

Post-quantum Asynchronous Deniable Key Exchange and the Signal Handshake

Version 1.1, August 2021*

Jacqueline Brendel
CISPA Helmholtz Center for
Information Security
jacqueline.brendel@cispa.de

Rune Fiedler
Technische Universität Darmstadt
rune.fiedler@cryptoplexity.de

Felix Günther
ETH Zürich
mail@felixguenther.info

Christian Janson
Technische Universität Darmstadt
christian.janson@cryptoplexity.de

Douglas Stebila
University of Waterloo
dstebila@uwaterloo.ca

ABSTRACT

The key exchange protocol that establishes initial shared secrets in the handshake of the Signal end-to-end encrypted messaging protocol has several important characteristics: (1) it runs asynchronously (without both parties needing to be simultaneously online), (2) it provides implicit mutual authentication while retaining deniability (transcripts cannot be used to prove either party participated in the protocol), and (3) it retains security even if some keys are compromised (forward secrecy and beyond). All of these properties emerge from clever use of the highly flexible Diffie–Hellman protocol.

While quantum-resistant key encapsulation mechanisms (KEMs) can replace Diffie–Hellman key exchange in some settings, there is no KEM-based replacement for the Signal handshake that achieves all three aforementioned properties, in part due to the inherent asymmetry of KEM operations. In this paper, we show how to construct asynchronous deniable key exchange by combining KEMs and designated verifier signature (DVS) schemes. There are several candidates for post-quantum DVS schemes, either direct constructions or via ring signatures. This yields a template for an efficient post-quantum realization of the Signal handshake with the same asynchronicity and security properties as the original Signal protocol.

KEYWORDS

authenticated key exchange, deniability, asynchronous, Signal protocol, post-quantum, designated verifier signatures

1 INTRODUCTION

The Signal protocol [69, 70], designed by Marlinspike and Perrin, has enabled mass adoption of end-to-end encrypted messaging in consumer applications such as WhatsApp, Signal, Facebook Messenger, Skype, and more. From a cryptographic perspective, the Signal protocol consists of an initial handshake and key exchange (called “X3DH” [70], a simplified version of which is shown in Figure 1), asymmetric and symmetric key exchange “ratchets” that establish new keys for every new chat message sent (called the

“double ratchet” algorithm [69]), and symmetric authenticated encryption for application data. Each of these components contributes to Signal’s interesting and useful security features:

- *Implicit mutual authentication in the handshake*: The session key K established in the handshake can only be computed by the intended peer. This comes from the terms involving the long-term secret keys a and b in Figure 1.
- *Forward secrecy in the handshake*: The session key K established in the handshake remains secret even if long-term keys are later compromised. This comes from the terms involving the ephemeral keys x and y in Figure 1.
- *Offline deniability [30, 55] of the handshake*: A judge seeing a transcript of an honest communication session cannot be convinced that a particular party was actually involved in the session. This comes from the use of Diffie–Hellman for authentication rather than signatures; all of the DH shared secrets input to the key derivation function in Figure 1 could have been computed unilaterally either by Alice or by Bob (e.g., both Alice and Bob can compute g^{as} , using a and s respectively). See [84] for a detailed analysis of the deniability of X3DH.
- *Asynchronicity*: The two communicating parties need never be online simultaneously, and can leave packets at an untrusted relay server until the other party comes back online. The handshake is made asynchronous by allowing each party to upload a *pre-key bundle* to an untrusted server in advance, consisting of long-term, medium-term, and ephemeral public keys, and an initiator can start sending text messages before their peer comes online. The restrictions on communication flow in an asynchronous protocol are weaker than those of non-interactive key exchange [43].
- *Forward secrecy and post-compromise security [23] in long-lived conversations*: Keys are updated using a new DH key exchange with each chat message via the asymmetric ratchet, enabling secrecy of past and future messages after a compromise.

1.1 Making Signal Post-quantum

Since the Diffie–Hellman problem upon which much of Signal relies is not secure against quantum adversaries, it is important to have a post-quantum alternative available.

The symmetric ratchet and authenticated encryption components of Signal are built on symmetric primitives, and thus are not

*Please see Appendix C for an overview of major changes from earlier versions of this work.




Figure 1: Simplified version of Signal’s X3DH handshake. Long-term keys a and b ; semi-static key s ; ephemeral keys x and y .

in immediate danger from quantum algorithms. The asymmetric ratchet was phrased by Marlinspike and Perrin [69] and analyzed by Cohn-Gordon, Cremers, Dowling, Garratt, and Stebila [22] in terms of Diffie–Hellman. Alwen, Coretti, and Dodis [1] generalized it into a primitive called continuous key agreement that can be built from KEMs, yielding post-quantum security. Hence, our focus in the rest of this paper is thus on the handshake.

The post-quantum primitives to be standardized by the United States National Institute of Standards and Technology (NIST) post-quantum standardization project are signatures and key encapsulation mechanisms (KEMs), so these would be most preferable to employ. It is certainly possible to generically construct an authenticated key exchange protocol from signatures and KEMs, but it is not possible to use *only* KEMs and signatures in a generic way to create a post-quantum replacement for Signal with all of the properties listed above. Suppose one tried to use KEMs instead of Diffie–Hellman in Figure 1. Recall that, to use a KEM for key exchange, one party uses the key generation algorithm to create a public-key/secret-key pair and transmits the public key to their peer; the peer encapsulates against that public key, producing a ciphertext and a shared secret, then transmits the ciphertext, which the first party decapsulates using their secret key to compute the shared secret. In the Signal handshake, one could try using KEM public keys to replace the Diffie–Hellman shares in Alice and Bob’s pre-key bundles. We can still obtain ephemeral key exchange (by having Alice encapsulate against Bob’s ephemeral public key) and implicit Bob-to-Alice authentication (by having Alice encapsulate against Bob’s long-term public key). However, we cannot obtain Alice-to-Bob authentication using KEMs without adding an extra flow: Bob cannot produce a ciphertext for Alice to decapsulate without knowing Alice’s public key first, so he cannot asynchronously produce a pre-key bundle for Alice to immediately use. This highlights the difference between Diffie–Hellman and KEMs: in DH, both parties’ shares are objects of the same type and can be generated independently, but in generic KEMs, public keys and ciphertexts are in principle objects of differing types and a ciphertext is generated with respect to a given public key. To obtain Alice-to-Bob authentication without adding an extra communication round, Alice could of course produce a signature for Bob to verify, but this undermines deniability.

The problem, in a nutshell, is to create an *asynchronous deniable authenticated key exchange protocol* that can be instantiated in the post-quantum setting, preferably with an efficient construction based on standardized primitives or at least cryptographic assumptions used in standardized primitives.

1.2 Options for PQ Asynchronous DAKE

There are several examples of authenticated key exchange protocols built generically from KEMs which have the potential for deniability [12, 13, 28, 44, 79] but do not have the desired asynchronicity property for reasons similar to the discussion above.

One post-quantum option that avoids the problem with KEMs described above is to use CSIDH [20], a primitive based on supersingular isogenies that yields a commutative group action which enables non-interactive key exchange. CSIDH could be used to achieve implicit Alice-to-Bob authentication while maintaining asynchronicity and deniability; indeed several key exchange protocols from CSIDH have been proposed [27, 54]. Unfortunately, there are several reasons CSIDH may not be a fully satisfactory solution: it is much more computationally expensive than most other forms of post-quantum cryptography; there is ongoing debate about the security of its concrete parameters [11, 73]; and the decisional form of a related problem [21] is not hard.

Most other post-quantum assumptions used in KEMs, including SIDH [52] and learning-with-errors (LWE) [76], are insecure against key reuse attacks without additional protection such as the Fujisaki–Okamoto transform [45] that leaves them unable to be used for non-interactive key exchange (since the ciphertext must be generated with respect to a given public key). There have been several attempts at SIDH-based non-interactive key exchange which have ended up being insecure [2, 32, 33, 36], and one attempt relying on an additional novel assumption [10] the security of which is unknown.

Brendel, Fischlin, Günther, Janson, and Stebila [15] previously considered the question of building a post-quantum version of the Signal handshake, highlighting many of these problems. They proposed decomposing the three operations of a KEM into a 4-operation “split KEM”, and showed how a Signal-like handshake could be built from a split KEM meeting a suitably strong security notion. They showed how CSIDH and LWE could be used to build split KEMs meeting a weaker security notion, but these constructions did not achieve the strong security notion required for their Signal-like handshake, effectively leaving the overall problem unsolved.

Unger and Goldberg [82, 83] also consider deniable authenticated key exchange (DAKE) protocols for secure messaging. Their protocol permits the optional use of a PQ KEM for ephemeral key exchange to achieve forward secrecy against future-quantum adversaries. To achieve deniability, they employ ring signatures with classical security and further rely on dual receiver encryption, which does not yet appear to have a PQ instantiation in the literature.

The recent work by Hashimoto, Katsumata, Kwiakowski, and Prest [48] is closest to ours. Their core protocol is meant to replace the Signal handshake based on (post-quantum) KEMs and signatures. It achieves security against exposure of long-term keys and session state and a weaker deniability level by encrypting the (regular) signature exchanged. Unlike Signal (and our proposed protocol), it does however not provide security against randomness exposure and lacks support for semi-static keys to mitigate the exhaustion of ephemeral pre-keys. Hashimoto et al. provide an implementation for their weakly-deniable protocol and further discuss a theoretical variant achieving stronger deniability based on ring signatures, strong knowledge-type assumptions for plaintext-aware [4] KEMs,

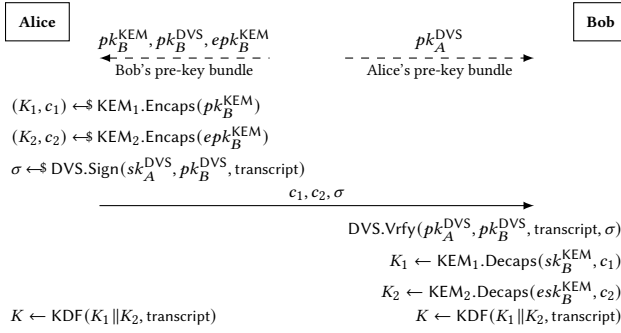


Figure 2: Our core asynchronous DAKE protocol, combining static and ephemeral key encapsulation schemes KEM_1 and KEM_2 , and a designated verifier signature DVS.

and non-interactive zero-knowledge arguments; the post-quantum security of these components is yet to be assessed.

1.3 Our Contributions

We show how to construct an asynchronous deniable authenticated key exchange protocol generically from designated verifier signature schemes and key encapsulation mechanisms.

Introduced by Jakobsson, Sako, and Impagliazzo [51], a *designated verifier signature (DVS) scheme* allows a signer to convince a chosen recipient, called the designated verifier, of the authenticity of a message, but in such a way that the designated verifier cannot convince any other party of the authenticity. In a DVS scheme, both the signer and the verifier have a public-key/secret-key pair; signing requires both the signer’s secret key and the verifier’s public key, and verification uses both parties’ public keys. To achieve the non-transferability property (called “source hiding”), a DVS scheme is accompanied by an additional simulation algorithm with which the designated verifier can, using its own secret key, construct a signature indistinguishable from one generated by the signer.

Asynchronous DAKE construction. We combine a DVS with a KEM to achieve an asynchronous deniable authenticated key exchange as shown in Figure 2. As expected, Bob-to-Alice authentication comes from an implicitly authenticated key exchange in which Alice encapsulates to Bob’s long-term KEM key (KEM_1 with long-term public key pk_B^{KEM} and ciphertext c_1 in Figure 2), and forward secrecy comes from a key exchange using an ephemeral KEM key (KEM_2 with public key epk_B^{KEM} and ciphertext c_2). Alice-to-Bob authentication comes from Alice using the designated verifier signature scheme to sign a transcript with Bob as the designated verifier; she can obtain Bob’s DVS verification key (pk_B^{DVS}) from his pre-key bundle. Since the source hiding property of the DVS scheme enables Bob to also have created a valid-looking signature from Alice with himself as the designated verifier, the transcript of the key exchange protocol could have been constructed by either Alice or Bob, yielding the desired deniability property.

Post-quantum designated verifier signatures. To achieve our goal of post-quantum asynchronous DAKE, we thus need a post-quantum designated verifier signature scheme. While there is a long line of research on DVS schemes from pre-quantum assumptions (including [16, 26, 51, 59, 62, 78, 80, 87]), comparatively little is available in the

literature on post-quantum DVS schemes. An isogeny-based DVS scheme was proposed in [81] but is insecure due to key reuse attacks identified in [46]. There are several lattice-based DVS schemes which may fit the bill [61, 71, 85, 86, 89], but these have not received much scrutiny in the mainstream cryptographic literature; we summarize this literature in Section 3.1.1. These lattice-based DVS schemes are direct constructions not based on any NIST candidates, so they would require their own thorough analysis.

DVS from ring signatures. Rather than constructing DVS schemes directly, it is possible to use a *ring signature scheme* [77] as a designated verifier signature scheme. In a ring signature scheme, one signer can sign a message intended to verify under a *ring* of public keys, only one of which is theirs; yet no one should be able to determine which signer produced such a signature. Following ideas sketched in [7, 77], we show in Section 3.2 how to use a 2-user ring signature scheme to build a DVS scheme: the ring used by the signer consists of the public keys of themselves and the one designated verifier. There are several candidates for post-quantum ring signatures whose properties we discuss in Section 3.2.3.

Doing the reverse—constructing a ring-signature scheme from a DVS scheme—seems more difficult, and at this point we have neither an implication nor separation. Should there indeed be a separation, it makes sense to try to build a DAKE protocol out of the weaker primitive of designated verifier signatures to open up the most options for construction.

Application to Signal handshake. We present a version of the Signal X3DH handshake which we call SPQR—Signal in a Post-Quantum Regime—based on our asynchronous DAKE design that uses KEMs and a designated verifier signature scheme. We show that the SPQR handshake achieves strong (“maximal-exposure”) session key security in a variant of the security model of [22] covering compromises of long- and medium-term keys and ephemeral randomness, as well as deniability.

Outline of the paper. In Section 2 we introduce preliminaries. Section 3 focuses on the security properties of designated verifier schemes and how to construct these in a post-quantum setting, including existing direct constructions as well as via ring signatures, and gives a discussion of our failed attempts at constructing DVS from chameleon hash functions in an earlier version of this work. In Section 4 we present a security model for key exchange that captures session key indistinguishability with implicit mutual authentication and weak forward secrecy, as well as offline deniability. In Section 5 we show that our core asynchronous deniable authenticated key exchange protocol from Figure 2 fulfills these security notions; in particular, offline deniability is based on the source hiding property of the DVS scheme. In Section 6 we introduce a complete post-quantum version of the Signal handshake that extends our core protocol to include additional components present in the Signal handshake (e.g., semi-static keys). In Section 7 we provide a security model for our full protocol and prove, in Section 8, its session key indistinguishability and deniability. In Section 9, we conclude with a discussion of the results and some limitations.

2 PRELIMINARIES

We begin by introducing notation and recapping some basic components.

2.1 Notation

To sample an element x uniformly at random from a set \mathcal{S} (or a distribution on an underlying set) we write $x \leftarrow \mathcal{S}$. For deterministic algorithms A we denote by $y \leftarrow A(x)$ the execution of A on input x with output y . Similarly, $y \leftarrow_s A(x)$ denotes the probabilistic execution of A , and $y \leftarrow A(x; r)$ the deterministic execution of a probabilistic algorithm A with its random coins fixed to r . Adversaries are typically denoted by \mathcal{A} and we write $\mathcal{A}^{\text{ORACLE}}$ to indicate that \mathcal{A} has access to the oracle ORACLE . Adversaries can have local quantum computation power but their oracle access and outputs are still classical. For an integer n , we denote by $[n]$ the set $\{1, \dots, n\}$. Double square brackets $\llbracket \cdot \rrbracket$ that enclose a boolean statement return the bit 1 if the statement is true, and 0 otherwise.

2.2 Key Encapsulation Mechanisms

The main building block for our post-quantum secure initial key agreement of Signal are so-called *key encapsulation mechanisms* that allow an *encapsulator* to transfer a shared secret key K via a ciphertext c to the *decapsulator*.

Definition 2.1 (Key Encapsulation Mechanisms). A *key encapsulation mechanism* KEM is a triple of algorithms $\text{KEM} = (\text{KGen}, \text{Encaps}, \text{Decaps})$. In more detail:

- $\text{KGen}() \mapsto (pk, sk)$: A probabilistic algorithm that outputs a public-key/secret-key pair with $(pk, sk) \in \mathcal{PK} \times \mathcal{SK}$.
- $\text{Encaps}(pk) \mapsto (c, K)$: A probabilistic algorithm taking as input a public key $pk \in \mathcal{PK}$ and outputs a ciphertext $c \in \mathcal{C}$ and the therein encapsulated key $K \in \mathcal{K}$.
- $\text{Decaps}(sk, c) \rightarrow K'$: A deterministic algorithm taking as input a ciphertext $c \in \mathcal{C}$ and secret key sk and outputs $K' \in \mathcal{K} \cup \{\perp\}$, where \perp indicates an error.

We say that a KEM $\text{KEM} = (\text{KGen}, \text{Encaps}, \text{Decaps})$ is δ -correct if, for every key pair $(pk, sk) \leftarrow_s \text{KGen}()$, and every encapsulation $(c, K) \leftarrow_s \text{Encaps}(pk)$, we have

$$\Pr[K' \neq K \mid K' \leftarrow \text{Decaps}(sk, c)] \leq \delta.$$

We call KEM (*perfectly*) *correct* if $\delta = 0$.

Security of KEMs is defined in terms of indistinguishability of encapsulated keys from random given the decapsulator's public key and the encapsulating ciphertext:

Definition 2.2 (IND-ATK Security of KEMs). Let $\text{KEM} = (\text{KGen}, \text{Encaps}, \text{Decaps})$ be a KEM with key space \mathcal{K} . We say that KEM is (t, ϵ) -IND-CPA-secure, resp. (t, ϵ, Q_D) -IND-CCA-secure, if for any adversary \mathcal{A} with running time at most t and (for IND-CCA) making at most Q_D queries to the DECAPS oracle, we have that

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

where $\mathcal{G}_{\text{KEM}}^{\text{indatk}}(\mathcal{A})$ (with $\text{atk} = \text{cpa}$, resp. $\text{atk} = \text{cca}$) is defined in Figure 3.

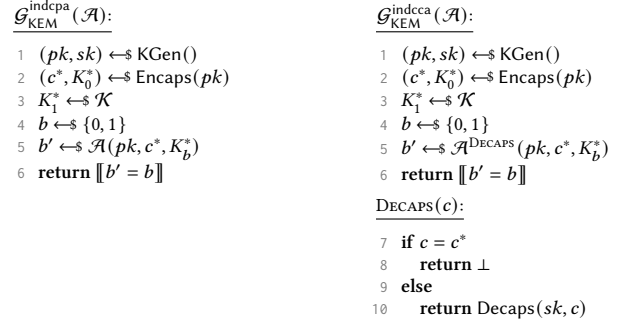


Figure 3: IND-CPA and IND-CCA security for KEM = $(\text{KGen}, \text{Encaps}, \text{Decaps})$ with key space \mathcal{K} .

