\mathcal{G}_{\text{C.0}}}(\mathcal{A}) = \text{Adv}_{\text{Kyber}}^{\mathcal{G}_0[R_i \wedge A_j \wedge pk_i \neq pk_j]}(\mathcal{A}).$$

If $ss'_i \neq z_j$ (for $ss' = \mathcal{G}_{256}(m'_i \| h_i)$), we can build a reduction \mathcal{B}_4 that returns $ss'_i \| H(ct_i), z_j \| H(ct_j)$ to break $(t, \epsilon_{\text{KDF}})$ -collision-resistance of KDF. Otherwise, we can build a reduction \mathcal{B}_5 that guesses $j \in [n]$ and embeds an image of \mathcal{G}_{256} in z_j , breaking the (t, ϵ'_G) -preimage resistance of \mathcal{G}_{256} for randomly sampled images (PR-r) if it guesses j correctly. Hence,

$$\text{Adv}_{\text{Kyber}}^{\mathcal{G}_{\text{C.0}}}(\mathcal{A}) \leq \epsilon_{\text{KDF}} + n \cdot \epsilon'_G.$$

Case D (Both decapsulations accept, distinct public keys: $A_i \wedge A_j \wedge pk_i \neq pk_j$).

Here, we bound the probability of the adversary producing ciphertexts that both get accepted during decapsulation and decapsulate to the same shared secret under distinct public keys.

Game D.0. This is the game conditioned on $A_i \wedge A_j \wedge pk_i \neq pk_j$ being satisfied.

$$\text{Adv}_{\text{Kyber}}^{\mathcal{G}_{\text{D.0}}}(\mathcal{A}) = \text{Adv}_{\text{Kyber}}^{\mathcal{G}_0[A_i \wedge A_j \wedge pk_i \neq pk_j]}(\mathcal{A}).$$

Now, $\text{KDF}(\mathbb{G}_{256}(m_i \| h_i) \| \text{H}(ct_i)) = ss_i = ss_j = \text{KDF}(\mathbb{G}_{256}(m_j \| h_j) \| \text{H}(ct_j))$ for $pk_i \neq pk_j$. It must be that \mathcal{A} provides us with a collision either in KDF, \mathbb{G}_{256} , or H. Once more, we can distinguish the cases:

1. In case $\mathbb{G}_{256}(m_i \| h_i) \| \text{H}(ct_i) \neq \mathbb{G}_{256}(m_j \| h_j) \| \text{H}(ct_j)$, then we can build a reduction \mathcal{B}_6 that outputs $\mathbb{G}_{256}(m_i \| h_i) \| \text{H}(ct_i), \mathbb{G}_{256}(m_j \| h_j) \| \text{H}(ct_j)$, breaking $(t, \epsilon_{\text{KDF}})$ -collision-resistance of KDF.
2. Otherwise, if $m_i \| h_i \neq m_j \| h_j$, then we can build a reduction \mathcal{B}_7 that outputs $m_i \| h_i, m_j \| h_j$, breaking $(t, \epsilon_{\mathbb{G}})$ -collision-resistance of \mathbb{G}_{256} .
3. Otherwise, we must have $m_i \| h_i = m_j \| h_j$ and $\text{H}(pk_i) = h_i = h_j = \text{H}(pk_j)$. By $\mathcal{G}_{\text{C.0}}$, $pk_i \neq pk_j$ and we can build a reduction \mathcal{B}_8 that outputs pk_i, pk_j , breaking (t, ϵ_{H}) -collision-resistance of H.

Jointly,

$$\text{Adv}_{\text{Kyber}}^{\mathcal{G}_{\text{C.0}}}(\mathcal{A}) \leq \epsilon_{\text{KDF}} + \epsilon_{\mathbb{G}} + \epsilon_{\text{H}}.$$

Collecting the bounds yields the claim. \square

4 Security Model

We analyze security of the PQXDH protocol in a computational, game-based security model for authenticated key exchange protocols, following the tradition of Bellare and Rogaway (BR) [BR94], but refined towards the maximum-exposure security that Signal’s handshakes (both X3DH and PQXDH) aim at. Our model is based on prior Signal and Signal-like security models by CCDGS [CCD⁺17] and particularly BFGJS [BFG⁺22, BFG⁺21], the latter introducing syntax and modeling for the particular class of asynchronous key exchange protocols that the Signal handshakes belong to.

In BR-style models, a computational adversary interacts with multiple users across multiple sessions of the key exchange protocol, fully controlling the network and being able to reroute, modify, inject, and drop messages at will (through a SEND oracle). For maximum-exposure, the model allows the adversary to compromise a user’s long-term key (via a CORRUPTLTKEY oracle) and semi-static keys (CORRUPTSSKEY), as well as reveal a session’s key (REVEALSESSKEY) and ephemeral randomness (REVEALRAND). The targeted security property finally is key indistinguishability, asking that keys established in so-called “fresh” sessions are indistinguishable from random keys. “Freshness” here encodes that a session is not trivially compromised and that the protocol under analysis aims to protect against the involved key compromises (defined through a set of “clean” predicates). We will detail those technical bits below.

For the analysis of PQXDH, we extend the prior models in several ways:

- Prior models [CCD⁺17, BFG⁺22, BFG⁺21] assumed semi-static keys being authentically distributed. However, only long-term keys can be verified out-of-band in Signal, and semi-static keys are signed with those keys instead. We are the first to capture these *signed semi-static keys* in a game-based model for a computational security analysis. In both the model and the security analysis (cf. Section 5), capturing these signatures leads to notable added complexity.
- In contrast to Signal’s classical X3DH handshake, PQXDH also involves *signed ephemeral (KEM) keys*. We capture this, too, expanding on the set of attack vectors our model covers.
- PQXDH adding KEM keys means that *multiple semi-static keys* (DH and KEM) can now be involved in a handshake. We capture this by accordingly modified identification of semi-static keys.

Note that CCDGS [CCD⁺17] also analyzed the ratcheting protocol, working in a multi-stage key exchange model [FG14]; PQXDH does not affect the ratcheting part, hence we focus on the initial handshake. BFGJS [BFG⁺22, BFG⁺21] in turn also studied deniability, which we do not consider in this work.

In the following, we will first review the syntax required to express both the PQXDH protocol and the security model formally; we mostly follow BFGJS [BFG⁺22, BFG⁺21] here, in parts verbatim, and focus on highlighting (with light-gray background) the main differences introduced in this work to account for signed semi-static and ephemeral keys as well as multiple identifiers for semi-static keys.

4.1 Syntax and Notation

Key exchange syntax. We define a two-party key exchange protocol via the following probabilistic algorithms:

- $\text{KGenLT}() \text{ } \text{s} \rightarrow (ltpk, ltsk)$: The *long-term key generation* algorithm that outputs a party’s public-key/secret-key pair.
- $\text{KGenSS}(ltsk) \text{ } \text{s} \rightarrow ((sspk_1, sssk_1, \sigma_1), \dots)$: The *semi-static key generation* algorithm that takes as input a long-term secret key $ltsk$ and outputs a vector of public-key/secret-key pairs and corresponding signatures.⁶
- $\text{Run}(ltsk, sssk, \vec{ltpk}, \pi, m) \text{ } \text{s} \rightarrow (\pi', m')$: 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 \vec{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 $m = (\text{create}, (\vec{ssid}, \text{type}))$, where \vec{ssid} indicates the semi-static keys⁷ to be used and type whether a full ($\text{type} = \text{full}$) or reduced ($\text{type} = \text{reduced}$) handshake should be performed.

Note that an explicit ephemeral key generation algorithm is not mandatory and can happen inside of Run .

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

- $\text{oid} \in [n_p]$: The identity of the session owner.
- $\text{pid} \in [n_p] \cup \{\star\}$: The identity of the intended peer, which may initially be unknown (indicated by \star).
- $\text{role} \in \{\text{initiator}, \text{responder}\}$: The role of the party.
- $\text{st}_{\text{exec}} \in \{\perp, \text{running}, \text{accepted}, \text{rejected}\}$: The status of this session’s execution.
- $\text{sid} \in \{0, 1\}^* \cup \{\perp\}$: A session identifier defining partnering.
- $\text{cid} \in \{0, 1\}^* \cup \{\perp\}$: 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 [DFGS15].
- $\text{K} \in \mathcal{K}_{\text{KE}} \cup \{\perp\}$: The session key established in this session, initialized to \perp .

⁶The length of this vector is protocol-dependent; e.g., in X3DH, semi-static keys consist of one DH key pair, and PQXDH adds a second semi-static key (for a KEM scheme). For simplicity, we have a set of semi-static keys be generated in one operation, yet in the protocol each key can be used independently.

⁷In the model, semi-static keys are identified by some value $\text{ssid} = "s, n"$ where s indicates the KGenSS call through which they were generated and n the key’s position in that call’s output. In practice, the latter signifies the type of semi-static key if there are several; for PQXDH, $n = 1$ for DH keys and $n = 2$ for KEM keys.

- $\text{type} \in \{\text{full}, \text{reduced}\}$: Indicator whether an ephemeral pre-key was used ($\text{type} = \text{full}$) for key establishment, or not ($\text{type} = \text{reduced}$).
- $\text{coins} \in \mathcal{R}_{\text{KE}}$: The random coins from the randomness space \mathcal{R}_{KE} used in the execution of Run; set by the game and read-only thereafter.
- $\text{sspk} \in (\{0, 1\}^*)^* \cup \{\perp\}$: The semi-static public keys used in this session.

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

- $\text{revrand} \in \{\text{true}, \text{false}\}$ indicates whether the random coins $\pi.\text{coins}$ have been revealed via a REVEALRAND query. The default value is `false`.
- $\text{pcorr} \in \{\text{true}, \text{false}\}$ 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 two sessions π_U^i and π_V^j are *partnered* if they agree on the session identifier: $\pi_U^i.\text{sid} = \pi_V^j.\text{sid} \neq \perp$. *Contributive* identifiers (*cid*) indicate when sessions may eventually derive the same key but are not fully partnered (yet); we use these to model security of initiator’s keys in incomplete handshakes.

4.2 Security Game

The security property of an authenticated key exchange protocol KE we consider in this work is *indistinguishability of session keys* (KI), formalized through the game $\mathcal{G}_{\text{KE}}^{\text{KI}}(\mathcal{A})$ in Figure 6 played by an adversary \mathcal{A} . At the start of the game, a random challenge bit $b_{\text{test}} \leftarrow_{\$} \{0, 1\}$ is fixed and long-term public-key/secret-key pairs are generated for all n_p honest parties and their public keys \vec{ltpk} provided to the adversary. The adversary is then able to interact with honest parties via the following queries:

- $\text{SEND}(U, i, m)$: Sends message m to session π_U^i , which corresponds to executing $\text{Run}(ltsk_U, \vec{ssk}_U, \vec{ltpk}, \pi_U^i, m)$, saving the updated session state π' as π_U^i , and returning the outgoing message m' to the adversary.
- $\text{CORRUPTLTKEY}(U)$: Returns party U ’s long-term secret key $ltsk_U$ to the adversary; recorded through the flag `corrltkU`.
- $\text{CORRUPTSSKEY}(U, \text{ssid})$: Returns party U ’s semi-static secret key ssk_U^{ssid} to the adversary; recorded through the flag `corrssksspk`, where $sspk$ is the public key corresponding to ssk_U^{ssid} .⁸
- $\text{REVEALRAND}(U, i)$: Returns the random coins of session π_U^i to the adversary, recorded in the session through the flag `revrand`.
- $\text{REVEALSESSKEY}(U, i)$: If session π_U^i has accepted, return its session key $\pi_U^i.\mathcal{K}$ to the adversary.
- $\text{TEST}(U, i)$: If the TEST query has been called before or session π_U^i has not accepted, then return \perp . Otherwise; if $b_{\text{test}} = 0$, return $\pi_U^i.\mathcal{K}$, otherwise return a randomly sampled session key from the protocol’s key space \mathcal{K}_{KE} . Record the test session as $\pi^* \leftarrow \pi_U^i$.

⁸Since we model that semi-static keys are not authentically distributed (but signed), sessions only know the public key received ($sspk$), but not the game label `ssid`, which was used in prior models [CCD⁺17, BFG⁺22, BFG⁺21] to track compromise more easily.

At the end of the game, the adversary outputs a bit b' . The adversary is said to win if $b' = b_{\text{test}}$ and the test session π^* is fresh. Formally, if the test session is fresh, the experiment outputs 1 if $b' = b_{\text{test}}$ and 0 otherwise; if the test session is not fresh, then the experiment outputs a random bit. The adversary's advantage in the key indistinguishability game measured as the experiment outputting 1 minus the adversary's guessing chance, $\frac{1}{2}$.

Definition 4.1 (Key indistinguishability). *Let KE be a key exchange protocol and \mathcal{A} an adversary against the key indistinguishability (KI) game $\mathcal{G}_{\text{KE}}^{\text{KI}}(\mathcal{A})$ in Figure 6. We say that KE achieves $(t, \epsilon, (q_{\text{Snd}}, q_{\text{CorrLT}}, q_{\text{CorrSS}}, q_{\text{RevR}}, q_{\text{RevSK}}))$ -key indistinguishability, if for any adversary \mathcal{A} against $\mathcal{G}_{\text{KE}}^{\text{KI}}(\mathcal{A})$ with running time at most t and making at most $q_{\text{Snd}}, q_{\text{CorrLT}}, q_{\text{CorrSS}}, q_{\text{RevR}}$, resp. q_{RevSK} queries to its SEND, CORRUPTLTKEY, CORRUPTSSKEY, REVEALRAND, resp. REVEALSESSKEY oracles, we have that*

$$\text{Adv}_{\text{KE}}^{\text{KI}}(\mathcal{A}) = \Pr \left[\mathcal{G}_{\text{KE}}^{\text{KI}}(\mathcal{A}) \right] - \frac{1}{2} \leq \epsilon.$$

Note that the model restricts the adversary to a single query to the TEST oracle.

Soundness. The model also captures soundness (via the predicate `sound`), i.e., that session identifiers appropriately reflect correct protocol executions. Concretely, soundness demands that an adversary cannot create one of the following situations (or else it will win the game immediately):

- (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. 6, line 18).
- (ii) Two initiator sessions accept with the same session identifier (Fig. 6, line 19).
- (iii) Three sessions accept with the same session identifier in full handshake type (Fig. 6, line 20).

