

Security Analysis of Signal’s PQXDH Handshake*

Version 2.1[†]

Rune Fiedler¹

Felix Günther²

¹ Cryptoplexity, Technische Universität Darmstadt, Darmstadt, Germany

² IBM Research Europe – Zurich, Rüschlikon, 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 against both classical and quantum adversaries, 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 fully parameterized, concrete security bounds for the classical and post-quantum session key security of PQXDH, and discuss how design choices in PQXDH make a KEM binding property necessary and how a lack of domain separation reduces the achievable security.

Our discussion of KEM binding and domain separation complements the concurrent tool-based analysis of PQXDH by Bhargavan, Jacomme, Kiefer, and Schmidt (USENIX Security 2024), which pointed out a potential re-encapsulation attack if the KEM shared secret does not bind the public key. In contrast to the tool-based analysis, we analyze all protocol modes of PQXDH and its “maximum-exposure” security. We further show that both Kyber (used in PQXDH) and the NIST standard ML-KEM (expected to replace Kyber) 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 [MP16b] as initial handshake protocol. In 2023, Signal released PQXDH [KS23], 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.

*A preliminary version of this paper appears in the proceedings of the *28th International Conference on Practice and Theory of Public-Key Cryptography (PKC 2025)*, © IACR 2025. DOI: [10.1007/978-3-031-91823-0_5](https://doi.org/10.1007/978-3-031-91823-0_5). This is the full version.

[†]We provide an overview of major changes between versions in Appendix A.

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 [KS23] adds a post-quantum 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.

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], subsequently (and concurrently to this work) expanded and published in [BJKS24]. 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 the BJKS analysis [BJKS24] of PQXDH and the CCDGS analysis [CCD⁺17] for X3DH is that while the tool-based analysis of BJKS provide machine-checked assurance, their ProVerif results remain on the symbolic, protocol-logic level assuming perfect cryptography and their 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 focuses on a reduced set of simpler properties and, for example, only captures the compromise of long-term user keys (but not of semi-static and ephemeral keys), does not model the full security contributed by certain key combinations, and separately considers attacks only against the classical DH components and only against the post-quantum KEM component. Such fine-grained compromise, however, is a main reason for combining multiple secrets in Signal’s X3DH and PQXDH handshakes, and combined classical/post-quantum hybrid guarantees a main reason behind the PQXDH design specifically. Furthermore, BJKS only model the semi-static KEM keys in PQXDH but not the ephemeral ones, sidestepping a domain separation issue in the protocol that allows signed ephemeral and semi-static KEM public keys to be confused. Hence, there remains a notable gap between the security guarantees formally confirmed so far and those aimed at by PQXDH.

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].

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

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 components in PQXDH. In particular, we model that the KEM public key and the semi-static DH shares are signed in PQXDH (instead of assuming them to be authentically distributed like in 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. Similarly, we specifically capture the key confusion of ephemeral and semi-static KEM public keys possible in PQXDH and when we can still expect security, enabling us to capture the ephemeral KEM public keys omitted in the model of BJKS [BJKS24].

Concrete security bound for PQXDH. We then analyze the PQXDH handshake with respect to its security against both classical and quantum adversaries. We establish classical 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 (EUF-CMA), 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 analyses in BJKS [BJKS24]. Against quantum adversaries, we establish security in the standard model, based on the KEM providing indistinguishability under chosen-ciphertext attacks (IND-CCA) and the same binding and correctness properties as before, and the key derivation function being a secure PRF when keyed by the KEM shared secret. Our resulting security theorems for PQXDH are 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 BJKS [BJKS24]. For example, our model and proofs pinpoint the stronger security achievable if PQXDH were to properly domain-separate ephemeral and semi-static KEM public keys. Furthermore, we argue how the key derivation approach taken by Signal’s X3DH and PQXDH leads to requiring a certain *binding* property from the KEM in our analysis, and how this requirement could be easily avoided.

For background, BJKS in their analysis discovered the following potential “KEM re-encapsulation attack”, a class of attacks introduced by Cremers, Dax, and Medinger (CDM) [CDM24], 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. As a 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 initial BJKS analysis [BJKS23, BJK23] excluded this attack by modeling that the KEM public key is added to the associated data when AEAD-encrypting messages, a suggestion that the PQXDH description however only follows optionally: it argues that Kyber [SAB⁺], 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. Capturing binding under revealed randomness is necessary to treat the “maximum-exposure” security of PQXDH. We prove that both Kyber [SAB⁺] (currently deployed in PQXDH) and the new NIST standard ML-KEM [Nat24] (provisioned to replace Kyber in PQXDH) satisfy this binding property.

²Security modeling still always requires trade-offs in abstraction, see our discussion in Section 6.

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 CDM [CDM24], yet their notion that would fit the PQXDH setting is *not* satisfied by ML-KEM [Sch24]. Furthermore, BJKS expanded their initial analysis in [BJKS24], published after this work, to also include a CryptoVerif analysis assuming a KEM binding property and proving it for Kyber. We discuss and relate the binding notion we introduce and require in our PQXDH analysis to the framework of CDM and to the notion used by BJKS in more detail in Section 3. 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 [CHN⁺24] propose K-Waay, based on adapting split KEMs [BFG⁺20]. In this work, we focus on Signal’s deployed PQXDH protocol.

2 Preliminaries

2.1 Notation

We refer to the i th element of a list or vector l with $l[i]$. 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 Pseudorandom Functions

We recall the definition of pseudorandom functions.

Definition 2.1. Let $\text{PRF}: \{0, 1\}^{l_k} \times \{0, 1\}^{l_i} \rightarrow \{0, 1\}^{l_o}$ be an efficient keyed function with key length l_k , input length l_i , and output length l_o . We say that PRF is $(t, \epsilon, q_{\text{PRF}})$ -PRF-secure, if for any adversary \mathcal{A} against $\mathcal{G}_{\text{PRF}}^{\text{prfsec}}(\mathcal{A})$ defined in Figure 1 with running time at most t and making at most q_{PRF} queries to the PRF oracle, we have that

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

2.3 Signatures

Definition 2.2 (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 .

$\mathcal{G}_{\text{PRF}}^{\text{prfsec}}(\mathcal{A}):$ 1 $K \leftarrow_{\$} \{0, 1\}^{l_k}$ 2 $g \leftarrow_{\$} \{\text{functions } f: \{0, 1\}^{l_i} \rightarrow \{0, 1\}^{l_o}\}$ 3 $b \leftarrow_{\$} \{0, 1\}$ 4 $b' \leftarrow_{\$} \mathcal{A}^{\text{PRF}}()$ 5 return $\llbracket b' = b \rrbracket$	$\text{PRF}(x):$ 6 if $b = 0$ 7 return $\text{PRF}(K, x)$ 8 else 9 return $g(x)$
---	--

Figure 1: Security game for PRF security for a function PRF defined in Definition 2.1.

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

Figure 2: Security game for EUF-CMA for a signature scheme $\text{SIG} = (\text{KGen}, \text{Sig}, \text{Vf})$ defined in Definition 2.3.

- $\text{Vf}(pk, m, \sigma) \rightarrow d$: 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 2.3 (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 2 with running time at most t and making at most q_{SIG} queries to the SIGN oracle, we have that*

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

2.4 Hash Functions

We follow the “human ignorance” approach of Rogaway [Rog06] to capture practical, unkeyed hash functions, giving explicit constructions for adversaries against, e.g., the collision resistance of a hash function in our results.

Definition 2.4 (Preimage resistance of hash functions). *Let $H: \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a (hash) function. We define the advantage of an adversary \mathcal{A} against the preimage resistance of H , distinguishing whether challenges are sampled from (a subset of) the domain $\mathcal{M} \subseteq \{0, 1\}^*$ (PR-d) or from the range $\{0, 1\}^n$ (PR-r), as*

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

Definition 2.5 (Collision resistance of hash functions). *Let $H: \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a (hash) function. We define the advantage of an adversary \mathcal{A} against the collision resistance (CR) of H as*

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

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 [MvOV97, 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} .

$$\begin{array}{ll}
\mathcal{G}_{(\mathbb{G},g,q)}^{\text{GDH}}(\mathcal{A}): & \text{DDH}(g^x, g^y, g^z): \\
1 \ a, b \leftarrow_{\$} \mathbb{Z}_q & 4 \ \mathbf{return} \ [\text{CDH}(g^x, g^y) = g^z] \\
2 \ g^z \leftarrow_{\$} \mathcal{A}^{\text{DDH}}((\mathbb{G}, g, q), g^a, g^b) & \\
3 \ \mathbf{return} \ [g^z = g^{ab}] &
\end{array}$$

Figure 3: Security game for the GapDH problem defined in Definition 2.7.

SHA3. Bertoni, Daemen, Peters, and Van Assche [BDPV08] show that SHA3 is indifferentiable from a random oracle when the underlying permutation is modeled as a random permutation. Under this assumption, 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.

2.5 Diffie–Hellman Key Exchange

Definition 2.6 (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 2.7 (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 3.

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:*

- $\text{KGen}() \xrightarrow{\$} (pk, sk)$: *The probabilistic key generation with randomness space $\mathcal{R}_{\text{KEM.KGen}}$ outputs a public-key/secret-key pair with $(pk, sk) \in \mathcal{PK} \times \mathcal{SK}$.*

$\mathcal{G}_{\text{KEM}}^{\text{IND-CCA}}(\mathcal{A}):$ 1 $b \leftarrow_{\$} \{0, 1\}$ 2 $(pk, sk) \leftarrow_{\$} \text{KEM.KGen}()$ 3 $(ct^*, ss_0^*) \leftarrow_{\$} \text{KEM.Enc}(pk)$ 4 $ss_1^* \leftarrow_{\$} \mathcal{SS}$ 5 $b' \leftarrow_{\$} \mathcal{A}^{\text{DECAPS}}(pk, ct^*, ss_b^*)$ 6 return $\llbracket b' = b \rrbracket$	$\mathcal{G}_{\text{KEM}}^{\text{OW-CCA}}(\mathcal{A}):$ 7 $(pk, sk) \leftarrow_{\$} \text{KGen}()$ 8 $(ct^*, ss^*) \leftarrow_{\$} \text{Enc}(pk)$ 9 $ss' \leftarrow_{\$} \mathcal{A}^{\text{DECAPS}}(pk, ct^*)$ 10 return $\llbracket ss' = ss^* \rrbracket$	$\text{DECAPS}(ct):$ 11 if $ct = ct^*$ 12 return \perp 13 else 14 return $\text{Dec}(sk, ct)$
---	---	---

Figure 4: Security games for IND-CCA security (left) and OW-CCA security (middle) of $\text{KEM} = (\text{KGen}, \text{Enc}, \text{Dec})$ defined in Definition 3.3 and Definition 3.4, respectively.

- $\text{Enc}(pk) \rightarrow (ct, ss)$: 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 $(pk, sk) \leftarrow_{\$} \text{KGen}()$, and every encapsulation $(ct, ss) \leftarrow_{\$} \text{Enc}(pk)$, 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 [Nat24], 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:

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

In our analysis, we require *indistinguishability under chosen-ciphertext attacks* (IND-CCA) against quantum adversaries and *one-way security under chosen-ciphertext attacks* (OW-CCA) against classical adversaries of the KEM scheme, where the former notion implies the latter.