2.3 (Twisted) Pseudorandom Functions

Beyond classical pseudorandom functions for key derivation, another crucial component for our SPQR protocol are special pseudorandom functions called *twisted pseudorandom functions* [44, 58]. In the following we recall the respective definitions and security games.

Definition 2.3. Let $F : \{0, 1\}^\kappa \times \{0, 1\}^t \rightarrow \{0, 1\}^\omega$ be an efficient keyed function with key length κ , input length t , and output length ω .

Let $\mathcal{G}_F^{\text{prfsec}}(\mathcal{A})$ be defined as in the top of Figure 4. We call F a (t, ϵ, Q_F) -pseudorandom function (or simply (t, ϵ, Q_F) -PRFSEC), if for any adversary \mathcal{A} with running time at most t and making at most Q_F queries to the PRFCHALLENGE oracle, we have that

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

Let $\mathcal{G}_F^{\text{tprfsec}}(\mathcal{A})$ be defined as in the bottom of Figure 4. We call F a (t, ϵ, q) -twisted pseudorandom function (or simply (t, ϵ, q) -tPRFSEC), if for any adversary \mathcal{A} with running time at most t , we have that

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

Note that one can easily build a twisted PRF tPRF from a PRF F in the standard model. Following Kurosawa and Furukawa [58], a secure construction doubling the key and label lengths is:

$$\text{tPRF}((k, k'), (e, e')) = F(k, e) \oplus F(e', k').$$

3 DESIGNATED VERIFIER SIGNATURES

Designated verifier signature (DVS) schemes were introduced by Jakobsson, Sako, and Impagliazzo [51]. Their goal is for a signer to convince a chosen recipient (the “designated verifier”) that a message is authentic but in such a way that the designated verifier cannot convince any other party of the authenticity of the message. This property is typically modeled by requiring that the designated verifier can efficiently simulate signatures that are indistinguishable from signatures produced by the signer.

<p>$\mathcal{G}_F^{\text{prfsec}}(\mathcal{A})$:</p> <ol style="list-style-type: none"> 1 $K \leftarrow \{0, 1\}^\kappa$ 2 $g \leftarrow \{\text{functions } f : \{0, 1\}^t \rightarrow \{0, 1\}^\omega\}$ 3 $b \leftarrow \{0, 1\}$ 4 $b' \leftarrow \mathcal{A}^{\text{PRFCHALLENGE}}()$ 5 return $\llbracket b' = b \rrbracket$ <p>$\text{PRFCHALLENGE}(x)$:</p> <ol style="list-style-type: none"> 6 if $b = 0$ 7 return $F(K, x)$ 8 else 9 return $g(x)$ <hr/> <p>$\mathcal{G}_{F,g}^{\text{trfsec}}(\mathcal{A})$:</p> <ol style="list-style-type: none"> 1 $g \leftarrow \{\text{functions } f : \{0, 1\}^t \rightarrow \{0, 1\}^\omega\}$ 2 $g' \leftarrow \{\text{functions } f : \{0, 1\}^\kappa \rightarrow \{0, 1\}^\omega\}$ 3 $K, K' \leftarrow \{0, 1\}^{2\kappa}$ 4 $b \leftarrow \{0, 1\}$ 5 $x, x_1, x_2, \dots, x_q \leftarrow \{0, 1\}^{(q+1)t}$ 6 $s_0 \leftarrow \{(x_1, F(K, x_1)), (x_2, F(K, x_2)), \dots, (x_q, F(K, x_q)), (K', F(K', x))\}$ 7 $s_1 \leftarrow \{(x_1, g(x_1)), (x_2, g(x_2)), \dots, (x_q, g(x_q)), (K', g'(K'))\}$ 8 $b' \leftarrow \mathcal{A}(s_b)$ 9 return $\llbracket b' = b \rrbracket$ 	<p>$\mathcal{G}_{\text{DVS}}^{\text{uf}}(\mathcal{A})$:</p> <ol style="list-style-type: none"> 1 $Q \leftarrow \emptyset$ 2 $\mathcal{L} \leftarrow \emptyset$ 3 $(pk_S, sk_S) \leftarrow \text{DVS.SKGen}()$ 4 $(pk_D, sk_D) \leftarrow \text{DVS.VKGen}()$ 5 for $i \in [n]$ 6 $(pk_i, sk_i) \leftarrow \text{DVS.VKGen}()$ 7 $\mathcal{L} \leftarrow \mathcal{L} \cup \{(pk_i, sk_i)\}$ 8 $(m^*, \sigma^*) \leftarrow \mathcal{A}^{\text{SIGN}}(pk_S, pk_D, \mathcal{L})$ 9 $d \leftarrow \text{DVS.Vrfy}(pk_S, pk_D, m^*, \sigma^*)$ 10 return $\llbracket d = \text{true} \wedge m^* \notin Q \rrbracket$ <hr/> <p>$\mathcal{G}_{\text{DVS}}^{\text{srchid}}(\mathcal{A})$:</p> <ol style="list-style-type: none"> 1 $(pk_S, sk_S) \leftarrow \text{DVS.SKGen}()$ 2 $(pk_D, sk_D) \leftarrow \text{DVS.VKGen}()$ 3 $b \leftarrow \{0, 1\}$ 4 $b' \leftarrow \mathcal{A}^{\text{CHALL}}(pk_S, sk_S, pk_D, sk_D)$ 5 return $\llbracket b' = b \rrbracket$
	<p>$\text{SIGN}(pk, m)$:</p> <ol style="list-style-type: none"> 11 if $pk = pk_D$ 12 $Q \leftarrow Q \cup \{m\}$ 13 else if $(pk, \cdot) \notin \mathcal{L}$ 14 return \perp 15 $\sigma \leftarrow \text{DVS.Sign}(sk_S, pk, m)$ 16 return σ <hr/> <p>$\text{CHALL}(m)$:</p> <ol style="list-style-type: none"> 6 if $b = 0$ 7 $\sigma \leftarrow \text{DVS.Sign}(sk_S, pk_D, m)$ 8 else 9 $\sigma \leftarrow \text{DVS.Sim}(pk_S, sk_D, m)$ 10 return σ

Figure 5: Existential unforgeability under chosen-message attacks (top) and source hiding (bottom) of a designated verifier signature scheme DVS.

Figure 4: Pseudorandomness ($\mathcal{G}_F^{\text{prfsec}}(\mathcal{A})$, top) and twisted pseudorandomness ($\mathcal{G}_F^{\text{trfsec}}(\mathcal{A})$, bottom) of a function F .

Definition 3.1. A designated verifier signature scheme (DVS) is a tuple of algorithms $\text{DVS} = (\text{SKGen}, \text{VKGen}, \text{Sign}, \text{Vrfy}, \text{Sim})$ along with a message space \mathcal{M} .

- $\text{SKGen}() \rightarrow (pk_S, sk_S)$: A probabilistic key generation algorithm that outputs a public-/secret-key pair for the signer.
- $\text{VKGen}() \rightarrow (pk_D, sk_D)$: A probabilistic key generation algorithm that outputs a public-/secret-key pair for the verifier.
- $\text{Sign}(sk_S, pk_D, m) \rightarrow \sigma$: A probabilistic signing algorithm that uses a signer secret key sk_S to produce a signature σ for a message $m \in \mathcal{M}$ for a designated verifier with public key pk_D .
- $\text{Vrfy}(pk_S, pk_D, m, \sigma) \rightarrow \text{true/false}$: A deterministic verification algorithm that checks a message m and signature σ against a signer public key pk_S and verifier public key pk_D .
- $\text{Sim}(pk_S, sk_D, m) \rightarrow \sigma$: A probabilistic signature simulation algorithm that uses the verifier's secret key sk_D to produce a signature σ on message m for signer public key pk_S .

A DVS scheme DVS is *correct*, if, for any honestly generated key pairs $(pk_S, sk_S), (pk_D, sk_D)$ and every message $m \in \mathcal{M}$, it holds that

$$\Pr[\text{Vrfy}(pk_S, pk_D, m, \text{Sign}(sk_S, pk_D, m)) = \text{true}] = 1.$$

We follow Laguillaumie and Vergnaud [59] in defining separate key generation algorithms for signers and designated verifiers; in some cases these two algorithms may be identical. Some DVS schemes, called *strong* DVS, are written with a verification algorithm that requires the designated verifier's secret key instead of the public key.

A long line of research has scrutinized the security of DVS schemes in different settings, including [16, 26, 51, 59, 62, 78, 80, 87]. For the purpose of this paper, it suffices to define the security notions of *unforgeability* and *source hiding*. Unforgeability for DVS

schemes is similar to that for standard signature schemes, providing the adversary with a signing oracle and asking it to forge a signature on a (fresh) message of its choice. Prior work restricts the signing oracle to the challenge designated verifier key. In contrast, and to account for settings where a signer's key is used with many other users' verifier keys (cf. Section 5), we allow the adversary to pick the designated verifier key to be used in the signing oracle from a set of additional, honestly generated key pairs.

Definition 3.2. A designated verifier signature scheme DVS is (t, ϵ, n, Q_S) -*unforgeable* if, for any adversary \mathcal{A} with running time at most t , having access to n additional DVS verifier key pairs beyond the challenge keys, and making at most Q_S queries to the SIGN oracle, we have that

$$\text{Adv}_{\text{DVS}}^{\text{uf}}(\mathcal{A}) = \Pr[\mathcal{G}_{\text{DVS}}^{\text{uf}}(\mathcal{A}) = 1] \leq \epsilon,$$

where $\mathcal{G}_{\text{DVS}}^{\text{uf}}(\mathcal{A})$ is as in Figure 5.

The second property we consider is called source hiding [59], demanding that it should be infeasible for an adversary to determine whether a given signature has been generated by the signer (using Sign) or by the designated verifier (using Sim), even if the adversary learns the secret keys of both parties.

Definition 3.3. A designated verifier signature scheme DVS is (t, ϵ, Q_{Ch}) -*source hiding* if, for any adversary \mathcal{A} with running time at most t and making at most Q_{Ch} to the CHALL oracle, we have that

$$\text{Adv}_{\text{DVS}}^{\text{srchid}}(\mathcal{A}) = \left| \Pr[\mathcal{G}_{\text{DVS}}^{\text{srchid}}(\mathcal{A}) = 1] - \frac{1}{2} \right| \leq \epsilon,$$

where $\mathcal{G}_{\text{DVS}}^{\text{srchid}}(\mathcal{A})$ is defined in Figure 5.

The property of source hiding also appears under different terms in the literature such as the designated verifier property [51, 78], non-transferability [80], source deniable [42], untransferability [16], and recently off-the-record [26]. While all these definitions share the intuition that the sender can blame another party (in particular, the designated receiver) as the originator of a signature, they vary in

the adversary capabilities, i.e., whether the adversary is unbounded or whether it gets access to the secret keys.

3.1 Post-quantum DVS Schemes: Prior Work and Failed Attempts

For this work, we are interested in DVS constructions that promise post-quantum security. Despite the long line of research on DVS schemes, there are only a few candidate post-quantum constructions available in the literature; furthermore, most of those have not received much scrutiny in the mainstream cryptographic literature. This led us to attempt building, in a prior version of this paper (see Appendix C), a generic construction of post-quantum-secure DVS schemes from chameleon hash functions through both full-domain-hash and Fiat–Shamir-style signature schemes, drawing from post-quantum building blocks much closer to schemes involved in NIST standardization. In the following, we summarize prior direct constructions and our own attempts, which ultimately failed, before turning to generic constructions from ring signatures in Section 3.2.

3.1.1 Post-quantum DVS constructions in the literature. An isogeny-based DVS scheme was proposed by Sun, Tian, and Wang [81] which turned out to be insecure due to key reuse attacks identified by Galbraith et al. [46].

Wang, Hu, and Wang [85] construct a DVS scheme directly from lattice assumptions (LWE and SIS) by combining the Bonsai tree lattice trapdoor of [19] with the GPV lattice-based signature scheme [47]; a subsequent paper of theirs [86] extends this to the identity-based setting.

Noh and Jeong [71] improve on [85, 86] by giving direct constructions from lattices that can be proven without relying on random oracles; they do so by replacing the random oracle with a chameleon hash function.

Li, Liu, and Yang [61] construct a DVS scheme directly from ideal lattice assumptions (ring-SIS) by combining a ring version of the GPV signature scheme [67] with a ring chameleon hash function [35] and adding a Fiat–Shamir-with-aborts technique [65, 66].

Zhang, Liu, Tang, and Tian [89] also give a DVS constructed directly from SIS by adapting the Lyubashevsky signature scheme [66].

3.1.2 Construction attempts: GPV and Fiat–Shamir. In a previous version of this paper (see Appendix C) we gave two attempted generic DVS constructions, to be instantiated from post-quantum building blocks close to schemes involved in NIST standardization.

- Our first DVS construction was based on the full-domain-hash signature scheme [5], although following the variant by Gentry, Peikert, and Vaikuntanathan [47] which uses a trapdoor function rather than a trapdoor permutation as in [5].
- Our second DVS construction was based on the method of Fiat and Shamir [41] for constructing a signature scheme from an honest-verifier zero-knowledge canonical identification protocol.

In both of these signature schemes, signatures were constructed in the normal “forward” direction by the signer using the hashing and signing algorithms in the normal way. One can attempt to construct signatures in the “backward” direction without the secret key by

applying the permutation (for the full-domain hash scheme) or generating an identification protocol transcript (in the Fiat–Shamir case), but a forger will get stuck without a way to make the hash of the message match the hash digest picked during the backwards signature generation. The key idea in both of our constructions was to replace the standard hash function with a *chameleon hash function* (CHF) [19, 57], which allows preimages of the hash function to be found with knowledge of a trapdoor, which will be held by the verifier.

The security proof for these constructions falsely modeled the CHF as random oracle, which is not faithful. For transparency and educational purposes we describe our insights in Appendix B.

3.2 Building Post-quantum DVS Schemes from Ring Signatures

We now turn to building DVS schemes generically from ring signatures, show which properties are required to obtain a post-quantum-secure instantiation and evaluate several ring signature candidates. Our constructions draws from the idea sketched in [7, 77], with syntax and security closely following the exposition of Bender, Katz, and Morselli [7].

Definition 3.4. A ring signature scheme is a tuple of algorithms $\text{Ring} = (\text{KGen}, \text{Sign}, \text{Vrfy})$ along with a message space \mathcal{M} .

- $\text{KGen}() \mapsto (pk, sk)$: A probabilistic key generation algorithm that outputs a public-/secret-key pair.
- $\text{Sign}(sk_s, m, R) \mapsto \sigma$: A probabilistic signing algorithm that uses a secret key sk_s to produce a signature σ for a message $m \in \mathcal{M}$ w.r.t. to a list of distinct public keys R , where (pk_s, sk_s) is an honestly generated key pair and $pk_s \in R$.
- $\text{Vrfy}(R, m, \sigma) \rightarrow \text{true/false}$: A deterministic verification algorithm that checks a message m and signature σ against a ring R .

A *2-user ring signature* is a ring signature fixed to rings of size 2. A ring signature scheme Ring is *correct*, if, for honestly generated key pairs $\{(pk_i, sk_i)\}_{i=1}^n$, any $s \in [n]$, and any message $m \in \mathcal{M}$, it holds that

$$\Pr [\text{Vrfy}(\{pk_i\}_{i=1}^n, m, \text{Sign}(sk_s, m, \{(pk_i)\}_{i=1}^n)) = \text{true}] = 1.$$

The unforgeability and anonymity property we require for ring signatures are subtly different from prior literature. Like in the unforgeability notion w.r.t. insider corruption defined in [7], we consider an unforgeability adversary with access to a corruption oracle CORR . However, our unforgeability adversary is limited to rings consisting of honestly generated public keys for both its final forgery as well as the queries to the signing oracle (like the unforgeability against chosen-subring attacks defined in [7]). It is easy to see that unforgeability w.r.t. insider corruption implies our unforgeability notion; a possibly similar notion is discussed informally in [49].

Definition 3.5. A ring signature scheme Ring is $(t, \epsilon, n, Q_S, Q_{Co})$ -unforgeable w.r.t. honest-ring insider corruption if, for any adversary \mathcal{A} with running time at most t , having access to n ring keys, and making at most Q_S queries to the SIGN oracle and Q_{Co} queries to the CORR oracle, we have that

$$\text{Adv}_{\text{Ring}}^{\text{uf}}(\mathcal{A}) = \Pr [\mathcal{G}_{\text{Ring}}^{\text{uf}}(\mathcal{A}) = 1] \leq \epsilon,$$

$\mathcal{G}_{\text{Ring}}^{\text{uf}}(\mathcal{A}):$ 1 $Q_S \leftarrow \emptyset$ 2 $Q_{Co} \leftarrow \emptyset$ 3 $\mathcal{L} \leftarrow \emptyset$ 4 for $i \in [n]$ 5 $(pk_i, sk_i) \leftarrow \text{Ring.KGen}()$ 6 $\mathcal{L} \leftarrow \mathcal{L} \cup \{pk_i\}$ 7 $(R^*, m^*, \sigma^*) \leftarrow \mathcal{A}^{\text{SIGN,CORR}}(\mathcal{L})$ 8 $d_1 \leftarrow \text{Ring.Vrfy}(R^*, m^*, \sigma^*)$ 9 $d_2 \leftarrow \llbracket (m^*, R^*) \notin Q_S \rrbracket$ 10 $d_3 \leftarrow \llbracket R^* \subseteq \mathcal{L} \setminus Q_{Co} \rrbracket$ 11 return $\llbracket d_1 \wedge d_2 \wedge d_3 \rrbracket$	$\text{SIGN}(s, m, R):$ 12 if $pk_s \notin R \vee s \notin [n]$ /sign wrt honest key 13 return \perp 14 if $R \not\subseteq \mathcal{L}$ /sign wrt honest ring 15 return \perp 16 $Q_S \leftarrow Q_S \cup \{(m, R)\}$ 17 $\sigma \leftarrow \text{Ring.Sign}(sk_s, m, R)$ 18 return σ <hr/> $\text{CORR}(i):$ 19 $Q_{Co} \leftarrow Q_{Co} \cup \{pk_i\}$ 20 return sk_i
$\mathcal{G}_{\text{Ring}}^{\text{anon}}(\mathcal{A}):$ 1 $\mathcal{L} \leftarrow \emptyset$ 2 for $i \in [n]$ 3 $(pk_i, sk_i) \leftarrow \text{Ring.KGen}()$ 4 $\mathcal{L} \leftarrow \mathcal{L} \cup \{pk_i, sk_i\}$ 5 $b \leftarrow \{0, 1\}$ 6 $b' \leftarrow \mathcal{A}^{\text{CHALL}}(\mathcal{L})$ 7 return $\llbracket b' = b \rrbracket$	$\text{CHALL}(m, i_0, i_1, R):$ 8 if $\{pk_{i_0}, pk_{i_1}\} \not\subseteq R$ /challenge signers in ring 9 return \perp 10 if $\{i_0, i_1\} \not\subseteq [n]$ /sign with honest keys only 11 return \perp 12 $\sigma \leftarrow \text{Ring.Sign}(sk_{i_b}, m, R)$ 13 return σ

Figure 6: Unforgeability w.r.t. insider corruption (top) and anonymity against key exposure (bottom) of a ring signature scheme Ring. The latter game is specialized for the ring size 2.

where $\mathcal{G}_{\text{Ring}}^{\text{uf}}(\mathcal{A})$ is as in Figure 6.