Freshness. Unrestricted access to the game oracles allows for trivial wins, e.g., by testing a session key and also revealing it (or its partner). The freshness predicate `fresh` rules out such trivial wins, and further encodes that the adversary has not obtained sufficiently many secrets, via CORRUPTLTKEY and/or CORRUPTSSKEY and/or REVEALRAND queries, to derive the session key of the test session itself and/or substituted signed keys in its execution; this is encoded through a set of so-called clean predicates (Figure 6, line 16 and lines 46 ff.).

Clean predicates. Following the terminology of CCDGS [CCD⁺17], the set of clean predicates capture the maximum-exposure properties of the protocol being analyzed. Figure 6, lines 46 ff. formally capture these for PQXDH. The core properties closely follow those for the classic X3DH handshake [CCD⁺17, BFG⁺22]: in handshakes where the long-term/semi-static, ephemeral/long-term, or ephemeral/semi-static (or, for type = full handshakes, ephemeral/ephemeral) secrets combination between initiator/responder is not revealed, we expect security of the derived session key. Jumping ahead, for those combinations where PQXDH employs both DH and KEM keys, we will obtain hybrid/combiner security terms in our analysis (see Section 5), capturing that an adversary has to break *both* primitives to successfully attack the protocol via these compromise angles.

In addition to prior work [CCD⁺17, BFG⁺22, BFG⁺21], our clean predicates further model the signatures on the responder's (semi-static and ephemeral public) keys and allow the adversary to replace semi-static and ephemeral keys sent over the network. Instead of assuming authentic distribution of semi-static public keys, we merely preclude the adversary from compromising the (long-term) signing key prior

$\mathcal{G}_{KE}^{KI}(\mathcal{A})$:

```

1   $b_{\text{test}} \leftarrow \{0, 1\}$  // sample challenge bit
2   $\pi^* \leftarrow \perp$  // variable for test session
3  for  $U \in [n_p]$  // generate long-term keys
4     $(pk_U, sk_U) \leftarrow \text{KGenLT}()$ 
5    for  $s \in [n_{ss}]$  // generate semi-static keys
6       $((ssp_{k_U}^{s,1}, sss_{k_U}^{s,1}, \sigma_U^{s,1}), \dots) \leftarrow \text{KGenSS}(ltsk_U)$ 
7       $sss_{k_U} \leftarrow \{sss_{k_U}^{s,1}, \dots\}^{s \in [n_{ss}]}$ 
8       $ltpk_U \leftarrow \{ltpk_U\}_{U \in [n_p]}$ 
9       $ssp_{k_U} \leftarrow \{(ssp_{k_U}^{s,1}, \sigma_U^{s,1}), \dots\}^{s \in [n_{ss}]}$ 
10  $b' \leftarrow \mathcal{A}(ltpk_U, ssp_{k_U})$  // run adversary
11 if  $\text{sound}() = \text{false}$ : return true // adversary wins if it breaks soundness
12 if  $\text{fresh}(\pi^*) = \text{false}$ :  $b' \leftarrow 0$  // attack invalid if test session is not fresh
13 return  $[b' = b_{\text{test}}]$  // determine win or loss

```

$\text{SEND}(U, i, m)$:

```

21 if  $\pi_U^i = \perp$  // initiate session: for responders, we have  $m = (\text{create}, (\text{ssid}, \text{type}))$ 
22    $\pi_U^i.\text{oid} \leftarrow U$  // set owner identity
23   if  $m = (\text{create}, \dots)$ :  $\pi_U^i.\text{role} \leftarrow \text{responder}$  // set responder role
24   else  $\pi_U^i.\text{role} \leftarrow \text{initiator}$  // set initiator role ( $m$  is first message)
25    $\pi_U^i.\text{coins} \leftarrow \mathcal{R}_{KE}$  // sample session randomness
26    $\pi_U^i.\text{st}_{\text{exec}} \leftarrow \text{running}$ 
27    $(\pi_U^i, m') \leftarrow \text{Run}(sk_U, sss_{k_U}, ltpk_U, \pi_U^i, m)$ 
   // run session, random coins in  $\pi_U^i$ 
28   if  $\pi_U^i.\text{st}_{\text{exec}} = \text{accepted}$ : // flag if peer was corrupted upon acceptance
29      $\pi_U^i.\text{pcorr} \leftarrow \text{corrtk}_{\pi_U^i.\text{pid}}$ 
30   return  $(m', \pi_U^i.\text{st}_{\text{exec}})$  // return message and session state

```

$\text{TEST}(U, i)$:

```

31 if  $\pi_U^i = \perp$  or  $\pi_U^i.\text{st}_{\text{exec}} \neq \text{accepted}$  or  $\pi^* \neq \perp$ : return  $\perp$ 
   // session does not exist, has not accepted yet, or test already asked
32  $\pi^* \leftarrow \pi_U^i$  // record test session
33  $K_0 \leftarrow \pi_U^i.K$ 
34  $K_1 \leftarrow \mathcal{K}_{KE}$ 
35 return  $K_{b_{\text{test}}}$  // return real-or-random challenge key

```

$\text{clean}_{\text{full}}(\pi^*)$:

```

46 return  $\text{clean}_{\text{reduced}}(\pi^*)$  or  $\text{clean}_{\text{EE}}(\pi^*)$ 

```

$\text{clean}_{\text{reduced}}(\pi^*)$:

```

47 return  $\text{clean}_{\text{LTSS}}(\pi^*)$  or  $\text{clean}_{\text{ELT}}(\pi^*)$  or  $\text{clean}_{\text{ESS}}(\pi^*)$ 

```

$\text{clean}_{\text{EE}}(\pi^*)$:

```

48 return  $\neg \pi^*.\text{revrand}$  and  $(\text{clean}_{\text{peerE}}(\pi^*)$  or  $\text{clean}_{\text{sigE}}(\pi^*))$ 
   // test session randomness is unrevealed and peer's ephemeral contribution is clean or cleanly signed

```

$\text{clean}_{\text{peerE}}(\pi^*)$:

```

49 return
50    $(\pi^*.\text{role} = \text{initiator} \text{ and } \exists \pi \neq \pi^* :$ 
    $(\pi.\text{role} = \text{responder} \text{ and } \pi^*.\text{cid} = \pi.\text{cid} \text{ and } \neg \pi.\text{revrand}))$ 
   // contributively-partnered responder session's randomness is unrevealed
51 or  $(\pi^*.\text{role} = \text{responder} \text{ and } \exists \pi \neq \pi^* :$ 
    $(\pi.\text{role} = \text{initiator} \text{ and } \pi^*.\text{sid} = \pi.\text{sid} \text{ and } \neg \pi.\text{revrand}))$ 
   // partnered initiator session's randomness is unrevealed

```

$\text{clean}_{\text{sigE}}(\pi^*)$:

```

52 return  $(\pi^*.\text{role} = \text{initiator} \text{ and } \neg \pi^*.\text{pcorr} \text{ and } \forall \pi :$ 
    $((\pi.\text{role} = \text{responder} \text{ and } \pi^*.\text{cid} = \pi.\text{cid}) \implies \neg \pi.\text{revrand}))$ 
   // long-term (signing) secret of the responder peer was uncompromised upon
   acceptance, and if a partner exists, then that partner's randomness is unrevealed

```

$\text{fresh}(\pi^*)$:

```

14 if  $\pi^*.\text{revealed} = \text{true}$ : return false // test session is revealed
15 if  $\exists \pi_V^i \neq \pi^* : (\pi_V^i.\text{sid} = \pi^*.\text{sid} \wedge \pi_V^i.\text{revealed} = \text{true})$ : return false
   // test session's partner is revealed
16 return  $\text{clean}_{\pi^*.\text{type}}(\pi^*)$ 
   // test session is clean wrt. to its handshake type (full resp. reduced)

```

$\text{sound}()$:

```

17 return  $\forall$  distinct  $\pi, \pi', \pi''$  (
18    $(\pi.\text{sid} = \pi'.\text{sid} \neq \perp \implies \pi.K = \pi'.K \wedge \pi.\text{type} = \pi'.\text{type} \wedge \pi.\text{cid} = \pi'.\text{cid})$ 
   // same session identifiers imply same key, type, contributive identifiers
19   and  $(\pi.\text{sid} = \pi'.\text{sid} \neq \perp \wedge \pi.\text{role} = \text{initiator} \implies \pi'.\text{role} = \text{responder})$ 
   // session identifiers of two initiator sessions never collide
20   and  $(\pi.\text{sid} = \pi'.\text{sid} = \pi''.\text{sid} \neq \perp \implies \pi.\text{type} = \text{reduced})$ )
   // session identifiers of three sessions only collide in reduced mode

```

$\text{CORRUPTLTKEY}(U)$:

```

36  $\text{corrtk}_U \leftarrow \text{true}$  // mark long-term key corrupted
37 return  $sk_U$  // return long-term secret key

```

$\text{CORRUPTSSKEY}(U, \text{ssid})$:

```

38  $\text{corrssk}_{\text{ssid}} \leftarrow \text{true}$  // mark semi-static key corrupted
39 return  $sss_{k_U}^{\text{ssid}}$  // return semi-static secret key

```

$\text{REVEALRAND}(U, i)$:

```

40 if  $\pi_U^i = \perp$ : return  $\perp$  // session does not exist
41  $\pi_U^i.\text{revrand} \leftarrow \text{true}$  // mark randomness revealed
42 return  $\pi_U^i.\text{coins}$  // return session's random coins

```

$\text{REVEALSESSKEY}(U, i)$:

```

43 if  $\pi_U^i = \perp$  or  $\pi_U^i.\text{st}_{\text{exec}} \neq \text{accepted}$ : return  $\perp$ 
   // session does not exist or has not yet derived session key
44  $\pi_U^i.\text{revealed} \leftarrow \text{true}$  // mark session key revealed
45 return  $\pi_U^i.K$  // return session key

```

$\text{clean}_{\text{LTSS}}(\pi^*)$:

```

53 return
54    $(\pi^*.\text{role} = \text{initiator} \text{ and } \neg \text{corrtk}_{\pi^*.\text{oid}} \text{ and}$ 
    $\forall sspk \in \pi^*.\text{sspks} : (\neg \text{corrssk}_{sspk} \text{ and } \neg \pi^*.\text{pcorr}))$ 
   // tested initiator's long-term and responder's semi-static secrets are uncompromised
   // and responder long-term was uncompromised upon acceptance
55 or  $(\pi^*.\text{role} = \text{responder} \text{ and } \neg \text{corrtk}_{\pi^*.\text{pid}} \text{ and}$ 
    $\forall sspk \in \pi^*.\text{sspks} : (\neg \text{corrssk}_{sspk}))$ 
   // initiator long-term and tested responder's semi-static secrets are uncompromised

```

$\text{clean}_{\text{ELT}}(\pi^*)$:

```

56 return
57    $(\pi^*.\text{role} = \text{initiator} \text{ and } \neg \pi^*.\text{revrand} \text{ and } \neg \text{corrtk}_{\pi^*.\text{pid}})$ 
   // tested initiator's randomness is unrevealed and responder long-term secret is uncompromised
58 or  $(\pi^*.\text{role} = \text{responder} \text{ and } \text{clean}_{\text{peerE}}(\pi^*) \text{ and } \neg \text{corrtk}_{\pi^*.\text{oid}})$ 
   // initiator's ephemeral contribution is clean and tested responder's long-term secret is uncompromised

```

$\text{clean}_{\text{ESS}}(\pi^*)$:

```

59 return
60    $(\pi^*.\text{role} = \text{initiator} \text{ and } \neg \pi^*.\text{revrand} \text{ and}$ 
    $\forall sspk \in \pi^*.\text{sspks} : (\neg \text{corrssk}_{sspk} \text{ and } \neg \pi^*.\text{pcorr}))$ 
   // tested initiator's randomness is unrevealed and responder semi-static secrets
   are uncompromised and responder long-term was uncompromised upon acceptance
61 or  $(\pi^*.\text{role} = \text{responder} \text{ and } \text{clean}_{\text{peerE}}(\pi^*) \text{ and}$ 
    $\forall sspk \in \pi^*.\text{sspks} : (\neg \text{corrssk}_{sspk}))$ 
   // initiator's ephemeral contribution is clean and tested responder's semi-static secrets are uncompromised

```

Figure 6: 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 (middle), and wrt. to clean predicates for PQXDH (bottom). Highlighted code reflects the main changes compared to the model of BFGJS [BFG⁺22, BFG⁺21].

to the targeted session having accepted (cf. Figure 6, lines 54 and 60). Treating signatures explicitly, we also capture the additional guarantees that signing KEM keys give: for the ephemeral/ephemeral combination to contribute security, it is now (also) sufficient to have the peer’s contribution be cleanly signed (cf. Figure 6, line 48). To that end we set a flag `pcorr` in a session if its peer was corrupted upon acceptance (Figure 6, line 29). For PQXDH, this part of the model ensures that a quantum-later adversary cannot, in addition to a “harvest now, decrypt later” attack, replace the ephemeral KEM public key with its own without being noticed.

In more formal detail, let π^* denote the test session. Depending on whether an ephemeral pre-key was used in the key derivation of π^* or not, we apply either the `cleanfull` or the `cleanreduced` predicate to π^* .

Since `cleanreduced` is part of the description of `cleanfull`, we first discuss the case $\pi^*.type = reduced$. Intuitively, a session key derived in such a session remains unknown to the adversary, if one of the four keys (i.e., excluding the unused DH_4) that constitute the master secret is “clean”, i.e., cannot be computed by the adversary. This is the case if one of the following three clean predicates holds for the test session π^* :

`cleanLTSS`: This predicate indicates whether the combination of the long-term key of the initiator and the semi-static keys of the responder are unknown to the adversary. If the test session is an initiator, the responder’s signing key must further be uncompromised upon acceptance.

`cleanELT`: This predicate indicates whether the combination of the ephemeral contribution of the initiator and the long-term key of the responder is unknown to the adversary.

`cleanESS`: This predicate indicates whether the combination of the ephemeral contribution of the initiator and the semi-static keys of the responder are unknown to the adversary. If the test session is an initiator, the responder’s signing key must further be uncompromised upon acceptance.

If the test session π^* is a responder session, the evaluation of `cleanELT` and `cleanESS` necessitates a further predicate called `cleanpeerE` (in all other cases, it is sufficient to consider the compromise of keys and randomness via the flags `corrltk`, `corrsk`, `revrand`, and `pcorr`).

`cleanpeerE`: This (sub)predicate indicates that the randomness used within any (sid- or cid-)partnered session is unknown to the adversary.