Definition 3.3 (IND-CCA Security of KEMs). Let $\text{KEM} = (\text{KGen}, \text{Enc}, \text{Dec})$ be a KEM. We say that KEM is $(t, \epsilon, q_{\text{Dec}})$ -IND-CCA-secure, if for any adversary \mathcal{A} against Game $\mathcal{G}_{\text{KEM}}^{\text{IND-CCA}}(\mathcal{A})$ defined in Figure 4 with running time at most t and making at most q_{Dec} queries to the DECAPS oracle, we have that

$$\text{Adv}_{\text{KEM}}^{\text{IND-CCA}}(\mathcal{A}) := \left| \Pr[\mathcal{G}_{\text{KEM}}^{\text{IND-CCA}}(\mathcal{A})] - \frac{1}{2} \right| \leq \epsilon.$$

Definition 3.4 (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 $\mathcal{G}_{\text{KEM}}^{\text{OW-CCA}}(\mathcal{A})$ defined in Figure 4 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. Revealing also the key generation randomness to the adversary is necessary to treat the “maximum-exposure” security of PQXDH, in which an adversary can 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*.

Aligning with language of Cremers, Dax, and Medinger (CDM) [CDM24] 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 $LEAK^{+r}$ -BIND-SS- $\{CT, PK\}$. In general, their notion X -BIND- P - 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 ($X = MAL$), honestly generated and leaked to the adversary ($X = LEAK$), or honestly generated without leakage of secrets ($X = HON$). We reproduce their $LEAK$ and MAL notions for ss simultaneously binding pk and ct in our syntax³ on the right-hand side of Figure 5.

In the following, we formalize our new $LEAK^{+r}$ -BIND-SS- $\{CT, PK\}$ binding property and relate it to the notions in the CDM framework [CDM24]. In particular, we show in Section 3.3 below that both Kyber and the NIST standard ML-KEM satisfy the $LEAK^{+r}$ -BIND-SS- $\{CT, PK\}$ property, while the notion from [CDM24] that implies ours, MAL -BIND-SS- $\{CT, PK\}$, is not satisfied by ML-KEM [Sch24]. While Signal’s current implementation of PQXDH uses Kyber [SAB⁺], it is already prepared⁴ to transition to the new NIST standard ML-KEM [Nat24].

Definition 3.5 ($LEAK^{+r}$ -BIND-SS- $\{CT, PK\}$). *Let $KEM = (KGen, Enc, Dec)$ be a KEM and $n \geq 2$. We say that KEM is (t, ϵ, n) - $LEAK^{+r}$ -BIND-SS- $\{CT, PK\}$ -secure, if for any adversary \mathcal{A} against Game $\mathcal{G}_{KEM, n}^{LEAK^{+r}\text{-BIND-SS-}\{CT, PK\}}$ defined in Figure 5 with running time at most t , we have that*

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

Figure 6 visualizes the relations between $LEAK^{+r}$ -BIND-SS- $\{CT, PK\}$ and the corresponding notions from [CDM24]: The MAL variant is strictly stronger than the $LEAK^{+r}$ variant, which in turn is strictly stronger than the $LEAK$ variant. Notably, ML-KEM separates the MAL and our $LEAK^{+r}$ variants; we show the other implications and separations below, incl. relating our multi-user notion to the 2-user case.

Relation to SH-CR [BJKS24]. The tool-based analysis in BJKS [BJKS24] introduces a similar KEM property to the binding notions we discuss here; dubbed “semi-honest collision resistance” (SH-CR): There, one party decapsulates a malicious ciphertext for an honestly generated but compromised key pair, and another party encapsulates against an adversarially-chosen public key. The adversary wins if the shared secrets from decapsulation and encapsulation match and either the ciphertexts or the public keys differ.

³Note that in the syntax of [CDM24] 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 MAL -BIND-SS- CT corresponding to their notion MAL -BIND- K - CT .

⁴<https://github.com/signalapp/libsignal/commit/0670f0d>




Figure 5: Security games for our (multi-user) $\text{LEAK}^{+r}\text{-BIND-SS-}\{CT,PK\}$ notion (left) and the $\text{LEAK-BIND-SS-}\{CT,PK\}$ and $\text{MAL-BIND-SS-}\{CT,PK\}$ notion from [CDM24] (right). In all games, the adversary’s goal is to produce colliding shared secrets ss under distinct ciphertexts ct or public keys pk .




Figure 6: Relations between our $\text{LEAK}^{+r}\text{-BIND-SS-}\{CT,PK\}$ binding notion and the corresponding MAL , LEAK notions from [CDM24] for KEMs. Solid arrows indicate implications, dashed ones loose implications, and crossed-out ones separations. Annotations indicate the respective theorems.

Intuitively, this notion is incomparable to ours, since it allows malicious key pairs but does not give the key generation randomness to the adversary. More formally, we show that a KEM can fulfill SH-CR but not $\text{LEAK}^{+r}\text{-BIND-SS-}\{CT,PK\}$, and vice-versa:

- $\text{LEAK}^{+r}\text{-BIND-SS-}\{CT,PK\}$ but not SH-CR: Given a $\text{LEAK}^{+r}\text{-BIND-SS-}\{CT,PK\}$ -secure scheme, modify key generation to add a constant 0 bit to the secret key and decapsulation returns an extra bitstring appended to the secret key as shared secret if the extra bit in the secret key is set to 1. This change does not affect $\text{LEAK}^{+r}\text{-BIND-SS-}\{CT,PK\}$ since the special case is never triggered for honestly generated key pairs, but it allows breaking SH-CR.
- SH-CR but not $\text{LEAK}^{+r}\text{-BIND-SS-}\{CT,PK\}$: given a SH-CR-secure scheme, modify key generation to add a trapdoor in the secret key (as we will do for the separation of $\text{LEAK-BIND-SS-}\{CT,PK\}$ and $\text{LEAK}^{+r}\text{-BIND-SS-}\{CT,PK\}$ in Theorem 3.9 in Section 3.2), for which decapsulation returns a special shared secret. This does not help the SH-CR adversary, who does not learn the key generation randomness, but it allows breaking $\text{LEAK}^{+r}\text{-BIND-SS-}\{CT,PK\}$ as in Theorem 3.9.

3.2 Relations of Binding Properties

First, we show the implication from MAL to LEAK^{+r} (Theorem 3.6) and from LEAK^{+r} to LEAK (Theorem 3.7). Second, we show that our notion is equivalent for n and 2 users, up to a factor of n^2

(Theorem 3.8). Third, we give a separation between $LEAK$ and $LEAK^{+r}$ (Theorem 3.9). Finally, we show that ML-KEM and Kyber are $LEAK^{+r}$ (Theorem 3.10 and Theorem 3.11), while ML-KEM is known to not be MAL [Sch24], giving us a separation for $LEAK^{+r}$ and MAL .

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

Theorem 3.6 ($MAL\text{-}BIND\text{-}SS\text{-}\{CT, PK\} \implies LEAK^{+r}\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$). *If a KEM \mathcal{KEM} is (t', ϵ') - $MAL\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure, then it is also (t, ϵ, n) - $LEAK^{+r}\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure with $t \approx t' + n$ and $\epsilon \leq \epsilon'$.*

Proof. An adversary \mathcal{A} against $LEAK^{+r}\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ 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 the adversary \mathcal{B} against $MAL\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ truthfully simulate $\mathcal{G}_{\mathcal{KEM}, n}^{LEAK^{+r}\text{-}BIND\text{-}SS\text{-}\{CT, PK\}}(\mathcal{A})$ 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, i.e., $\epsilon \leq \epsilon'$. The adversary \mathcal{B} runs KGen n times and \mathcal{A} once, resulting in $t \approx t' + n$. \square

Theorem 3.7 ($LEAK^{+r}\text{-}BIND\text{-}SS\text{-}\{CT, PK\} \implies LEAK\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$). *If a KEM \mathcal{KEM} is (t', ϵ', n) - $LEAK^{+r}\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure, then it is also (t, ϵ) - $LEAK\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure with $t \approx t'$ and $\epsilon \leq \epsilon'$ for any n .*

Proof. The adversary \mathcal{B} against $LEAK^{+r}\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ runs the adversary \mathcal{A} against $LEAK\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ on the first two key pairs (excluding key generation randomness) of its own input, 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} , i.e., $\epsilon \leq \epsilon'$. The adversary \mathcal{B} runs \mathcal{A} once, resulting in $t \approx t'$. \square

Theorem 3.8 ($LEAK^{+r}\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$: n users \iff 2 users). *Let $n \geq 2$. If a KEM \mathcal{KEM} is (t, ϵ, n) - $LEAK^{+r}\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure, then it is also $(t, \epsilon, 2)$ - $LEAK^{+r}\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure, and if a KEM \mathcal{KEM} is $(t', \epsilon', 2)$ - $LEAK^{+r}\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure, then it is also $(t, n^2 \cdot \epsilon', n)$ - $LEAK^{+r}\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure for $t \approx t' + n$.*

Proof. The “ \implies ” direction straightforwardly holds by letting the adversary against the n -user variant ignore the remaining $n - 2$ keys.

The “ \impliedby ” direction holds via a guessing argument: Let the adversary \mathcal{B} against the 2-user variant 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 the adversary \mathcal{A} against the n -user variant uses i (and possibly j) for its attack, \mathcal{B} is successful if \mathcal{A} is. The adversary \mathcal{B} runs \mathcal{A} , guesses two indices, and runs KGen $n - 2$ times, resulting in $t \approx t' + n$. The chance of \mathcal{B} guessing correctly is $\frac{1}{n^2}$, establishing the claim. \square

Theorem 3.9 ($LEAK\text{-}BIND\text{-}SS\text{-}\{CT, PK\} \not\implies LEAK^{+r}\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$). *Assuming a $LEAK\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure KEM \mathcal{KEM}' and a preimage-resistant (PR-d) hash function H (when sampling from $\{0, 1\}^{256}$), there exists a KEM \mathcal{KEM} which is $LEAK\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure but not $LEAK^{+r}\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure. Concretely, for any $LEAK\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ adversary \mathcal{A} against \mathcal{KEM}' we construct efficient adversaries $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ such that*

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

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

Figure 7: Algorithmic description of ML-KEM and Kyber using functions H, G, and J and a public key encryption scheme PKE. Solid boxes are exclusive to Kyber, dashed boxes are exclusive to ML-KEM.

Proof. Let $\text{KEM}' = (\text{KGen}', \text{Enc}', \text{Dec}')$ with shared secret space \mathcal{SS}' and randomness space $\mathcal{R}_{\text{KGen}'}$ for key generation be $\text{LEAK-BIND-SS-}\{CT, PK\}$ -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 $\{0, 1\}^{256} \times \mathcal{R}_{\text{KGen}'}$ for key generation.

<u>KGen(; r):</u> 1 $(x, r_{KG}) \leftarrow r$ 2 $y \leftarrow \text{H}(x)$ 3 $(pk', sk') \leftarrow \text{KGen}'(; r_{KG})$ 4 return $(pk', (sk', y))$	<u>Enc(pk):</u> 5 return $(\text{Enc}'(pk), \perp)$	<u>Dec(sk, pk, ct):</u> 6 $(sk', y) \leftarrow sk$ 7 $(ct', z) \leftarrow ct$ 8 if $y = \text{H}(z)$ 9 return ss^* 10 return $\text{Dec}'(sk', pk, ct')$
--	---	--

KEM is not $\text{LEAK}^{+r}\text{-BIND-SS-}\{CT, PK\}$ -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 $\text{LEAK-BIND-SS-}\{CT, PK\}$ -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 $\text{LEAK-BIND-SS-}\{CT, PK\}$ of KEM implies breaking $\text{LEAK-BIND-SS-}\{CT, PK\}$ of KEM' , which we have excluded by assumption. \square

3.3 Kyber and ML-KEM

Figure 7 gives an algorithmic description of Kyber [SAB⁺] and the new NIST standard ML-KEM [Nat24]. 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 (where J is called KDF in 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 $\text{J}(ss' \parallel \text{H}(ct))$

as shared secret while ML-KEM outputs ss' directly. Otherwise, both implicitly reject by computing the shared secret as $J(z\|H(ct))$ 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 $LEAK^{+r}$ -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 3.10 (ML-KEM is $LEAK^{+r}$ -BIND-SS- $\{CT, PK\}$ -secure). *Let \mathcal{A} be an adversary against the (t, ϵ, n) - $LEAK^{+r}$ -BIND-SS- $\{CT, PK\}$ security of ML-KEM and assume H, J behave like independent random oracles. Then we can construct adversaries $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ (given in the proof) with running time $\approx t$ such that*

$$\epsilon \leq \frac{n^2}{2^{256}} + 2 \cdot \text{Adv}_J^{\text{CR}}(\mathcal{B}_1) + \frac{q_{\text{RO}}^2}{2^{256}} + \text{Adv}_H^{\text{CR}}(\mathcal{B}_2) + 2 \cdot \text{Adv}_{G_{256}}^{\text{CR}}(\mathcal{B}_3),$$

where q_{RO} is the number of random oracle queries made by \mathcal{A} .