We consider an anonymity notion based on anonymity against full key exposure [7]. The first difference is that we directly give all secret keys to the adversary instead of providing a signing and a corruption oracle to the adversary, where the latter in [7] returns the key generation randomness. The other difference is that we parameterize the game in the number of queries Q_{Ch} allowed to the challenge oracle. As a result, anonymity against full key exposure hence implies our anonymity notion with $Q_{Ch} = 1$. Similarly, the anonymity notions of [63] and [40], where the attacker has access to a key generation oracle, imply our anonymity notion with $Q_{Ch} = 1$.

Definition 3.6. A ring signature scheme Ring is (t, ϵ, n, Q_{Ch}) -anonymous against key exposure if, for any adversary \mathcal{A} with running time at most t , having access to n ring keys, and making at most Q_{Ch} queries to the CHALL oracle, we have that

$$\text{Adv}_{\text{Ring}}^{\text{anon}}(\mathcal{A}) = \Pr \left[\mathcal{G}_{\text{Ring}}^{\text{anon}}(\mathcal{A}) = 1 \right] \leq \epsilon,$$

where $\mathcal{G}_{\text{Ring}}^{\text{anon}}(\mathcal{A})$ is as in Figure 6.

It is easy to see that one can transform any $(t, \epsilon, n, 1)$ -anonymous (as per Definition 3.6) ring signature scheme into a $(t, \epsilon \cdot Q_{Ch}, n, Q_{Ch})$ -anonymous scheme via a hybrid argument.

3.2.1 The construction. Our construction, denoted RingDVS, is a straightforward adaption of a 2-user ring signature Ring to the DVS setting as detailed in Figure 7. The security of the resulting DVS scheme hinges on the unforgeability and anonymity of the ring signature as per Definitions 3.5 and 3.6.

THEOREM 3.7 (UNFORGEABILITY OF RingDVS). *If Ring is a $(t, \epsilon, n + 2, Q_S, Q_{Co})$ -unforgeable w.r.t. honest-ring insider corruption 2-user ring signature scheme, then RingDVS defined in Figure 7 is*

$\text{RingDVS.SKGen}():$ 1 $(pk_S, sk_S) \leftarrow \text{Ring.KGen}()$ 2 return (pk_S, sk_S) <hr/> $\text{RingDVS.VKGen}():$ 3 $(pk_D, sk_D) \leftarrow \text{Ring.KGen}()$ 4 return (pk_D, sk_D)	$\text{RingDVS.Sign}(sk_S, pk_D, m):$ 5 $\sigma \leftarrow \text{Ring.Sign}(sk_S, m, \{pk_S, pk_D\})$ 6 return σ <hr/> $\text{RingDVS.Sim}(pk_S, sk_D, m):$ 7 $\sigma \leftarrow \text{Ring.Sign}(sk_D, m, \{pk_S, pk_D\})$ 8 return σ <hr/> $\text{RingDVS.Vrfy}(pk_S, pk_D, m, \sigma):$ 9 return $\text{Ring.Vrfy}(\{pk_S, pk_D\}, m, \sigma)$
--	---

Figure 7: Designated-verifier signature scheme RingDVS = RingDVS[Ring] constructed from a 2-user ring signature scheme Ring.

(t', ϵ, n, Q_S) -existentially unforgeable under chosen message attacks, with $t' \approx t$.

PROOF. We reduce the existential unforgeability of RingDVS to the unforgeability w.r.t. honest-ring insider corruption of Ring.

Initialization of \mathcal{A} . The adversary \mathcal{B} against unforgeability of the ring signature receives as input a list \mathcal{L} of honestly generated public keys $\{pk_i\}_{i=1}^{n+2}$. Next, \mathcal{B} corrupts all keys except the first two via its CORR oracle. It sets the first two public keys as challenge keys for \mathcal{A} as $pk_S \leftarrow pk_1$ and $pk_D \leftarrow pk_2$. (Observe that we choose these two indices wlog. for easier bookkeeping.) The reduction then initializes the adversary \mathcal{A} against unforgeability of the DVS on input $(pk_S, pk_D, \{(pk_i, sk_i)\}_{i=3}^{n+2})$.

Queries to SIGN. Queries of \mathcal{A} to the SIGN oracle are of the form (pk, m) . If pk is not one of the honestly generated keys that the reduction gave to \mathcal{A} , return \perp . For each query, \mathcal{B} queries its own signing oracle on $(1, m, \{pk_1, pk\})$ and returns the answer directly to \mathcal{A} . If $pk = pk_2$, record m in Q .

Existential Forgery. At some point, \mathcal{A} outputs a DVS forgery (m^*, σ^*) wrt. pk_S and pk_D . The reduction outputs $(m^*, \sigma^*, pk_1, pk_2)$ as its own forgery.

The reduction soundly simulates the unforgeability game against RingDVS. It simulates the signing oracle truthfully by using its own signing oracle.

If \mathcal{A} outputs a valid DVS forgery wrt. sender key $pk_S = pk_1$ and verifier key $pk_D = pk_2$, the output of \mathcal{B} is a valid ring forgery wrt. the ring $\{pk_1, pk_2\}$ by construction of RingDVS. Furthermore, since $m \notin Q$, \mathcal{A} has not queried its SIGN oracle on m and pk_D . Thus, the message-ring pair $(m, \{pk_1, pk_2\})$ was not queried by \mathcal{B} to its oracle either. Lastly, the forgery is wrt. the keys $\{pk_1, pk_2\}$ that \mathcal{B} did not corrupt. Hence, all winning conditions for the ring unforgeability game are met.

The running time t of \mathcal{B} is dominated by the running time t' of \mathcal{A} and we write $t \approx t'$; simulating the signing oracle and querying the corruption oracle n times are not expensive. If \mathcal{A} outputs a successful DVS forgery with probability ϵ , then \mathcal{B} is able to produce a valid ring forgery with the same probability. \square

THEOREM 3.8 (SOURCE HIDING OF RingDVS). *If Ring is a (t, ϵ, n, Q_{Ch}) -anonymous against key exposure 2-user ring signature for $n \geq 2$, then RingDVS as shown in Figure 7 is (t', ϵ, Q_{Ch}) -source hiding, with $t' \approx t$.*

PROOF. We reduce the source hiding of RingDVS to the anonymity against key exposure of Ring.

Initialization of \mathcal{A} . The adversary \mathcal{B} against anonymity of the ring signature receives as input a list of honestly generated key pairs $\{(pk_i, sk_i)\}_{i=1}^n$. It sets the first two public keys as challenge keys for \mathcal{A} as $pk_S \leftarrow pk_1$ and $pk_D \leftarrow pk_2$. The reduction then initializes the source hiding adversary \mathcal{A} on input (sk_S, pk_S, sk_D, pk_D) .

Queries to CHALL. \mathcal{A} 's queries to the CHALL oracle are of the form m . For each of the Q_{Ch} queries, \mathcal{B} forwards the query to its own CHALL oracle as $(m, 1, 2, \{pk_1, pk_2\})$ and returns the answer it gets directly to \mathcal{A} .

Output. When \mathcal{A} outputs its guess b' , the reduction outputs b' .

The reduction soundly simulates the source hiding game against RingDVS for \mathcal{A} . The runtime of \mathcal{B} is essentially the runtime of \mathcal{A} plus the runtime to forward the challenge queries and responses and we write $t \approx t'$.

Adversary \mathcal{A} distinguishing between outputs of RingDVS.Sign and RingDVS.Sim amounts to distinguishing between Ring signatures under the two signing keys sk_1 and sk_2 in the ring $\{pk_1, pk_2\}$. Hence, \mathcal{B} inherits \mathcal{A} 's winning probability ϵ . \square

3.2.2 Implications and the inverse direction. Our construction above establishes that DVS schemes with the security properties needed for this work (i.e., existential unforgeability under chosen-message attacks and source hiding) can be generically constructed from 2-user ring signatures that provide unforgeability w.r.t. honest insider corruption and anonymity against key exposure. We note that the latter security properties are weaker than those put forward by Bender, Katz, and Morselli [7].

It is not immediately obvious to us whether, conversely, 2-user ring signatures can generically be constructed from DVS schemes. In a DVS scheme, a signature of a signer Alice designated to a verifier Bob may be distinguishable from a signature of Bob designated to Alice. This appears to pose a challenge in creating a 2-user ring signature from DVS without breaking the anonymity property. As a result, it might be the case that DVS schemes are, conceptually, easier to construct than ring signature schemes, enabling potential avenues for tailored, post-quantum DVS constructions.

3.2.3 Post-quantum ring signature candidates. Several post-quantum ring signature schemes were suggested in the literature. In the following, we list a selection of schemes having concrete instantiations and report on the signature sizes and other practical parameters provided in the corresponding works to illustrate their practicality. All schemes except Raptor (listed first) come with security proofs for unforgeability and anonymity definitions that imply our notions.

Lu, Au, and Zhang [64] introduce Raptor, which uses a chameleon hash function based on the NIST finalist FALCON [75], producing signatures of size approximately 5 KB for a 2-user ring. However, they argue that the best-known attack is inefficient instead of proving unforgeability and anonymity.

The following two schemes use zero-knowledge proofs based on symmetric primitives, akin to the NIST alternate candidate Picnic [88]: Derler, Ramacher, and Slamanig [29] provide a scheme using NIZK proofs and accumulators. For their smallest reported

ring size 2^5 , signatures can have a size of 719 KB. Katz, Kolesnikov, and Wang [53] use NIZKPoK with the MPC-in-the-head paradigm. For their smallest ring size 2^7 , signing takes 2 seconds and produces signatures of size 285 KB.

In terms of lattice-based constructions, a series of works [38–40] by Esgin et al. provide constructions relying on the hardness of M-LWE and M-SIS. The most recent candidate has a signature size of 18 KB for a 2-user ring. A construction by Lyubashevsky, Nguyen, and Seiler [68] relies on (variants of) M-LWE and M-SIS and their smallest signature for rings of size 2^5 is 16 KB. Beullens, Katsumata, and Pintore [8] introduce Falafel that also relies on M-LWE and M-SIS and produces signatures of size 29 KB in less than 100 milliseconds.

4 SECURITY MODEL FOR ASYNCHRONOUS DENIABLE KEY EXCHANGE

From a formal perspective, an asynchronous authenticated key exchange protocol is just a traditional authenticated key exchange protocol with a specific type of message flow. In particular, asynchronicity allows one party to post pre-key bundles containing long-term and possibly ephemeral public keys, provided that they can be constructed without knowing the intended partner. We will formalize security for this setting based on a Bellare–Rogaway-type model [3] with implicit authentication and (weak) forward secrecy using post-specified peers [18, 56]. The model presented in this section is simplified to deal with basic Bellare–Rogaway-type security with only long-term keys as a warm-up; in Section 7 we present a more granular model that accommodates the complex characteristics found in the Signal protocol handshake, including semi-static keys and stronger security against maximal exposure.

Parties and sessions. Let \mathcal{P} be the set of n_p parties, each of whom has a long-term public-key/secret-key pair generated by an algorithm KGenLT. Each party may run multiple instances of the protocol simultaneously or sequentially, each of which is called a session. The i th session at party P is denoted π_P^i . For each session, the party maintains the following collection of session-specific information:

- $\text{oid} \in \mathcal{P}$: The identity of the session owner.
- $\text{pid} \in \mathcal{P} \cup \{\star\}$: 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, indicating cryptographically relevant material for key derivation.
- $K \in \mathcal{K}_{\text{KE}} \cup \{\perp\}$: The session key established in this session.
- Any additional protocol-specific data used during execution.

Protocol specification. A 2-party key exchange protocol consists of the following algorithms:

- $\text{KGenLT}() \mapsto (pk, sk)$: A probabilistic long-term key generation algorithm that outputs a public-key/secret-key pair.
- $\text{Run}(sk, \hat{pk}, \pi, m) \mapsto (\pi', m')$: A probabilistic session execution algorithm that takes as input a party's long-term secret key sk , a list of long-term public keys for all other honest parties \hat{pk} , a session state π , and an incoming message m , and outputs

an updated session state π' and a (possibly empty) outgoing message m' . To set up the session sending the first message, Run is called with a distinguished message $m = \text{create}$.

In a deniable key exchange protocol, we will demand the existence of an additional algorithm:

- $\text{Fake}(pk_U, sk_V) \xrightarrow{s} T$: A probabilistic transcript simulation algorithm that takes as input one party's public key and the other party's secret key and generates a transcript T of a protocol interaction between them.

Asynchronous key exchange. In principle, a key exchange protocol can have an arbitrary number of message flows, which correspond to multiple calls to Run for a single session. In normal execution of an asynchronous authenticated key exchange protocol, the following three calls to Run occur: 1) a call to Run at the responder (Bob)¹ with $m = \text{create}$, which sets up the responder session and outputs the responder's pre-key bundle, including an ephemeral public key; 2) a call to Run at the initiator with the responder's pre-key bundle (long-term public and ephemeral public keys) which generates a session key and outputs a key exchange message; and 3) a call to Run at the responder with the initiator's long-term public key and key exchange message which generates a session key and has no output message.

Partnering. Two sessions π_U^i and π_V^j are said to be *partners* if they agree on the session identifier ($\pi_U^i.\text{sid} = \pi_V^j.\text{sid}$).

Session key indistinguishability. The first security property we want of an authenticated key exchange protocol is indistinguishability of session keys. At the start of the security experiment, long-term public-key/secret-key pairs are generated for all n_p honest parties and the public keys pk are provided to the adversary, as well as a random challenge bit b_{test} fixed for the duration of the experiment. 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 $\text{Run}(sk_U, \vec{pk}, \pi_U^i, m)$, 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 sk_U to the adversary.
- $\text{REVEALSESSKEY}(U, i)$: If session π_U^i has accepted, return its session key $\pi_U^i.K$ to the adversary.
- $\text{TEST}(U, i)$: If the TEST query has been called before or session π_U^i has not accepted, then return \perp . Otherwise, if $b_{\text{test}} = 0$, return $\pi_U^i.K$, otherwise return an element of \mathcal{K}_{KE} chosen uniformly at random. Record $\pi^* \leftarrow \pi_U^i$.

The test session $\pi^* = \pi_{U^*}^{i^*}$ is called *fresh* if the following all hold:

- (1) $\text{REVEALSESSKEY}(U^*, i^*)$ was never called.
- (2) $\text{REVEALSESSKEY}(V, j)$ was never called for any V, j such that $\pi^*.\text{sid} = \pi_V^j.\text{sid}$.
- (3) Either

- (a) there exists an honest partner session π_p^* ($\pi_p^*.\text{sid} = \pi^*.\text{sid}$ if π^* is a responder, and $\pi_p^*.\text{cid} = \pi^*.\text{cid}$ if π^* is an initiator), covering weak forward secrecy, or
- (b) $\text{CORRUPTLTKEY}(\pi^*.\text{oid})$ and $\text{CORRUPTLTKEY}(\pi^*.\text{pid})$ was never called, covering implicit authentication.

At the end of the experiment, 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 is the absolute value of the difference between $\frac{1}{2}$ and the probability that the experiment outputs 1.

Deniability. The second security property we want is deniability. At the start of this experiment, long-term public-key/secret-key pairs are generated for all n_p honest parties and the public and secret keys are provided to the adversary. A random challenge bit b is fixed for the duration of the experiment. The adversary is given repeated access to a CHALL oracle which takes as input two party identifiers U and V . If b is 0, then CHALL will generate an honest transcript of an interaction between U and V using the Run algorithm and each party's secret keys. If b is 1, then CHALL will generate a simulated transcript of an interaction between U and V using the Fake algorithm. At the end of the experiment, the adversary outputs a guess b' of b . The experiment outputs 1 if $b' = b$ and 0 otherwise. The adversary's advantage in the deniability game is the absolute value of the difference between $\frac{1}{2}$ and the probability the experiment outputs 1.

There are several prior works giving definitions of offline deniability for key exchange [24, 25, 31, 82, 83]. Our definition differs from previous ones in that it gives access to secret keys to the Fake algorithm (corresponding to the simulator in simulation-based definitions) and to the adversary (i.e., the judge). This models the informal deniability requirement from the Signal specification [70, §4.4]. See Appendix A for a more detailed discussion.

5 SECURITY OF THE CORE PROTOCOL

We now show that our core protocol Π from Figure 2 achieves the security properties defined in Section 4. Key indistinguishability of Π depends on the IND-CCA security of the two KEMs, the unforgeability of the DVS, and the security of the KDF; deniability of Π depends on the source hiding of the DVS. Both proofs are in the standard model.

To formally capture Π in the security model of Section 4, we need to specify a few more details:

- Alice takes the initiator role, Bob the responder role.
- The transcript in Figure 2 corresponds to the session identifier and consists of the parties' identities and long-term public keys, the responder's ephemeral public key, and the KEM ciphertexts; the contributive identifier corresponds to the pre-key bundle part of the transcript, received by Alice from Bob:

$$\begin{aligned} \text{transcript} = \text{sid} &= (A, B, pk_A^{\text{DVS}}, pk_B^{\text{KEM}}, pk_B^{\text{DVS}}, epk_B^{\text{KEM}}, c_1, c_2), \\ \text{cid} &= (B, pk_B^{\text{KEM}}, pk_B^{\text{DVS}}, epk_B^{\text{KEM}}). \end{aligned}$$

¹Note that we call Bob the *responder* in our model despite Bob outputting the first, asynchronous key exchange message. Based on the high-level protocol interaction, we deem it more natural to call Alice, who decides to *initiate* a Signal session with Bob, the initiator (in contrast to, e.g., [22, 82, 83]).

5.1 Key Indistinguishability

THEOREM 5.1 (KEY INDISTINGUISHABILITY OF Π). *Let DVS be a $(t, \epsilon_{\text{DVS}}, n_p, Q_S)$ -unforgeable DVS scheme, KEM_1 be a $(t, \epsilon_{\text{KEM}_1}, n_s)$ -IND-CCA-secure KEM, KEM_2 be a $(t, \epsilon_{\text{KEM}_2}, 1)$ -IND-CCA-secure KEM, and KDF be a $(t, \epsilon_{\text{KDF}}, n_s)$ -PRF-secure key derivation function when keyed through either of the key components K_1 and K_2 . Then the asynchronous DAKE protocol Π from Figure 2 provides key indistinguishability (as defined in Section 4) in that the advantage ϵ' of any adversary \mathcal{A} running in time $t' \approx t$ is upper bounded as*

$$\epsilon' \leq n_s \cdot \begin{pmatrix} n_s \cdot (\epsilon_{\text{KEM}_2} + \epsilon_{\text{KDF}}) \\ + n_p \cdot (\epsilon_{\text{KEM}_1} + \epsilon_{\text{KDF}}) \\ + n_p^2 \cdot (\epsilon_{\text{DVS}} + n_s \cdot (\epsilon_{\text{KEM}_2} + \epsilon_{\text{KDF}})) \end{pmatrix},$$

where $n_s \leq Q_{\text{Snd}}$ is the maximum number of sessions (upper bounded by the number Q_{Snd} of SEND queries) and n_p the number of parties.