For test sessions in full handshake mode, i.e., where $\pi^*.type = full$, it must either hold that `cleanreduced` is true or that the additional input to the master secret computation is clean. The latter is captured by the following predicate:

`cleanEE`: This predicate indicates that the ephemeral contribution of the test session is unknown to the adversary and that the ephemeral contribution of the peer is unknown to the adversary (captured by `cleanpeerE`) or was “cleanly” signed (captured by the new `cleansigE` predicate we introduce). It is the second part of this “or” statement by which our model captures that KEM ephemeral keys are protected by signatures against replacement.

Again, the predicate `cleanpeerE` helps to determine within `cleanEE` whether the randomness of the test session’s (contributive) partners is unrevealed. The following predicate precludes that the adversary substitutes the responder’s ephemeral key and signs it with the signing key obtained from corrupting the long-term key.

`cleansigE`: This (sub)predicate indicates that for initiator test sessions the peer was not corrupted at the time of acceptance and, if the test session has a contributively partnered session, then the ephemeral contribution of that session is not known to the adversary.

5 PQXDH Analysis

We formalize the PQXDH handshake in Figure 7, highlighting the changes compared to X3DH to achieve post-quantum security in **dark gray**. Ours is the first reductionist analysis to model the signatures on public keys, which we accent with **light gray**. Long-term key pairs consist of a DH key pair and a signing key pair.⁹ Semi-static key pairs consist of a DH key pair and a KEM key pair, each signed under the long-term signing key. The protocol then works as follows.

First, Bob produces a pre-key bundle with semi-static and ephemeral keys as indicated by the $\vec{\text{ssid}}$ vector and **type**: The pre-key bundle always includes the semi-static DH key and for reduced handshakes a semi-static KEM key. Only for full handshakes, the pre-key bundle includes ephemeral DH and KEM keys.¹⁰ Signatures on the semi-static DH key and the KEM key (regardless if it is semi-static or ephemeral) are always included.

Second, Alice verifies both signatures and samples an ephemeral DH key pair of her own. She computes several DH secret combinations, namely long-term/semi-static, ephemeral/long-term, ephemeral/semi-static, and ephemeral/ephemeral (the last one only for full handshakes), and encapsulates against Bob's KEM public key. She derives the session key via a key derivation function KDF on input all three/four DH shared secrets and the KEM shared secret and sends her sampled DH ephemeral public key and the KEM ciphertext to Bob.

Third, Bob computes the same DH shared secrets and decapsulates the KEM ciphertext. He derives the session key in the same manner and accepts.

Security

We show that PQXDH achieves key indistinguishability wrt. to the maximum-exposure attack vectors formalized through the clean predicates in our model (cf. Section 4). As for the prior analysis of X3DH [CCD⁺17] whose proof structure we follow in parts, the bound is highly non-tight; we still give it in concrete terms for clarity.

Theorem 5.1 (Key indistinguishability of PQXDH). *The PQXDH protocol given in Figure 7 with randomness space $\mathcal{R}_{\text{KE}} = \mathcal{R}_{\text{DH.KGen}} \times \mathcal{R}_{\text{KEM.KGen}} \times \mathcal{R}_{\text{KEM.Enc}} \times \mathcal{R}_{\text{SIG.Sig}}$ achieves $(t, \epsilon, (q_{\text{Snd}}, q_{\text{CorrLT}}, q_{\text{CorrSS}}, q_{\text{RevR}}, q_{\text{RevSK}}))$ -key indistinguishability, assuming (for $t' \approx t$) the GapDH problem in the group (\mathbb{G}, g, q) is $(t', \epsilon_{\text{GDH}}, 1.5(q_{\text{RO}} + q_{\text{Snd}})^2)$ -hard, KEM is a $(t', \epsilon_{\text{CCA}}, n_s)$ -OW-CCA-secure and $(t', \epsilon_{\text{LEAK}^+}, n_p \cdot n_{\text{ss}} + n_s)$ -LEAK⁺-BIND-SS- $\{CT, PK\}$ -secure KEM with public-key collision probability γ_{coll} and correctness error δ_{corr} , SIG is a $(t', \epsilon_{\text{SIG}}, 2n_{\text{ss}} + n_s)$ -unforgeable signature scheme, and KDF behaves like a (programmable) random oracle. Concretely, for any efficient adversary \mathcal{A} we have*

$$\text{Adv}_{\text{PQXDH}}^{\text{KI}}(\mathcal{A}) \leq \frac{(n_p + n_p \cdot n_{\text{ss}} + n_s)^2}{q} + \gamma_{\text{coll}}(n_p \cdot n_{\text{ss}} + n_s) + n_s \cdot \delta_{\text{corr}} + \epsilon_{\text{LEAK}^+} + \left(\begin{array}{ll} (n_p \cdot (\epsilon_{\text{SIG}} + n_p \cdot n_{\text{ss}} \cdot \epsilon_{\text{GDH}})) & // \text{clean}_{\text{LTSS}} \\ + (n_s \cdot n_p \cdot \epsilon_{\text{GDH}}) & // \text{clean}_{\text{ELT}} \\ + (n_p \cdot (\epsilon_{\text{SIG}} + n_{\text{ss}} \cdot n_s \cdot \epsilon_{\text{GDH}})) & // \text{clean}_{\text{ESS}} \wedge \text{type} = \text{full} \\ + (n_p \cdot (\epsilon_{\text{SIG}} + n_{\text{ss}} \cdot n_s \cdot \min(\epsilon_{\text{GDH}}, \epsilon_{\text{CCA}}))) & // \text{clean}_{\text{ESS}} \wedge \text{type} = \text{reduced} \\ + (n_s^2 \cdot \min(\epsilon_{\text{GDH}}, \epsilon_{\text{CCA}})) & // \text{clean}_{\text{EE}} \wedge \text{clean}_{\text{peerE}} \\ + (n_p \cdot (\epsilon_{\text{SIG}} + n_s^2 \cdot q_{\text{RO}} \cdot \epsilon_{\text{CCA}})) & // \text{clean}_{\text{EE}} \wedge \text{clean}_{\text{sigE}} \end{array} \right).$$

For easier accessibility, we first give a summary of the proof; the full proof follows below.

⁹The PQXDH description [KS24] suggests to use a single key pair for both DH and the signature scheme XEdDSA [Per16], modeling which we leave to future work.

¹⁰Ephemeral DH and KEM pre-keys are used as long as there are some left on the Signal server. In practice, ephemeral DH and KEM pre-keys may run out at different points in time. We do not model this to improve accessibility of the proof.

KGenLT():

```

1 ( $ltpk^{DH}, ltsk^{DH}$ )  $\leftarrow$  DH.KGen()
2 ( $ltpk^{SIG}, ltsk^{SIG}$ )  $\leftarrow$  SIG.KGen()
3 return ( $(ltpk^{DH}, ltpk^{SIG}), (ltsk^{DH}, ltsk^{SIG})$ )

```

KGenSS($ltsk$):

```

4 ( $sspk^{DH}, sssk^{DH}$ )  $\leftarrow$  DH.KGen(); ( $sspk^{KEM}, sssk^{KEM}$ )  $\leftarrow$  KEM.KGen()
5 ( $ltsk^{DH}, ltsk^{SIG}$ )  $\leftarrow$   $ltsk$ 
6  $\sigma^{DH} \leftarrow$  SIG.Sig( $ltsk^{SIG}, sspk^{DH}$ );  $\sigma^{KEM} \leftarrow$  SIG.Sig( $ltsk^{SIG}, sspk^{KEM}$ )
7  $sssk^{DH*} \leftarrow (sssk^{DH}, sspk^{DH}, \sigma^{DH})$ ;  $sssk^{KEM*} \leftarrow (sssk^{KEM}, sspk^{KEM}, \sigma^{KEM})$ 
8 return ( $(sspk^{DH}, sssk^{DH*}, \sigma^{DH}), (sspk^{KEM}, sssk^{KEM*}, \sigma^{KEM})$ )

```

Alice

```

Run( $ltsk_A, sssk_A, ltpk, \pi_A, m$ ; coins)
( $r_4, \perp, \perp, r_5$ )  $\leftarrow$   $\pi_A$ .coins
( $ltsk_A^{DH}, ltsk_A^{SIG}$ )  $\leftarrow$   $ltsk_A$ 
( $epk_A^{DH}, esk_A^{DH}$ )  $\leftarrow$  DH.KGen( $r_4$ )
( $B, sspk, epk_B, \sigma_B^{DH}, \sigma_B^{KEM}$ )  $\leftarrow$   $m$ 
( $ltpk_B^{DH}, ltpk_B^{SIG}$ )  $\leftarrow$   $ltpk_B$ 
( $sspk_B^{DH}, sspk_B^{KEM}$ )  $\leftarrow$   $sspk$ 
if SIG.Vf( $ltpk_B^{SIG}, sspk_B^{DH}, \sigma_B^{DH}$ ) = false
  return ( $\pi_A, \epsilon$ )
 $DH_1 \leftarrow$  DH( $ltsk_A^{DH}, sspk_B^{DH}$ )
 $DH_2 \leftarrow$  DH( $esk_A^{DH}, ltpk_B^{DH}$ )
 $DH_3 \leftarrow$  DH( $esk_A^{DH}, sspk_B^{DH}$ )
if  $epk_B \neq \perp$  // full handshake
  ( $epk_B^{DH}, epk_B^{KEM}$ )  $\leftarrow$   $epk_B$ 
   $DH_4 \leftarrow$  DH( $esk_A^{DH}, epk_B^{DH}$ )
  if SIG.Vf( $ltpk_B^{SIG}, epk_B^{KEM}, \sigma_B^{KEM}$ ) = false
    return ( $\pi_A, \epsilon$ )
  ( $ct, ss$ )  $\leftarrow$  KEM.Enc( $epk_B^{KEM}; r_5$ )
   $\pi_A$ .type  $\leftarrow$  full
else // reduced handshake
   $DH_4 \leftarrow \epsilon$ 
  if SIG.Vf( $ltpk_B^{SIG}, sspk_B^{KEM}, \sigma_B^{KEM}$ ) = false
    return ( $\pi_A, \epsilon$ )
  ( $ct, ss$ )  $\leftarrow$  KEM.Enc( $sspk_B^{KEM}; r_5$ )
   $\pi_A$ .type  $\leftarrow$  reduced
 $\pi_A$ .K  $\leftarrow$  KDF( $DH_1 || DH_2 || DH_3 || DH_4 || ss$ )
 $\pi_A$ .sid  $\leftarrow$  ( $A, B, ltpk_A, ltpk_B,$ 
   $sspk, epk_B, epk_A^{DH}, ct$ )
 $\pi_A$ .pid  $\leftarrow$  B
 $\pi_A$ .sspks  $\leftarrow$   $sspk$ 
 $\pi_A$ .cid  $\leftarrow$  ( $B, ltpk_B, sspk, epk_B$ )
 $\pi_A$ .stexec  $\leftarrow$  accepted
return ( $\pi_A, m' = (A, epk_A^{DH}, ct)$ )

```

Bob

Run($ltsk_B, ltpk, sssk_B, \pi_B, (create, (ssid, type))$)

```

( $r_1, r_2, r_3, \perp$ )  $\leftarrow$   $\pi_B$ .coins
( $sssk_B^{DH}, sspk_B^{DH}, \sigma_B^{DH}$ )  $\leftarrow$   $sssk_B^{ssid[1]}$ 
if type = full // full handshake
  ( $epk_B^{DH}, esk_B^{DH}$ )  $\leftarrow$  DH.KGen( $r_1$ )
  ( $epk_B^{KEM}, esk_B^{KEM}$ )  $\leftarrow$  KEM.KGen( $r_2$ )
  ( $ltsk_B^{DH}, ltsk_B^{SIG}$ )  $\leftarrow$   $ltsk_B$ 
   $\sigma_B^{KEM} \leftarrow$  SIG.Sig( $ltsk_B^{SIG}, epk_B^{KEM}; r_3$ )
   $epk_B \leftarrow (epk_B^{DH}, epk_B^{KEM})$ 
   $sspk \leftarrow (sspk_B^{DH}, \perp)$ 
else // reduced handshake
  ( $sspk_B^{KEM}, sssk_B^{KEM}, \sigma_B^{KEM}$ )  $\leftarrow$   $sssk_B^{ssid[2]}$ 
   $epk_B \leftarrow \perp$ 
   $sspk \leftarrow (sspk_B^{DH}, sspk_B^{KEM})$ 
   $\sigma_B^{KEM} \leftarrow \sigma_B^{KEM}$ 
 $\pi_B$ .pid  $\leftarrow$  *
 $\pi_B$ .sspks  $\leftarrow$   $sspk$ 
 $\pi_B$ .type  $\leftarrow$  type
 $\pi_B$ .cid  $\leftarrow$  ( $B, ltpk_B, sspk, epk_B$ )
return ( $\pi_B, m = (B, sspk, epk_B, \sigma_B^{DH}, \sigma_B^{KEM})$ )

```

Run($ltsk_B, sssk_B, ltpk, \pi_B, m'$; coins)

```

( $A, epk_A^{DH}, ct$ )  $\leftarrow$   $m'$ 
( $ltpk_A^{DH}, ltpk_A^{SIG}$ )  $\leftarrow$   $ltpk_A$ 
 $DH_1 \leftarrow$  DH( $ltpk_A^{DH}, sssk_B^{DH}$ )
 $DH_2 \leftarrow$  DH( $epk_A^{DH}, ltsk_B^{DH}$ )
 $DH_3 \leftarrow$  DH( $epk_A^{DH}, sssk_B^{DH}$ )
if  $\pi_B$ .type = full // full handshake
   $DH_4 \leftarrow$  DH( $epk_A^{DH}, esk_B^{DH}$ )
   $ss \leftarrow$  KEM.Dec( $esk_B^{KEM}, ct$ )
else // reduced handshake
   $DH_4 \leftarrow \epsilon$ 
   $ss \leftarrow$  KEM.Dec( $sssk_B^{KEM}, ct$ )
 $\pi_B$ .K  $\leftarrow$  KDF( $DH_1 || DH_2 || DH_3 || DH_4 || ss$ )
 $\pi_B$ .sid  $\leftarrow$  ( $A, B, ltpk_A, ltpk_B,$ 
   $sspk, epk_B, epk_A^{DH}, ct$ )
 $\pi_B$ .pid  $\leftarrow$  A
 $\pi_B$ .stexec  $\leftarrow$  accepted
return ( $\pi_B, \epsilon$ )

```

Figure 7: The PQXDH protocol. Darker gray code lines are additions compared to the original X3DH protocol to add post-quantum security. Lighter gray code lines are signature elements of the X3DH and PQXDH not covered in the prior analyses of [CCD⁺17].