When treating G, J as *quantum-accessible* (independent) random oracles, following Zhandry [Zha15] we get a query complexity of $\Theta(2^{256/3})$ for finding a collision between the two instead of the $\frac{q_{\text{RO}}^2}{2^{256}}$ term in the theorem bound.

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

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

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

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 7 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.

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

Game A.1 (Colliding z values). We let $\mathcal{G}_{\text{A.0}}$ 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:

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

Game A.2 (Collision under 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 $z_i \| ct_i, z_j \| ct_j$ as a collision under J:

$$\text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_{\text{A.1}}}(\mathcal{A}) \leq \text{Adv}_{\text{J}}^{\text{CR}}(\mathcal{B}_1).$$

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.

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

Since $ct_i \neq ct_j$, we also have $z_i \| ct_i \neq z_j \| ct_j$, yet $J(z_i \| ct_i) = J(z_j \| ct_j)$. As in the prior case, the reduction \mathcal{B}_1 that outputs $z_i \| ct_i, z_j \| ct_j$ produces a collision under J:

$$\text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_{\text{B.0}}}(\mathcal{A}) \leq \text{Adv}_{\text{J}}^{\text{CR}}(\mathcal{B}_1).$$

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 $J(z_i \| ct_i) = ss_i = ss_j = \mathbf{G}_{256}(m'_j \| h_j)$ 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 $\mathcal{G}_{256}(m_i \| \mathbf{H}(pk_i)) = ss_i = ss_j = \mathcal{G}_{256}(m_j \| \mathbf{H}(pk_j))$, it must be that \mathcal{A} provides us with a collision either in \mathbf{H} or in \mathcal{G}_{256} . Once more, we can distinguish two cases:

1. The first case is $m_i \| \mathbf{H}(pk_i) = m_j \| \mathbf{H}(pk_j)$. Then we can build a reduction \mathcal{B}_2 that outputs pk_i, pk_j as collision under \mathbf{H} .
2. Otherwise, $m_i \| \mathbf{H}(pk_i) \neq m_j \| \mathbf{H}(pk_j)$, and we can build reduction \mathcal{B}_3 that outputs $m_i \| \mathbf{H}(pk_i), m_j \| \mathbf{H}(pk_j)$ as collision under \mathcal{G}_{256} .

Jointly,

$$\text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_{\text{D.0}}}(\mathcal{A}) \leq \text{Adv}_{\mathbf{H}}^{\text{CR}}(\mathcal{B}_2) + \text{Adv}_{\mathcal{G}_{256}}^{\text{CR}}(\mathcal{B}_3).$$

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 $h_i = \mathbf{H}(pk_i) = \mathbf{H}(pk_j) = h_j$ in $\mathcal{G}_{256}(m_i \| h_i) = ss_i = ss_j = \mathcal{G}_{256}(m_j \| h_j)$. 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**.

As in the prior case, the reduction \mathcal{B}_3 that outputs $m_i \| \mathbf{H}(pk_i), m_j \| \mathbf{H}(pk_j)$ produces a collision under \mathcal{G}_{256} :

$$\text{Adv}_{\text{ML-KEM}}^{\mathcal{G}_{\text{E.0}}}(\mathcal{A}) \leq \text{Adv}_{\mathcal{G}_{256}}^{\text{CR}}(\mathcal{B}_3).$$

Collecting the bounds yields the claim. □

The Kyber proof does not need to consider collisions *across* hash functions, allowing us to avoid relying on the random oracle. Instead, we additionally rely on the hardness of finding preimages under \mathcal{G}_{256} for randomly sampled images. Overall, the Kyber proof otherwise follows a similar strategy as the one for ML-KEM.

Theorem 3.11 (Kyber is $LEAK^{+r}$ -BIND-SS- $\{CT, PK\}$ -secure). *Let \mathcal{A} be an adversary against the (t, ϵ, n) - $LEAK^{+r}$ -BIND-SS- $\{CT, PK\}$ security of Kyber. Then we can construct adversaries $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3, \mathcal{B}_4, \mathcal{B}_5$ (given in the proof) with running time $\approx t$ such that*

$$\epsilon \leq \text{Adv}_{\mathbf{H}}^{\text{CR}}(\mathcal{B}_1) + \text{Adv}_{\mathbf{H}}^{\text{CR}}(\mathcal{B}_5) + \text{Adv}_{\mathcal{G}_{256}}^{\text{CR}}(\mathcal{B}_4) + n \cdot \text{Adv}_{\mathcal{G}_{256}}^{\text{PR-r}}(\mathcal{B}_3) + 4 \cdot \text{Adv}_{\mathbf{J}}^{\text{CR}}(\mathcal{B}_2) + \frac{n^2}{2^{256}}.$$

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

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

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

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 7 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 $J(x_i \| H(ct_i)) = ss_i = ss_j = J(x_j \| H(ct_j))$ 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 as collision under H , or a reduction \mathcal{B}_2 that outputs $x_i \| H(ct_i), x_j \| H(ct_j)$ as collision under J :

$$\text{Adv}_{\text{Kyber}}^{\mathcal{G}_{A,0}}(\mathcal{A}) \leq \text{Adv}_H^{\text{CR}}(\mathcal{B}_1) + \text{Adv}_J^{\text{CR}}(\mathcal{B}_2).$$

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.

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

Game B.1 (Colliding z values). We let $\mathcal{G}_{B.0}$ 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:

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

Game B.2 (Collision under J). Now, we have established $z_i \parallel \text{H}(ct_i) \neq z_j \parallel \text{H}(ct_j)$ (since $z_i \neq z_j$), yet $\text{J}(z_i \parallel \text{H}(ct_i)) = \text{J}(z_j \parallel \text{H}(ct_j))$. The reduction \mathcal{B}_2 that outputs $z_i \parallel \text{H}(ct_i), z_j \parallel \text{H}(ct_j)$ ⁵ produces a collision under J:

$$\text{Adv}_{\text{Kyber}}^{\mathcal{G}_{B.1}}(\mathcal{A}) \leq \text{Adv}_{\text{J}}^{\text{CR}}(\mathcal{B}_2).$$

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.

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

If $ss'_i \neq z_j$ (for $ss'_i = \text{G}_{256}(m'_i \parallel h_i)$), again the reduction \mathcal{B}_2 from Game $\mathcal{G}_{A.0}$ that outputs $ss'_i \parallel \text{H}(ct_i)$ and $z_j \parallel \text{H}(ct_j)$ produces a collision under J. Otherwise, we can build a reduction \mathcal{B}_3 that guesses $j \in [n]$ and embeds an image of G_{256} in z_j , producing a preimage under G_{256} for randomly sampled images (PR-r) if it guesses j correctly. Hence,

$$\text{Adv}_{\text{Kyber}}^{\mathcal{G}_{C.0}}(\mathcal{A}) \leq \text{Adv}_{\text{J}}^{\text{CR}}(\mathcal{B}_2) + n \cdot \text{Adv}_{\text{G}_{256}}^{\text{PR-r}}(\mathcal{B}_3).$$

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.

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

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

1. In case $\text{G}_{256}(m_i \parallel h_i) \parallel \text{H}(ct_i) \neq \text{G}_{256}(m_j \parallel h_j) \parallel \text{H}(ct_j)$, then reduction \mathcal{B}_2 outputs $\text{G}_{256}(m_i \parallel h_i) \parallel \text{H}(ct_i), \text{G}_{256}(m_j \parallel h_j) \parallel \text{H}(ct_j)$ as collision under J.
2. Otherwise, if $m_i \parallel h_i \neq m_j \parallel h_j$, then we can build a reduction \mathcal{B}_4 that outputs $m_i \parallel h_i, m_j \parallel h_j$ as collision under G_{256} .
3. Otherwise, we must have $m_i \parallel h_i = m_j \parallel h_j$ and $\text{H}(pk_i) = h_i = h_j = \text{H}(pk_j)$. By event D, $pk_i \neq pk_j$ and we can build a reduction \mathcal{B}_5 that outputs pk_i, pk_j as collision under H.

Jointly,

$$\text{Adv}_{\text{Kyber}}^{\mathcal{G}_{D.0}}(\mathcal{A}) \leq \text{Adv}_{\text{J}}^{\text{CR}}(\mathcal{B}_2) + \text{Adv}_{\text{G}_{256}}^{\text{CR}}(\mathcal{B}_4) + \text{Adv}_{\text{H}}^{\text{CR}}(\mathcal{B}_5).$$

Collecting the bounds yields the claim. □

⁵This is the same reduction \mathcal{B}_2 as in Game $\mathcal{G}_{A.0}$, which outputs $x_i \parallel \text{H}(ct_i), x_j \parallel \text{H}(ct_j)$, where x_i, x_j are z_i, z_j for implicitly rejecting decapsulations that we consider here.

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).⁶ 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 traditional 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 Double Ratchet 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.

⁶This in particular covers forward secrecy (compromise of long-term and semi-static keys after a handshake took place) and a form of post-compromise security [CCG16] (compromise of long-term and semi-static keys before a handshake with uncompromised ephemeral randomness takes place) for the handshake. Note that our model only considers the initial handshake; Signal’s subsequent Double Ratchet protocol [MP16a] provides further (classical) forward security and post-compromise security.

- $\text{KGenSS}(ltsk) \mapsto ((sspk_1, sssk_1, \sigma_1), \dots)$: 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.⁷
- $\text{Run}(ltsk, \vec{ssk}, \vec{lpk}, \pi, m) \mapsto (\pi', m')$: The *session execution algorithm* that takes as input a party's long-term secret key $ltsk$, that party's semi-static secret keys \vec{ssk} , all parties' long-term public keys \vec{lpk} , 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 keys⁸ 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, each *party* $P \in [n_p]$ holds a long-term public-key/secret-key pair generated by KGenLT and may run multiple instances of the protocol (simultaneously or sequentially); we denote the i th such *session* of party P by π_P^i . The security game maintains the following information for each session:

- $\text{oid} \in [n_p]$: 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].
- $K \in \mathcal{K}_{\text{KE}} \cup \{\perp\}$: The session key established in this session, initialized to \perp .
- $\text{type} \in \{\text{full}, \text{reduced}\}$: 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.
- $\text{sspks} \in (\{0, 1\}^*)^* \cup \{\perp\}$: 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:

⁷The 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.

⁸In 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.