PROOF. We proceed via a sequence of game hops starting from the key indistinguishability game for an adversary \mathcal{A} . We bound the difference between each hop until we reach a game where the adversary's advantage is 0.

Game 0. The initial key indistinguishability game, denoted \mathcal{G}_0 , letting $\epsilon' := \text{Adv}_{\text{ADAKE}}^{\mathcal{G}_0}(\mathcal{A}) = |\Pr[\mathcal{G}_0 = 1] - \frac{1}{2}|$.

Game 1 (Guess test session π^*). We first guess the tested session π^* and “invalidate” the game by overwriting the adversary's bit guess with 0 if the adversary calls TEST on a different session. Guessing among the n_s many sessions (where n_s is at most the number Q_{Snd} of calls to the SEND oracle),

$$\text{Adv}_{\text{ADAKE}}^{\mathcal{G}_0}(\mathcal{A}) \leq n_s \cdot \text{Adv}_{\text{ADAKE}}^{\mathcal{G}_1}(\mathcal{A}).$$

For the remaining proof, we distinguish the following three cases for the test session being fresh:

- There exists an honest partner session π_p^* ($\pi_p^*.\text{sid} = \pi^*.\text{sid}$ if π^* is a responder, and $\pi_p^*.\text{cid} = \pi^*.\text{cid}$ if π^* is an initiator).
- The tested session is an initiator (“Alice”) session and $\text{CORRUPTLTKEY}(\pi^*.\text{pid})$ was never called.
- The tested session is a responder (“Bob”) session and neither $\text{CORRUPTLTKEY}(\pi^*.\text{oid})$ nor $\text{CORRUPTLTKEY}(\pi^*.\text{pid})$ was ever called.²

Treating these cases as events in \mathcal{G}_1 , and writing $\mathcal{G}_1[X]$ to indicate that event X occurs, by the union bound we have:

$$\text{Adv}_{\text{ADAKE}}^{\mathcal{G}_1}(\mathcal{A}) \leq \text{Adv}_{\text{ADAKE}}^{\mathcal{G}_1[\text{A}]}(\mathcal{A}) + \text{Adv}_{\text{ADAKE}}^{\mathcal{G}_1[\text{B}]}(\mathcal{A}) + \text{Adv}_{\text{ADAKE}}^{\mathcal{G}_1[\text{C}]}(\mathcal{A}).$$

Case A (Honest partner). In the first proof case, there exists a session π_p^* that agrees with the tested session π^* on the responder's ephemeral KEM public key epk^{KEM} used. We will leverage this to embed a challenge into the ephemeral KEM ciphertext c_2 .

Game A.1 (Guess partnered session). We first guess a session π_p^* which is partnered via sid (if π^* is a responder) or cid (if π^* is an initiator) to the test session π^* , and let the adversary lose if the guess is incorrect. By this case's prerequisites, (at least) one

²In our full SPQR protocol (see Section 6), we will strengthen this case by having Bob use semi-static DVS keys. This limits the time window for a key-compromise impersonation (KCI) attack [9] against Bob, as in the Signal handshake [70, §4.6].

partner session exists and is guessed with probability at least $1/n_s$, hence

$$\text{Adv}_{\text{ADAKE}}^{\mathcal{G}_1[\text{A}]}(\mathcal{A}) \leq n_s \cdot \text{Adv}_{\text{ADAKE}}^{\mathcal{G}_{\text{A.1}}}(\mathcal{A}).$$

Game A.2 (Ephemeral KEM). We now replace the KEM key K_2 with a random key \widetilde{K}_2 in π^* and also in π_p^* (unless the latter is a responder and receives a different ciphertext c_2 than sent by π^*).

We bound the difference introduced by this step through a reduction to the IND-CCA security of the KEM_2 scheme, which simulates $\mathcal{G}_{\text{A.1}}$ truthfully except for the following changes and runs in time $t \approx t'$. It embeds the obtained challenge public key pk into the ephemeral KEM public key epk of the responder session among π^* and π_p^* , the challenge ciphertext c^* as c_2 of the initiator session (among π^* and π_p^*), and the challenge (real-or-random) key K_b^* as K_2 into both π^* and π_p^* . If π^* is an initiator session, it uses its DECAPS oracle (at most once, i.e., $Q_D \leq 1$) to decrypt a potentially different ciphertext $c'_2 \neq c_2 = c^*$ received by π_p^* . Depending on the IND-CCA KEM challenge bit, the reduction perfectly simulates $\mathcal{G}_{\text{A.1}}$ or $\mathcal{G}_{\text{A.2}}$, hence

$$\text{Adv}_{\text{ADAKE}}^{\mathcal{G}_{\text{A.1}}}(\mathcal{A}) \leq \epsilon_{\text{KEM}_2} + \text{Adv}_{\text{ADAKE}}^{\mathcal{G}_{\text{A.2}}}(\mathcal{A}).$$

Game A.3 (KDF). We finally replace the key derivation function KDF in both π^* and π_p^* (in the latter only if it uses \widetilde{K}_2) with a random function, in particular replacing the session key K of π^* with a randomly sampled key \widetilde{K} .

We bound the introduced advantage difference via a reduction to the pseudorandomness of the key derivation function KDF, treated as a PRF keyed through the second key component K_2 and taking $(K_1, \text{transcript})$ as label. The reduction runs in time $t \approx t'$ and simulates Game $\mathcal{G}_{\text{A.2}}$ truthfully, except that it does not sample \widetilde{K}_2 itself but instead uses its oracle PRFCHALLENGE to compute the session key values derived from \widetilde{K}_2 . It calls its oracle at most twice, once for π^* and possibly once for π_p^* on a different label, hence $Q_{\text{PRF}} \leq n_s$. Depending on whether its oracle output is the true KDF evaluation or that of a random function, the reduction perfectly simulates $\mathcal{G}_{\text{A.2}}$ or $\mathcal{G}_{\text{A.3}}$, thus

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

In Game $\mathcal{G}_{\text{A.3}}$, the challenge key K_{test} for π^* is a uniformly random key, independent of b_{test} . Furthermore, by the first two freshness conditions, \mathcal{A} cannot reveal K_{test} via a REVEALSESSKEY query on π^* or any partnered session who might hold the same key. Thus, in $\mathcal{G}_{\text{A.3}}$, \mathcal{A} cannot do better than guessing, leaving it with advantage $\text{Adv}_{\text{ADAKE}}^{\mathcal{G}_{\text{A.3}}}(\mathcal{A}) = 0$.

Case B (Initiator tested, peer uncorrupted). In the second proof case, we have that the tested initiator session π^* has an uncorrupted intended peer. We will leverage this to embed a challenge into the static KEM ciphertext c_1 .

Game B.1 (Guess responder identity). We first guess the test session's intended peer, $V = \pi^*.\text{pid}$, among the n_p many parties in the game and let the adversary lose if we guess incorrectly. This reduces the adversary's advantage by a factor at most n_p :

$$\text{Adv}_{\text{ADAKE}}^{\mathcal{G}_1[\text{B}]}(\mathcal{A}) \leq n_p \cdot \text{Adv}_{\text{ADAKE}}^{\mathcal{G}_{\text{B.1}}}(\mathcal{A}).$$

Game B.2 (Static KEM). We can now replace the KEM key K_1 in π^* (and any responder session of V receiving the same ciphertext c_1) with a random key \widetilde{K}_1 .

We bound the advantage difference introduced by this step through a reduction to the IND-CCA security of the KEM₁ scheme. The reduction runs in time $t \approx t'$ and simulates $\mathcal{G}_{B.1}$ truthfully, but embeds the obtained challenge public key pk as V 's public KEM key pk_V^{KEM} at the outset of the game. It further embeds the challenge ciphertext c^* as c_1 sent by π^* and the challenge (real-or-random) key K_b^* as K_1 into π^* (and any responder session of V receiving c^*). The reduction uses the DECAPS oracle to decapsulate any ciphertexts $c_1 \neq c^*$ received by sessions of V (calling the oracle at most n_s times), and never has to respond to $\text{CORRUPTLTKEY}(V)$ queries as otherwise π^* would not be fresh. Depending on the IND-CCA KEM challenge bit, the reduction perfectly simulates $\mathcal{G}_{B.1}$ or $\mathcal{G}_{B.2}$, hence

$$\text{Adv}_{\text{ADAKE}}^{\mathcal{G}_{B.1}}(\mathcal{A}) \leq \epsilon_{\text{KEM}_1} + \text{Adv}_{\text{ADAKE}}^{\mathcal{G}_{B.2}}(\mathcal{A}).$$

Game B.3 (KDF). We finally replace the key derivation function KDF in π^* (and any other session using \widetilde{K}_1) with a random function, in particular replacing the session key K of π^* with a randomly sampled key \widetilde{K} .

Analogous to Game $\mathcal{G}_{A.3}$, we can bound the introduced advantage difference by the pseudorandomness of KDF when keyed through the first key component K_1 and taking $(K_2, \text{transcript})$ as label. The challenge static KEM key \widetilde{K}_1 may possibly be decapsulated in many responder sessions of V , who use distinct transcript labels unless they are partnered with π^* ; the PRF reduction, running in time $t \approx t'$, may hence make up to n_s queries to its PRFCHALLENGE oracle. Simulating either of the two games in the reduction, we get

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

At this point, the challenge key K_{test} for π^* is uniformly random and independent, as only partnered sessions will use the same transcript label to derive their session keys, but for π^* to be fresh those cannot be revealed. Thus $\text{Adv}_{\text{ADAKE}}^{\mathcal{G}_{B.3}}(\mathcal{A}) = 0$.

Case C (Responder tested, both parties uncorrupted). In the final proof case, we know that the tested responder session π^* has an uncorrupted intended peer. We will leverage this to ensure that there is a partnered initiator session (which signed the transcript) and then embed a challenge into the ephemeral KEM ciphertext c_2 between these two sessions.

Game C.1 (Guess initiator and responder identities). We first guess the (responder) test session's owner $V = \pi^*.oid$ and intended (initiator) peer $U = \pi^*.pid$ among the n_p many parties in the game and "invalidate" the game (overwriting \mathcal{A} 's bit guess by 0) if we guess incorrectly. Guessing both parties induces a quadratic loss in n_p :

$$\text{Adv}_{\text{ADAKE}}^{\mathcal{G}_1[C]}(\mathcal{A}) \leq n_p^2 \cdot \text{Adv}_{\text{ADAKE}}^{\mathcal{G}_{C.1}}(\mathcal{A}).$$

Game C.2 (Signature unforgeability). We now "invalidate" the game (overwriting \mathcal{A} 's bit guess by 0) if the test session π^* accepts a DVS signature σ on a transcript that no session of U has issued.

We bound this event by a reduction against the existential unforgeability of DVS, running in time $t \approx t'$ and simulating $\mathcal{G}_{C.1}$ with

the following modification: Instead of generating parties' DVS keys itself, the reduction embeds the unforgeability game's challenge public keys as $pk_U = pk_S$ and $pk_V = pk_D$, and assigns the additional DVS public-secret key pairs from the unforgeability game's list \mathcal{L} to the remaining parties. (Note that the reduction obtains the secret keys for the latter keys, allowing it to fully simulate those parties.) The reduction uses its signing oracle to compute signatures under $pk_U = pk_S$ (and for any peer public key pk). As U and V remain uncorrupted in this proof case, the reduction never has to answer a $\text{CORRUPTLTKEY}(U)$ or $\text{CORRUPTLTKEY}(V)$ query. In the case that π^* receives a valid DVS transcript-signature pair (transcript, σ) that no session of U sent (and hence transcript was not queried to the DVS SIGN oracle), the reduction outputs this pair as its forgery and wins. Therefore,

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

Game C.3 (Guess partnered session). As of $\mathcal{G}_{C.2}$, we know that π^* receives a DVS signature on a transcript value transcript = $\pi^*.sid$ sent by some session of U . We now guess this (sid-partnered) session π_p^* (among the n_s many sessions) and, invalidating the game (overwriting \mathcal{A} 's bit guess by 0) upon wrong guess, get

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

Game C.4 (Ephemeral KEM). We next replace the KEM key K_2 with a random key \widetilde{K}_2 in π^* and π_p^* .

As in Game $\mathcal{G}_{A.2}$, we bound the introduced advantage difference by the IND-CCA security of the KEM₂ scheme. The reduction runs in time $t \approx t'$, embeds the challenge pk and c^* into π^* 's ephemeral KEM public key, resp. π_p^* 's c_2 ciphertext, and uses the challenge key K_b^* in place of K_2 in both sessions. It does not need to use its DECAPS oracle (i.e., $Q_D = 0$), since pk is not used in another session and we are at this point guaranteed that π^* receives π_p^* 's ephemeral ciphertext. (So in fact we only need IND-CPA security of KEM₂ here.) The reduction simulates the difference between $\mathcal{G}_{C.3}$ and $\mathcal{G}_{C.4}$, so

$$\text{Adv}_{\text{ADAKE}}^{\mathcal{G}_{C.3}}(\mathcal{A}) \leq \epsilon_{\text{KEM}_2} + \text{Adv}_{\text{ADAKE}}^{\mathcal{G}_{C.4}}(\mathcal{A}).$$

Game C.5 (KDF). In the final game hop, we replace KDF in both π^* and π_p^* with a random function, replacing the session key K of π^* with a randomly sampled key \widetilde{K} .

As in Game $\mathcal{G}_{A.3}$, this is bounded by the pseudorandomness of KDF with key K_2 and label $(K_1, \text{transcript})$. Due to π^* and π_p^* agreeing on the transcript input to KDF, the corresponding reduction only makes one query, $Q_{\text{PRF}} = 1 \leq n_s$, running in time $t \approx t'$. Simulating the game difference through this reduction, we get

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

This completes the last proof case, as the challenge key K_{test} for π^* is now uniformly random and independent (beyond partnered sessions), leaving \mathcal{A} with advantage $\text{Adv}_{\text{ADAKE}}^{\mathcal{G}_{C.5}}(\mathcal{A}) = 0$. \square

5.2 Deniability

THEOREM 5.2 (DENIABILITY OF Π). *Let $DVS = (\text{SKGen}, \text{VKGen}, \text{Sign}, \text{Vrfy}, \text{Sim})$ be a $(t, \epsilon_{\text{srchid}}, Q_{Ch})$ -source hiding DVS scheme. Then the asynchronous DAKE protocol Π from Figure 2 provides deniability (as defined in Section 4) in that the advantage ϵ' of any adversary \mathcal{A} running in time $t' \approx t$ and making up to Q_{Ch} challenge queries is upper bounded as $\epsilon' \leq n_p^2 \cdot \epsilon_{\text{srchid}}$, where n_p is the number of parties.*

PROOF. The proof follows by a standard hybrid argument. Let \mathcal{A} be a successful adversary against deniability of Π , then we can construct a reduction \mathcal{B} against the source hiding property of DVS. Observe that \mathcal{B} computes for each of the n_p parties a long-term key pair. It randomly guesses the identifiers of two parties $\text{id}^*, \text{rid}^* \in [n_p]$ for which \mathcal{A} can distinguish between Run and Fake. Let a number $i \in [n_p]$ uniquely denote two independent values id, rid in a query (e.g., encoded as $(\text{id} - 1) \cdot n_p + \text{rid}$) and let $i^* \in [n_p]$ denote the specific guess $\text{id}^*, \text{rid}^*$ of \mathcal{B} . For party id^* , \mathcal{B} replaces the sampled long-term key with its challenge key pair (pk_S, sk_S) and similarly it replaces for party rid^* with (pk_D, sk_D) .

In case \mathcal{A} makes a query i for $1 \leq i < i^*$, then \mathcal{B} answers as if $b = 0$, i.e., it runs $DVS.\text{Sign}$. For all $i^* < i \leq n_p^2$, if \mathcal{A} makes a query, then \mathcal{B} answers as if $b = 1$, i.e., it runs $DVS.\text{Sim}$. If \mathcal{A} queries $i = i^*$, then \mathcal{B} passes it to its own oracle. In all cases \mathcal{B} returns the transcript and the session key K to \mathcal{A} . Finally, when \mathcal{A} returns its guess bit b' , \mathcal{B} returns b' as its guess.

Observe that \mathcal{B} faithfully simulates the deniability game for \mathcal{A} . Moreover, the runtime of \mathcal{B} is essentially the runtime of \mathcal{A} plus the runtime to generate the keys and answer the oracle queries.

Now we analyze the winning probability of \mathcal{A} against deniability. For this, we define the hybrids $H_0, \dots, H_{n_p^2}$ with H_i being the hybrid that answers all challenge queries for indices $1, \dots, i$ with Run and the challenge queries for indices $i+1, \dots, n_p^2$ with Fake. The extreme hybrids are $H_{n_p^2}$, which answers all the challenge queries with Run, and H_0 , which answers all queries with Fake. Observe that H_{i-1} and H_i only differ in an execution of Run or Fake. Hence, the probability of distinguishing between H_{i-1} and H_i is bounded by ϵ_{srchid} . Since there are n_p^2 many hybrids, we overall obtain that \mathcal{A} 's probability of winning the deniability game is bounded by $\epsilon' \leq n_p^2 \cdot \epsilon_{\text{srchid}}$. \square

6 SIGNAL IN A POST-QUANTUM REGIME

We now extend our core protocol Π from Figure 2 to capture all the characteristics of the Signal handshake. The core protocol already captures implicit mutual authentication, forward secrecy, offline deniability, and asynchronicity. Signal's X3DH has a few more subtle aspects and security features to consider, which we address in our extended asynchronous DAKE protocol: SPQR (Signal in a Post-Quantum Regime), depicted in Figure 8.

Semi-static keys. In Signal, asynchronicity is facilitated by a central, untrusted server which stores the users' pre-key bundles. To enable multiple users to asynchronously contact some responder user, say Bob, the latter uploads multiple ephemeral public pre-keys to the Signal server, of which one is handed to any initiator session that wants to contact Bob (along with the other pre-key bundle elements) and then deleted from the Signal server.