Proof summary. The proof proceeds via a series of game hops. First, we exclude collisions in the DH and KEM public keys, yielding the terms $\frac{(n_p + n_p \cdot n_{ss} + n_s)^2}{q} + \gamma_{\text{coll}}(n_p \cdot n_{ss} + n_s)$. Second, we exclude KEM correctness errors when decapsulating in sessions ($n_s \cdot \delta_{\text{corr}}$). Now, soundness holds, as the session identifiers carrying non-colliding DH/KEM keys make them non-colliding themselves and sessions agreeing on session identifiers in particular derive the same session keys. Third, we ensure that KEM shared secrets bind the public keys and ciphertexts, introducing the $\epsilon_{\text{LEAK+}}$ term; this in particular rules out re-encapsulation attacks where sessions that are not partnered possibly derive the same session key.¹¹

At this stage we separate the proof into several cases based on the `clean` predicate (cf. the annotated sum in the bound of Theorem 5.1). Let us illustrate the fourth case (`cleanESS ∧ type = reduced`) here since it covers the most interesting proof techniques, including handling the signatures and a hybrid DH/KEM argument; the other cases follow in a similar fashion. In this case, the `cleanESS` predicate is satisfied and the adversary is testing a session running a reduced (`type = reduced`) handshake. In the proof, we first guess the responder identity involved in the test session, introducing a n_p loss. Second, we ensure that the guessed responder’s semi-static KEM and DH keys are not tampered with in the test session, reducing to the unforgeability of the signature scheme (ϵ_{SIG}). Third and fourth, we guess the two semi-static key identifiers of the responder as well as the initiator session, with a loss of $n_{ss}^2 \cdot n_s$. Fifth, we abort if the adversary queries the key derivation function KDF, modeled as random oracle, on the master secret of the test session. To bound this step, we embed *both* a GapDH challenge *and* a OW-CCA challenge into the test session (concretely, into the ephemeral DH and KEM encapsulation contribution of the initiator and the semi-static DH and KEM keys of the responder, which we both guessed before). If the adversary detects this change, the reduction can win *both* games, yielding the minimum of the advantages as the *hybrid* bound, i.e., $\min(\epsilon_{\text{GDH}}, \epsilon_{\text{CCA}})$. Finally, we can replace the session key of the test session with a random key, making it independent of the challenge bit b_{test} . Concluding that the session key of the test session is now independent of b_{test} since non-partnered sessions derive distinct keys (in particular due to the KEM binding hop in the beginning), this completes the proof case. \square

We get hybrid guarantees for the fourth and fifth proof cases (`cleanESS` in reduced mode and `cleanEE` in full mode with a clean peer, i.e., `cleanpeerE`), showing that for these attack vectors, an adversary would need to break *both* GapDH in the DH group and OW-CCA security of the KEM. The other attack vectors either involve only DH secrets (and hence do not provide post-quantum security), or, for the sixth case (`cleanEE` in full mode with clean signatures), the adversary can manipulate the (unsigned) ephemeral DH key, hence only the KEM secret ensures security here. All in all, PQXDH *maintains* the classical guarantees of X3DH for the first three cases (as intended), extends the guarantees for the fourth and fifth case to provide *hybrid* guarantees, and *adds* a post-quantum guarantee for the last case due to the signed ephemeral KEM key (compared to no guarantee under this attack vector for X3DH).

We now give the full proof for Theorem 5.1.

Proof. Game 0. We start with Game \mathcal{G}_0 being the original key indistinguishability game $\mathcal{G}_{\text{PQXDH}}^{\text{KI}}(\mathcal{A})$, i.e.,

$$\text{Adv}_{\text{PQXDH}}^{\text{KI}}(\mathcal{A}) = \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_0}(\mathcal{A}).$$

Game 1 (DH and KEM key collisions). We let \mathcal{G}_0 abort (overwriting the adversary’s output with 0) if any two honestly generated DH keys or KEM keys coincide. There are n_p many long-term DH keys, $n_p \cdot n_{ss}$ many semi-static DH keys, and at most n_s many ephemeral DH keys, so by the birthday bound

¹¹As we will detail in the proof (Game 3) and discuss in Section 6, if PQXDH included the session context (e.g., the values in the session identifier) in the key derivation, as is good practice, this problem would not occur: different KEM public keys or ciphertexts would lead to different session keys (by being part of the KDF input as session context).

the probability for any two DH keys colliding can be upper-bounded by $\frac{(n_p + n_p \cdot n_{ss} + n_s)^2}{q}$. As for KEM keys, there are $n_p \cdot n_{ss}$ semi-static ones and at most n_s ephemeral ones, which collide with probability at most $\gamma_{\text{coll}}(n_p \cdot n_{ss} + n_s)$; cf. Definition 3.1. Put together, we have

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_0}(\mathcal{A}) \leq \frac{(n_p + n_p \cdot n_{ss} + n_s)^2}{q} + \gamma_{\text{coll}}(n_p \cdot n_{ss} + n_s) + \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_1}(\mathcal{A}).$$

Game 2 (KEM correctness). We next let \mathcal{G}_1 abort if any honestly created KEM encapsulation fails to decapsulate correctly. This introduces a KEM correctness error term for each of the n_s many sessions:

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_1}(\mathcal{A}) \leq n_s \cdot \delta_{\text{corr}} + \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_2}(\mathcal{A}).$$

Soundness. At this point soundness holds unconditionally; we consider the three sub-conditions of the sound predicate:

- *Agreement on shared key, type, contributive identifiers* (Figure 6, line 18): The DH keys and the KEM key and the KEM ciphertext in the session identifier determine all inputs to the key derivation function KDF. By \mathcal{G}_2 , KEM ciphertexts decapsulate correctly and both parties derive the same session key. The session identifier further determines the type by epk_B being empty (type = reduced) or not (type = full). Finally, the contributive identifier is a subset of the session identifier, hence matching if the latter match.
- *No initiator session identifiers collide* (Figure 6, line 19): Every initiator samples a fresh ephemeral DH key epk_A^{DH} included in the session identifier. These DH keys do not collide by Game \mathcal{G}_1 . Hence, initiator session identifiers do not collide.
- *No three session identifiers collide in full mode* (Figure 6, line 20): Three colliding session identifiers implies colliding identifiers for two initiator sessions or for two responder sessions. Collisions of initiator session are ruled out above already. Collisions of responder sessions imply colliding ephemeral DH keys epk_B^{DH} , which are ruled out by Game \mathcal{G}_1 .

Game 3 (KEM shared secret collisions). We will need later that two sessions that use distinct (semi-static or ephemeral) KEM public keys or distinct ciphertexts do not end up using the same shared secret ss (and hence possibly the same session key, despite not being partnered).¹² In this game, we abort if any two sessions derive the same shared secret ss while using different KEM public keys (either key may be an ephemeral key in a full session or a semi-static key in a reduced session) or different KEM ciphertexts.

We bound the probability of this abort happening by the advantage of a reduction \mathcal{B}_1 in breaking KCR security of the KEM scheme for $n = n_p \cdot n_{ss} + n_s$. Let \mathcal{B}_1 use the first $n_p \cdot n_{ss}$ challenge keys as semi-static KEM keys in KGenSS instead of generating them itself, and the remaining challenge keys in place of the (at most n_s) ephemeral KEM keys generated by responder sessions (letting the KEM.KGen randomness obtained in the KCR game replace r_2 in the session's random coins). If any two sessions π_i, π_j obtain the same KEM shared secret $ss_i = ss_j$ (as the output of encapsulation for initiator sessions, or as the

¹²If PQXDH included the session context (e.g., the values in the session identifier) in the key derivation, we would not need this game hop, since different KEM public keys or ciphertexts would lead to different session keys. See also the discussion in Section 6.

output of decapsulation for responder sessions) involving two distinct public keys pk_i, pk_j or two distinct ciphertexts ct_i, ct_j , \mathcal{B}_1 outputs (i, ct_i, j, ct_j) and wins. Hence,

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_2}(\mathcal{A}) \leq \epsilon_{\text{LEAK}^+} + \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_3}(\mathcal{A}).$$

Separating the clean cases. At this point, we will divide the proof into six sub-cases following the structure of the $\text{clean}_{\text{full}}$ resp. $\text{clean}_{\text{reduced}}$ predicates evaluated on the test session π^* . We can bound the advantage of the adversary in Game \mathcal{G}_3 by the sum of its advantages in Games $\mathcal{G}_3[\mathbf{c}]$ which are \mathcal{G}_3 conditioned on a sub-predicate \mathbf{c} :

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_3}(\mathcal{A}) \leq \sum_{\mathbf{c} \in \left\{ \begin{array}{l} \text{clean}_{\text{LTSS}}(\pi^*), \text{clean}_{\text{ELT}}(\pi^*), \\ \text{clean}_{\text{ESS}}(\pi^*) \wedge \pi^*.type=\text{full}, \text{clean}_{\text{ESS}}(\pi^*) \wedge \pi^*.type=\text{reduced}, \\ \text{clean}_{\text{EE}}(\pi^*) \wedge \text{clean}_{\text{peerE}}(\pi^*) \wedge \pi^*.type=\text{full}, \\ \text{clean}_{\text{EE}}(\pi^*) \wedge \text{clean}_{\text{sigE}}(\pi^*) \wedge \pi^*.type=\text{full} \end{array} \right\}} \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_3[\mathbf{c}]}(\mathcal{A}).$$

In the remainder of the proof, we will bound each of these cases, numbered **A–F**, separately.

Case A ($\text{clean}_{\text{LTSS}}(\pi^*)$).

In this proof case, the $\text{clean}_{\text{LTSS}}$ predicate ensures for the test session π^* that

1. the initiator’s long-term key is uncompromised,
2. the responder’s semi-static key is uncompromised, and
3. if π^* is an initiator, the responder’s long-term (signing) key was uncompromised upon acceptance.

Via the last point, we can guarantee that initiator and responder indeed agree on the semi-static key (identifier), given signatures are unforgeable. Then, similar to the classical Signal proof [CCD⁺17], the uncompromised long-term/semi-static Diffie–Hellman combination then ensures key indistinguishability for the test session.

Game A.0. This case begins with Game \mathcal{G}_3 conditioned on $\text{clean}_{\text{LTSS}}(\pi^*)$ being satisfied.

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{A.0}}}(\mathcal{A}) = \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_3[\text{clean}_{\text{LTSS}}(\pi^*)]}(\mathcal{A}).$$

Game A.1 (Guess responder identity V^*). We first guess the identity V^* of the responder involved in the test session, overwriting the adversary’s bit guess with 0 if this guess was incorrect. This step loses at most a factor of the number of users n_p :

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{A.0}}}(\mathcal{A}) \leq n_p \cdot \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{A.1}}}(\mathcal{A}).$$

Game A.2 (Signature unforgeability). We now abort the game (again, returning 0 as the adversary’s bit guess) in the event that the test session π^* is an initiator session and accepts using a semi-static DH public key $\text{sspk}_{V^*}^{\text{DH}}$ that was not generated through a KGenSS run for V^* . This ensures that the test session accepts with $\pi^*.sspk = (\text{sspk}_{V^*}^{\text{DH}}, \cdot)$ corresponding to a DH key pair of which the adversary does

not know the secret key, since $\text{clean}_{\text{LTSS}}(\pi^*)$ guarantees that the semi-static secret belong to $\text{sspk}_{V^*}^{\text{DH}}$ is not compromised. The probability of such an abort can be bounded by the advantage of the following reduction \mathcal{B}_2 against the $(t, \epsilon_{\text{SIG}}, 2n_{ss} + n_s)$ -unforgeability of SIG.

The reduction \mathcal{B}_2 samples all key components itself except for the signature key of V^* : In place of the long-term public signature key $\text{ltpk}_{V^*}^{\text{SIG}}$ of V^* it uses the public key pk obtained in its unforgeability game. In its simulation of Game $\mathcal{G}_{A.1}$, \mathcal{B}_2 uses its signing oracle to obtain signatures under $\text{ltsk}_{V^*}^{\text{SIG}}$: two per semi-static key (for DH and KEM public keys) and up to n_s signatures on ephemeral KEM keys. Since $\text{clean}_{\text{LTSS}}(\pi^*) = \text{true}$, we know that $\pi^*. \text{pcorr} = \text{false}$, i.e., the long-term key of V^* was not corrupted when the test session accepted and so \mathcal{B}_2 does not have to answer a $\text{CORRUPTLTKEY}(V^*)$ query prior to the abort event. Hence, \mathcal{B}_2 provides a perfect simulation of Game $\mathcal{G}_{A.1}$ up to when the abort would happen, and if π^* receives a signature σ_{V^*} on a semi-static DH public key that V^* did not generate, then \mathcal{B}_2 can output this as its forgery and wins. Thus,

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{A.1}}(\mathcal{A}) \leq \epsilon_{\text{SIG}} + \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{A.2}}(\mathcal{A}).$$

Game A.3 (Guess initiator identity U^*). We guess the identity U^* of the initiator to the test session, overwriting the adversary's bit guess with 0 if this guess was incorrect. This step again loses at most a factor of the number of users n_p :

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{A.2}}(\mathcal{A}) \leq n_p \cdot \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{A.3}}(\mathcal{A}).$$

Game A.4 (Guess semi-static key identifier ssid^* of V^*). We now guess the identifier ssid of the responder V^* 's (uncorrupted) semi-static DH key $\text{sspk}_{V^*}^{\text{DH}}$. Note that depending on the role of π^* this is either the test session's own key (if $\pi^*. \text{role} = \text{responder}$), or of the intended peer (if $\pi^*. \text{role} = \text{initiator}$). We denote the guessed identifier by ssid^* , and abort, setting the adversary's output bit to 0, if this guess is incorrect, losing at most a factor of the number of semi-static keys per user n_{ss} :

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{A.3}}(\mathcal{A}) \leq n_{ss} \cdot \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{A.4}}(\mathcal{A}).$$

Game A.5 (GapDH). We now abort the game (again, returning 0 as the adversary's bit guess) in the event that the adversary queries the random oracle on a value formatted like a master secret and beginning with $\text{CDH}(\text{ltpk}_{U^*}^{\text{DH}}, \text{sspk}_{V^*}^{\text{DH}})$, i.e., the shared DH key between $\text{ltpk}_{U^*}^{\text{DH}}$ and $\text{sspk}_{V^*}^{\text{DH}}$, where $\text{sspk}_{V^*}^{\text{DH}}$ is the semi-static DH key with identifier ssid^* . The probability of such an abort can be bounded by the advantage of the following reduction \mathcal{B}_3 against the $(t, \epsilon_{\text{GDH}}, q_{\text{DDH}})$ -hardness of the GapDH problem in (\mathbb{G}, g, q) .