- $\text{revrand} \in \{\text{true}, \text{false}\}$ 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 8 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} and \vec{sspk} 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 $\text{Run}(\vec{ltsk}_U, \vec{sssk}_U, \vec{ltpk}, \pi_U^i, m)$, 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 \vec{ltsk}_U to the adversary; recorded through the flag `corrltkU`.
- $\text{CORRUPTSSKEY}(U, \text{ssid})$: Returns party U ’s semi-static secret key $\vec{sssk}_U^{\text{ssid}}$ to the adversary; recorded through the flag `corrssksspk`, where \vec{sspk} is the public key corresponding to $\vec{sssk}_U^{\text{ssid}}$.⁹
- $\text{REVEALRAND}(U, i)$: Returns the randomness $\pi_U^i.\text{coins}$ of session π_U^i to the adversary, recorded in the session through the flag `revrand`.
- $\text{REVEALSESSKEY}(U, i)$: Returns the session key $\pi_U^i.\mathcal{K}$ of session π_U^i to the adversary, recorded in the session through the flag `revealed`.
- $\text{TEST}(U, i)$: If a TEST query has been made before or session π_U^i has not accepted, then return \perp . Otherwise; if $b_{\text{test}} = 0$, return $\pi_U^i.\mathcal{K}$, otherwise return a randomly sampled session key from the protocol’s key space \mathcal{K}_{KE} . Record the test session as $\pi^* \leftarrow \pi_U^i$.

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 8. We say that KE achieves $(t, \epsilon, (q_{\text{Snd}}, q_{\text{CorrLT}})$,*

⁹Since we model that semi-static keys are not authentically distributed (but signed), sessions only know the public key received (\vec{sspk}), but not the game label `ssid`, which was used in prior models [CCD⁺17, BFG⁺22, BFG⁺21] to track compromise more easily.

```

 $\mathcal{G}_{KE}^{KI}(\mathcal{A})$ :
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    $(ltpk_U, ltsk_U) \leftarrow \text{KGenLT}()$ 
5   for  $s \in [n_{ss}]$  // generate semi-static keys
6      $((ssp_{U,s}^{s,1}, sss_{U,s}^{s,1}, \sigma_U^{s,1}), \dots) \leftarrow \text{KGenSS}(ltsk_U)$ 
7      $ss\vec{s}k_U \leftarrow \{ss_{U,s}^{s,1}, \dots\}^{s \in [n_{ss}]}$ 
8      $l\vec{p}k \leftarrow \{ltpk_U\}_{U \in [n_p]}$ 
9      $ss\vec{p}k \leftarrow \{(ssp_{U,s}^{s,1}, \sigma_U^{s,1}), \dots\}^{s \in [n_{ss}]}$ 
10   $b' \leftarrow \mathcal{A}(l\vec{p}k, ss\vec{p}k)$  // 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

SEND( $U, i, m$ ):
21 if  $\pi_U^i = \perp$ : // initiate session: for responders, we have  $m = (\text{create}, (ssid, \text{type}))$ 
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    $(\pi_U^i, m') \leftarrow \text{Run}(ltsk_U, ss\vec{s}k_U, l\vec{p}k, \pi_U^i, m)$ 
   // 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    $\pi_U^i.\text{pcorr} \leftarrow \text{corrltk}_{\pi_U^i.\text{pid}}$ 
30  return  $(m', \pi_U^i.\text{st}_{\text{exec}})$  // return message and session state

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

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

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

CORRUPTLTKEY( $U$ ):
36  $\text{corrltk}_U \leftarrow \text{true}$  // mark long-term key corrupted
37 return  $ltsk_U$  // return long-term secret key

CORRUPTSSKEY( $U, ssid$ ):
38  $\text{corrssk}_{ssp_{U,ssid}} \leftarrow \text{true}$  // mark semi-static key corrupted
39 return  $ss\vec{s}k_U^{ssid}$  // return semi-static secret key

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

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

```

Figure 8: Key indistinguishability (KI) game for key exchange protocol KE (top), in which adversary \mathcal{A} has access to oracles SEND, TEST, CORRUPTLTKEY, CORRUPTSSKEY, REVEALRAND, and REVEALSESSKEY (bottom), and wrt. to (protocol-specific) clean predicates clean^{KE} for KE. Highlighted code reflects the main changes compared to the model of BFGJS [BFG⁺22, BFG⁺21]. The clean predicates for PQXDH protocol capturing classical security ($\text{clean}^{\text{PQXDH}}$) and post-quantum security ($\text{clean}^{\text{PQXDH-Q}}$) are given in Figure 9.

$(q_{\text{CorrSS}}, q_{\text{RevR}}, q_{\text{RevSK}})$ –key indistinguishability, if for any adversary \mathcal{A} against $\mathcal{G}_{KE}^{KI}(\mathcal{A})$ 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

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

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):

Predicates $\text{clean}^{\text{PQXDH}}$ capturing classical security

```

cleanfullPQXDH( $\pi^*$ ):
46 return cleanreducedPQXDH( $\pi^*$ ) or cleanEE( $\pi^*$ )

cleanreducedPQXDH( $\pi^*$ ):
47 return cleanLTSS( $\pi^*$ ) or cleanELT( $\pi^*$ ) or cleanESS( $\pi^*$ )

cleanEE( $\pi^*$ ):
48 return  $\neg\pi^*.revrand$  and (cleanpeerE( $\pi^*$ ) or cleansigE( $\pi^*$ ))
    // test session randomness is unrevealed and peer's ephemeral contribution is clean or cleanly signed

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

cleansigE( $\pi^*$ ):
52 return ( $\pi^*.role = \text{initiator}$  and  $\neg\pi^*.pcorr$ 
    and  $\forall\pi : ((\pi.role = \text{responder}$  and  $\pi^*.pid = \pi.oid$ 
    and  $\pi^*.cid[4][2] = \pi.cid[4][2]) \implies \neg\pi.revrand$ )
    and ( $\pi^*.cid[4][2], \cdot \notin \text{sspks}$ )
    // long-term (signing) secret key of the responder peer was uncompromised upon acceptance,
    // and if a partner with the same ephemeral KEM key exists
    // then that partner's randomness is unrevealed,
    // and the received ephemeral KEM public key was not replaced with a semi-static one)

```

Predicates $\text{clean}^{\text{PQXDH-Q}}$ capturing post-quantum security

```

cleanfullPQXDH-Q( $\pi^*$ ):
62 return cleanEEQ( $\pi^*$ )

cleanreducedPQXDH-Q( $\pi^*$ ):
63 return cleanESSQ( $\pi^*$ )

cleanEEQ( $\pi^*$ ):
64 return  $\neg\pi^*.revrand$  and cleanpeerE( $\pi^*$ )
    // test session randomness is unrevealed and peer's ephemeral contribution is clean

```

```

cleanLTSS( $\pi^*$ ):
53 return
54 ( $\pi^*.role = \text{initiator}$  and  $\neg\text{corrltk}_{\pi^*.oid}$  and  $\neg\text{corrsk}_{\pi^*.sspks[1]}$  and  $\neg\pi^*.pcorr$ )
    // tested initiator's long-term and responder's semi-static DH secret keys are uncompromised
    // and responder long-term (signing) secret key was uncompromised upon acceptance
55 or ( $\pi^*.role = \text{responder}$  and  $\neg\text{corrltk}_{\pi^*.pid}$  and  $\neg\text{corrsk}_{\pi^*.sspks[1]}$ )
    // initiator long-term and tested responder's semi-static DH secret keys are uncompromised

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

cleanESS( $\pi^*$ ):
59 return
60 ( $\pi^*.role = \text{initiator}$  and  $\neg\pi^*.revrand$ 
    and  $\forall \text{sspk} \in \pi^*.sspks : (\neg\text{corrsk}_{\text{sspk}}$  and  $\neg\pi^*.pcorr$ 
    and  $\pi^*.type = \text{reduced} \implies (\pi^*.sspks[2], \cdot \in \text{sspks}')$ )
    // tested initiator's randomness is unrevealed
    // and responder semi-static secret keys are uncompromised
    // and responder long-term (signing) secret key was uncompromised upon acceptance
    // and in reduced mode the received semi-static KEM public key
    // was honestly generated as semi-static key)
61 or ( $\pi^*.role = \text{responder}$  and cleanpeerE( $\pi^*$ )
    and  $\forall \text{sspk} \in \pi^*.sspks : (\neg\text{corrsk}_{\text{sspk}})$ )
    // initiator's ephemeral contribution is clean and tested responder's semi-static secret keys are uncompromised

```

```

cleanESSQ( $\pi^*$ ):
65 return
66 ( $\pi^*.role = \text{initiator}$  and  $\neg\pi^*.revrand$  and
     $\exists s \in [n_{ss}] : \pi^*.sspks[2] = \text{sspk}_{\pi^*.pid}^{s,2}$  and  $\neg\text{corrsk}_{\pi^*.sspks[2]}$ )
    // tested initiator's randomness is unrevealed and
    // responder's semi-static KEM secret key was honestly generated and is uncompromised
67 or ( $\pi^*.role = \text{responder}$  and cleanpeerE( $\pi^*$ ) and  $\neg\text{corrsk}_{\pi^*.sspks[2]}$ )
    // initiator's ephemeral contribution is clean and
    // tested responder's semi-static KEM secret key is uncompromised

```

Figure 9: Clean predicates for the PQXDH protocol capturing classical security ($\text{clean}^{\text{PQXDH}}$, top) and post-quantum security ($\text{clean}^{\text{PQXDH-Q}}$, bottom). For post-quantum security, the elliptic curve-related clean predicates ($\text{clean}_{\text{LTSS}}$, $\text{clean}_{\text{ELT}}$, $\text{clean}_{\text{sigE}}$) become moot; security is expected from the ephemeral ($\text{clean}_{\text{EE}}^{\text{Q}}$) resp. semi-static ($\text{clean}_{\text{ESS}}^{\text{Q}}$) KEM component for full resp. reduced handshakes, upon non-tampered transmission. Lighter gray code reflects the main changes in the classical security predicates compared to the model of BFGJS [BFG⁺22, BFG⁺21]. Highlighted code reflects changes in the post-quantum security predicate $\text{clean}_{\text{ESS}}^{\text{Q}}$ compared to the classical-security version $\text{clean}_{\text{ESS}}$. Dashed-boxed code prevents trivial confusion attacks between ephemeral and semi-static KEM public keys in PQXDH, since their signatures lack domain separation (cf. discussion in Section 4.3).

- (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. 8, line 18).
- (ii) Two initiator sessions accept with the same session identifier (Fig. 8, line 19).
- (iii) Three sessions accept with the same session identifier in full handshake type (Fig. 8, 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 8, line 16, and Figure 9).

4.3 Clean Predicates for PQXDH: Capturing Classic and Post-quantum Security

Following the terminology of CCDGS [CCD⁺17], the specific security properties targeted by a key exchange protocol KE are expressed through a set of *clean predicates* clean^{KE} . Figure 9 formalizes the maximum-exposure security properties for PQXDH, both for classical security ($\text{clean}^{\text{PQXDH}}$) as well as for post-quantum security ($\text{clean}^{\text{PQXDH-Q}}$).

Classic security. The core classical security properties in $\text{clean}^{\text{PQXDH}}$ (Figure 9, top) closely follow those for the classical 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 5.1), 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 model the signatures on the responder’s (ephemeral KEM and semi-static public) keys and allow the adversary to replace semi-static keys (on top of replacing 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 to the targeted session having accepted (cf. Figure 8, lines 54 and 60). Treating signatures explicitly, we also capture the additional guarantee that signing KEM keys gives: for the ephemeral/ephemeral combination to contribute security, it is now (also) sufficient to have the peer’s contribution be cleanly signed (cf. Figure 8, line 48). To that end we set a flag `pcorr` in a session if its peer was corrupted upon acceptance (Figure 8, 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.