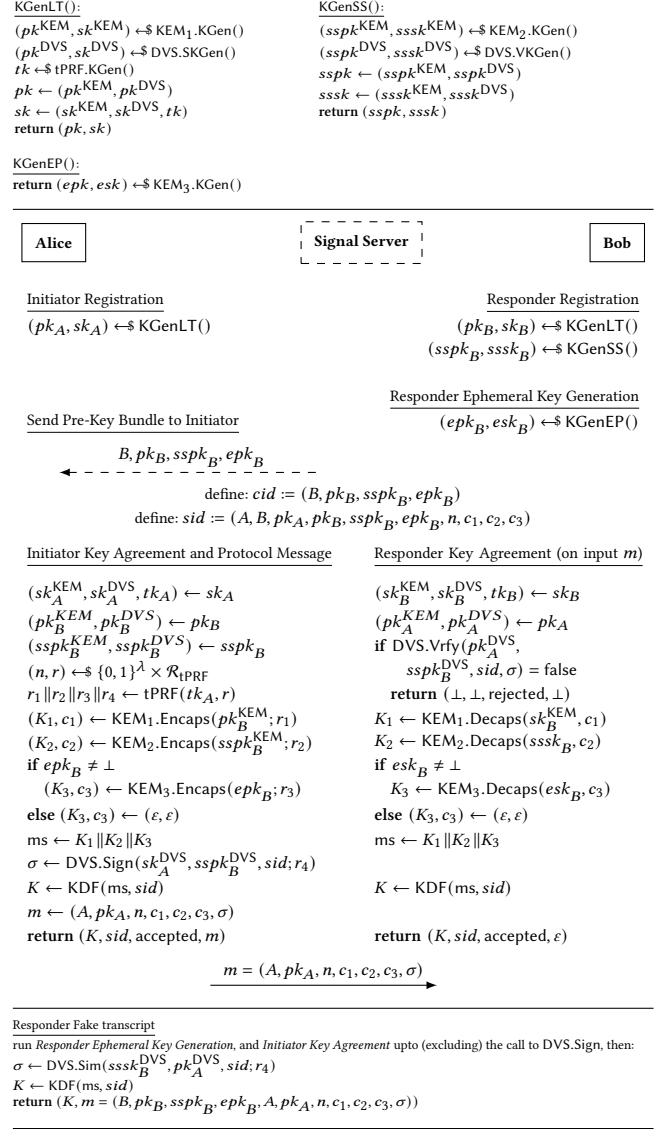


Figure 8: The SPQR protocol (top: key generation, middle: protocol flow, bottom: fake transcript generation), combining static, semi-static and ephemeral key encapsulation schemes KEM_1 , KEM_2 , and KEM_3 , a designated verifier signature DVS, and a twisted pseudorandom function tPRF.

Bob will periodically upload new ephemeral pre-keys; however, if Bob has been offline for a long time, those pre-keys may run out. Therefore, the Signal protocol also includes a *semi-static* key in user pre-key bundles, and always includes key derivations based on that semi-static key. If the Signal server runs out of ephemeral pre-keys, the corresponding key share is not derived and left out; in that case the semi-static key share still provides delayed forward secrecy [14]. We capture this similarly in SPQR by encapsulating a key-ciphertext pair (K_3, c_3) against Bob's ephemeral KEM public key epk_B only if the latter is present.

Maximal-exposure security. Signal aims for very strong security guarantees, considering beyond long-term and session key compromise and also compromise of semi-static and ephemeral keys (via the randomness of sessions) [17, 22, 60]. We model this in an accordingly strong key exchange model (in Section 7) and prove (in Section 8) that SPQR achieves equivalent security in the post-quantum setting as Signal does in the classical setting. In particular, we show that session keys remain secret, as long as any of the (Alice–Bob) secret combinations ephemeral–ephemeral, ephemeral–semi-static, ephemeral–long-term, and long-term–semi-static are uncompromised. Secrecy from the first three is straightforwardly achieved via encapsulations against the corresponding ephemeral, semi-static, and long-term KEM keys of Bob. To achieve secrecy from the last one (i.e., when all initiator randomness is compromised), beyond relying on the DVS scheme for initiator authentication, we apply a NAXOS-like [60] trick to extract randomness from Alice’s long-term secrets via a twisted PRF [44, 58]. Twisted PRFs can be generically instantiated from regular PRFs (see Section 2.3) and yield output indistinguishable from random as long as a session’s long-term secret *or* randomness is uncompromised.

We present our formal security results for SPQR in Section 8 after introducing the full security model next.

7 FULL SECURITY MODEL

In this section we present the extensions to the core ADAKE model that we will use to prove our post-quantum Signal construction SPQR depicted in Figure 8 secure. The key-indistinguishability game is fully specified in Figure 9 and the deniability game in Figure 10. The main differences to the core security models are as follows

- Signal employs *semi-static* keys; these keys are authenticated via signatures using the long-term key of the respective party, reused, and updated regularly. We thus establish multiple of these keys for each party, identifying each key pair uniquely via an identifier $\text{ssid} \in [n_{\text{ss}}]$. Sessions receive semi-static keys in an authenticated manner in the model (just like long-term keys). The adversary is able to corrupt semi-static keys of a user U via the $\text{CORRUPTSSKEY}(U, \text{ssid})$ oracle, similar to the corruption of long-term keys via $\text{CORRUPTLTKEY}(U)$.
- The usage of ephemeral pre-keys is optional in Signal (as the pre-keys stored on the Signal server may run out). We model this by introducing two types of sessions, full and reduced, depending on whether an ephemeral pre-key is received by the initiator in the pre-key bundle or not.
- The adversary is now granted maximal-exposure capabilities by also revealing the randomness used in a party’s execution Run. To this end, we make the used randomness explicit in syntax via the session state variable coins , which during setup of the session samples random coins from the appropriate randomness spaces. The adversary then has access to a $\text{REVEALRANDOM}(U, i)$ oracle that returns the coins sampled in session π_U^i , and marks them as revealed via a flag revrand .

7.1 Key Indistinguishability

Definition 7.1. An asynchronous DAKE key exchange protocol ADAKE is $(t, \epsilon, (Q_{\text{Snd}}, Q_{\text{CorrLT}}, Q_{\text{CorrSS}}, Q_{\text{RevR}}, Q_{\text{RevSK}}))$ -key-indistinguishable if for any adversary \mathcal{A} with running time at most t , we have that

$$\text{Adv}_{\text{ADAKE}}^{\text{ke-kind}}(\mathcal{A}) = \left| \Pr \left[\mathcal{G}_{\text{ADAKE}}^{\text{ke-kind}}(\mathcal{A}) = 1 \right] - \frac{1}{2} \right| \leq \epsilon,$$

where $\mathcal{G}_{\text{ADAKE}}^{\text{ke-kind}}(\mathcal{A})$ is defined in Figure 9 and Q_x for $x \in \{\text{Snd}, \text{CorrLT}, \text{CorrSS}, \text{RevR}, \text{RevSK}\}$ denotes the number of queries to the oracles SEND, CORRUPTLTKEY, CORRUPTSSKEY, REVEALRANDOM and REVEALSESSKEY, respectively. The model restricts the adversary to a single query to the TEST oracle.

In addition to the state variables given for the core protocol in Section 4, the following protocol-specific variables are introduced:

- $\text{ssid} \in [n_{\text{ss}}]$ denotes the identifier of the responder’s semi-static key used in this session. If $\pi.\text{role} = \text{initiator}$ this refers to $\text{sspk}_{\text{pid}}^{\text{ssid}}$, if $\pi.\text{role} = \text{responder}$ this refers to $\text{sspk}_{\text{oid}}^{\text{ssid}}$.
- $\text{type} \in \{\text{full}, \text{reduced}\}$ indicates whether an ephemeral pre-key was included in the pre-key bundle, or not. Setting $\text{type} = \text{full}$ indicates that an ephemeral pre-key has been received and used by the initiator, whereas $\text{type} = \text{reduced}$ means that no ephemeral pre-key has been received resp. used by the initiator.
- $\text{coins} \in \mathcal{R}_{\text{KE}}$ denotes the random coins from the randomness space \mathcal{R}_{KE} used in the execution of Run.
- $\text{revrand} \in \{\text{true}, \text{false}\}$ indicates whether the random coins $\pi.\text{coins}$ have been revealed via a REVEALRANDOM query. The default value is false.

In order to fully describe the security game $\mathcal{G}_{\text{KE}}^{\text{ke-kind}}(\mathcal{A})$ that is played between the adversary and the challenger, we introduce the following game-specific flags associated with a user $U \in [n_p]$. They indicate whether a party U ’s long-term or one of its semi-static secret keys have been compromised by the adversary:

- $\text{corrltk}_U \in \{\text{true}, \text{false}\}$ indicates whether the long-term secret key of party U has been compromised by the adversary via a $\text{CORRUPTLTKEY}(U)$ query. The default value is false.
- $\text{corrssk}_{\text{ssid}}^U \in \{\text{true}, \text{false}\}$ indicates whether the semi-static secret key with index ssid of party U has been compromised by the adversary via a CORRUPTSSKEY query. The default value is false.

7.1.1 Session Partnering and Correctness. As in the core model, (full) session partnering is defined via session identifiers: We say that a session π_U^i owned by U is *partnered* with a session π_V^j owned by V if they agree on the session identifier, i.e., $\pi_U^i.\text{sid} = \pi_V^j.\text{sid} \neq \perp$. In order to identify sessions which may eventually derive the same key but are not fully partnered (yet), we have introduced the concept of *contributive identifiers* cid . More precisely, we say that a session π_U^i owned by U is *contributively partnered* with a session π_V^j owned by V , if they agree on their contributive session identifier, i.e., whenever $\pi_U^i.\text{cid} = \pi_V^j.\text{cid} \neq \perp$.

We say that an asynchronous authenticated key exchange protocol $\text{KE} = (\text{KGenLT}, \text{KGenSS}, \text{KGenEP}, \text{Run})$ with randomness

<p>$\mathcal{G}_{KE}^{\text{ke-kind}}(\mathcal{A})$:</p> <pre> 1 $b_{\text{test}} \leftarrow \{0, 1\}$ #sample test bit 2 $\pi^* \leftarrow \perp$ #variable for test session 3 for $U \in [n_p]$ 4 $(pk_U, sk_U) \leftarrow \text{KGenLT}()$ #long-term key generation 5 for $ssid \in [n_{ss}]$ 6 $(sspk_U^{ssid}, sssk_U^{ssid}) \leftarrow \text{KGenSS}()$ #semi-static key generation 7 $\vec{pk} \leftarrow \{pk_U\}_{U \in [n_p]}$; $\vec{sspk} \leftarrow \{sspk_U^{ssid}\}_{U \in [n_p]}$ 8 $b' \leftarrow \mathcal{A}^{\text{SEND, TEST, CORRUPTLTKEY, CORRUPTSSKEY, REVEALRANDOM, REVEALSESSKEY}}(\vec{pk}, \vec{sspk})$ #run adversary 9 if $\text{sound}() = \text{false}$ #adversary wins if it breaks soundness 10 return 1 11 if $\text{fresh}(\pi^*) = \text{false}$ #attack invalid if test session is not fresh 12 $b' \leftarrow 0$ 13 return $\llbracket b' = b_{\text{test}} \rrbracket$ #determine win or loss</pre>	<p>$\text{fresh}(\pi^*)$:</p> <pre> 14 if $\pi^*.revealed = \text{true}$ then return false #test session is revealed 15 if $\exists \pi_V^i \neq \pi^* : (\pi_V^i.sid = \pi^*.sid \wedge \pi_V^i.revealed = \text{true})$ then return false #test session's partner is revealed 16 return $(\pi^*.type = \text{full and clean}_{\text{full}}(\pi^*))$ #test session in full handshake mode and test session key is clean or $(\pi^*.type = \text{reduced and clean}_{\text{reduced}}(\pi^*))$ #test session in reduced handshake mode and test session key is clean</pre> <p>$\text{sound}()$:</p> <pre> 17 return \forall distinct π, π', π'' 18 $(\pi.sid = \pi'.sid \neq \perp \implies \pi.K = \pi'.K \wedge \pi.type = \pi'.type \wedge \pi.cid = \pi'.cid)$ #same session identifier imply same shared key, type, and contributive identifiers 19 and $(\pi.sid = \pi'.sid \neq \perp \wedge \pi.role = \text{initiator} \implies \pi'.role = \text{responder})$ #session identifiers of two initiator sessions never collide 20 and $(\pi.sid = \pi'.sid = \pi''.sid \neq \perp \implies \pi.type = \text{reduced})$ #session identifiers of three sessions only collide in reduced mode</pre>	
<p>$\text{SEND}(U, i, m)$:</p> <pre> 21 if $\pi_U^i = \perp$ #initiate session: for responders, $m = \text{create}$ carries semi-static key identifier 22 $\pi_U^i.oid \leftarrow U$ #set owner identity 23 if $m = \text{create}$ then 24 $\pi_U^i.role \leftarrow \text{responder}$; $\pi_U^i.ssid \leftarrow m.ssid$ #set responder role and semi-static key identifier (carried in m) 25 else $\pi_U^i.role \leftarrow \text{initiator}$ #set initiator role (m is first protocol message) 26 $\pi_U^i.coins \leftarrow \mathcal{R}_{KE}$ #sample session randomness 27 $\pi_U^i.st_{\text{exec}} \leftarrow \text{running}$ 28 $(\pi_U^i, m') \leftarrow \text{Run}(sk_U, \vec{pk}, \vec{sspk}, \pi_U^i, m; \pi_U^i.coins)$ #run session, with explicit random coins 29 return $(m', \pi_U^i.st_{\text{exec}})$ #return message and session state</pre>	<p>$\text{CORRUPTLTKEY}(U)$:</p> <pre> 38 $\text{corrltk}_U \leftarrow \text{true}$ #mark long-term key as corrupted 39 return sk_U #return long-term secret key</pre> <p>$\text{CORRUPTSSKEY}(U, ssid)$:</p> <pre> 40 $\text{corrssk}_U^{ssid} \leftarrow \text{true}$ #mark semi-static key as corrupted 41 return $sssk_U^{ssid}$ #return semi-static secret key</pre> <p>$\text{REVEALRANDOM}(U, i)$:</p> <pre> 42 if $\pi_U^i = \perp$ then return \perp #session does not exist 43 $\pi_U^i.revrand \leftarrow \text{true}$ #mark randomness as revealed 44 return $\pi_U^i.coins$ #return session's random coins</pre> <p>$\text{REVEALSESSKEY}(U, i)$:</p> <pre> 45 if $\pi_U^i = \perp$ or $\pi_U^i.st_{\text{exec}} \neq \text{accepted}$ then return \perp #session does not exist or has not derived key yet 46 $\pi_U^i.revealed \leftarrow \text{true}$ #mark session key as revealed 47 return $\pi_U^i.K$ #return session key</pre>	
<p>$\text{TEST}(U, i)$:</p> <pre> 30 if $\pi_U^i = \perp$ or $\pi_U^i.st_{\text{exec}} \neq \text{accepted}$ or $\pi^* \neq \perp$ #session does not exist, has not accepted yet, or test already asked 31 return \perp 32 $\pi^* \leftarrow \pi_U^i$ #record test session 33 if $b_{\text{test}} = 0$ 34 $K_{\text{test}} \leftarrow \pi_U^i.K$ #real session key 35 else 36 $K_{\text{test}} \leftarrow \mathcal{K}_{KE}$ #random key from key space 37 return K_{test} #return challenge key</pre>	<p>$\text{clean}_{\text{full}}(\pi^*)$:</p> <pre> 48 return $\text{clean}_{\text{reduced}}(\pi^*)$ or $\text{clean}_{EE}(\pi^*)$</pre> <p>$\text{clean}_{\text{reduced}}(\pi^*)$:</p> <pre> 49 return $\text{clean}_{\text{LTSS}}(\pi^*)$ or $\text{clean}_{\text{ELT}}(\pi^*)$ or $\text{clean}_{\text{ESS}}(\pi^*)$</pre> <p>$\text{clean}_{EE}(\pi^*)$:</p> <pre> 50 return $\pi^*.revrand = \text{false}$ and $\text{clean}_{\text{peerE}}(\pi^*)$ #randomness of test session is unrevealed and ephemeral contribution of peer is clean</pre> <p>$\text{clean}_{\text{peerE}}(\pi^*)$:</p> <pre> 51 return 52 $(\pi^*.role = \text{initiator and } \exists \pi \neq \pi^* :$ $(\pi.role = \text{responder and } \pi.cid = \pi.cid \text{ and } \pi.revrand = \text{false}))$ #there exists a contributively partnered responder session exists whose randomness is unrevealed 53 or $(\pi^*.role = \text{responder and } \exists \pi \neq \pi^* :$ $(\pi.role = \text{initiator and } \pi.sid = \pi.sid \text{ and } \pi.revrand = \text{false}))$ #there exists a partnered initiator session (which is unique by sound) whose randomness is unrevealed</pre>	<p>$\text{clean}_{\text{LTSS}}(\pi^*)$:</p> <pre> 54 return 55 $(\pi^*.role = \text{initiator and corrltk}_{\pi^*.oid} = \text{false and corrssk}_{\pi^*.ssid} = \text{false})$ #long-term secret of initiator test session and semi-static key of responder peer are uncorrupted 56 or $(\pi^*.role = \text{responder and corrltk}_{\pi^*.pid} = \text{false and corrssk}_{\pi^*.ssid} = \text{false})$ #long-term secret of responder peer and semi-static key of initiator test session are uncorrupted</pre> <p>$\text{clean}_{\text{ELT}}(\pi^*)$:</p> <pre> 57 return 58 $(\pi^*.role = \text{initiator and } \pi^*.revrand = \text{false and corrltk}_{\pi^*.pid} = \text{false})$ #randomness of initiator test session is unrevealed and long-term secret of responder peer is uncorrupted 59 or $(\pi^*.role = \text{responder and clean}_{\text{peerE}}(\pi^*) \text{ and corrltk}_{\pi^*.oid} = \text{false})$ #long-term secret of responder test session is uncorrupted and ephemeral contribution of initiator peer is clean</pre> <p>$\text{clean}_{\text{ESS}}(\pi^*)$:</p> <pre> 60 return 61 $(\pi^*.role = \text{initiator and } \pi^*.revrand = \text{false and corrssk}_{\pi^*.ssid} = \text{false})$ #randomness of initiator test session is unrevealed and semi-static secret of responder peer is uncorrupted 62 or $(\pi^*.role = \text{responder and clean}_{\text{peerE}}(\pi^*) \text{ and corrssk}_{\pi^*.ssid} = \text{false})$ #semi-static secret of responder test session is uncorrupted and ephemeral contribution of initiator peer is clean</pre>

Figure 9: Key indistinguishability game $\mathcal{G}_{KE}^{\text{ke-kind}}(\mathcal{A})$ for key exchange protocol KE against adversary \mathcal{A} ; composed of the main game (top section), oracles (middle section), and clean predicates defining freshness (bottom section). Without loss of generality, we assume that all queries that the adversary makes to the oracles are *well-defined* and *valid*, i.e., of the expected type and in the appropriate ranges.

space \mathcal{R}_{KE} is *correct* if any protocol execution between honest parties without interference by the adversary results in two sessions which accept with the same session key and session identifier.

7.1.2 Soundness. Soundness, captured in the predicate *sound*, describes the behavior with respect to a correct protocol execution. If an adversary \mathcal{A} manages to create one of the following situations, it will win the game immediately:

- (i) Two sessions accept with the same session identifier, but derive different session keys, indicate different handshake types (full vs. reduced), or do not agree on their contributive identifiers (Fig. 9, Line 18).
- (ii) Two initiator sessions accept with the same session identifier (Fig. 9, Line 19).
- (iii) Three sessions accept with the same session identifier in full handshake type (Fig. 9, Line 20).