Reduction \mathcal{B}_3 samples all key components itself except for replacing the long-term DH key $\text{ltpk}_{U^*}^{\text{DH}}$ of U^* and the semi-static DH key $\text{sspk}_{V^*}^{\text{DH}}$ of V^* with its own challenge values g^a and g^b . Since $\text{clean}_{\text{LTSS}}(\pi^*) = \text{true}$, \mathcal{B}_3 never has to answer the queries $\text{CORRUPTLTKEY}(U^*)$ or $\text{CORRUPTSSKEY}(V^*, \text{ssid}^*)$ in its simulation of Game $\mathcal{G}_{A.4}$. Though, \mathcal{B}_3 may have to answer for REVEALSESSKEY and SEND queries that involve the substituted keys. Hence, the reduction patches the random oracle and the SEND oracle to generate the session keys in a consistent manner. In consequence, the reduction answers queries to the REVEALSESSKEY oracle consistently as well.

Figure 8 gives an algorithmic description of the changes. In particular, the reduction gives special treatment to SEND queries that involve any of the challenge secrets, i.e., for SEND queries where the initiator is the initiator of the test session U^* partnered with the responder of the test session V^* using the semi-static key ssid^* of the test session, the responder is U^* , or the responder is V^* with semi-static key ssid^* and the initiator is not U^* . In these cases the reduction cannot compute all DH shared secrets

RO'(x):

```

1 if  $DH_1 \| DH_2 \| DH_3 \| DH_4 \| ss = x$  // if we can parse the query string  $x$  as three or four DH shared secrets ( $DH_4$  may be empty) and a KEM shared
   secret...
2   if  $DDH(ltpk_{U^*}^{DH}, sspk_{V^*}^{DH}, DH_1)$  // referring to  $V^*$ 's key from  $ssid^*$ 
3      $\mathcal{B}_3$  halts and returns  $DH_1$  as its GapDH solution // change for Game  $\mathcal{G}_{A.5}$ 
4   return  $patch(DH_1, DH_2, DH_3, DH_4, ss)$  // extra routine to ensure consistency
5 return  $RO(x)$ 

```

patch($PoS_1, PoS_2, PoS_3, DH_4, ss$):

```

// the first three arguments may pairs of DH public keys or DH shared secrets
6 foreach  $(e_1, e_2, e_3, e_4, e_5, y)$  in  $L$  // iterate over list entries
7   if  $DDH\text{-}eq(PoS_1, e_1) \wedge DDH\text{-}eq(PoS_2, e_2) \wedge DDH\text{-}eq(PoS_3, e_3) \wedge DH_4 = e_4 \wedge ss = e_5$ 
8     return  $y$  // query matches a prior one, respond with recorded value  $y$ 
9    $y \leftarrow \{0, 1\}^{256}$  // new query; sample value at random
10   $L \leftarrow L \cup \{(PoS_1, PoS_2, PoS_3, DH_4, ss, y)\}$  // record response  $y$  on this query
11 return  $y$ 

```

DDH-eq(PoS_1, PoS_2):

```

12 if  $(pk_1, pk_2) = PoS_1 \wedge \text{if } (pk_3, pk_4) = PoS_2$  // if both  $PoS_1$  and  $PoS_2$  parse as two public keys...
13   return  $\{pk_1, pk_2\} = \{pk_3, pk_4\}$ 
14 if  $(pk_1, pk_2) = PoS_1$  // if  $PoS_1$  parses as two public keys...
15   return  $DDH(pk_1, pk_2, PoS_2)$ 
16 if  $(pk_1, pk_2) = PoS_2$  // if  $PoS_2$  parses as two public keys...
17   return  $DDH(pk_1, pk_2, PoS_1)$ 
18 return  $PoS_1 = PoS_2$  // else, both are DH shared secrets

```

SEND(W, i, m) substitutes KDF call in Run where GapDH challenge keys are involved:

```

19 ...
20  $U \leftarrow \pi_W^i.\text{initiator}$     $V \leftarrow \pi_W^i.\text{responder}$     $ssid \leftarrow \pi_W^i.\text{ssid}$ 
21  $(\dots, epk_U^{DH}, \dots) \leftarrow \pi_W^i.\text{sid}$ 
22 if  $U = U^* \wedge V = V^* \wedge ssid = ssid^*$ 
23   compute  $DH_2, DH_4, ss$  with the corresponding secret keys as specified in Run
24    $\pi_W^i.K \leftarrow patch((ltpk_{U^*}^{DH}, sspk_{V^*}^{DH}), DH_2, (epk_U^{DH}, sspk_{V^*}^{DH}), DH_4, ss)$ 
25 else if  $V = U^*$ 
26   compute  $DH_1, DH_3, DH_4, ss$  with the corresponding secret keys as specified in Run
27    $\pi_W^i.K \leftarrow patch(DH_1, (epk_U^{DH}, ltpk_{U^*}^{DH}), DH_3, DH_4, ss)$ 
28 else if  $U \neq U^* \wedge V = V^* \wedge ssid = ssid^*$ 
29   compute  $DH_1, DH_2, DH_4, ss$  with the corresponding secret keys as specified in Run
30    $\pi_W^i.K \leftarrow patch(DH_1, DH_2, (epk_U^{DH}, sspk_{V^*}^{DH}), DH_4, ss)$ 
31 else //  $\mathcal{B}_3$  can compute all secrets
32   compute  $DH_1, DH_2, DH_3, DH_4, ss$  with the corresponding secret keys as specified in Run
33    $\pi_W^i.K \leftarrow patch(DH_1, DH_2, DH_3, DH_4, ss)$ 
34 ...

```

Figure 8: Algorithmic description of how \mathcal{B}_3 simulates the random oracle and the SEND oracle in game $\mathcal{G}_{A.5}$. To ensure consistency of queries we use an additional (global) list L , which is initialized as empty. This list L contains six-tuples consisting of three pairs of DH public keys or DH shared secrets (PoS), a DH shared secret, a KEM shared secret, and the associated RO output. The `patch` routine checks if incoming queries match list entries with the `DDH-eq` routine: `DDH-eq` returns true if both PoS argument are the same two public keys, if the two PoS arguments are a valid DH tuple according to the DDH oracle, or if the two PoS arguments are identical DH shared secrets. We patch the SEND oracle to use the `patch` routine in lieu of computing the KDF, whenever there are challenge secret keys involved that \mathcal{B}_3 does not know.

entering the key derivation function and chooses a session key via the `patch` routine. Specifically, the `patch` routine takes as arguments three pairs of DH public keys or DH shared secrets PoS (since DH_1, DH_2, DH_3 are the DH shared secrets that the reduction can potentially not compute), a DH shared secret, and a KEM shared secret, and returns a RO output that is consistent with previous queries. To this end, `patch` saves all queries together with their RO response in a list L and checks new queries against L . Note here that two PoS are equivalent if the pair of public keys in one of them forms a valid Diffie–Hellman triple with the shared secret in the other one according to the DDH oracle or if the PoS components are equivalent, as described by the DDH-eq routine. If a query does not match any previous query, the `patch` routine randomly chooses a new session key and records the query with the response in L .

Since the SEND oracle computes consistent session keys for each session, \mathcal{B}_3 also consistently answers to the REVEALSESSKEY oracle queries. To ensure consistency with queries to the random oracle, the reduction also uses the `patch` routine for RO queries. Finally, in case the adversary queries the random oracle on a value beginning with $CDH(ltpk_{U^*}^{DH}, sspk_{V^*}^{DH})$, the reduction \mathcal{B}_3 halts and returns this value to its own challenger.

Note that for each RO query the reduction makes up to $1 + 2|L| \leq 3|L|$ queries to the DDH oracle (line 2 and up to two times per entry in line 7). For each query to the SEND oracle,¹³ the reduction calls the `patch` routine once, which makes up to $3|L|$ queries to the DDH oracle. Each call to `patch` adds at most one entry to L , so $|L|$ grows from 1 to $q_{RO} + q_{Snd}$. Using the Gaussian sum we get

$$q_{DDH} \leq \sum_{i=1}^{q_{RO} + q_{Snd}} 3 \cdot (i - 1) = 3 \cdot \frac{(q_{RO} + q_{Snd} - 1) \cdot (q_{RO} + q_{Snd})}{2} \leq 1.5(q_{RO} + q_{Snd})^2.$$

Unless \mathcal{A} queries the random oracle on $CDH(ltpk_{U^*}^{DH}, sspk_{V^*}^{DH})$ in the DH_1 position, \mathcal{B}_3 provides a perfect simulation of game $\mathcal{G}_{A.4}$. If \mathcal{A} makes such a random oracle query, then \mathcal{B}_3 will detect this and win its GapDH game. Thus,

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{A.4}}(\mathcal{A}) \leq \epsilon_{\text{GDH}} + \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{A.5}}(\mathcal{A}).$$

Game A.6 (Replacing the session key). We now replace the session key of the test session π^* with a uniformly sampled key. Since Game $\mathcal{G}_{A.5}$ has ruled out that the adversary queries the RO on the master secret $DH_1 \| DH_2 \| DH_3 \| DH_4 \| ss$ of the test session π^* , the adversary has no chance of detecting this change. Thus,

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{A.5}}(\mathcal{A}) = \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{A.6}}(\mathcal{A}).$$

To conclude this proof case, observe that in Game $\mathcal{G}_{A.6}$ the session key of the test session π^* is now a uniformly random key, independent of b_{test} . Furthermore, \mathcal{A} cannot reveal the session key of the test session π^* via a REVEALSESSKEY query on π^* or any partnered session which might hold the same key. Finally, any non-partnered session will derive a different session key: any difference in the identities A, B yields different DH public keys by Game \mathcal{G}_1 , any difference in the DH shares implies a difference in the DH[1]–DH[4] inputs, while differing KEM public keys or ciphertexts yield a different ss input by Game \mathcal{G}_3 . Thus, \mathcal{A} cannot gain any information about the test bit b_{test} and can do no better than to guess:

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{A.6}}(\mathcal{A}) \leq 0.$$

¹³If PQXDH would include the session context (e.g., the values in the session identifier) in the key derivation, as is good practice, the reduction would need less DDH oracle queries: Using the context, the reduction can detect when a challenge DH key is combined with a maliciously generated (ephemeral) key. In the modified SEND oracle in Figure 8, only the first PoS in line 24 remains a pair of public keys, while the other three occurrences are solved via context. In consequence, at maximum PoS per entry in L contains two public keys and we get a maximum of $q_{RO} + q_{Snd}$ DDH queries.

Case B ($\text{clean}_{\text{ELT}}(\pi^*)$).

In this proof case, the $\text{clean}_{\text{ELT}}$ predicate ensures for the test session π^* that

1. the initiator's randomness is not revealed, and
2. the responder's long-term key is uncompromised.

Similar to the classical Signal proof [CCD⁺17], the uncompromised ephemeral/long-term Diffie–Hellman combination ensures key indistinguishability for the test session.

Game B.0. This case begins with Game \mathcal{G}_3 conditioned on $\text{clean}_{\text{ELT}}(\pi^*)$ being satisfied.

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{B.0}}}(\mathcal{A}) = \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_3[\text{clean}_{\text{ELT}}(\pi^*)]}(\mathcal{A}).$$

Game B.1 (Guess initiator session). We guess the initiator session π_i^* contributing to the test session π^* (i.e., either the test session itself if it is an initiator session, or the initiator session partnered to the test session), overwriting the adversary's bit guess with 0 if this guess was incorrect. This step loses at most a factor of the number of sessions n_s :

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{B.0}}}(\mathcal{A}) \leq n_s \cdot \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{B.1}}}(\mathcal{A}).$$

Game B.2 (Guess responder identity V^*). We guess the identity V^* of the responder involved in the test session, overwriting the adversary's bit guess with 0 if this guess was incorrect. This step loses at most a factor of the number of users n_p :

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{B.1}}}(\mathcal{A}) \leq n_p \cdot \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{B.2}}}(\mathcal{A}).$$

Game B.3 (GapDH). We now abort the game (again, returning 0 as the adversary's bit guess) in the event that the adversary queries the random oracle on a value formatted like a master secret and beginning with $\text{CDH}(epk^{\text{DH}}, ltpk_{V^*}^{\text{DH}})$, i.e., the shared DH key between the ephemeral key of the test session's initiator epk^{DH} and $ltpk_{V^*}^{\text{DH}}$. The probability of such an abort can be bounded by the advantage of the following reduction \mathcal{B}_4 against the $(t, \epsilon_{\text{GDH}}, q_{\text{DDH}})$ -hardness of the GapDH problem in (\mathbb{G}, g, q) .

The reduction follows the idea of Game $\mathcal{G}_{\text{A.5}}$, only that now we embed the GapDH challenge keys as the ephemeral key in the initiator session π_i^* and as the long-term key of V^* . The reduction does not need to answer corresponding REVEALRAND queries on π_i^* or CORRUPTLTKEY queries on V^* due to $\text{clean}_{\text{ELT}}$. It can use a similar strategy to the one in Game $\mathcal{G}_{\text{A.5}}$ to patch the random oracle, yet it has an easier time ensuring consistency of the SEND oracle: It can always compute DH_1 and DH_3 by using the responder's semi-static secret key, and DH_4 (if it is a full handshake) with the honest party's ephemeral secret key. Outside of the test session \mathcal{B}_4 can use the initiator's ephemeral secret key to compute DH_2 if the ephemeral key was honestly generated. Hence, \mathcal{B}_4 can compute all inputs to the KDF for all SEND queries itself, except for the SEND queries to V^* with a malicious ephemeral key on the initiator side. In consequence, \mathcal{B}_4 queries the DDH oracle at most once per SEND query and once when checking RO queries for the DH_2 position. All in all, \mathcal{B}_4 needs a maximum of $q_{\text{Snd}} + q_{\text{RO}}$ queries to the DDH oracle, which is well below the number in Game $\mathcal{G}_{\text{A.5}}$.