Trivial KEM key confusion due to lack of domain separation. Unfortunately, PQXDH signs ephemeral and semi-static KEM public keys without any domain separation (e.g., labels indicating the key type). This leads to a trivial *KEM key confusion attack*: an active adversary can swap out an ephemeral for a semi-static KEM public key (or vice versa), without Alice being able to distinguish the two types, resulting in reduced forward security guarantees. Note that this is different from the key confusion attack

noticed by BJKS [BJKS24], where the adversary exchanges the semi-static DH and the semi-static KEM key.¹⁰

The formal analysis of BJKS [BJKS24] treated all KEM public keys as semi-static, side-stepping the issue. Here, we model both ephemeral and semi-static KEM keys and capture their distinct security properties *if* the adversary does not mount such trivial key confusion attack. We highlight the code ruling out this trivial attack with [dashed boxes] in Figure 9 (lines 52 and 60). We emphasize that this attack would be easily avoided by domain separating PQXDH’s signature, allowing stronger security guarantees (cf. our proof details), a fix suggested also by BJKS [BJKS24] and envisioned by the Signal team for the next protocol upgrade.

Post-quantum security. Via a separate set $\text{clean}^{\text{PQXDH-Q}}$ (Figure 9, bottom) of clean predicates for post-quantum security, we capture the security that remains against a quantum adversary: PQXDH is expected to provide key indistinguishability through a **non-tampered KEM encapsulation**, either ephemeral (in a full handshake) or semi-static (in a reduced handshake). The authentic transmission of KEM keys captures an adversary that can *not yet* forge signatures (i.e., the signatures still protect the authenticity of the KEM keys) or a quantum adversary that does not tamper with the transmission of the KEM keys. Note that since a quantum adversary can efficiently break discrete logarithms, the DH shares cannot contribute to security anymore.

Clean predicates details. Let us go through the clean predicates in formal detail now. We start with the classic security predicates, which are in parts also used for defining post-quantum security. 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 $\text{clean}_{\text{full}}$ or the $\text{clean}_{\text{reduced}}$ predicate to π^* .

Since $\text{clean}_{\text{reduced}}$ is part of the description of $\text{clean}_{\text{full}}$, we first discuss the case $\pi^*.\text{type} = \text{reduced}$. Intuitively, a session key derived in such a session remains unknown to the adversary, if one of the four shared secrets (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^* :

clean_{LTS}: This predicate indicates whether the combination of the long-term key of the initiator and the semi-static DH key 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.**

clean_{ELT}: 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.

clean_{ESS}: 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 [and in reduced mode the received semi-static KEM public key was honestly generated as semi-static key].** The latter, [dashed-boxed] part rules out the trivial attack due to lack of domain separation in PQXDH’s signatures discussed above. If such domain separation were added, PQXDH could be proven secure wrt. stronger clean predicates not ruling out such confusion as trivial.

If the test session π^* is a responder session, the evaluation of $\text{clean}_{\text{ELT}}$ and $\text{clean}_{\text{ESS}}$ necessitates a further predicate called $\text{clean}_{\text{peerE}}$ (in all other cases, it is sufficient to consider the compromise of keys and randomness via the flags `corrltk`, `corrsk`, `revrand`, and `pcorr`).

¹⁰The latter confusion is not a trivial attack against Signal’s implementation of PQXDH since the sizes of DH and KEM public keys are distinct. Hence, we do not exclude it in the model but explicitly rely on this distinction in our analysis.

$\text{clean}_{\text{peerE}}$: 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 $\pi^*.\text{type} = \text{full}$, it must either hold that $\text{clean}_{\text{reduced}}$ is true or that the additional input to the master secret computation is clean. The latter is captured by the following predicate:

clean_{EE} : 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 $\text{clean}_{\text{peerE}}$) or was “cleanly” signed (captured by the new $\text{clean}_{\text{sigE}}$ 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 $\text{clean}_{\text{peerE}}$ helps to determine within clean_{EE} 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.

$\text{clean}_{\text{sigE}}$: This (sub)predicate indicates that for initiator test sessions the peer was not corrupted at the time of acceptance and that, if the test session’s peer has issued the ephemeral KEM public key used in the test session (that is, $\pi^*.\text{cid}[4][2]$)¹¹, then the randomness for generating that ephemeral KEM key was not revealed, [and that the received ephemeral KEM public key was not replaced with] [a semi-static one]. The latter, [dashed-boxed] part again rules out the trivial key confusion attack due to lack of domain separation in PQXDH’s signatures discussed above.

Finally, the post-quantum security relies solely on the ephemeral/ephemeral KEM contribution for full handshakes resp. on the ephemeral/semi-static KEM contribution for reduced handshakes, captured via the following two predicates:

$\text{clean}_{\text{EE}}^Q$: As for classic security, this predicate requires that the ephemeral contribution of both the test session and its peer session are unknown to the adversary. Note that since PQXDH’s signatures are only classically secure, they cannot contribute to post-quantum security and the $\text{clean}_{\text{sigE}}$ sub-predicate is accordingly omitted.

$\text{clean}_{\text{ESS}}^Q$: Similar to classic security, this predicate requires that the ephemeral contribution of the initiator and the semi-static contribution of the responder are unknown to the adversary. Just that here, [only the KEM secret key] is relevant for security and [must have been honestly generated], since again we cannot hope for tampering protection from the (classically-secure) signatures. Requiring honestly generated keys for post-quantum security also excludes the trivial KEM key confusion attack.

5 PQXDH Analysis

We formalize the PQXDH handshake in Figure 10, 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

¹¹Note that other elements in cid might differ. A prior version incorrectly compared the full cid; we acknowledge Hashimoto, Katsumata, and Wiggers for pointing out this oversight.

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( $ltsk_A, sssk_A, \vec{ltpk}, \pi_A, m$ )
( $r_4, \perp, \perp, r_5$ )  $\leftarrow$   $\pi_A$ .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( $ltsk_B, \vec{ltpk}, sssk_B, \pi_B, (create, (ssid, type))$ )
( $r_1, r_2, r_3, \perp$ )  $\leftarrow$   $\pi_B$ .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( $ltsk_B, sssk_B, \vec{ltpk}, \pi_B, m'; coins$ )
( $A, epk_A^{DH}, ct$ )  $\leftarrow$   $m'$ 
( $ltpk_A^{DH}, ltpk_A^{SIG}$ )  $\leftarrow$   $ltpk_A$ 
 $DH_1 \leftarrow$  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 10: The PQXDH protocol. Darker gray code are additions compared to the original X3DH protocol to add post-quantum security. Lighter gray code are signature elements of the X3DH and PQXDH protocols not covered in the prior analysis of [CCD⁺17].

key pair.¹² 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.¹³ 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.

5.1 Classic Security of PQXDH

We first show that PQXDH achieves key indistinguishability against classical adversaries wrt. the maximum-exposure attack vectors formalized through the clean predicates $\text{clean}^{\text{PQXDH}}$ in our model (cf. Figure 9). This result relies on the GapDH assumption in the DH group, the KEM providing one-way CCA and $LEAK\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ security as well as low key-collision probability and correctness errors, the signature scheme being existentially unforgeable, and the key derivation function behaving like a random oracle. Note that we further explicitly assume the representations of KEM and Diffie–Hellman public keys to be disjoint. This is the case for Signal’s PQXDH implementation, employing Kyber as KEM and X25519 or X448 for elliptic-curve DH [LHT16].

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 (Classic key indistinguishability of PQXDH). *The PQXDH protocol given in Figure 10 with randomness space $\mathcal{R}_{KE} = \mathcal{R}_{DH.KGen} \times \mathcal{R}_{KEM.KGen} \times \mathcal{R}_{KEM.Enc} \times \mathcal{R}_{SIG.Sig}$ achieves $(t, \epsilon, (q_{Snd}, q_{CorrLT}, q_{CorrSS}, q_{RevR}, q_{RevSK}))$ -key indistinguishability wrt. the classical security clean predicates $\text{clean}^{\text{PQXDH}}$ (cf. Figure 9), assuming (for $t' \approx t$) the GapDH problem in the group (\mathbb{G}, g, q) is $(t', \epsilon_{GDH}, 1.5(q_{RO} + q_{Snd})^2)$ -hard, KEM is a $(t', \epsilon_{CCA}, n_s)$ -OW-CCA-secure and $(t', \epsilon_{LEAK+r}, n_p \cdot n_{ss} + n_s)$ - $LEAK^{+r}\text{-}BIND\text{-}SS\text{-}\{CT, PK\}$ -secure KEM with public-key collision probability γ_{coll} and correctness error δ_{corr} , SIG is a $(t', \epsilon_{SIG}, 2n_{ss} + n_s)$ -unforgeable signature scheme, KDF behaves like a (programmable) random oracle, and the group (\mathbb{G}, g, q) and the public key space of KEM have disjoint representations. Concretely, for any efficient classical*

¹²The 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.

¹³Ephemeral 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.

adversary \mathcal{A} we have

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

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

Proof summary. The proof proceeds via a series of game hops.

Game 1 (DH/KEM pk collisions). First, we exclude collisions in the DH and KEM public keys, yielding the term $\frac{(n_p + n_p \cdot n_{ss} + n_s)^2}{q} + \gamma_{\text{coll}}(n_p \cdot n_{ss} + n_s)$.

Game 2 (KEM correctness). Second, we exclude KEM correctness errors when decapsulating in sessions, yielding the term $n_s \cdot \delta_{\text{corr}}$. 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 $\epsilon_{\text{LEAK}+r}$ term; this in particular rules out re-encapsulation attacks where sessions that are not partnered possibly derive the same session key.¹⁴

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

¹⁴As 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).

involve only DH secrets (and hence do not provide post-quantum security), or, for the sixth case (clean_{EE} 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 public key collisions). We let \mathcal{G}_0 abort (overwriting the adversary's output with 0) if any two honestly generated DH public keys or KEM public 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 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 3.2. 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 8, 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 8, 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 8, 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).¹⁵ 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 $LEAK^{+r}$ -BIND-SS- $\{CT, PK\}$ 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 $LEAK^{+r}$ -BIND-SS- $\{CT, PK\}$ 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 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,

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

Separating the clean cases. At this point, we will divide the proof into six sub-cases following the structure of the $\text{clean}_{\text{full}}^{\text{PQXDH}}$ resp. $\text{clean}_{\text{reduced}}^{\text{PQXDH}}$ 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 $\mathcal{G}_3[c]$ which are \mathcal{G}_3 conditioned on a sub-predicate c :

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

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, 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 the initiator uses a semi-static key that was honestly generated by the responder, 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.

¹⁵If 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.

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^* . Hence, the test session accepts with $\pi^*.\text{sspks} = (\text{sspk}_{V^*}^{\text{DH}}, \cdot)$ 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 corresponding 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 $\mathcal{G}_{\text{A.1}}$, \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 $\mathcal{G}_{\text{A.1}}$ 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. Note that the adversary cannot masquerade a signed KEM public key as signed DH public key since public keys are distinct for DH and KEM by assumption. Thus,

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

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 :

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

Game A.4 (Guess semi-static DH key identifier $\text{ssid}_{\text{DH}}^*$ 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 $\text{ssid}_{\text{DH}}^*$, 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} :

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

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

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  $DDH(ltpk_{U^*}^{DH}, sspk_{V^*}^{DH}, DH_1)$ 
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  $U \leftarrow \pi_W^i.\text{initiator}; V \leftarrow \pi_W^i.\text{responder}; sspk_V^{DH} \leftarrow \pi_W^i.\text{sspks}[1]$ 
21  $(\dots, epk_U^{DH}, \dots) \leftarrow \pi_W^i.\text{sid}$ 
22 if  $U = U^* \wedge sspk_V^{DH} = sspk_{V^*}^{DH}$ 
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  $U \neq U^* \wedge sspk_V^{DH} = sspk_{V^*}^{DH}$ 
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 11: Algorithmic description of how \mathcal{B}_3 simulates the random oracle and the SEND oracle in game $\mathcal{G}_{A.5}$. We ensure query consistency via an additional global, initially empty list L . 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. The SEND oracle uses patch in lieu of computing KDF whenever there are challenge secret keys involved that \mathcal{B}_3 does not know.

beginning with $\text{CDH}(ltpk_{U^*}^{\text{DH}}, sspk_{V^*}^{\text{DH}})$, i.e., the shared DH key between $ltpk_{U^*}^{\text{DH}}$ and $sspk_{V^*}^{\text{DH}}$, where $sspk_{V^*}^{\text{DH}}$ is the semi-static DH key with identifier $\text{ssid}_{\text{DH}}^*$. 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) for $q_{\text{DDH}} \leq 1.5(q_{\text{RO}} + q_{\text{Snd}})^2$.