7.1.3 Freshness. Granting the adversary \mathcal{A} access to the oracles described in Figure 9 without restriction would allow the adversary to trivially win the game, e.g., by testing a session key and then revealing it or corrupting all secrets used in the key derivation of a session. Therefore, as in the core model, we require the tested session to be *fresh* and for this introduce the predicate *fresh* (cf. Figure 9) which takes as input the test session π^* and prohibits all “trivial wins”. On a high level, the session key derived in the test session is considered to be *fresh* if the following criteria hold:

- (i) The session key of the test session has not been revealed to \mathcal{A} via a REVEALSESSKEY query (Fig. 9, Line 14).
- (ii) The session key of any partnered session (i.e., any session with the same session identifier as the test session) has not been revealed to \mathcal{A} via a REVEALSESSKEY query (Fig. 9, Line 15).
- (iii) \mathcal{A} has not obtained sufficiently many secrets to derive the session key of the test session itself via CORRUPTLTKEY and/or CORRUPTSSKEY and/or REVEALRANDOM queries (Fig. 9, Line 16).

Clean keys. Following the terminology of Cohn-Gordon, Cremers, Dowling, Garratt, and Stebila [22] the last criterion of freshness is captured by a series of so-called clean predicates, which we discuss next. The formal description can also be found in Figure 9. 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 assume that $\pi^*.\text{type} = \text{reduced}$. Intuitively, a session key derived in such a session remains unknown to the adversary, if one of the three keys that constitute the master secret, is “clean”, i.e., cannot be computed by the adversary. This is the case if either of the following three clean predicates holds for the test session π^* :

$\text{clean}_{\text{LTSS}}$: This predicate indicates whether the combination of the long-term key of the initiator and the semi-static key of the responder is unknown to the adversary.

$\text{clean}_{\text{ELT}}$: This predicate indicates whether the combination of the ephemeral contribution³ of the initiator and the long-term key of the responder is unknown to the adversary.

³Recall that the ephemeral contribution in initiator sessions is determined by the session specific randomness coins $\in \mathcal{R}_{KE}$.

$\text{clean}_{\text{ESS}}$: This predicate indicates whether the combination of the ephemeral contribution of the initiator and the semi-static key of the responder is unknown to the adversary.

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 flags corrltk , corrssk , and revrand , respectively). For responder test sessions, $\text{clean}_{\text{peerE}}$ indicates whether the randomness used in any partnered initiator session π_p^* (if test session responder) 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 whether the combination of the ephemeral contribution of the initiator and the ephemeral pre-key of the responder is unknown to the adversary.

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 or uncorrupted, respectively.

7.1.4 Main differences to the model in [22]. Our authenticated key exchange model closely follows the one used in the original Signal analysis of Cohn-Gordon et al. [22]. We make the following modifications: Since we are only concerned about the initial key agreement and not the subsequent symmetric and asymmetric ratcheting stages, we can forgo the notion of *multi-stage* AKE security, where multiple sessions keys are derived in a series of stages. Lastly, we explicitly take the *deniability* feature of Signal into account in a separate notion to avoid establishing a post-quantum solution that forgoes a key requirement of the specification.

7.2 Deniability

Definition 7.2. An asynchronous DAKE protocol ADAKE is (t, ϵ, Q_{Ch}) -deniable if for any adversary \mathcal{A} with running time at most t and making at most Q_{Ch} many queries to its CHALL oracle, we have that

$$\text{Adv}_{\text{ADAKE}}^{\text{adake-den}}(\mathcal{A}) = \left| \Pr \left[\mathcal{G}_{\text{ADAKE}}^{\text{adake-den}}(\mathcal{A}) = 1 \right] - \frac{1}{2} \right| \leq \epsilon,$$

where $\mathcal{G}_{\text{ADAKE}}^{\text{adake-den}}(\mathcal{A})$ is defined in Figure 10.

The main difference between the textual description of the deniability game in Section 5.2 and Figure 10 is the use of semi-static keys. Here, we generate n_{ss} many semi-static keys per party, all of which are given to the adversary. When querying the challenge oracle, the adversary may choose the semi-static key that the responder uses. The pre-key bundle of the responder may depend on the semi-static key. The Initiator key agreement and the Fake algorithm also use the semi-static key as specified by ADAKE.

8 SPQR SECURITY PROOFS

In this section we present the security results for our SPQR protocol (Figure 8) via theorem statements and detailed proofs of both key-indistinguishability and deniability in the previously described security model (Section 7).

```

 $\mathcal{G}_{\text{ADAKE}}^{\text{adake-den}}(\mathcal{A})$ :
1  $\mathcal{L} \leftarrow \emptyset$  //list of keys for the adversary
2 for  $U \in [n_p]$ 
3    $(pk_U, sk_U) \leftarrow \text{KGenLT}()$  //long-term key generation
4    $\mathcal{L} \leftarrow \mathcal{L} \cup \{(pk_U, sk_U)\}$ 
5   for  $ssid \in [n_{ss}]$ 
6      $(sspk_U^{ssid}, sssk_U^{ssid}) \leftarrow \text{KGenSS}()$  //semi-static key generation
7      $\mathcal{L} \leftarrow \mathcal{L} \cup \{(sspk_U^{ssid}, sssk_U^{ssid})\}$ 
8    $\vec{pk} \leftarrow \{pk_U\}_{U \in [n_p]}$ ;  $sspk \leftarrow \{sspk_U^{ssid}\}_{U \in [n_p]}$ 
9    $b \leftarrow \{0, 1\}$ 
10   $b' \leftarrow \mathcal{A}^{\text{CHALL}}(\mathcal{L})$ 
11  return  $\llbracket b' = b \rrbracket$ 
 $\text{CHALL}(\text{iid}, \text{rid}, \text{ssid})$ :
12 if  $b = 0$ 
13    $\pi_{\text{rid}}.\text{role} \leftarrow \text{responder}$ ;  $\pi_{\text{rid}}.\text{st}_{\text{exec}} \leftarrow \text{running}$  //initialize session variables
14    $\pi_{\text{iid}}.\text{role} \leftarrow \text{initiator}$ ;  $\pi_{\text{iid}}.\text{st}_{\text{exec}} \leftarrow \text{running}$ 
15    $(\pi_{\text{rid}}, m) \leftarrow \text{Run}(sk_{\text{rid}}, \vec{pk}, sspk, \pi_{\text{rid}}, (\text{create}, \text{ssid}))$  //build pre-key bundle
16    $(\pi_{\text{iid}}, m') \leftarrow \text{Run}(sk_{\text{iid}}, \vec{pk}, sspk, \pi_{\text{iid}}, m)$  //initiator sends message
17    $(K, T) \leftarrow (\pi_{\text{iid}}, K, (m, m'))$  //save session key and transcript
18 else
19    $(K, T) \leftarrow \text{Fake}(pk_{\text{iid}}, sk_{\text{rid}}, sspk, \text{ssid})$ 
20 return  $(K, T)$ 

```

Figure 10: Security game for deniability of an asynchronous DAKE protocol ADAKE against an adversary \mathcal{A} .

Before we start, we translate the informal protocol description of SPQR given in Figure 8 into the syntax of our model. The resulting protocol flow is depicted in Figure 11.

8.1 Key Indistinguishability

THEOREM 8.1 (KEY INDISTINGUISHABILITY OF SPQR). *Let DVS be a $(t, \epsilon_{\text{DVS}}, n_p \cdot n_{ss}, Q_S)$ -unforgeable DVS scheme.*

Let KEM_1 be a $(t, \epsilon_{\text{KEM}_1}, n_s)$ -IND-CCA-secure KEM, KEM_2 be a $(t, \epsilon_{\text{KEM}_2}, n_s)$ -IND-CCA-secure KEM, KEM_3 be a $(t, \epsilon_{\text{KEM}_3}, 1)$ -IND-CCA-secure KEM with randomness space $\mathcal{R}_{\text{KEM}_3}$, and δ_{corr} be the maximal correctness error among KEM_1 , KEM_2 , and KEM_3 .

Let KDF be a $(t, \epsilon_{\text{KDF}}, n_s)$ -PRF-secure key derivation function when keyed through any key component K_1, K_2, K_3 , and tPRF a $(t, \epsilon_{\text{tPRF}}, n_s)$ -secure twisted pseudorandom function with label space $\mathcal{R}_{\text{tPRF}}$.

Then the SPQR protocol with randomness space $\mathcal{R}_{\text{KE}} = \{0, 1\}^\lambda \times \mathcal{R}_{\text{tPRF}} \times \mathcal{R}_{\text{KEM}_3}$ as shown in Figure 8 and formalized in Figure 11 provides $(t', \epsilon', (Q_{\text{Snd}}, Q_{\text{CorrLT}}, Q_{\text{CorrSS}}, Q_{\text{RevR}}, Q_{\text{RevSK}}))$ -key indistinguishability for $t \approx t'$ and

$$\epsilon' \leq \frac{n_s^2}{2^\lambda} + \frac{n_s^2}{2^{|\mathcal{R}_{\text{tPRF}}|}} + \frac{n_s^2}{2^{|\mathcal{R}_{\text{KEM}_3}|}} + 3n_s \cdot \delta_{\text{corr}} + n_s \cdot n_p^2 \cdot \left(\frac{n_{ss} \cdot (\epsilon_{\text{DVS}} + 2n_s \cdot (\epsilon_{\text{tPRF}} + \epsilon_{\text{KEM}_2} + \epsilon_{\text{KDF}}))}{+ n_s \cdot (2\epsilon_{\text{tPRF}} + \epsilon_{\text{KEM}_1} + \epsilon_{\text{KEM}_3} + 2\epsilon_{\text{KDF}})} \right),$$

where $n_s \leq Q_{\text{Snd}}$ is the maximum number of sessions (upper bounded by the number Q_{Snd} of SEND queries), n_p the number of parties, and n_{ss} the number of semi-static keys per party.

PROOF. We proceed via a sequence of game hops starting from $\mathcal{G}_{\text{SPQR}}^{\text{ke-kind}}(\mathcal{A})$ (cf. Figure 9), branching off into the cleanness predicates. In the final games, we will have that the adversary \mathcal{A} has probability exactly $\frac{1}{2}$ in guessing the test challenge bit b . Along the way, we will further establish that the soundness predicate sound is satisfied.

Game 0. The initial game, Game \mathcal{G}_0 , is the key indistinguishability game $\mathcal{G}_{\text{SPQR}}^{\text{ke-kind}}(\mathcal{A})$ for SPQR played by \mathcal{A} . By definition,

$$\epsilon' := \text{Adv}_{\text{SPQR}}^{\text{ke-kind}}(\mathcal{A}) = \text{Adv}_{\text{SPQR}}^{\mathcal{G}_0}(\mathcal{A}) = \left| \Pr[\mathcal{G}_0 = 1] - \frac{1}{2} \right|.$$

Game 1 (Nonce and randomness collisions). We modify \mathcal{G}_0 to overwrite the adversary's output with 0, if any two initiator sessions hold the same nonce n or the same randomness value r , or if two responder sessions pick the same ephemeral KEM key pair. As initiator nonces are uniformly random λ -bit strings, the initiator randomness is a uniformly random element from tPRF's label space $\mathcal{R}_{\text{tPRF}}$, and the responder randomness is a uniformly random element from KEM_3 's randomness space $\mathcal{R}_{\text{KEM}_3}$, we can upper-bound the probability of this happening across the at most n_s sessions by the birthday bound:

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_0}(\mathcal{A}) \leq \frac{n_s^2}{2^\lambda} + \frac{n_s^2}{2^{|\mathcal{R}_{\text{tPRF}}|}} + \frac{n_s^2}{2^{|\mathcal{R}_{\text{KEM}_3}|}} + \text{Adv}_{\text{SPQR}}^{\mathcal{G}_1}(\mathcal{A}).$$

Game 2 (KEM correctness). We modify \mathcal{G}_1 to overwrite the adversary's output with 0 if for any key pair $(pk, sk) \leftarrow \text{KGen}_l()$ and encapsulation $(c, K) \leftarrow \text{Encaps}_l(pk)$ used in any of the sessions, where $l \in \{1, 2, 3\}$, we have that $K \neq K' \leftarrow \text{Decaps}_l(sk, c)$. The probability of this happening for any tuple (pk, sk, c) is upper-bounded by the maximal correctness error δ_{corr} among KEM_1 , KEM_2 , and KEM_3 . As there are at most three such tuples per session, we can bound any correctness errors happening by:

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_1}(\mathcal{A}) \leq 3n_s \cdot \delta_{\text{corr}} + \text{Adv}_{\text{SPQR}}^{\mathcal{G}_2}(\mathcal{A}).$$

Soundness. At this point, soundness (i.e., $\text{sound}() = \text{true}$) holds unconditionally and this will not change in any of the subsequent game hops. Consider the three sub-conditions of the sound predicate:

- *Shared key, type, contributive identifier* (Figure 9, Line 18): Session identifiers fix the KEM keys and ciphertexts involved in key derivation, hence by game \mathcal{G}_2 KEM correctness implies agreement on K_1, K_2 , and (if type = full) K_3 and thus also on K under deterministic key derivation KDF. Session identifiers have distinct entries depending, and ensuring agreement, on the session type ($\text{epk}_B = \perp$ if and only if type = reduced). Since the entries in the contributive identifier are a (proper) subset of the entries of session identifiers, agreement on session identifiers also yields agreement on contributive identifiers.
- *No initiator session identifiers collide* (Figure 9, Line 19): As of Game \mathcal{G}_1 , each initiator session picks a unique nonce. This nonce is part of the session identifier and thus ensures uniqueness of initiator session identifiers.
- *No three session identifiers collide in full mode* (Figure 9, Line 20): We ruled out initiator collisions above already. For session identifiers to collide in two responder sessions in full mode, the two sessions would need to use the same ephemeral pre-key (epk, esk) . Since we ruled out collisions in the randomness space $\mathcal{R}_{\text{KEM}_3}$ sampled in responder sessions, the ephemeral pre-keys derived via $\text{KGenEP}()$ are unique.



Figure 11: Formal specification of the Run algorithm of asynchronous DAKE SPQR wrt. the model given in Section 7. Note that the generation of long-term and semi-static keys during registration happens at the outset of the game $\mathcal{G}_{\text{KE}}^{\text{ke-kind}}(\mathcal{A})$ and the Signal Server is abstracted away. Thus these elements are not included explicitly in the above description.

Game 3 (Guess test session π^*). Next, we guess the tested session π^* among the at most n_s sessions total at the outset of the game, and “invalidate” the game by overwriting the adversary’s bit guess with 0 if the adversary calls TEST on a different session. With probability $1/n_s$, the guess is correct and this change goes unnoticed, so

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_2}(\mathcal{A}) \leq n_s \cdot \text{Adv}_{\text{SPQR}}^{\mathcal{G}_3}(\mathcal{A}).$$

Game 4 (Guess initiator identity U). We first guess the test session’s own identity if it is an initiator session, or the test session’s peer identity if it is a responder session. Note that since the test session has necessarily accepted, the peer in a responder session is also set to a valid identity in $[n_p]$, i.e., is not set to \star anymore. We denote the guessed initiator identity by U and overwrite 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 :

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_3}(\mathcal{A}) \leq n_p \cdot \text{Adv}_{\text{SPQR}}^{\mathcal{G}_4}(\mathcal{A}).$$

Game 5 (Guess responder identity V). Next, we guess the identity of the involved (intended) responder. This is π^* ’s own identity if it is a responder session, or its intended peer identity if π^* is an initiator session. We denote the guessed responder identity by V and again overwrite 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 :

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_4}(\mathcal{A}) \leq n_p \cdot \text{Adv}_{\text{SPQR}}^{\mathcal{G}_5}(\mathcal{A}).$$

Recall that the adversary’s bit guess at the end of the game is considered only if $\text{fresh}(\pi^*)$ holds for the tested session π^* . Freshness requires that the session key in neither π^* nor in a partnered session was revealed and that one of these four cleanness conditions is satisfied: $\text{clean}_{\text{LTSS}}(\pi^*)$ or $\text{clean}_{\text{ELT}}(\pi^*)$ or $\text{clean}_{\text{ESS}}(\pi^*)$ or $(\pi^*.\text{type} = \text{full} \text{ and } \text{clean}_{\text{EE}}(\pi^*))$.

We will now branch out into four sub-cases following the structure of the cleanness predicates, bounding the adversary’s winning advantage $\text{Adv}_{\text{SPQR}}^{\mathcal{G}_5}(\mathcal{A})$ by the sum of its advantages when conditioning the adversary on each of the cleanness sub-conditions being satisfied (which we write as $\mathcal{G}_5[c]$ for predicate c). Via the union bound:

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_5}(\mathcal{A}) \leq \sum_{c \in \{\text{clean}_{\text{LTSS}}(\pi^*), \text{clean}_{\text{ELT}}(\pi^*), \text{clean}_{\text{ESS}}(\pi^*), \pi^*.\text{type}=\text{full} \wedge \text{clean}_{\text{EE}}(\pi^*)\}} \text{Adv}_{\text{SPQR}}^{\mathcal{G}_5[c]}(\mathcal{A}).$$

Case A ($\text{clean}_{\text{LTSS}}(\pi^*)$). In this proof case, we are guaranteed that either

- (1) π^* is an initiator session owned by U for which both its own long-term key and its intended peer V ’s semi-static key are uncorrupted or
- (2) π^* is a responder session owned by V whose own semi-static and intended peer U ’s long-term keys are both uncorrupted.

We will leverage this to show that the KEM ciphertext c_2 exchanged with the test session π^* was generated for an uncorrupted KEM

key with good randomness, bootstrapping key indistinguishability from the corresponding encapsulated key K_2 .

Game A.0. This is the game conditioned on $\text{clean}_{\text{LTSS}}(\pi^*)$ being satisfied.

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_{A.0}}(\mathcal{A}) = \text{Adv}_{\text{SPQR}}^{\mathcal{G}_5[\text{clean}_{\text{LTSS}}(\pi^*)]}(\mathcal{A}).$$

Game A.1 (Guess semi-static key of V). We now guess the identifier ssid of the responder V ’s (uncorrupted) semi-static key $\text{sspk}_V^{\text{ssid}}$. Note that depending on the role of π^* this is either the test session’s own key (if $\pi^*.\text{role} = \text{responder}$), or of the intended peer (if $\pi^*.\text{role} = \text{initiator}$). We denote the guessed identifier by ssid^* , and abort, setting the adversary’s output bit to 0, if this guess is incorrect, losing at most a factor n_{ss} of the number of semi-static keys per user:

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

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 a responder session and accepts having received a DVS signature σ that no session of U has issued. The probability of such an abort can be bounded by the advantage of the following reduction \mathcal{B}_1 against the $(t, \epsilon_{\text{DVS}}, Q_S)$ -existential unforgeability of DVS.