Unless \mathcal{A} queries the random oracle on $\text{CDH}(epk^{\text{DH}}, ltpk_{V^*}^{\text{DH}})$ in the DH_2 position, \mathcal{B}_4 provides a perfect simulation of $\mathcal{G}_{\text{B.2}}$. If \mathcal{A} does make such a random oracle query, then \mathcal{B}_4 will detect this and win its GapDH game. Thus,

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{B.2}}}(\mathcal{A}) \leq \epsilon_{\text{GDH}} + \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{B.3}}}(\mathcal{A}).$$

Game B.4 (Replacing the session key). We now replace the session key of the test session π^* with a uniformly sampled key. Since Game $\mathcal{G}_{B.3}$ has ruled out that the adversary queries the RO on the master secret $DH_1 \| DH_2 \| DH_3 \| DH_4 \| ss$ of the test session π^* , the adversary has no chance of detecting this change. Thus,

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{B.3}}(\mathcal{A}) = \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{B.4}}(\mathcal{A}).$$

To conclude this proof case, observe that in Game $\mathcal{G}_{B.4}$ the session key of the test session π^* is now a uniformly random key, independent of b_{test} . As before, \mathcal{A} can neither reveal the session key of the test session π^* nor of any partnered session, nor does any other session derive the same session key by Game \mathcal{G}_3 . Thus, \mathcal{A} cannot gain any information about the test bit b_{test} and can do no better than to guess:

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{B.4}}(\mathcal{A}) \leq 0.$$

Case C ($\text{clean}_{\text{ESS}}(\pi^*) \wedge \pi^*.type = \text{full}$).

In this proof case, the test session performs a full handshake and the $\text{clean}_{\text{ESS}}$ predicate ensures for the test session π^* that

1. the initiator's randomness is not revealed, and
2. the responder's semi-static key is uncompromised, and
3. if π^* is an initiator, the responder's long-term (signing) key was uncompromised upon acceptance.

Via the last point, similarly to Case A, we can guarantee that initiator and responder indeed agree on the semi-static key (identifier), given signatures are unforgeable. Then, similar to the classical Signal proof [CCD⁺17], the uncompromised ephemeral/semi-static Diffie–Hellman combination ensures key indistinguishability for the test session.

Game C.0. This case begins with Game \mathcal{G}_3 conditioned on $\text{clean}_{\text{ESS}}(\pi^*) \wedge \pi^*.type = \text{full}$ being satisfied.

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{C.0}}(\mathcal{A}) = \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_3[\text{clean}_{\text{ESS}}(\pi^*) \wedge \pi^*.type = \text{full}]}(\mathcal{A}).$$

Game C.1 (Guess responder identity V^*). We first guess the identity V^* of the responder to the test session, overwriting the adversary's bit guess with 0 if this guess was incorrect. This step loses at most a factor of the number of users n_p :

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{C.0}}(\mathcal{A}) \leq n_p \cdot \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{C.1}}(\mathcal{A}).$$

Game C.2 (Signature unforgeability). We now abort the game (again, returning 0 as the adversary's bit guess) in the event that the test session π^* is an initiator session and accepts using a semi-static DH public key $ssp k_{V^*}^{\text{DH}}$ that was not generated through a KGenSS run for V^* . This ensures that the test session accepts with $\pi^*.sspk s = (ssp k_{V^*}^{\text{DH}}, \cdot)$ corresponding to a DH key pair of which the adversary does not know the secret key, since $\text{clean}_{\text{ESS}}(\pi^*)$ guarantees that the semi-static secret belong to $ssp k_{V^*}^{\text{DH}}$ is not compromised. Similar to Game $\mathcal{G}_{A.2}$, the probability of such an abort can be bounded by the advantage of a reduction against the $(t, \epsilon_{\text{SIG}}, 2n_{ss} + n_s)$ -unforgeability of SIG . The argument is identical to the Game $\mathcal{G}_{A.2}$. Thus,

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{C.1}}(\mathcal{A}) \leq \epsilon_{\text{SIG}} + \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{C.2}}(\mathcal{A}).$$

Game C.3 (Guess semi-static key identifier ssid^* of V^*). We now guess the identifier ssid of the responder V^* 's (uncorrupted) semi-static key $\text{sspk}_{V^*}^{\text{DH}}$. Note that depending on the role of π^* this is either the test session's own key (if $\pi^*.\text{role} = \text{responder}$), or of the intended peer (if $\pi^*.\text{role} = \text{initiator}$). We denote the guessed identifier by ssid^* , and abort, setting the adversary's output bit to 0, if this guess is incorrect, losing at most a factor n_{ss} of the number of semi-static keys per user:

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{C.2}}}(\mathcal{A}) \leq n_{ss} \cdot \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{C.3}}}(\mathcal{A}).$$

Game C.4 (Guess initiator session). We guess the initiator session π_i^* involved in the test session π^* (i.e., either the test session itself if it is an initiator session, or the initiator session partnered to the test session), overwriting the adversary's bit guess with 0 if this guess was incorrect. This step loses at most a factor of the number of sessions n_s :

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{C.3}}}(\mathcal{A}) \leq n_s \cdot \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{C.4}}}(\mathcal{A}).$$

Game C.5 (GapDH). We now abort the game (again, returning 0 as the adversary's bit guess) in the event that the adversary queries the random oracle on a value formatted like a master secret, beginning with $\text{CDH}(\text{epk}^{\text{DH}}, \text{sspk}_{V^*}^{\text{DH}})$, i.e., the shared DH key between the ephemeral key of the test session's initiator epk^{DH} and $\text{sspk}_{V^*}^{\text{DH}}$, where $\text{sspk}_{V^*}^{\text{DH}}$ is the semi-static DH key with identifier ssid^* . The probability of such an abort can be bounded by the advantage of the following reduction \mathcal{B}_5 against the $(t, \epsilon_{\text{GDH}}, q_{\text{DDH}})$ -hardness of the GapDH problem in (\mathbb{G}, g, q) .

The proof follows the idea of Games $\mathcal{G}_{\text{A.5}}$ and $\mathcal{G}_{\text{B.3}}$. Though, now we embed the GapDH challenge keys as the ephemeral key in the initiator session π_i^* and as semi-static key of V^* with semi-static id ssid^* . The reduction does not need to answer corresponding REVEALRAND queries on π_i^* or CORRUPTSSKEY(V^*, ssid^*) queries due to $\text{clean}_{\text{ESS}}$. It can use a similar strategy to patch the random oracle: It can always compute DH_1 and DH_2 by using the initiator's and responder's long-term secret key, respectively, and DH_4 with the honest party's ephemeral secret key. Outside of the test session \mathcal{B}_5 can use the initiator's ephemeral secret key to compute DH_3 if the ephemeral key was honestly generated. Hence, \mathcal{B}_5 can compute all inputs to the KDF for all SEND queries itself, except for the SEND queries to the test session or to V^* with ssid^* and a malicious ephemeral key on the initiator side. In consequence, \mathcal{B}_5 queries the DDH oracle at most once per SEND query and once when checking RO queries for the DH_3 position. All in all, \mathcal{B}_5 needs a maximum of $q_{\text{Snd}} + q_{\text{RO}}$ queries to the DDH oracle, which is well below the number in Game $\mathcal{G}_{\text{A.5}}$.

Unless \mathcal{A} queries the random oracle on $\text{CDH}(\text{epk}^{\text{DH}}, \text{ltpk}_{V^*}^{\text{DH}})$ in the DH_2 position, \mathcal{B}_5 provides a perfect simulation of $\mathcal{G}_{\text{C.4}}$. If \mathcal{A} does make such a random oracle query, then \mathcal{B}_5 will detect this and win its GapDH game. Thus,

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{C.4}}}(\mathcal{A}) \leq \epsilon_{\text{GDH}} + \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{C.5}}}(\mathcal{A}).$$

Game C.6 (Replacing the session key). We now replace the session key of the test session π^* with a uniformly sampled key. Since Game $\mathcal{G}_{\text{C.5}}$ has ruled out that the adversary queries the RO on the master secret $DH_1 \| DH_2 \| DH_3 \| DH_4 \| ss$ of the test session π^* , the adversary has no chance of detecting this change. Thus,

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{C.5}}}(\mathcal{A}) = \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{C.6}}}(\mathcal{A}).$$

To conclude this proof case, observe that in Game $\mathcal{G}_{\text{C.6}}$ the session key of the test session π^* is now a uniformly random key, independent of b_{test} . As before, \mathcal{A} can neither reveal the session key of the test session π^* nor of any partnered session, nor does any other session derive the same session key by Game \mathcal{G}_3 . Thus, \mathcal{A} cannot gain any information about the test bit b_{test} and can do no better than to guess:

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{C.6}}}(\mathcal{A}) \leq 0.$$

Case D ($\text{clean}_{\text{ESS}}(\pi^*) \wedge \pi^*.type = \text{reduced}$).

In this proof case, the test session performs a reduced handshake and the $\text{clean}_{\text{ESS}}$ predicate ensures for the test session π^* that

1. the initiator's randomness is not revealed, and
2. the responder's semi-static key is uncompromised, and
3. if π^* is an initiator, the responder's long-term (signing) key was uncompromised upon acceptance.

Via the last point, we can guarantee that initiator and responder indeed agree on the semi-static key (identifier), given signatures are unforgeable. The difference to Case C is that, in a reduced handshake, the semi-static KEM public key of the responder is used (and there is no ephemeral/ephemeral Diffie–Hellman combination). This means we get the following *hybrid* guarantee: both the ephemeral/semi-static Diffie–Hellman combination being uncompromised *and* the semi-static KEM key being uncompromised are, on their own, sufficient to ensure key indistinguishability for the test session.

Game D.0. This case begins with Game \mathcal{G}_3 conditioned on $\text{clean}_{\text{ESS}}(\pi^*) \wedge \pi^*.type = \text{reduced}$ being satisfied.

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{D.0}}}(\mathcal{A}) = \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_3[\text{clean}_{\text{ESS}}(\pi^*) \wedge \pi^*.type=\text{reduced}]}(\mathcal{A}).$$

Game D.1 (Guess responder identity V^*). We first guess the identity V^* of the responder to the test session, overwriting the adversary's bit guess with 0 if this guess was incorrect. This step loses at most a factor of the number of users n_p :

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{D.0}}}(\mathcal{A}) \leq n_p \cdot \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{D.1}}}(\mathcal{A}).$$

Game D.2 (Signature unforgeability). In this game case, we are interested in the authenticity of *both* semi-static DH and KEM keys, and hence now abort the game (again, returning 0 as the adversary's bit guess) in the event that the test session π^* is an initiator session and accepts using a semi-static DH public key $\text{sspk}_{V^*}^{\text{DH}}$ or a semi-static KEM public key $\text{sspk}_{V^*}^{\text{KEM}}$ that was not generated through a KGenSS run for V^* . Since for reduced handshakes both public keys are recorded in $\pi^*.\text{sspks} = (\text{sspk}_{V^*}^{\text{DH}}, \text{sspk}_{V^*}^{\text{KEM}})$, this ensures that the test session accepts with semi-static DH and KEM public keys of which the adversary does not know the secret, since $\text{clean}_{\text{ESS}}(\pi^*)$ guarantees that the semi-static secrets corresponding to $\text{sspk}_{V^*}^{\text{DH}}$ or to $\text{sspk}_{V^*}^{\text{KEM}}$ are not compromised.

Again similar to Game $\mathcal{G}_{\text{A.2}}$, we can reduce this hop to the $(t, \epsilon_{\text{SIG}}, 2n_{ss} + n_s)$ -unforgeability of SIG:

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{D.1}}}(\mathcal{A}) \leq \epsilon_{\text{SIG}} + \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{D.2}}}(\mathcal{A}).$$

Game D.3 (Guess semi-static key identifiers ssid_1^* , ssid_2^* of V^*). We now guess the ssid identifiers of the responder V^* 's (uncompromised) semi-static DH and KEM keys, $\text{sspk}_{V^*}^{\text{DH}}$ resp. $\text{sspk}_{V^*}^{\text{KEM}}$. Note that depending on the role of π^* this is either the test session's own key (if $\pi^*.role = \text{responder}$), or of the intended peer (if $\pi^*.role = \text{initiator}$). We denote the guessed identifiers by ssid_1^* (for the DH key) and ssid_2^* (for the KEM key), and abort, setting the adversary's output bit to 0, if this guess is incorrect, losing at most a factor n_{ss} of the number of semi-static keys of each type per user, for each of the two guesses:

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{D.2}}}(\mathcal{A}) \leq n_{ss}^2 \cdot \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{D.3}}}(\mathcal{A}).$$

Game D.4 (Guess initiator session). We guess the initiator session π_i^* involved in the test session π^* (i.e., either the test session itself if it is an initiator session, or the initiator session partnered to the test session), overwriting the adversary’s bit guess with 0 if this guess was incorrect. This step loses at most a factor of the number of sessions n_s :

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{D.3}}}(\mathcal{A}) \leq n_s \cdot \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{D.4}}}(\mathcal{A}).$$

Game D.5 (GapDH + OW-CCA). We now abort the game (again, returning 0 as the adversary’s bit guess) in the event that the adversary queries the random oracle on a value formatted like a master secret, beginning with $\text{CDH}(epk^{\text{DH}}, sspk_{V^*}^{\text{DH}})$, i.e., the shared DH key between the ephemeral key of the test session’s initiator epk^{DH} and $sspk_{V^*}^{\text{DH}}$, and ending with the KEM shared secret ss^* resulting from π_i^* encapsulating against $sspk_{V^*}^{\text{KEM}}$, where $sspk_{V^*}^{\text{DH}}$ and $sspk_{V^*}^{\text{KEM}}$ are the semi-static DH and KEM keys in $sspk_{V^*}^{\text{ssid}^*}$. The probability of such an abort can be bounded by the *minimum* of the advantages of the following reduction \mathcal{B}_6 playing *simultaneously* against the $(t, \epsilon_{\text{GDH}}, q_{\text{DDH}})$ -hardness of the GapDH problem in (\mathbb{G}, g, q) and $(t, \epsilon_{\text{CCA}}, n_s)$ -OW-CCA security of KEM.

The reduction embeds the GapDH challenge as the ephemeral DH share in the initiator session π_i^* and as semi-static DH key of V^* with semi-static id ssid_1^* (as in Game $\mathcal{G}_{\text{C.5}}$), and the KEM public key pk_{KEM} from the OW-CCA game as the semi-static key $sspk_{V^*}^{\text{KEM}}$ of V^* with semi-static id ssid_2^* . Furthermore, it uses the KEM challenge ciphertext ct^* in π_i^* . Similar to Game $\mathcal{G}_{\text{C.5}}$, the reduction does not need to answer corresponding REVEALRAND on π_i^* or CORRUPTSSKEY($V^*, \text{ssid}_1^*/\text{ssid}_2^*$) queries due to $\text{clean}_{\text{ESS}}$. Furthermore, as for Game $\mathcal{G}_{\text{C.5}}$, patching the SEND oracle is relatively easy: Except for queries to the test session or to V^* using $sspk_{V^*}^{\text{DH}}$ and a malicious ephemeral DH key on the initiator side, the reduction can compute all DH shared secrets with the DH secret keys and it can decapsulate the KEM shared secret with the KEM secret key or, for sessions using $sspk_{V^*}^{\text{KEM}}$, the DECAPS oracle of the OW-CCA game. Otherwise (and in particular in the test session), the reduction uses a freshly sampled key as session key. If the DDH oracle accepts a query, then the reduction returns the DH_3 and ss^* from the corresponding RO query to its GapDH and OW-CCA games after finishing the simulation for \mathcal{A} .