Reduction \mathcal{B}_3 samples all key components itself except for replacing the long-term DH key $ltpk_{U^*}^{\text{DH}}$ of U^* and the semi-static DH key $sspk_{V^*}^{\text{DH}}$ 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 $\text{CORRUPTSSKEY}(V^*, \text{ssid}_{\text{DH}}^*)$ in its simulation of Game $\mathcal{G}_{\text{A.4}}$. 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 11 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 $sspk_{V^*}^{\text{DH}}$ of the test session, the responder is U^* , or the responder is V^* with semi-static key $sspk_{V^*}^{\text{DH}}$ and the initiator is not U^* . In these cases the reduction cannot compute all DH shared secrets 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 $\text{CDH}(ltpk_{U^*}^{\text{DH}}, sspk_{V^*}^{\text{DH}})$, 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,¹⁶ 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 $q_{\text{RO}} + q_{\text{Snd}}$. Using the Gaussian sum we get

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

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

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

¹⁶If 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 11, 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 $q_{\text{RO}} + q_{\text{Snd}}$ DDH queries.

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:

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

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 session’s randomness is not revealed, and
2. the responder’s long-term key is uncompromised, and
3. if π^* is a responder, then there exists a partnered initiator session.

Via the last point, we can guarantee that the responder uses an ephemeral key that was honestly generated by the initiator. 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.

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

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 :

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

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 :

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

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 $(t', \epsilon_{\text{GDH}}, q_{\text{Snd}} + q_{\text{RO}})$ -hardness of the GapDH problem in (\mathbb{G}, g, q) .

The reduction follows the idea of Game $\mathcal{G}_{A.5}$, 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 $\mathcal{G}_{A.5}$ 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 $\mathcal{G}_{A.5}$.

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 $\mathcal{G}_{B.2}$. If \mathcal{A} does make such a random oracle query, then \mathcal{B}_4 will detect this and win its GapDH game. Thus,

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

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

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 session’s randomness is not revealed, and
2. the responder’s semi-static key is uncompromised, and
3. if π^* is an initiator, then the responder’s long-term (signing) key was uncompromised upon acceptance, and
4. if π^* is a responder, then there exists a partnered initiator session.

Via the second last point, similarly to Case A, we can guarantee that the initiator uses a semi-static key that was honestly generated by the responder, given signatures are unforgeable. Via the last point, similarly to Case B, we can guarantee that the responder uses an ephemeral key that was honestly generated by the initiator. 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 $\text{clean}_{\text{ESS}}(\pi^*) \wedge \pi^*.\text{type} = \text{full}$ being satisfied.

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

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 :

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

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 $\text{sspk}_{V^*}^{\text{DH}}$ that was not generated through a KGenSS run for V^* . Note that DH public keys and KEM public keys are distinct by assumption and cannot be confused with each other. Hence, the test session accepts with $\pi^*.\text{sspk} = (\text{sspk}_{V^*}^{\text{DH}}, \cdot)$ corresponding to a DH key pair of which the adversary does not know the secret key, since $\text{clean}_{\text{ESS}}(\pi^*)$ guarantees that the semi-static secret belong to $\text{sspk}_{V^*}^{\text{DH}}$ is not compromised. Similar to Game $\mathcal{G}_{\text{A.2}}$, the probability of such an abort can be bounded by the advantage of a reduction against the $(t', \epsilon_{\text{SIG}}, 2n_{\text{ss}} + n_s)$ -unforgeability of SIG. The argument is identical to the Game $\mathcal{G}_{\text{A.2}}$. Thus,

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

Game C.3 (Guess semi-static DH key identifier $\text{ssid}_{\text{DH}}^*$ 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 $\text{ssid}_{\text{DH}}^*$, 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:

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

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 $\text{ssid}_{\text{DH}}^*$. The probability of such

an abort can be bounded by the advantage of the following reduction \mathcal{B}_5 against the $(t', \epsilon_{\text{GDH}}, q_{\text{Snd}} + q_{\text{RO}})$ -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 $\text{ssid}_{\text{DH}}^*$. The reduction does not need to answer corresponding REVEALRAND queries on π_i^* or CORRUPTSSKEY(V^* , $\text{ssid}_{\text{DH}}^*$) 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 $\text{ssid}_{\text{DH}}^*$ 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 $\text{CDH}(epk^{\text{DH}}, ltpk_{V^*}^{\text{DH}})$ 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 $\mathcal{G}_{\text{C.3}}$. 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 session'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
4. and the received semi-static KEM public key was honestly generated as semi-static key (ensured via the `[dashed-box trivial attack check]`), and
5. if π^* is a responder, then there exists a partnered initiator session.

Via points 3 and 4, we can guarantee that the initiator uses a semi-static key that was honestly generated by the responder, given signatures are unforgeable. Via the last point, we can guarantee that the responder uses an ephemeral key that was honestly generated by the initiator. 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 as well as the ephemeral encapsulation randomness 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 $\text{clean}_{\text{ESS}}(\pi^*) \wedge \pi^*.\text{type} = \text{reduced}$ being satisfied.

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

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 $(t', \epsilon_{\text{SIG}}, 2n_{\text{ss}} + n_s)$ -unforgeability of SIG. Note that DH public keys and KEM public keys have disjoint representations by assumption and cannot be confused with each other. In addition, the `[dashed-box trivial attack check]` (Figure 9, line 60) ensures that the KEM public key must be an honestly generated semi-static key from sspk and hence cannot be confused with a signed ephemeral KEM public key (note that keys do not collide by Game \mathcal{G}_1).¹⁷

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

Game D.3 (Guess semi-static key identifiers $\text{ssid}_{\text{DH}}^*$, $\text{ssid}_{\text{KEM}}^*$ of V^*). We now guess the ssid identifiers of the responder V^* ’s (uncorrupted) 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 $\pi^*.\text{role} = \text{responder}$), or of the intended peer (if $\pi^*.\text{role} = \text{initiator}$). We denote the guessed identifiers by $\text{ssid}_{\text{DH}}^*$ (for the DH key) and $\text{ssid}_{\text{KEM}}^*$ (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:

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

¹⁷If signatures in PQXDH were properly domain-separated, this trivial key confusion attack against PQXDH would be avoided, enabling a proof in a stronger model not forbidding the trivial key confusion attack; cf. the discussion in Section 4.3.

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 :

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

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 with identifiers $\text{ssid}_{\text{DH}}^*$ and $\text{ssid}_{\text{KEM}}^*$, respectively. 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 $(t', \epsilon_{\text{GDH}}, q_{\text{Snd}} + q_{\text{RO}})$ -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 $\text{ssid}_{\text{DH}}^*$ (as in Game $\mathcal{G}_{C.5}$), 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 $\text{ssid}_{\text{KEM}}^*$. Furthermore, it uses the KEM challenge ciphertext ct^* in π_i^* . Similar to Game $\mathcal{G}_{C.5}$, the reduction does not need to answer corresponding REVEALRAND on π_i^* or CORRUPTSSKEY($V^*, \text{ssid}_{\text{DH}}^*/\text{ssid}_{\text{KEM}}^*$) queries due to $\text{clean}_{\text{ESS}}$. Furthermore, as for Game $\mathcal{G}_{C.5}$, 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 $\mathcal{G}_{A.5}$; 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 $\mathcal{G}_{D.4}$. 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:

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

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 $\mathcal{G}_{D.5}$ 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,

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

To conclude this proof case, observe that in Game $\mathcal{G}_{D.6}$ 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 randomness is not revealed, and
2. there exists a (contributively) partnered session whose randomness is not revealed.

Via the last point, we can guarantee that initiator and responder agree on the ephemeral contributions. Similar to Case D, we get a *hybrid* guarantee here: both the ephemeral/ephemeral Diffie–Hellman combination being uncompromised *and* the ephemeral KEM key as well as the ephemeral encapsulation randomness 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 $(t', \epsilon_{\text{GDH}}, q_{\text{RO}} + 1)$ -hardness of the GapDH problem in (\mathbb{G}, g, q) and $(t', \epsilon_{\text{CCA}}, 1)$ -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,¹⁸ 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

¹⁸This 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 $q_{RO} + 1$ 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 $CDH(epk_{U^*}^{DH}, epk_{V^*}^{DH})$ 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. π^* is an initiator, and
2. the test session's randomness is not revealed, and
3. the responder's long-term (signing) key was uncompromised upon acceptance of π^* , and
4. the randomness of a potential partner session that sent the ephemeral KEM public key used by π^* was not revealed¹⁹, and
5. the received ephemeral KEM public key was not replaced with a semi-static one (ensured via the dashed-box trivial attack check).

Via points 3 and 5, we can guarantee that the initiator uses an ephemeral KEM key that was honestly generated by the responder, given signatures are unforgeable. The KEM encapsulation then ensures key indistinguishability for the test session.

¹⁹Note that other elements in cid might be replaced, cf. the check in the $\text{clean}_{\text{sigE}}$ predicate, Figure 9 line 52. A prior version incorrectly compared the full cid ; we acknowledge Hashimoto, Katsumata, and Wiggers for pointing out this oversight.

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.

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

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 :

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

Game F.2 (Signature unforgeability). We now abort the game (again, returning 0 as the adversary's bit guess) in the event that the initiator test session π^* accepts using an ephemeral KEM public key $\text{epk}_{V^*}^{\text{KEM}}$ that was not generated by V^* . Note again that DH public keys and KEM public keys are distinct by assumption and cannot be confused with each other. Similar to Game $\mathcal{G}_{\text{D.2}}$, the `[dashed-box trivial attack check]` (Figure 9, line 52) ensures that the received ephemeral KEM public key cannot be confused with a signed semi-static KEM public key from sspk (note that keys do not collide by Game \mathcal{G}_1).²⁰ 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 per $\text{clean}_{\text{sigE}}$ (recall that the ephemeral KEM public key is part of the sessions' cid, namely $\text{cid}[4][2]$). Similar to Game $\mathcal{G}_{\text{A.2}}$, the probability of such an abort can be bounded by the advantage of a reduction against the $(t', \epsilon_{\text{SIG}}, 2n_{\text{ss}} + n_s)$ -unforgeability of SIG. The proof is identical to the one for Game $\mathcal{G}_{\text{A.2}}$. Thus,

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

Game F.3 (Guess initiator and responder sessions). We guess the initiator test session and the partnered responder sessions π^* resp. π_r^* , 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:

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

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^* encapsulating against $\text{epk}_{V^*}^{\text{KEM}}$. The probability of such an abort can be bounded by the advantage of the following reduction \mathcal{B}_8 playing against $(t', \epsilon_{\text{CCA}}, 1)$ -OW-CCA security of KEM.