Reduction \mathcal{B}_1 samples all key components itself except for the DVS keys: In place of the long-term public DVS key pk_U^{DVS} of U and the semi-static public DVS key $\text{sspk}_V^{\text{DVS}}$ of V it uses the public keys pk_S and pk_D , respectively, obtained in its unforgeability game. For all semi-static DVS (verifier) keys, it uses the public-secret key pairs obtained through the list \mathcal{L} in the unforgeability game, while sampling all static DVS keys itself. (Hence, it knows the secret key component for all DVS keys but pk_U^{DVS} and $\text{sspk}_V^{\text{DVS}}$). In its simulation of Game $\mathcal{G}_{A.2}$, \mathcal{B}_1 uses its signing oracle to compute signatures under sk_U^{DVS} (and for any peer semi-static public key sspk). Since $\text{clean}_{\text{LTSS}}(\pi^*) = \text{true}$, \mathcal{B}_1 never has to answer the query $\text{CORRUPTLTKEY}(U)$ or $\text{CORRUPTSSKEY}(V, \text{ssid})$. Hence \mathcal{B}_1 can provide a perfect simulation of $\mathcal{G}_{A.2}$, and if π^* as a responder receives a signature σ on a session-identifier message sid that no session of U has issued, \mathcal{B}_1 can output this as its forgery and wins. Thus,

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

Game A.3 (Guess partnered initiator session). By Game $\mathcal{G}_{A.2}$, we are now ensured that a responder test session does not accept unless an honest session π_p^* has sent the ciphertext c_2 that π^* received, as c_2 is signed under σ . We now guess this initiator session π_p^* (if π^* is a responder), aborting and setting \mathcal{A} ’s output to 0, if the test session is a responder and the guess was incorrect. This reduces the adversary’s advantage by a factor of at most the number of sessions n_s :

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

Game A.4 (Twisted PRF randomness). Next, we replace all tPRF evaluations involving U ’s long-term secret tk_U by the evaluation of a randomly chosen function. This, in particular, replaces the

value r_2 in π^* (if π^* is an initiator) or in π_p^* (if π^* is a responder) with an independent random value \tilde{r}_2 (recall that the randomness value r is unique per session as of Game \mathcal{G}_1).

We bound the advantage difference introduced by this step based on the $(t, \epsilon_{\text{tPRF}}, n_s)$ -twisted pseudorandomness of tPRF via the following reduction \mathcal{B}_2 . The reduction \mathcal{B}_2 receives a sequence of n_s tuples (x_i, y_i) (and a tuple (K', z) but this is not relevant for our purposes here), which is either $((x_1, \text{tPRF}(K, x_1)), \dots, (x_q, \text{tPRF}(K, x_q)))$ or $((x_1, g(x_1)), \dots, (x_q, g(x_q)))$ for random values K, x_1, \dots, x_q and a randomly chosen function g .

During the reduction, instead of sampling the tPRF key tk_U itself, \mathcal{B}_2 will simply use y_i as the expanded randomness in the i -th initiator session of U (there are at most n_s such sessions), setting $r_1 \| r_2 \| r_3 \| r_4 \leftarrow y_i$. (\mathcal{B}_2 simulates the rest of the game as usual, in particular generating the tPRF keys for all other users itself.) As its bit guess, \mathcal{B}_2 outputs 1 if \mathcal{A} wins the game and 0 otherwise. Depending on which sequence \mathcal{B}_2 is given, it either simulates $\mathcal{G}_{A.3}$ or $\mathcal{G}_{A.4}$, thus

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

Game A.5 (Semi-static KEM). In the following, let

$$(c_2, K_2) \leftarrow \text{KEM}_2.\text{Encaps}(sspk_V^{\text{ssid}^*}; \tilde{r}_2)$$

be the encapsulation computed in the initiator session between π^* and π_p^* under the semi-static key identified by ssid^* of V . Recall that by the previous game, \tilde{r}_2 is an independent random value, unknown to the adversary. This allows us to now replace the key K_2 encapsulated in c_2 with a randomly sampled key \tilde{K}_2 in π^* and its partnered session(s), if existent. Furthermore, we replace K_2 with \tilde{K}_2 in any session of V using ssid^* that has received the same encapsulating ciphertext c_2 .⁴

We can now bound \mathcal{A} 's difference in advantage by the advantage of a reduction \mathcal{B}_3 in winning the $(t, \epsilon_{\text{KEM}_2}, n_s)$ -IND-CCA security game for KEM_2 . The reduction \mathcal{B}_3 obtains the IND-CCA challenge (pk, c^*, K_b^*) and simulates the game for \mathcal{A} as follows: It samples the test bit b_{test} itself and generates all long-term, semi-static, and ephemeral pre-keys itself, except for the key pair $(sspk_V^{\text{ssid}^*}, sssk_V^{\text{ssid}^*})$ of the previously guessed responder identity V and identifier ssid^* . The reduction embeds its received challenge public key pk by setting $sspk_V^{\text{ssid}^*} = pk$. As predicate $\text{clean}_{\text{LTSS}}$ holds, \mathcal{A} never asks the query $\text{CORRUPTSSKEY}(V, \text{ssid}^*)$ and the reduction thus never needs to output the secret key sk corresponding to pk .

Whenever a decapsulation of some ciphertext $c \neq c^*$ using sk is necessary to faithfully simulate the game for \mathcal{A} , \mathcal{B}_3 simply forwards this ciphertext to its decapsulation oracle DECAPS (making at most n_s queries as claimed). In both π^* and its partnered session(s) π_p^* (if existent), \mathcal{B}_3 embeds K_b^* wherever K_2 would be used and c^* wherever c_2 would be used. The same replacement is employed in responder sessions of party V that receive c_2 as an encapsulation under $sspk_V^{\text{ssid}^*} = pk$. When \mathcal{A} stop with output b' , the reduction \mathcal{B}_3 returns $\llbracket b_{\text{test}} = b' \rrbracket$.

⁴Note that we know the involved initiator session of U (it is either the test session π^* itself or its partnered session π_p^*) and the identity V of the owner of the involved semi-static key pair with id ssid^* . This allows us to precompute c_2 at the outset of the game and thus easily identify responder sessions that receive c_2 .

Observe that \mathcal{B}_3 perfectly simulates $\mathcal{G}_{A.4}$ if $b = 0$ in $\mathcal{G}_{\text{KEM}_2}^{\text{indcca}}(\mathcal{B}_3)$ and $\mathcal{G}_{A.5}$ otherwise. Hence, any difference in \mathcal{A} 's advantage in the two games is bounded by the distinguishing advantage of \mathcal{B}_3 against the IND-CCA security of KEM_2 :

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_{A.4}}(\mathcal{A}) \leq \epsilon_{\text{KEM}_2} + \text{Adv}_{\text{SPQR}}^{\mathcal{G}_{A.5}}(\mathcal{A}).$$

Game A.6 (Session key KDF). Lastly, we replace the output of the session key derivation $K \leftarrow \text{KDF}(K_1 \| \tilde{K}_2 \| K_3, \text{sid})$ in the test session and its partnered session(s), as well as in any other session using \tilde{K}_2 , by the output of a random function; in particular replacing K with a uniformly random key \tilde{K} . We show that any adversary that can efficiently distinguish Game $\mathcal{G}_{A.6}$ from Game $\mathcal{G}_{A.5}$ can be turned into an efficient adversary \mathcal{B}_4 against the $(t, \epsilon_{\text{KDF}}, n_s)$ -pseudorandomness of the key derivation function KDF , treated as a PRF keyed through the second key component K_2 and taking (K_1, K_3, sid) as label.

The reduction \mathcal{B}_4 generates all key pairs itself and initializes \mathcal{A} as usual. In particular, \mathcal{B}_4 samples the test bit b_{test} itself and can answer all CORRUPTSSKEY , CORRUPTLTKEY queries truthfully. Similarly, the reduction \mathcal{B}_4 can execute all SEND queries. Furthermore, \mathcal{B}_4 can reveal the randomness and the session keys of sessions, with the exception of session keys in the test session and its partnered session(s) (which is unproblematic since these queries would trigger an immediate loss for the adversary when checking $\text{fresh}(\pi^*)$).

In any session using \tilde{K}_2 as of Game $\mathcal{G}_{A.5}$, and in particular in the test session π^* and its partner(s), \mathcal{B}_4 queries (K_1, K_3, sid) to its PRFCHALLENGE oracle to compute the session key, where $\text{sid} = (U, V, pk_U, pk_V, sspk_V, epk_V, n, c_1, c_2, c_3)$; this amount to at most n_s oracle queries, as claimed. The returned values are either $\text{KDF}(K_1 \| \tilde{K}_2 \| K_3, \text{sid})$ for a uniformly random key \tilde{K}_2 if $b = 0$, or the outputs of a uniformly random function g if $b = 1$. When \mathcal{A} terminates with output b' , \mathcal{B}_4 returns $\llbracket b_{\text{test}} = b' \rrbracket$.

Note that \mathcal{B}_4 perfectly simulates $\mathcal{G}_{A.5}$ if $b = 0$ and $\mathcal{G}_{A.6}$ if $b = 1$. Hence, if \mathcal{A} can distinguish the two games, the reduction can win the PRF security game against KDF with the same advantage and we have

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_{A.5}}(\mathcal{A}) \leq \epsilon_{\text{KDF}} + \text{Adv}_{\text{SPQR}}^{\mathcal{G}_{A.6}}(\mathcal{A}).$$

Finalize. To conclude this proof case, observe that in Game $\mathcal{G}_{A.6}$ the challenge K_{test} for π^* is now a uniformly random key, independent of b_{test} . Furthermore, \mathcal{A} cannot reveal K_{test} via a REVEALSESSKEY query on π^* or any partnered session who might hold the same key. Thus, \mathcal{A} cannot gain any information about the test bit b_{test} and can do no better than to guess:

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

Case B (clean_{ELT}(π^)).* In this proof case, we are guaranteed that either

- (1) π^* is an initiator session owned by U whose randomness is unrevealed and whose intended peer V 's long-term key is uncorrupted or

- (2) π^* is a responder session owned by V and (via $\text{clean}_{\text{peerE}}$) there exists a unique partnered initiator session π_p^* whose randomness is unrevealed and which is unique (via sound). We further know that π_p^* is owned by U , as the matching session identifiers include the initiator identity guessed in Game \mathcal{G}_4 .

Game B.0. We now condition on $\text{clean}_{\text{ELT}}(\pi^*)$:

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_{B.0}}(\mathcal{A}) = \text{Adv}_{\text{SPQR}}^{\mathcal{G}_5[\text{clean}_{\text{ELT}}(\pi^*)]}(\mathcal{A}).$$

Game B.1 (Guess unique partnered initiator session). As mentioned above, if the test session π^* is a responder session, by $\text{clean}_{\text{peerE}}$ there exists an initiator partner session π_p^* to π^* which furthermore is unique by sound . We now guess this partnered initiator session π_p^* owned by party U ; if π^* is an initiator session we simply ignore the guess. The game is changed to overwrite \mathcal{A} 's output to 0 if the test session is a responder and the guess was incorrect. This reduces the adversary's advantage by a factor of at most the number of sessions n_s :

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

Game B.2 (Twisted PRF randomness). Next, we replace the tPRF evaluation of the initiator session π owned by U by the evaluation of a randomly chosen function (here $\pi = \pi^*$, if $\pi^*.role = \text{initiator}$, and $\pi = \pi_p^*$, if $\pi^*.role = \text{responder}$). In particular, we replace the value r_1 in π with an independent random value \tilde{r}_1 (recall the randomness value r is unique per session as of Game \mathcal{G}_1).

We bound the advantage difference introduced by this step by the $(t, \epsilon_{\text{tPRF}}, 0)$ -twisted pseudorandomness of tPRF via the following reduction \mathcal{B}_5 . The reduction receives (K', z) which is either $(K', \text{tPRF}(K', x))$ if $b = 0$, or $(K', g'(K'))$ if $b = 1$, where K', x are random values and g' is a random function.

The reduction \mathcal{B}_5 then generates all keys and parameters for the key exchange games itself, but sets $tk_U \leftarrow K'$. It uses tk_U in all sessions of U except for π , where instead of evaluating $\text{tPRF}(tk_U, r)$, \mathcal{B}_5 sets $r_1 \| r_2 \| r_3 \| r_4 \leftarrow z$. Upon a potential $\text{CORRUPTLTKEY}(U)$ query, \mathcal{B}_5 can hand out tk_U as part of U 's secret key (note that r remains hidden as $\text{REVEALRANDOM}(\pi)$ is never called). As its bit guess, \mathcal{B}_5 outputs 1 if \mathcal{A} wins the game and 0 otherwise. Depending on which sequence \mathcal{B}_5 is given, it either simulates $\mathcal{G}_{B.1}$ or $\mathcal{G}_{B.2}$, and thus:

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_{B.1}}(\mathcal{A}) \leq \epsilon_{\text{tPRF}} \cdot \text{Adv}_{\text{SPQR}}^{\mathcal{G}_{B.2}}(\mathcal{A}).$$

Game B.3 (Long-term KEM). In the following, let

$$(c_1, K_1) \leftarrow \text{KEM}_1.\text{Encaps}(pk_V^{\text{KEM}}; \tilde{r}_1)$$

be the encapsulation computed at session π , where again $\pi = \pi^*$, if $\pi^*.role = \text{initiator}$, and $\pi = \pi_p^*$, if $\pi^*.role = \text{responder}$. Recall that by the previous game, \tilde{r}_1 is an independent random value, unknown to the adversary. In Game $\mathcal{G}_{B.3}$, we now replace the encapsulated key K_1 with a randomly sampled key \tilde{K}_1 in the test session π^* and its partnered session(s), if existent. Furthermore, in any responder session of V that receives the same encapsulating ciphertext c_1 , we replace K_1 with \tilde{K}_1 , too. Observe that, knowing the involved initiator session π as well as the long-term key identity V in advance,

we can precompute c_1 at the outset of the game and then simply check when c_1 is received by responder sessions owned by V .

We bound the difference in \mathcal{A} 's advantage by the advantage of a reduction \mathcal{B}_6 against the $(t, \epsilon_{\text{KEM}_1}, n_s)$ -IND-CCA security of KEM_1 as follows. \mathcal{B}_6 obtains a challenge (pk, c^*, K_b^*) and simulates the game for \mathcal{A} as follows: It samples the test bit b_{test} itself and generates all key pairs to initialize \mathcal{A} itself, except for the long-term KEM public key of V , for which it only sets $pk_V^{\text{KEM}} = pk$.

Note that $\text{clean}_{\text{ELT}}$ ensures that \mathcal{A} never calls $\text{CORRUPTLTKEY}(V)$, so \mathcal{B}_6 never has to output sk_V^{KEM} . Whenever \mathcal{B}_6 would have to use sk_V^{KEM} to decapsulate some ciphertext $c \neq c^*$ in some responder session of V , it does so via its DECAPS oracle (querying the oracle at most n_s times as claimed). In the test session π^* and its potential initiator partner π_p^* , \mathcal{B}_6 embeds K_b^* in the place of K_1 and c^* in the place of c_1 . Also, in responder sessions of V receiving c_1 (recall, \mathcal{B}_6 knows c_1 from the start of the game), \mathcal{B}_6 uses K_b^* in the place of K_1 and c^* in the place of c_1 . When \mathcal{A} stops and outputs its bit guess b' , \mathcal{B}_6 returns $\llbracket b_{\text{test}} = b' \rrbracket$.

The simulation \mathcal{B}_6 provides for \mathcal{A} perfectly represents Game $\mathcal{G}_{B.2}$ if $b = 0$ in the IND-CCA game for KEM_1 , and Game $\mathcal{G}_{B.3}$ otherwise. Any difference in \mathcal{A} 's advantage between the two games hence translates into a distinguishing advantage of \mathcal{B}_6 in the IND-CCA game against KEM_1 :

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_{B.2}}(\mathcal{A}) \leq \epsilon_{\text{KEM}_1} + \text{Adv}_{\text{SPQR}}^{\mathcal{G}_{B.3}}(\mathcal{A}).$$

Game B.4 (Session key KDF). As the final step in this proof case, we replace in Game $\mathcal{G}_{B.4}$ the session key derived in the test and partnered session as $K \leftarrow \text{KDF}(K_1 \| K_2 \| K_3, \text{sid})$, as well as in any other session using \tilde{K}_1 , by the output of a random function; in particular replacing K with a randomly sampled key \tilde{K} . As in the previous case in Game $\mathcal{G}_{A.6}$ we can bound the advantage introduced by this change by the advantage of an adversary \mathcal{B}_7 against the $(t, \epsilon_{\text{KDF}}, n_s)$ -pseudorandomness property of KDF, treated as a PRF keyed through the first key component K_1 and taking (K_2, K_3, sid) as label:

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_{B.3}}(\mathcal{A}) \leq \epsilon_{\text{KDF}} + \text{Adv}_{\text{SPQR}}^{\mathcal{G}_{B.4}}(\mathcal{A}).$$

Finalize. To conclude the proof, we observe that in Game $\mathcal{G}_{B.4}$, the challenge session key is uniformly random independent of b_{test} and cannot be revealed by \mathcal{A} , hence \mathcal{A} cannot do better than guessing:

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

Case C ($\text{clean}_{\text{ESS}}(\pi^*)$). In this proof case, we are guaranteed that either

- (1) π^* is an initiator session owned by U whose session randomness is unrevealed and whose intended peer V 's semi-static key in question is uncorrupted or
- (2) π^* is a responder session owned by V whose semi-static key in question is uncorrupted and (via $\text{clean}_{\text{peerE}}$ and sound) there exists a unique partnered session π_p^* owned by U whose randomness is unrevealed.

Similarly to the cases before, we leverage this to show that the KEM ciphertext c_2 associated with the test session π^* was generated using an uncorrupted KEM key with good randomness, yielding key secrecy for the corresponding encapsulated key K_2 .

Game C.0. We now condition on $\text{clean}_{\text{ESS}}(\pi^*)$:

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_{C.0}}(\mathcal{A}) = \text{Adv}_{\text{SPQR}}^{\mathcal{G}_5[\text{clean}_{\text{ESS}}(\pi^*)]}(\mathcal{A}).$$

Game C.1 (Guess unique partnered initiator session). We guess the unique existing partner session π_p^* of π^* , if the test session is a responder session; if π^* is an initiator session, we simply ignore the guess. As before we set \mathcal{A} 's output to 0 if the guess was incorrect. We thus reduce the adversary's advantage by a factor of at most the number of sessions n_s :

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

Game C.2 (Guess semi-static key of V). Next, we guess the identifier ssid of the (uncorrupted) semi-static key $\text{sspk}_V^{\text{ssid}}$ of party V , which, depending on the role of π^* , is either the test session's own key (if $\pi^*.\text{role} = \text{responder}$) or that of its intended peer (if $\pi^*.\text{role} = \text{initiator}$). As before, we denote the guessed identifier by ssid^* , and abort with 0 if this guess is incorrect, losing at most a factor of the number of semi-static keys per user n_{ss} :

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_{C.1}}(\mathcal{A}) \leq n_{ss} \cdot \text{Adv}_{\text{SPQR}}^{\mathcal{G}_{C.2}}(\mathcal{A}).$$

Game C.3 (Twisted PRF randomness). Next, we replace the tPRF evaluation of the initiator session π owned by U by the evaluation of a randomly chosen function (here $\pi = \pi^*$, if $\pi^*.\text{role} = \text{initiator}$, and $\pi = \pi_p^*$, if $\pi^*.\text{role} = \text{responder}$).

In particular, we replace the value r_2 in π with an independent random value \tilde{r}_2 (recall the randomness value r is unique per session as of Game \mathcal{G}_1). This change thus assures that the randomness involved in the ensuing encapsulation under the semi-static public key of V is unknown to the adversary.