The reduction ensures consistency with the random oracle by checking RO queries against the above mentioned patches. In consequence, the remaining queries to the DDH oracle are once per SEND query and when checking RO queries for the DH_3 position. All in all, \mathcal{B}_6 needs a maximum of $q_{\text{Snd}} + q_{\text{RO}}$ queries to the DDH oracle, which is well below the number in $\mathcal{G}_{\text{A.5}}$; it further makes at most n_s many DECAPS query in the OW-CCA game.

Unless \mathcal{A} queries the random oracle on $\text{CDH}(epk^{\text{DH}}, sspk_{V^*}^{\text{DH}})$ in the DH_3 position *and* ss^* in the ss position, \mathcal{B}_6 provides a perfect simulation of $\mathcal{G}_{\text{D.4}}$. If \mathcal{A} does make such a random oracle query, then \mathcal{B}_6 will detect this and win *both* its GapDH game and its OW-CCA game. We can thus bound the game hop by the *minimum* advantage against the two games:

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{D.4}}}(\mathcal{A}) \leq \min(\epsilon_{\text{GDH}}, \epsilon_{\text{CCA}}) + \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{D.5}}}(\mathcal{A}).$$

Game D.6 (Replacing the session key). We now replace the session key of the test session π^* with a uniformly sampled key. Since Game $\mathcal{G}_{\text{D.5}}$ has ruled out that the adversary queries the RO on the master secret $DH_1 \| DH_2 \| DH_3 \| DH_4 \| ss$ of the test session π^* , the adversary has no chance of detecting this change. Thus,

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{D.5}}}(\mathcal{A}) = \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{\text{D.6}}}(\mathcal{A}).$$

To conclude this proof case, observe that in Game $\mathcal{G}_{\text{D.6}}$ the session key of the test session π^* is now a uniformly random key, independent of b_{test} . As before, \mathcal{A} can neither reveal the session key of the test

session π^* nor of any partnered session, nor does any other session derive the same session key by Game \mathcal{G}_3 . Thus, \mathcal{A} cannot gain any information about the test bit b_{test} and can do no better than to guess:

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{D.6}}(\mathcal{A}) \leq 0.$$

Case E ($\text{clean}_{\text{EE}}(\pi^*) \wedge \text{clean}_{\text{peerE}}(\pi^*) \wedge \text{type}(\pi^*) = \text{full}$).

In this proof case, the test session performs a full handshake and the clean_{EE} predicate and its $\text{clean}_{\text{peerE}}$ sub-predicate ensure for the test session π^* that

1. the test session's owner's randomness is not revealed, and
2. the randomness of the session (contributively) partnered to the test session is not revealed.

Similar to Case D, we get a *hybrid* guarantee here: both the ephemeral/ephemeral Diffie–Hellman combination being uncompromised *and* the ephemeral KEM key being uncompromised are, on their own, sufficient to ensure key indistinguishability for the test session.

Game E.0. This case begins with Game \mathcal{G}_3 conditioned on $\text{clean}_{\text{EE}}(\pi^*) \wedge \text{clean}_{\text{peerE}}(\pi^*) \wedge \pi^*.\text{type} = \text{full}$ being satisfied.

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{E.0}}(\mathcal{A}) = \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_3[\text{clean}_{\text{EE}}(\pi^*) \wedge \text{clean}_{\text{peerE}}(\pi^*) \wedge \pi^*.\text{type}=\text{full}]}(\mathcal{A}).$$

Game E.1 (Guess initiator and responder sessions). We guess the initiator and responder sessions π_i^* resp. π_r^* involved in the test session π^* (i.e., both the test session itself and its *sid-* resp. *cid-*partner), overwriting the adversary's bit guess with 0 if this guess was incorrect. This step loses at most a factor of the number of sessions n_s squared:

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{E.0}}(\mathcal{A}) \leq n_s^2 \cdot \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{E.1}}(\mathcal{A}).$$

Game E.2 (GapDH + OW-CCA). We now abort the game (again, returning 0 as the adversary's bit guess) in the event that the adversary queries the random oracle on a value formatted like a master secret, beginning with $\text{CDH}(epk_{U^*}^{\text{DH}}, epk_{V^*}^{\text{DH}})$, i.e., the shared DH key between the ephemeral keys $epk_{U^*}^{\text{DH}}$ and $epk_{V^*}^{\text{DH}}$ of the guessed initiator session π_i^* , resp. responder session π_r^* , and ending with the KEM shared secret ss^* resulting from π_i^* encapsulating against $epk_{V^*}^{\text{KEM}}$. The probability of such an abort can be bounded by the *minimum* of the advantages of the following reduction \mathcal{B}_7 playing *simultaneously* against the $(t, \epsilon_{\text{GDH}}, q_{\text{DDH}})$ -hardness of the GapDH problem in (\mathbb{G}, g, q) and $(t, \epsilon_{\text{CCA}}, n_s)$ -OW-CCA security of KEM.

The reduction embeds the GapDH challenge as the ephemeral keys in the sessions π_i^* and π_r^* , and the KEM public key pk_{KEM} from the OW-CCA game as ephemeral key $epk_{V^*}^{\text{KEM}}$ in π_r^* . It uses the KEM challenge ciphertext ct^* in π_i^* . Similar to Game $\mathcal{G}_{D.5}$, the reduction does not need to answer corresponding REVEALRAND queries on π_i^* or π_r^* due to $\text{clean}_{\text{EE}} \wedge \text{clean}_{\text{peerE}}$. Furthermore, as for Game $\mathcal{G}_{D.5}$, patching the SEND oracle is relatively easy: Except for queries to the test session or to π_r^* with a modified initiator ephemeral key,¹⁴ the reduction can compute all DH shared secrets itself. For the latter query, \mathcal{B}_7 uses a single call to the DDH oracle of GapDH resp. DECAPS oracle of OW-CCA to simulate. In the test session, the reduction uses a freshly sampled key as session key. If the DDH oracle accepts a query, then the

¹⁴This can happen if the test session is a responder session and the adversary mauls the initiator public key on the wire.

reduction returns the DH_3 and ss from the corresponding RO query to its GapDH and OW-CCA games after finishing the simulation for \mathcal{A} .

The reduction ensures consistency with the random oracle by checking RO queries against the above mentioned patches. It queries the DDH oracle possibly once for π_r^* and when checking RO queries for the DH_4 position, so overall a maximum of $q_{\text{RO}} + 1$ queries, which is well below the number in Game $\mathcal{G}_{A.5}$. It further makes at most one DECAPS query in the OW-CCA game.

Unless \mathcal{A} queries the random oracle on $\text{CDH}(epk_{U^*}^{\text{DH}}, epk_{V^*}^{\text{DH}})$ in the DH_4 position and ss^* in the ss position, \mathcal{B}_7 provides a perfect simulation of $\mathcal{G}_{E.1}$. If \mathcal{A} does make such a random oracle query, then \mathcal{B}_7 will detect this and win *both* its GapDH game and its OW-CCA game. We can thus bound the game hop by the *minimum* advantage against the two games:

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{E.1}}(\mathcal{A}) \leq \min(\epsilon_{\text{GDH}}, \epsilon_{\text{CCA}}) + \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{E.2}}(\mathcal{A}).$$

Game E.3 (Replacing the session key). We now replace the session key of the test session π^* with a uniformly sampled key. Since Game $\mathcal{G}_{E.2}$ has ruled out that the adversary queries the RO on the master secret $DH_1 \| DH_2 \| DH_3 \| DH_4 \| ss$ of the test session π^* , the adversary has no chance of detecting this change. Thus,

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{E.2}}(\mathcal{A}) = \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{E.3}}(\mathcal{A}).$$

To conclude this proof case, observe that in Game $\mathcal{G}_{E.3}$ the session key of the test session π^* is now a uniformly random key, independent of b_{test} . As before, \mathcal{A} can neither reveal the session key of the test session π^* nor of any partnered session, nor does any other session derive the same session key by Game \mathcal{G}_3 . Thus, \mathcal{A} cannot gain any information about the test bit b_{test} and can do no better than to guess:

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{E.3}}(\mathcal{A}) \leq 0.$$

Case F ($\text{clean}_{\text{EE}}(\pi^*) \wedge \text{clean}_{\text{sigE}}(\pi^*) \wedge \text{type}(\pi^*) = \text{full}$).

In this proof case, the test session performs a full handshake and the clean_{EE} predicate and its $\text{clean}_{\text{sigE}}$ sub-predicate ensure for the test session π^* that

1. the test session's owner's randomness is not revealed, and
2. if π^* is an initiator, the responder's long-term (signing) key was uncompromised upon acceptance and a potential responder partner session's randomness was not revealed.

Via the second point, we can guarantee that initiator and responder agree on the ephemeral KEM key, given signatures are unforgeable. The KEM encapsulation then ensures key indistinguishability for the test session.

Game F.0. This case begins with Game \mathcal{G}_3 conditioned on $\text{clean}_{\text{EE}}(\pi^*) \wedge \text{clean}_{\text{sigE}}(\pi^*) \wedge \text{type}(\pi^*) = \text{full}$ being satisfied.

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{F.0}}(\mathcal{A}) = \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_3[\text{clean}_{\text{EE}}(\pi^*) \wedge \text{clean}_{\text{sigE}}(\pi^*) \wedge \text{type}(\pi^*) = \text{full}]}(\mathcal{A}).$$

Game F.1 (Guess responder identity V^*). We first guess the identity V^* of the responder to the test session, overwriting the adversary's bit guess with 0 if this guess was incorrect. This step loses at most a factor of the number of users n_p :

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{F.0}}(\mathcal{A}) \leq n_p \cdot \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{F.1}}(\mathcal{A}).$$

Game F.2 (Signature unforgeability). We now abort the game (again, returning 0 as the adversary’s bit guess) in the event that the test session π^* is an initiator session and accepts using an ephemeral KEM public key $epk_{V^*}^{\text{KEM}}$ that was not generated by V^* . This ensures that the test session accepts with an ephemeral KEM public key of which the adversary does not know the secret key, since the adversary cannot reveal V^* ’s session randomness. Similar to Game $\mathcal{G}_{A.2}$, the probability of such an abort can be bounded by the advantage of a reduction against the $(t, \epsilon_{\text{SIG}}, 2n_{ss} + n_s)$ -unforgeability of SIG. The proof is identical to the one for Game $\mathcal{G}_{A.2}$. Thus,

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{F.1}}(\mathcal{A}) \leq \epsilon_{\text{SIG}} + \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{F.2}}(\mathcal{A}).$$

Game F.3 (Guess initiator and responder sessions). We guess the initiator and responder sessions π_i^* resp. π_r^* involved in the test session π^* (i.e., both the test session itself and its partner), overwriting the adversary’s bit guess with 0 if this guess was incorrect. This step loses at most a factor of the number of sessions n_s squared:

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{F.2}}(\mathcal{A}) \leq n_s^2 \cdot \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{F.3}}(\mathcal{A}).$$

Game F.4 (OW-CCA). We now abort the game (again, returning 0 as the adversary’s bit guess) in the event that the adversary queries the random oracle on a value formatted like a master secret and ending with the KEM shared secret ss^* resulting from π_i^* encapsulating against $epk_{V^*}^{\text{KEM}}$. The probability of such an abort can be bounded by the advantage of the following reduction \mathcal{B}_8 playing against $(t, \epsilon_{\text{CCA}}, n_s)$ -OW-CCA security of KEM.

The reduction embeds the challenge pk_{KEM}, ct^* from the OW-CCA game as ephemeral key $epk_{V^*}^{\text{KEM}}$ of V^* , and as KEM ciphertext in π_r^* . The reduction does not need to answer a REVEALRAND query on π_r^* due to clean_{EE} . When \mathcal{A} halts, the reduction guesses $i \in [q_{\text{RO}}]$ and returns the ss position of the i th RO query (assuming this query is formatted like a master secret) as the target shared secret.

Unless \mathcal{A} queries the random oracle on ss^* in the ss position, \mathcal{B}_8 provides a perfect simulation of Game $\mathcal{G}_{F.3}$. If \mathcal{A} does make such a random oracle query, then \mathcal{B}_8 wins its OW-CCA game if it guesses the i th RO query correctly. We can thus bound the game hop by the advantage against OW-CCA times a loss of q_{RO} :

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{F.3}}(\mathcal{A}) \leq q_{\text{RO}} \cdot \epsilon_{\text{CCA}} + \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{F.4}}(\mathcal{A}).$$

Game F.5 (Replacing the session key). We now replace the session key of the test session π^* with a uniformly sampled key. Since Game $\mathcal{G}_{F.4}$ has ruled out that the adversary queries the RO on the master secret $DH_1 \| DH_2 \| DH_3 \| DH_4 \| ss$ of the test session π^* , the adversary has no chance of detecting this change. Thus,

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{F.4}}(\mathcal{A}) = \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{F.5}}(\mathcal{A}).$$

To conclude this proof case, observe that in Game $\mathcal{G}_{F.5}$ the session key of the test session π^* is now a uniformly random key, independent of b_{test} . As before, \mathcal{A} can neither reveal the session key of the test session π^* nor of any partnered session, nor does any other session derive the same session key by Game \mathcal{G}_3 . Thus, \mathcal{A} cannot gain any information about the test bit b_{test} and can do no better than to guess:

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{F.5}}(\mathcal{A}) \leq 0.$$

Collecting the bounds from all cases yields the overall theorem bound. \square

6 Discussion and Conclusion

In this work, we provided a reductionist security analysis of Signal’s PQXDH in a “maximum-exposure”, game-based security model, augmented compared to prior versions [CCD⁺17, BFG⁺22, BFG⁺21] to capture the KEM component but also key signing, and gave a fully-parameterized, concrete security bound for PQXDH.

Our bound relies (among other things) on the $LEAK^+$ -BIND- $SS-\{CT, PK\}$ binding property of the KEM. While we show that this property is satisfied by both Kyber and ML-KEM, the current and future KEMs in PQXDH may not satisfy this property; so achieving PQXDH-like security without relying on a binding property is of general interest. We can indeed forgo this assumption (and the corresponding advantage term), if in the key derivation of PQXDH, we include the KEM public key and ciphertext (or, ideally, the whole session context as is good practice). This supports and complements a proposal discussed in the tool-based formal verification of PQXDH by BJKS [BJK23, BJKS23], to make sure the key agreement in the initial handshake does not rely on the follow-up AEAD encryption (which would indeed violate key indistinguishability as we prove it). As a side note, including context into the key derivation also supports tighter security proofs [CCG⁺19].