The reduction embeds the challenge pk_{KEM}, ct^* from the OW-CCA game as ephemeral key $\text{epk}_{V^*}^{\text{KEM}}$ of V^* , and as KEM ciphertext in π^* . The reduction does not need to answer a REVEALRAND query on π_r^* due to clean_{EE} . If π_r^* receives a different ciphertext than ct^* , the reduction uses its DECAPS oracle (once) to compute the required shared secret. 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 $\mathcal{G}_{\text{F.3}}$. If \mathcal{A} does make such a random oracle query, then \mathcal{B}_8 wins its OW-CCA game if it guesses the

²⁰Again, this trivial key confusion attack against PQXDH would be avoided if its signatures were properly domain-separated; cf. the discussion in Section 4.3.

i th RO query correctly.²¹ We can thus bound the game hop by the advantage against OW-CCA times a loss of q_{RO} :

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

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 $\mathcal{G}_{\text{F.4}}$ 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,

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

To conclude this proof case, observe that in Game $\mathcal{G}_{\text{F.5}}$ 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:

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

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

5.2 Post-quantum Security of PQXDH

Next, we show that PQXDH achieves key indistinguishability against quantum adversaries wrt. the maximum-exposure attack vectors formalized through the clean predicates $\text{clean}^{\text{PQXDH-Q}}$ in our model (cf. Figure 9). In particular, these clean predicates model authentic distribution of the KEM keys (e.g., due to the adversary being passive or *not yet* able to forge signatures) and do not rely on the contribution of the DH shared secrets. This result is in the standard model and relies on the KEM providing IND-CCA and $LEAK^{+r}$ -BIND-SS- $\{CT, PK\}$ security, as well as low key-collision probability and correctness errors, and the KDF providing PRF security. The proof follows the same structure as the previous proof, but is shorter following the fewer cases of the $\text{clean}^{\text{PQXDH-Q}}$ predicate in which one can expect security against a quantum adversary.

Theorem 5.2 (Post-quantum key indistinguishability of PQXDH). *The PQXDH protocol given in Figure 10 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 wrt. the post-quantum security clean predicates $\text{clean}^{\text{PQXDH-Q}}$ (cf. Figure 9), assuming (for $t' \approx t$) KEM is a $(t', \epsilon_{\text{CCA}}, n_s)$ -IND-CCA-secure and $(t', \epsilon_{\text{LEAK}^{+r}}, n_p \cdot n_{ss} + n_s)$ - $LEAK^{+r}$ -BIND-SS- $\{CT, PK\}$ -secure KEM with public-key collision probability γ_{coll} and correctness error δ_{corr} , and KDF is $(t', \epsilon_{\text{PRF}}, n_s)$ -PRF-secure when keyed via the last 256 input bits. Concretely, for any efficient quantum adversary \mathcal{A} we have*

$$\begin{aligned} \text{Adv}_{\text{PQXDH-Q}}^{\text{KI}}(\mathcal{A}) &\leq \frac{(n_p + n_p \cdot n_{ss} + n_s)^2}{q} + \gamma_{\text{coll}}(n_p \cdot n_{ss} + n_s) + n_s \cdot \delta_{\text{corr}} + \epsilon_{\text{LEAK}^{+r}} \\ &\quad + n_{ss} \cdot n_s \cdot (\epsilon_{\text{CCA}} + \epsilon_{\text{PRF}}) \quad // \text{clean}_{\text{ESS}}^{\text{Q}} \\ &\quad + n_s^2 \cdot (\epsilon_{\text{CCA}} + \epsilon_{\text{PRF}}), \quad // \text{clean}_{\text{EE}}^{\text{Q}} \end{aligned}$$

where we write PQXDH-Q to indicate that we assess the post-quantum security of PQXDH wrt. to the $\text{clean}^{\text{PQXDH-Q}}$ clean predicates.

²¹Observe that this guessing loss does not arise in the proof cases D (Game $\mathcal{G}_{\text{D.5}}$) and E (Game $\mathcal{G}_{\text{E.2}}$), because the reduction there can leverage the DDH oracle of the GapDH game to identify the target random oracle query.

Proof. The proof against quantum adversaries is structurally similar to the proof against classical adversaries and likewise proceeds via a series of game hops. We first exclude collisions in the public keys and KEM decapsulation errors, yielding the terms $\frac{(n_p + n_p \cdot n_{ss} + n_s)^2}{q} + \gamma_{\text{coll}}(n_p \cdot n_{ss} + n_s)$ and $n_s \cdot \delta_{\text{corr}}$. Now, soundness holds. Next, we preclude the KEM re-encapsulation attack with a $LEAK^{+r}$ game hop, yielding the term $\epsilon_{LEAK^{+r}}$. Again, we branch out over the cases of the clean predicate – now in the quantum version with only two cases. In either case we guess the involved sessions and (ephemeral or semi-static) KEM public keys, then embed an IND-CCA challenge in the test session’s KEM encapsulation, and conclude with a PRF security hop from the now random KEM shared secret. Note that all of these game hops rely on standard-model assumptions, i.e., we do not need to rely on a (quantum) random oracle.

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

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

We proceed with Games 1, 2, and 3 as in the proof of Theorem 5.1, excluding collisions in DH and KEM public keys, KEM decapsulation errors, and collisions of KEM shared secrets. Similarly to the proof of Theorem 5.1, these game hops ensure soundness. With these three game hops we have

$$\text{Adv}_{\text{PQXDH-}Q}^{\mathcal{G}_0}(\mathcal{A}) \leq \frac{(n_p + n_p \cdot n_{ss} + n_s)^2}{q} + \gamma_{\text{coll}}(n_p \cdot n_{ss} + n_s) + n_s \cdot \delta_{\text{corr}} + \epsilon_{LEAK^{+r}} + \text{Adv}_{\text{PQXDH-}Q}^{\mathcal{G}_3}(\mathcal{A}).$$

Separating the clean^Q cases. At this point, we will divide the proof into two sub-cases following the $\text{clean}_{\text{PQXDH-}Q}^Q$ predicate with its sub-predicates $\text{clean}_{\text{EE}}^Q$ for full sessions and $\text{clean}_{\text{ESS}}^Q$ for reduced sessions. We can bound the advantage of the adversary in Game \mathcal{G}_3 by the sum of its advantages in Games $\mathcal{G}_3[\text{c}]$ which are \mathcal{G}_3 conditioned on a sub-predicate c :

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

In the remainder of the proof, we will bound both of these cases, called **A** and **B**, separately.

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

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

1. the initiator session’s randomness is not revealed, and
2. the responder’s semi-static KEM key is uncompromised, and
3. if π^* is an initiator, then the semi-static KEM key was honestly generated, and
4. if π^* is a responder, then there exists a partnered initiator session.

Hence, the semi-static KEM shared secret ensures key indistinguishability for the test session.

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

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

Game A.1 (Guess semi-static KEM key identifier $\text{ssid}_{\text{KEM}}^*$ of V^*). We now guess the `ssid` identifier of the responder V^* 's (uncorrupted) semi-static KEM key, $\text{sspk}_{V^*}^{\text{KEM}}$. 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 $\pi^*.role = \text{initiator}$). In the latter case, the `highlighted code` of $\text{clean}_{\text{ESS}}^Q$ (Figure 9, line 66) ensures that $\text{sspk}_{V^*}^{\text{KEM}}$ is indeed an honestly generated key. We denote the guessed identifier by $\text{ssid}_{\text{KEM}}^*$ 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:

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

Game A.2 (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 :

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

Game A.3 (IND-CCA). We replace the KEM shared secret ss with a uniformly random value ss^* in the initiator session π_i^* and test session π^* , as well as any session of V^* using the same semi-static KEM key $\text{sspk}_{V^*}^{\text{KEM}}$ and receiving the same ciphertext. The probability of the adversary detecting this change can be bounded by the the advantage of the following reduction \mathcal{B}_1 playing against the $(t', \epsilon_{\text{CCA}}, n_s)$ -IND-CCA security of KEM.

The reduction embeds the KEM public key pk_{KEM} from the IND-CCA game as the semi-static key $\text{sspk}_{V^*}^{\text{KEM}}$ of V^* with semi-static id $\text{ssid}_{\text{KEM}}^*$. Furthermore, it uses the KEM challenge ct^* , ss_b^* in π_i^* and ss_b^* in π^* . The reduction does not need to answer corresponding REVEALRAND queries on π_i^* or CORRUPTSSKEY(V^* , $\text{ssid}_{\text{KEM}}^*$) queries due to $\text{clean}_{\text{ESS}}^Q$. For SEND queries to further sessions of V^* with $\text{ssid}_{\text{KEM}}^*$, the reduction can use ss_b^* as decapsulation of ct^* and queries the DECAPS oracle of the IND-CCA game otherwise. Hence, the reduction makes at most n_s many DECAPS queries in the IND-CCA game.

If the IND-CCA game provides the real KEM shared secret as challenge, then \mathcal{B}_1 perfectly simulates $\mathcal{G}_{A.2}$. Otherwise, the IND-CCA challenge is a randomly sampled KEM shared secret and \mathcal{B}_1 perfectly simulates $\mathcal{G}_{A.3}$. If \mathcal{A} can distinguish these two games, then \mathcal{B}_1 can win its IND-CCA game. Thus,

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

Game A.4 (KDF pseudorandomness). We now replace KDF, when keyed with ss^* as per $\mathcal{G}_{A.3}$ through the last 256 input bits, with a random function. This in particular effectively replaces the session key of the test session π^* with a uniformly sampled key. We bound the advantage difference of this change with a reduction to $(t', \epsilon_{\text{PRF}}, n_s)$ -PRF security of KDF, where KDF is keyed in its last 256-bit input ss and treats the remaining inputs DH_1, DH_2, DH_3, DH_4 as label.

The reduction \mathcal{B}_2 does not sample ss^* itself but instead queries its PRF oracle in the sessions where it would need to compute KDF on input ss^* ; in particular in the test session π^* .

Depending on the secret bit in the PRF game, the reduction perfectly simulates $\mathcal{G}_{A.3}$ or $\mathcal{G}_{A.4}$. If \mathcal{A} can distinguish these two games, then \mathcal{B}_2 can win its pseudorandomness game. Thus,

$$\text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{A.3}}(\mathcal{A}) \leq \epsilon_{\text{PRF}} + \text{Adv}_{\text{PQXDH}}^{\mathcal{G}_{A.4}}(\mathcal{A}).$$

To conclude this proof case, observe that in Game $\mathcal{G}_{A.4}$ 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:

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

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

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

1. the test session's randomness is not revealed, and
2. there exists a (contributively) partnered session whose randomness is not revealed.

Via the last point, we can guarantee that initiator and responder agree on the ephemeral KEM key. Neither the ephemeral KEM key nor the KEM encapsulation randomness is compromised. Hence, the ephemeral KEM shared secret ensures key indistinguishability for the test session.

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

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

Game B.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:

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

Game B.2 (IND-CCA). We replace the KEM shared secret ss with a uniformly random value ss^* in the test session π^* . Furthermore, we replace ss with ss^* in π_i^* if π^* is a responder, and in π_r^* if π^* is an initiator and π_r^* received the KEM ciphertext sent by π^* . The probability of the adversary detecting this change can be bounded by the the advantage of the following reduction \mathcal{B}_3 playing against the $(t', \epsilon_{\text{CCA}}, 1)$ -IND-CCA security of KEM.

The reduction embeds the KEM public key pk_{KEM} from the IND-CCA game as the ephemeral KEM key $epk_{V^*}^{\text{KEM}}$ of V^* in the responder test session π_r^* . Furthermore, it uses the KEM challenge ct^* , ss_b^* in π_i^* and ss_b^* in π_r^* . The reduction does not need to answer corresponding REVEALRAND queries on π_i^* or π_r^* due to $\text{clean}_{\text{EE}}^Q$. If π_r^* receives ct^* , the reduction uses ss^* there as well; otherwise, the reduction uses its DECAPS oracle (once) to compute the required shared secret. Hence, the reduction makes at most 1 DECAPS query in the IND-CCA game.