Our reduction for PQXDH has an additional case (compared to the X3DH analysis [CCD⁺17]), where key indistinguishability hinges solely on the signed ephemeral KEM key. This guarantees security against an *active* adversary who later gets quantum powers. Hence, PQXDH protects against an even stronger class of attacks than the “harvest now, decrypt later” attack which motivated the design.

Like prior Signal analyses [CCD⁺17, BJKS23, BJK23], we model the long-term DH and signing keys as separate ones; we leave it to future work to more accurately reflect the implementation which re-uses the same keys for both purposes.

Acknowledgments

We thank Ehren Kret and Rolfe Schmidt for discussions on the PQXDH protocol, and Franziskus Kiefer for insights on the KEM re-encapsulation attack. 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.

References

- [AHK⁺23] Joël Alwen, Dominik Hartmann, Eike Kiltz, Marta Mularczyk, and Peter Schwabe. Post-quantum multi-recipient public key encryption. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, CCS ’23*, page 1108–1122, New York, NY, USA, 2023. Association for Computing Machinery. (Cited on page 3.)
- [BC] Bruno Blanchet and Vincent Cheval. Proverif: Cryptographic protocol verifier in the formal model. <https://bblanche.gitlabpages.inria.fr/proverif/>. (Cited on page 2.)
- [BCD⁺24] Manuel Barbosa, Deirdre Connolly, João Diogo Duarte, Aaron Kaiser, Peter Schwabe, Karolin Varner, and Bas Westerbaan. X-wing. *IACR Communications in Cryptology*, 1(1), 2024. (Cited on page 3.)
- [BDPV08] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. On the indistinguishability of the sponge construction. In Nigel P. Smart, editor, *Advances in Cryptology – EUROCRYPT 2008*, volume 4965 of *Lecture Notes in Computer Science*, pages 181–197, Istanbul, Turkey, April 13–17, 2008. Springer, Heidelberg, Germany. (Cited on page 5.)

- [BFG⁺20] Jacqueline Brendel, Marc Fischlin, Felix Günther, Christian Janson, and Douglas Stebila. Towards post-quantum security for Signal’s X3DH handshake. In Orr Dunkelman, Michael J. Jacobson Jr., and Colin O’Flynn, editors, *SAC 2020: 27th Annual International Workshop on Selected Areas in Cryptography*, volume 12804 of *Lecture Notes in Computer Science*, pages 404–430, Halifax, NS, Canada (Virtual Event), October 21–23, 2020. Springer, Heidelberg, Germany. (Cited on page 3.)
- [BFG⁺21] Jacqueline Brendel, Rune Fiedler, Felix Günther, Christian Janson, and Douglas Stebila. Post-quantum asynchronous deniable key exchange and the Signal handshake. *Cryptology ePrint Archive*, Report 2021/769, 2021. <https://eprint.iacr.org/2021/769>. (Cited on pages 2, 15, 16, 17, 18, 19, and 37.)
- [BFG⁺22] Jacqueline Brendel, Rune Fiedler, Felix Günther, Christian Janson, and Douglas Stebila. Post-quantum asynchronous deniable key exchange and the Signal handshake. In Goichiro Hanaoka, Junji Shikata, and Yohei Watanabe, editors, *PKC 2022: 25th International Conference on Theory and Practice of Public Key Cryptography, Part II*, volume 13178 of *Lecture Notes in Computer Science*, pages 3–34, Virtual Event, March 8–11, 2022. Springer, Heidelberg, Germany. (Cited on pages 2, 3, 15, 16, 17, 18, 19, and 37.)
- [BJK23] Karthikeyan Barghavan, Charlie Jacomme, and Franziskus Kiefer. Formal analysis of the PQXDH protocol, 2023. <https://github.com/Inria-Prosecco/pqxdh-analysis>. (Cited on pages 2, 3, and 37.)
- [BJKS23] Karthikeyan Barghavan, Charlie Jacomme, Franziskus Kiefer, and Rolfe Schmidt. An analysis of Signal’s PQXDH, October 2023. <https://cryspen.com/post/pqxdh/>. (Cited on pages 2, 3, and 37.)
- [Bla] Bruno Blanchet. Cryptoverif: Cryptographic protocol verifier in the computational model. <https://bblanche.gitlabpages.inria.fr/CryptoVerif/>. (Cited on page 2.)
- [BR94] Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In Douglas R. Stinson, editor, *Advances in Cryptology – CRYPTO’93*, volume 773 of *Lecture Notes in Computer Science*, pages 232–249, Santa Barbara, CA, USA, August 22–26, 1994. Springer, Heidelberg, Germany. (Cited on page 15.)
- [CCD⁺17] Katriel Cohn-Gordon, Cas J. F. Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the Signal messaging protocol. In *IEEE European Symposium on Security and Privacy, EuroS&P 2017*, pages 451–466, 2017. (Cited on pages 1, 2, 5, 15, 17, 18, 21, 22, 25, 29, 30, and 37.)
- [CCG⁺19] Katriel Cohn-Gordon, Cas Cremers, Kristian Gjøsteen, Håkon Jacobsen, and Tibor Jager. Highly efficient key exchange protocols with optimal tightness. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part III*, volume 11694 of *Lecture Notes in Computer Science*, pages 767–797, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Heidelberg, Germany. (Cited on page 37.)
- [CD23] Wouter Castryck and Thomas Decru. An efficient key recovery attack on SIDH. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023, Part V*, volume 14008 of *Lecture Notes in Computer Science*, pages 423–447, Lyon, France, April 23–27, 2023. Springer, Heidelberg, Germany. (Cited on page 3.)

- [CDM23] Cas Cremers, Alexander Dax, and Niklas Medinger. Keeping up with the KEMs: Stronger security notions for KEMs. Cryptology ePrint Archive, Paper 2023/1933, 2023. Version 1.0.6 (April 3, 2024), <https://eprint.iacr.org/2023/1933>. (Cited on pages 3, 7, and 8.)
- [CHDN⁺24] Daniel Collins, Loïs Huguenin-Dumittan, Ngoc Khanh Nguyen, Nicolas Rolin, and Serge Vaudenay. K-Waay: Fast and deniable Post-Quantum X3DH without ring signatures. In *33rd USENIX Security Symposium (USENIX Security 24)*, Philadelphia, PA, August 2024. USENIX Association. (Cited on page 3.)
- [DFGS15] Benjamin Dowling, Marc Fischlin, Felix Günther, and Douglas Stebila. A cryptographic analysis of the TLS 1.3 handshake protocol candidates. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM CCS 2015: 22nd Conference on Computer and Communications Security*, pages 1197–1210, Denver, CO, USA, October 12–16, 2015. ACM Press. (Cited on page 16.)
- [DG22] Samuel Dobson and Steven D. Galbraith. Post-quantum signal key agreement from SIDH. In Jung Hee Cheon and Thomas Johansson, editors, *Post-Quantum Cryptography - 13th International Workshop, PQCrypto 2022, Virtual Event, September 28-30, 2022, Proceedings*, volume 13512 of *Lecture Notes in Computer Science*, pages 422–450. Springer, 2022. (Cited on page 3.)
- [FG14] Marc Fischlin and Felix Günther. Multi-stage key exchange and the case of Google’s QUIC protocol. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *ACM CCS 2014: 21st Conference on Computer and Communications Security*, pages 1193–1204, Scottsdale, AZ, USA, November 3–7, 2014. ACM Press. (Cited on page 15.)
- [FJ24] Rune Fiedler and Christian Janson. A deniability analysis of Signal’s initial handshake PQXDH. *Proc. Priv. Enhancing Technol.*, 2024(4), 2024. (Cited on page 2.)
- [FO99] Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In Michael J. Wiener, editor, *Advances in Cryptology – CRYPTO’99*, volume 1666 of *Lecture Notes in Computer Science*, pages 537–554, Santa Barbara, CA, USA, August 15–19, 1999. Springer, Heidelberg, Germany. (Cited on page 10.)
- [HKKP22] Keitaro Hashimoto, Shuichi Katsumata, Kris Kwiatkowski, and Thomas Prest. An efficient and generic construction for Signal’s handshake (X3DH): Post-quantum, state leakage secure, and deniable. *Journal of Cryptology*, 35(3):17, July 2022. (Cited on page 3.)
- [JK18] Tibor Jager and Rafael Kurek. Short digital signatures and ID-KEMs via truncation collision resistance. In Thomas Peyrin and Steven Galbraith, editors, *Advances in Cryptology – ASIACRYPT 2018, Part II*, volume 11273 of *Lecture Notes in Computer Science*, pages 221–250, Brisbane, Queensland, Australia, December 2–6, 2018. Springer, Heidelberg, Germany. (Cited on page 5.)
- [JKN21] Tibor Jager, Rafael Kurek, and David Niehues. Efficient adaptively-secure IB-KEMs and VRFs via near-collision resistance. In Juan Garay, editor, *PKC 2021: 24th International Conference on Theory and Practice of Public Key Cryptography, Part I*, volume 12710 of *Lecture Notes in Computer Science*, pages 596–626, Virtual Event, May 10–13, 2021. Springer, Heidelberg, Germany. (Cited on page 5.)
- [KS23] Ehren Kret and Rolfe Schmidt. The PQXDH key agreement protocol, September 2023. Revision 1, <https://signal.org/docs/specifications/pqxdh/>. (Cited on pages 1 and 2.)

- [KS24] Ehren Kret and Rolfe Schmidt. The PQXDH key agreement protocol, January 2024. Revision 3, <https://signal.org/docs/specifications/pqxdh/>. (Cited on pages 1, 2, 3, and 21.)
- [LT11] Gaëtan Leurent and Søren S. Thomsen. Practical near-collisions on the compression function of BMW. In Antoine Joux, editor, *Fast Software Encryption – FSE 2011*, volume 6733 of *Lecture Notes in Computer Science*, pages 238–251, Lyngby, Denmark, February 13–16, 2011. Springer, Heidelberg, Germany. (Cited on page 5.)
- [MMP⁺23] Luciano Maino, Chloe Martindale, Lorenz Panny, Giacomo Pope, and Benjamin Wesolowski. A direct key recovery attack on SIDH. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023, Part V*, volume 14008 of *Lecture Notes in Computer Science*, pages 448–471, Lyon, France, April 23–27, 2023. Springer, Heidelberg, Germany. (Cited on page 3.)
- [MP16] Moxie Marlinspike and Trevor Perrin. The X3DH key agreement protocol, November 2016. <https://signal.org/docs/specifications/x3dh/>. (Cited on page 1.)
- [MvV97] Alfred J. Menezes, Paul C. van Oorschot, and Scott A. Vanstone. *Handbook of Applied Cryptography*. The CRC Press series on discrete mathematics and its applications. CRC Press, 2000 N.W. Corporate Blvd., Boca Raton, FL 33431-9868, USA, 1997. (Cited on page 5.)
- [NIS23] NIST. Module-lattice-based key-encapsulation mechanism standard, August 2023. FIPS 203 (draft). <https://doi.org/10.6028/NIST.FIPS.203.ipd>. (Cited on pages 3, 6, 7, and 10.)
- [OP01] Tatsuaki Okamoto and David Pointcheval. The gap-problems: A new class of problems for the security of cryptographic schemes. In Kwangjo Kim, editor, *PKC 2001: 4th International Workshop on Theory and Practice in Public Key Cryptography*, volume 1992 of *Lecture Notes in Computer Science*, pages 104–118, Cheju Island, South Korea, February 13–15, 2001. Springer, Heidelberg, Germany. (Cited on page 5.)
- [Per16] Trevor Perrin. The XEdDSA and VXEdDSA signature schemes, October 2016. <https://signal.org/docs/specifications/xeddsa/>. (Cited on page 21.)
- [PS14] Inna Polak and Adi Shamir. Using random error correcting codes in near-collision attacks on generic hash-functions. In Willi Meier and Debdeep Mukhopadhyay, editors, *Progress in Cryptology - INDOCRYPT 2014: 15th International Conference in Cryptology in India*, volume 8885 of *Lecture Notes in Computer Science*, pages 219–236, New Delhi, India, December 14–17, 2014. Springer, Heidelberg, Germany. (Cited on page 5.)
- [Rob23] Damien Robert. Breaking SIDH in polynomial time. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023, Part V*, volume 14008 of *Lecture Notes in Computer Science*, pages 472–503, Lyon, France, April 23–27, 2023. Springer, Heidelberg, Germany. (Cited on page 3.)
- [SAB⁺] Peter Schwabe, Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Gregor Seiler, Damien Stehlé, and Jintai Ding. CRYSTALS-Kyber. <https://pq-crystals.org/kyber/>. (Cited on pages 3, 6, 7, and 10.)
- [Sch24] Sophie Schmieg. Unbindable kemmy schmidt: ML-KEM is neither MAL-BIND-K-CT nor MAL-BIND-K-PK. Cryptology ePrint Archive, Paper 2024/523, 2024. <https://eprint.iacr.org/2024/523>. (Cited on pages 3, 7, and 8.)

- [Sig] Signal: Technical information. <https://signal.org/docs/>. (Cited on page 1.)
- [VGIK20] Nihal Vatandas, Rosario Gennaro, Bertrand Ithurburn, and Hugo Krawczyk. On the cryptographic deniability of the Signal protocol. In Mauro Conti, Jianying Zhou, Emiliano Casalicchio, and Angelo Spognardi, editors, *ACNS 20: 18th International Conference on Applied Cryptography and Network Security, Part II*, volume 12147 of *Lecture Notes in Computer Science*, pages 188–209, Rome, Italy, October 19–22, 2020. Springer, Heidelberg, Germany. (Cited on page 2.)
- [WY05] Xiaoyun Wang and Hongbo Yu. How to break MD5 and other hash functions. In Ronald Cramer, editor, *Advances in Cryptology – EUROCRYPT 2005*, volume 3494 of *Lecture Notes in Computer Science*, pages 19–35, Aarhus, Denmark, May 22–26, 2005. Springer, Heidelberg, Germany. (Cited on page 5.)
- [YCW14] Hongbo Yu, Jiazhe Chen, and Xiaoyun Wang. Partial-collision attack on the round-reduced compression function of Skein-256. In Shiho Moriai, editor, *Fast Software Encryption – FSE 2013*, volume 8424 of *Lecture Notes in Computer Science*, pages 263–283, Singapore, March 11–13, 2014. Springer, Heidelberg, Germany. (Cited on page 5.)