If the IND-CCA game provides the real KEM shared secret as challenge, then \mathcal{B}_3 perfectly simulates $\mathcal{G}_{B.1}$. Otherwise, the IND-CCA challenge is a randomly sampled KEM shared secret and \mathcal{B}_3 perfectly simulates $\mathcal{G}_{B.2}$. If \mathcal{A} can distinguish these two games, then \mathcal{B}_3 can win its IND-CCA game. Thus,

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

Game B.3 (KDF pseudorandomness). We now replace KDF, when keyed with ss^* as per $\mathcal{G}_{B.2}$ through the last 256 input bits, with a random function. This in particular effectively replaces the session key of the test session π^* with a uniformly sampled key. We bound the advantage difference of this change with a reduction to $(t', \epsilon_{\text{PRF}}, 2)$ -PRF security of KDF, where KDF is keyed in its last 256-bit input ss and treats the remaining inputs DH_1, DH_2, DH_3, DH_4 as label.

The reduction \mathcal{B}_4 does not sample ss^* itself but instead queries its PRF oracle in the sessions where it would need to compute KDF on input ss^* ; i.e., in the test session π^* and possibly its partner session.

Depending on the secret bit in the PRF game, the reduction perfectly simulates $\mathcal{G}_{B.2}$ or $\mathcal{G}_{B.3}$. If \mathcal{A} can distinguish these two games, then \mathcal{B}_4 can win its pseudorandomness game. Thus,

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

To conclude this proof case, observe that in Game $\mathcal{G}_{B.3}$ 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:

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

Collecting the bounds from all cases yields the overall theorem bound. □

6 Conclusion and Discussion

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 fully-parameterized, concrete bounds for PQXDH’s classical and post-quantum security.

PQXDH design discussion. Our bound relies (among other things) on the $LEAK^{+r}$ -BIND- SS - $\{CT, PK\}$ binding property of the KEM. While we show that this property is satisfied by both Kyber and ML-KEM, future KEMs deployed in PQXDH or similar protocols 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 the key derivation of PQXDH takes the KEM public key and ciphertext (or, ideally, the whole session context as is good practice) as additional arguments. This supports and complements a proposal discussed in the tool-based formal verification of PQXDH by BJKS [BJKS24], 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].

The lack of domain separation in PQXDH when signing ephemeral and semi-static KEM public keys allows for unnecessary, trivial key confusion attacks. Our analysis further confirms and pinpoints the stronger security guarantees achievable if PQXDH were to properly domain-separate these keys.

Our classical security 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 or against a *passive* quantum adversary. Hence, PQXDH protects against an even stronger class of attacks than the “harvest now, decrypt later” attack which motivated the design.

Limitations and future work. Our analysis is focused on the PQXDH handshake, analyzing both (full and reduced) modes in isolation. We hence do not treat attacks across protocol modes or versions (X3DH vs. PQXDH), but note that BJKS [BJKS24] highlighted the importance of domain separating the KDF calls to avoid confusion between X3DH in full mode and PQXDH in reduced mode.

Prior works on Signal’s initial handshake have assumed authentic distribution of public keys instead of modelling signatures [CCD⁺17, BFG⁺22], treated the signing keys informally [VGIK20], or, like our work, assumed that long-term signing keys and long-term DH keys are distinct [FJ24, BJKS24, CHN⁺24]. We leave it to future work to more accurately reflect the implementation which re-uses the same keys for both purposes.

It remains open to find the minimal changes necessary to PQXDH that guarantee security against active quantum adversaries, i.e., showing security without assuming honest distribution of KEM keys. In Section 1 we discussed related work [HKKP22, BFG⁺22, DG22, CHN⁺24] that proposes post-quantum *replacements* for X3DH and PQXDH. Preferably, these minimal changes would also preserve deniability of the protocol.

Acknowledgments

We thank Ehren Kret and Rolfe Schmidt for discussions on the PQXDH protocol, Franziskus Kiefer for insights on the potential KEM re-encapsulation attack on PQXDH, Keitaro Hashimoto, Shuichi Katsumata, and Thom Wiggers for pointing out an oversight in a previous version of our `cleansigE` predicate, and the anonymous reviewers for their insightful comments. R.F. was supported by the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE.

References

- [AHK⁺23] Joël Alwen, Dominik Hartmann, Eike Kiltz, Marta Mularczyk, and Peter Schwabe. Post-quantum multi-recipient public key encryption. In Weizhi Meng, Christian Damsgaard Jensen, Cas Cremers, and Engin Kirda, editors, *ACM CCS 2023: 30th Conference on Computer and Communications Security*, pages 1108–1122, Copenhagen, Denmark, November 26–30, 2023. ACM Press. (Cited on page 4.)
- [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, Karoline Varner, and Bas Westerbaan. X-wing. *IACR Communications in Cryptology (CiC)*, 1(1):21, 2024. (Cited on page 4.)
- [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 – EURO-*

CRYPT 2008, volume 4965 of *Lecture Notes in Computer Science*, pages 181–197, Istanbul, Turkey, April 13–17, 2008. Springer Berlin Heidelberg, Germany. (Cited on page 6.)

- [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, Cham, Switzerland. (Cited on page 4.)
- [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. (Cited on pages 3, 17, 19, 20, 21, 22, and 46.)
- [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, Cham, Switzerland. (Cited on pages 3, 4, 17, 19, 20, 21, 22, 46, and 47.)
- [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 and 3.)
- [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 and 3.)
- [BJKS24] Karthikeyan Bhargavan, Charlie Jacomme, Franziskus Kiefer, and Rolfe Schmidt. Formal verification of the PQXDH post-quantum key agreement protocol for end-to-end secure messaging. In Davide Balzarotti and Wenyuan Xu, editors, *USENIX Security 2024: 33rd USENIX Security Symposium*, Philadelphia, PA, USA, August 14–16, 2024. USENIX Association. (Cited on pages 2, 3, 4, 8, 23, 46, and 47.)
- [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 Berlin Heidelberg, Germany. (Cited on page 17.)
- [CCD⁺17] Katriel Cohn-Gordon, Cas Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the Signal messaging protocol. In *2017 IEEE European Symposium on Security and Privacy*, pages 451–466, Paris, France, April 26–28, 2017. IEEE Computer Society Press. (Cited on pages 2, 3, 6, 17, 19, 22, 25, 26, 29, 33, 35, 46, and 47.)
- [CCG16] Katriel Cohn-Gordon, Cas J. F. Cremers, and Luke Garratt. On post-compromise security. In Michael Hicks and Boris Köpf, editors, *CSF 2016: IEEE 29th Computer Security Foundations Symposium*, pages 164–178, Lisbon, Portugal, June 27–1, 2016. IEEE Computer Society Press. (Cited on page 17.)

- [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, Cham, Switzerland. (Cited on page 46.)
- [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, Cham, Switzerland. (Cited on page 4.)
- [CDM24] Cas Cremers, Alexander Dax, and Niklas Medinger. Keeping up with the KEMs: Stronger security notions for KEMs and automated analysis of KEM-based protocols. In Bo Luo, Xiaojing Liao, Jun Xu, Engin Kirda, and David Lie, editors, *ACM CCS 2024: 31st Conference on Computer and Communications Security*, pages 1046–1060, Salt Lake City, UT, USA, October 14–18, 2024. ACM Press. (Cited on pages 3, 4, 8, and 9.)
- [CHN⁺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 Davide Balzarotti and Wenyuan Xu, editors, *USENIX Security 2024: 33rd USENIX Security Symposium*, Philadelphia, PA, USA, August 14–16, 2024. USENIX Association. (Cited on pages 4 and 47.)
- [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 18.)
- [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*, pages 422–450, Virtual Event, September 28–30, 2022. Springer, Cham, Switzerland. (Cited on pages 4 and 47.)
- [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 17.)
- [FJ24] Rune Fiedler and Christian Janson. A deniability analysis of Signal’s initial handshake PQXDH. *Proceedings on Privacy Enhancing Technologies*, 2024(4):907–928, October 2024. (Cited on pages 2 and 47.)
- [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 Berlin Heidelberg, Germany. (Cited on page 11.)
- [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 pages 4 and 47.)

- [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, Cham, Switzerland. (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, Cham, Switzerland. (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 2, 3, and 26.)
- [LHT16] A. Langley, M. Hamburg, and S. Turner. Elliptic Curves for Security. RFC 7748 (Informational), January 2016. (Cited on page 26.)
- [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 Berlin 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, Cham, Switzerland. (Cited on page 4.)
- [MP16a] Moxie Marlinspike and Trevor Perrin. The Double Ratchet algorithm, November 2016. <https://www.signal.org/docs/specifications/doubleratchet/>. (Cited on page 17.)
- [MP16b] Moxie Marlinspike and Trevor Perrin. The X3DH key agreement protocol, November 2016. <https://signal.org/docs/specifications/x3dh/>. (Cited on page 1.)
- [MvOV97] 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.)
- [Nat24] National Institute of Standards and Technology. Module-lattice-based key-encapsulation mechanism standard. Technical Report Federal Information Processing Standards Publications (FIPS PUBS) 203, U.S. Department of Commerce, Washington, D.C., 2024. (Cited on pages 3, 7, 8, and 11.)
- [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 Berlin Heidelberg, Germany. (Cited on page 6.)

- [Per16] Trevor Perrin. The XEdDSA and VEdDSA signature schemes, October 2016. <https://signal.org/docs/specifications/xeddsa/>. (Cited on page 26.)
- [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, Cham, Switzerland. (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, Cham, Switzerland. (Cited on page 4.)
- [Rog06] Phillip Rogaway. Formalizing human ignorance. In Phong Q. Nguyen, editor, *Progress in Cryptology - VIETCRYPT 06: 1st International Conference on Cryptology in Vietnam*, volume 4341 of *Lecture Notes in Computer Science*, pages 211–228, Hanoi, Vietnam, September 25–28, 2006. Springer Berlin Heidelberg, Germany. (Cited on pages 5 and 52.)
- [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, 7, 8, and 11.)
- [Sch24] Sophie Schmieg. Unbindable kemmy schmidt: ML-KEM is neither MAL-BIND-K-CT nor MAL-BIND-K-PK. Cryptology ePrint Archive, Report 2024/523, 2024. (Cited on pages 4, 8, and 10.)
- [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, Cham, Switzerland. (Cited on pages 2 and 47.)
- [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 Berlin 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 Berlin Heidelberg, Germany. (Cited on page 5.)
- [Zha15] Mark Zhandry. A note on the quantum collision and set equality problems. *Quantum Inf. Comput.*, 15(7&8):557–567, 2015. (Cited on page 12.)

A Summary of Major Changes

- **version 1.0 – May 2024:** Initial release
- **version 2.0 – August 2024:**
 - added post-quantum(-only) analysis (Section 5.2) via separate clean predicates (Figure 9) and theorem statement (Theorem 5.2) for quantum adversaries
 - discuss key confusion between DH and KEM keys as well as ephemeral and semi-static KEM keys in Sections 1 and 4.3 and incorporate them into Theorems 5.1 and 5.2
 - define and use hash functions with the “human ignorance” approach [Rog06] (cp. Definitions 2.4 and 2.5, Theorems 3.6 to 3.11)
 - rename the adversarial model $LEAK^+$ for the KEM binding notion to $LEAK^{+r}$
- **version 2.1 – May 2025 (corresponding to the PKC 2025 proceedings version):**
 - fix an oversight in the $\text{clean}_{\text{sigE}}$ predicate (Figure 9) pointed out by Keitaro Hashimoto, Shuichi Katsumata, and Thom Wiggers