We bound the advantage difference induced by this step by the twisted $(t, \epsilon_{\text{tPRF}}, 0)$ -twisted pseudorandomness of tPRF via the following reduction \mathcal{B}_8 . The reduction receives (K', z) which is either $(K', \text{tPRF}(K', x))$ if $b = 0$, or $(K', g'(K'))$ if $b = 1$, where K', x are random values and g' is a random function.

The reduction \mathcal{B}_8 then generates all keys and parameters for the key exchange games itself, in particular it sets $tk_U \leftarrow K'$. Instead of evaluating $\text{tPRF}(tk_U, r)$ for $(n, r) \leftarrow \pi.\text{coins}$, \mathcal{B}_8 sets $r_1 \| r_2 \| r_3 \| r_4 \leftarrow z$. Upon a potential $\text{CORRUPTLTKEY}(U)$ query, \mathcal{B}_8 can hand out tk_U as part of U 's secret key (while r remains hidden as $\text{REVEALRANDOM}(\pi)$ is never called). As its bit guess, \mathcal{B}_8 outputs 1 if \mathcal{A} wins the game and 0 otherwise. Depending on which sequence \mathcal{B}_8 is given, it either simulates $\mathcal{G}_{C.2}$ or $\mathcal{G}_{C.3}$, and thus:

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_{C.2}}(\mathcal{A}) \leq \epsilon_{\text{tPRF}} + \text{Adv}_{\text{SPQR}}^{\mathcal{G}_{C.3}}(\mathcal{A}).$$

Game C.4 (Semi-static KEM). In the following, let

$$(c_2, K_2) \leftarrow \text{KEM}_2.\text{Encaps}(\text{sspk}_V^{\text{ssid}^*}; r_2)$$

be the encapsulation computed at session π , where $\pi = \pi^*$, if $\pi^*.\text{role} = \text{initiator}$, and $\pi = \pi_p^*$, if $\pi^*.\text{role} = \text{responder}$.

We can now replace the key K_2 encapsulated in c_2 under the semi-static key of V with identifier ssid^* with a randomly sampled key \tilde{K}_2 in π^* and its partnered session(s), if existent. Furthermore,

we replace K_2 with \tilde{K}_2 in any session of V that has received the same encapsulating ciphertext c_2 .

As in previous cases, we can bound \mathcal{A} 's difference in advantage that was introduced by this change by the advantage of a reduction \mathcal{B}_9 in winning the $(t, \epsilon_{\text{KEM}_2}, n_s)$ -IND-CCA security game for KEM_2 , where the reduction \mathcal{B}_9 obtains its IND-CCA challenge (pk, c^*, K_b^*) and simulates the game \mathcal{A} by generating all parameters of the game itself, except for embedding its received challenge public key pk by setting $\text{sspk}_U^{\text{ssid}^*} = pk$. The predicate $\text{clean}_{\text{ESS}}$ holds, thus we know that \mathcal{A} never asks a $\text{CORRUPTSSKEY}(V, \text{ssid}^*)$ query and the reduction need never output the secret key sk corresponding to pk . Whenever a decapsulation of some ciphertext $c \neq c^*$ using sk is necessary to faithfully simulate the game for \mathcal{A} , \mathcal{B}_9 simply forwards this ciphertext to its decapsulation oracle DECAPS (querying its oracle at most n_s times as claimed). In both π^* and its partnered session(s) π_p^* (if existent), \mathcal{B}_9 embeds K_b^* wherever K_2 would be used and c^* , wherever c_2 would be used. The same replacement is employed in any responder sessions of party V that receive c_2 . At some point, \mathcal{A} will stop with output b' , and the reduction \mathcal{B}_9 returns 0 if $b' = b_{\text{test}}$ and 1 otherwise.

Observe that \mathcal{B}_9 perfectly simulates $\mathcal{G}_{C.3}$ if $b = 0$ in $\mathcal{G}_{\text{KEM}_2}^{\text{indcca}}(\mathcal{B}_9)$ and $\mathcal{G}_{C.4}$ otherwise. Hence, any difference in \mathcal{A} 's advantage in the two games is bounded by the distinguishing advantage of \mathcal{B}_9 against the IND-CCA security of KEM_2 :

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_{C.3}}(\mathcal{A}) \leq \epsilon_{\text{KEM}_2} + \text{Adv}_{\text{SPQR}}^{\mathcal{G}_{C.4}}(\mathcal{A}).$$

Game C.5 (Session key KDF). As the final step in this proof case, we replace in Game $\mathcal{G}_{C.5}$ the session key derived in the test and partnered session as $K \leftarrow \text{KDF}(K_1 \| K_2 \| K_3, \text{sid})$ by a randomly sampled key \tilde{K} . As in the previous cases we can bound the advantage introduced by this change by the advantage of an adversary \mathcal{B}_{10} against $(t, \epsilon_{\text{KDF}}, n_s)$ -PRFSEC property of KDF, this time keyed via K_2 :

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

Finalize. To conclude the proof, we observe that the adversary expects the challenge K_{test} to be the output of the key derivation function KDF applied to the master secret ms and session identifier sid if $b_{\text{test}} = 0$ or a uniformly random string if $b_{\text{test}} = 1$. In all of the above cases, this distinction cannot be made by \mathcal{A} anymore as both keys are now uniformly random. Thus, \mathcal{A} cannot gain any information about the test bit b_{test} and can do no better than to guess, causing us to arrive at the final bound

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

Case D ($\pi^*.\text{type} = \text{full}$ and $\text{clean}_{\text{EE}}(\pi^*)$). In this proof case, we are guaranteed

- (1) π^* is an initiator session owned by U whose session randomness is unrevealed and that has received an ephemeral pre-key that was generated using unrevealed randomness in a session of intended partner V , or
- (2) π^* is a responder session owned by V whose ephemeral pre-key generation was executed with unrevealed randomness and there exists a partnered initiator session π_p^* owned by U whose session randomness is unrevealed.

Similarly to the cases before, we leverage this to show that the KEM ciphertext c_3 associated with the test session π^* was generated using an uncorrupted KEM key with good randomness, yielding key indistinguishability for the corresponding encapsulated key K_3 .

Game D.0. We now condition on the test session running in full mode and $\text{clean}_{\text{EE}}(\pi^*)$ being satisfied:

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_{D.0}}(\mathcal{A}) = \text{Adv}_{\text{SPQR}}^{\mathcal{G}_5[\pi^*.\text{type}=\text{full} \wedge \text{clean}_{\text{EE}}(\pi^*)]}(\mathcal{A}).$$

Game D.1 (Guess unique (contributive) partner session). We first guess the unique existing (contributive) partner session π_p^* of π^* : If π^* is a responder session, π_p^* is its sid-partner, if π^* is an initiator session, π_p^* is its contributively partnered session via cid. (Recall that this unique contributive partner exists since we ruled out collisions in the ephemeral pre-keys, and $\pi^*.\text{type} = \text{full}$, so π^* received such ephemeral pre-key contained in its contributive identifier.) The game sets \mathcal{A} 's output bit to 0 if the guess was incorrect. We thus reduce the adversary's advantage by a factor of at most the number of sessions n_s :

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

Game D.2 (Twisted PRF randomness). Next, we replace the tPRF evaluation of the initiator session π owned by U by the evaluation of a randomly chosen function (here $\pi = \pi^*$, if $\pi^*.\text{role} = \text{initiator}$, and $\pi = \pi_p^*$, if $\pi^*.\text{role} = \text{responder}$).

In particular, we replace the value r_3 in π with an independent random value \tilde{r}_3 (recall the randomness value r is unique per session as of Game \mathcal{G}_1). This change thus assures that the randomness involved in the ensuing encapsulation under the ephemeral pre-key of V is unknown to the adversary.

As in previous cases, we bound the advantage difference induced by this step by the $(t, \epsilon_{\text{tPRF}}, 0)$ -twisted pseudorandomness of tPRF via a reduction \mathcal{B}_{11} . The reduction receives (K', z) which is either $(K', \text{tPRF}(K', x))$ if $b = 0$, or $(K', g'(K'))$ if $b = 1$, where K', x are random values and g' is a random function.

Instead of evaluating $\text{tPRF}(tk_U, r)$ for $(n, r) \leftarrow \pi.\text{coins}$, \mathcal{B}_{11} sets $r_1 \| r_2 \| r_3 \| r_4 \leftarrow z$; as in prior cases, \mathcal{B}_{11} can still answer a potential $\text{CORRUPTLTKEY}(U)$ query, but r remains hidden since $\text{REVEALRANDOM}(\pi)$ is never called. As its bit guess, \mathcal{B}_{11} outputs 1 if \mathcal{A} wins the game and 0 otherwise. Depending on which sequence \mathcal{B}_{11} is given, it either simulates $\mathcal{G}_{D.1}$ or $\mathcal{G}_{D.2}$, and thus:

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

Game D.3 (Ephemeral pre-key KEM). In the following, let

$$(c_3, K_3) \leftarrow \text{KEM}_3.\text{Encaps}(epk_V; r_3)$$

be the encapsulation computed at session π , where $\pi = \pi^*$, if $\pi^*.\text{role} = \text{initiator}$, and $\pi = \pi_p^*$, if $\pi^*.\text{role} = \text{responder}$.

We can now replace the key K_3 encapsulated in c_3 under the ephemeral pre-key of V with a randomly sampled key \tilde{K}_3 in π^* and its partnered session(s), if existent. Furthermore, we replace K_3 with \tilde{K}_3 in any session of V that has received the same encapsulating ciphertext c_3 .

Similar to previous cases, we can bound \mathcal{A} 's difference in advantage that was introduced by this change by the advantage of a

reduction \mathcal{B}_{12} in winning the $(t, \epsilon_{\text{KEM}_3}, 1)$ -IND-CCA security game for KEM_3 , which embeds the received challenge (pk, c^*, K_b^*) by setting $epk_V = pk$, $c_3 = c^*$, and $K_3 = K_b^*$. The predicate clean_{EE} holds, thus we know that \mathcal{A} never asks a $\text{REVEALRANDOM}(\tilde{\pi})$ query, where $\tilde{\pi} = \pi^*$ if the π^* is the responder, and $\tilde{\pi} = \pi_p^*$, if π^* is the initiator; hence \mathcal{B}_{12} need never output the secret key sk corresponding to pk . If session $\tilde{\pi}$ receives a different ciphertext than c_3 , \mathcal{B}_{12} uses (once) its DECAPS oracle to obtain the resulting key. At some point, \mathcal{A} will stop with output b' , and the reduction \mathcal{B}_{12} returns 0 if $b' = b_{\text{test}}$ and 1 otherwise.

Observe that \mathcal{B}_{12} perfectly simulates $\mathcal{G}_{D.2}$ if $b = 0$ in $\mathcal{G}_{\text{KEM}_3}^{\text{indcca}}(\mathcal{B}_{12})$ and $\mathcal{G}_{D.3}$ otherwise. Hence, any difference in \mathcal{A} 's advantage in the two games is bounded by the distinguishing advantage of \mathcal{B}_{12} against the IND-CCA security of KEM_3 :

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_{D.2}}(\mathcal{A}) \leq \epsilon_{\text{KEM}_3} + \text{Adv}_{\text{SPQR}}^{\mathcal{G}_{D.3}}(\mathcal{A}).$$

Game D.4 (Session key KDF). As the final step in this proof case, we replace in Game $\mathcal{G}_{D.4}$ the session key derived in the test and partnered session as $K \leftarrow \text{KDF}(K_1 \| K_2 \| K_3, \text{sid})$ by a randomly sampled key \tilde{K} . As in the previous cases we can bound the advantage introduced by this change by the advantage of an adversary \mathcal{B}_{13} against $(t, \epsilon_{\text{KDF}}, 2)$ -PRFSEC property of KDF (note that here, \tilde{K} is used at most in two sessions, π^* and π_p^*):

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_{D.3}}(\mathcal{A}) \leq \epsilon_{\text{KDF}} + \text{Adv}_{\text{SPQR}}^{\mathcal{G}_{D.4}}(\mathcal{A}).$$

Finalize. To conclude the proof, we observe that the adversary expects the challenge K_{test} to be the output of the key derivation function KDF applied to the master secret ms and session identifier sid if $b_{\text{test}} = 0$ or a uniformly random string if $b_{\text{test}} = 1$. In all of the above cases, this distinction cannot be made by \mathcal{A} anymore as both keys are now uniformly random. Thus, \mathcal{A} cannot gain any information about the test bit b_{test} and can do no better than to guess, causing us to arrive at the final bound

$$\text{Adv}_{\text{SPQR}}^{\mathcal{G}_{D.4}}(\mathcal{A}) \leq 0. \quad \square$$

8.2 Proof of Deniability

THEOREM 8.2 (DENIABILITY OF SPQR). *If DVS is a $(t, \epsilon_{\text{srchid}}, Q_{\text{Ch}})$ -source hiding designated verifier signature, then the SPQR protocol as shown in Figure 8 is $(t', \epsilon', Q'_{\text{Ch}})$ -deniable, where $t' \approx t$, $\epsilon' \leq n_p^2 n_{ss} \cdot \epsilon_{\text{srchid}}$, where n_p is the number of parties and n_{ss} the number of semi-static keys per party, and $Q'_{\text{Ch}} = Q_{\text{Ch}}$.*

PROOF. An adversary \mathcal{B} against the source hiding property of DVS can use a successful adversary \mathcal{A} against deniability of $\text{SPQR}[\text{KEM}_1, \text{KEM}_2, \text{KEM}_3, \text{DVS}, \text{KDF}, \text{tPRF}]$ to succeed in its own game. The challenger starts \mathcal{B} with two DVS key pairs who then simulates the asynchronous DAKE key exchange deniability game $\mathcal{G}_{\text{SPQR}}^{\text{adake-den}}(\mathcal{A})$ for \mathcal{A} as follows. For each of the n_p parties \mathcal{B} generates a long-term key pair and n_{ss} many semi-static keys. It randomly guesses the identifiers of two parties $\text{id}^*, \text{rid}^* \in [n_p]$ and the identifier of a semi-static key $\text{ssid}^* \in [n_{ss}]$ for which the deniability adversary can distinguish between Run and Fake. Let a number $i \in [n_p^2 n_{ss}]$ uniquely denote three independent values $\text{id}, \text{rid}, \text{ssid}$ in a query (e.g., as $(\text{id} - 1) \cdot n_p \cdot n_{ss} + (\text{rid} - 1) \cdot n_{ss} + \text{ssid}$) and let

$i^* \in [n_p^2 n_{ss}]$ denote the specific guess iid^* , rid^* , $ssid^*$ of the reduction. For the party iid^* , \mathcal{B} replaces the DVS sender key pair in the long-term key with its own challenge key pair (pk_S, sk_S) . For the party rid^* , \mathcal{B} replaces the DVS verifier key pair in the semi-static key with $id\ ssid^*$ with its own challenge key pair (pk_D, sk_D) . It starts \mathcal{A} with all key pairs.

\mathcal{B} answers the queries of \mathcal{A} to the CHALL oracle as follows: First, it runs the responder ephemeral key generation. Then, it runs the initiator key agreement until computing the DVS signature (i.e., it computes a nonce and the three KEM ciphertexts (n, c_1, c_2, c_3) and sets the master secret ms to the concatenation of the KEM encapsulations). In the next step, the reduction computes the DVS signature on the session identifier sid . Here \mathcal{B} distinguishes between three cases: The first case is that the query is for $1 \leq i < i^*$. Then the reduction behaves as if $b = 0$, i.e., it executes $DVS.Sign$. The second case is that the query is for $i = i^*$. In this case the reduction forwards the query to its own oracle to obtain a DVS signature or a simulated one depending on the outside challenge bit. The third case is that the query is for $i^* < i \leq n_p^2 n_{ss}$. Then the reduction behaves as if $b = 1$, i.e., it executes $DVS.Sim$. In all cases the reduction then proceeds to compute the session key K from the master secret and the session id. Finally, the reduction returns the transcript and the session key K to \mathcal{A} . Hence, the transcript and session key were computed either as specified by Run or as specified by $Fake$, depending on the query index i and the secret bit of the DVS challenger. Finally, when \mathcal{A} returns its guess bit b' , \mathcal{B} returns b' as its guess.

Observe that \mathcal{B} faithfully simulates the deniability game for \mathcal{A} . Moreover, the runtime of \mathcal{B} is essentially the runtime of \mathcal{A} plus the runtime to generate the keys and answer the oracle queries.

Now let us analyze the winning probability of \mathcal{A} against deniability. For this, we define the hybrids $H_0, \dots, H_{n_p^2 n_{ss}}$ with H_i being the hybrid that answers all challenge queries for indices $1, \dots, i$ by Run and all other challenge queries for indices $i + 1, \dots, n_p^2 n_{ss}$ are answered with $Fake$. The extreme hybrids are $H_{n_p^2 n_{ss}}$ which answers all the challenge queries with Run and H_0 which answers all queries by $Fake$. Observe that H_{i-1} and H_i only differ in an execution of Run or $Fake$ depending on the reduction \mathcal{B} 's challenge oracle. Hence, it is easy to see that the probability of distinguishing between H_{i-1} and H_i is bounded by ϵ_{srchid} .

Now let us analyze \mathcal{A} 's advantage in more detail and we denote by $H_i^{\mathcal{A}}$ the output of the adversary in the i th hybrid. Then it follows:

$$\begin{aligned} & \text{Adv}_{\text{SPQR}}^{\text{adake-den}}(\mathcal{A}) \\ &= \left| \Pr \left[\mathcal{G}_{\text{SPQR}}^{\text{adake-den}}(\mathcal{A}) = 1 \right] - \frac{1}{2} \right| \\ &= \left| \Pr \left[\mathcal{G}_{\text{SPQR}}^{\text{adake-den}}(\mathcal{A}) = 1 | b = 1 \right] \cdot \Pr[b = 1] \right. \\ &\quad \left. + \Pr \left[\mathcal{G}_{\text{SPQR}}^{\text{adake-den}}(\mathcal{A}) = 1 | b = 0 \right] \cdot \Pr[b = 0] - \frac{1}{2} \right| \\ &= \left| \frac{1}{2} \cdot \left(\Pr \left[\mathcal{G}_{\text{SPQR}}^{\text{adake-den}}(\mathcal{A}) = 1 | b = 1 \right] \right. \right. \\ &\quad \left. \left. + \Pr \left[\mathcal{G}_{\text{SPQR}}^{\text{adake-den}}(\mathcal{A}) = 1 | b = 0 \right] \right) - \frac{1}{2} \right| \end{aligned}$$

$$\begin{aligned} &= \left| \frac{1}{2} \cdot \left(\Pr \left[\mathcal{G}_{\text{SPQR}}^{\text{adake-den}}(\mathcal{A}) = 1 | b = 1 \right] \right. \right. \\ &\quad \left. \left. + 1 - \Pr \left[\mathcal{G}_{\text{SPQR}}^{\text{adake-den}}(\mathcal{A}) = 0 | b = 0 \right] \right) - \frac{1}{2} \right| \\ &= \left| \frac{1}{2} \cdot \left(\Pr[1 \leftarrow_{\$} \mathcal{A} | b = 1] - \Pr[1 \leftarrow_{\$} \mathcal{A} | b = 0] \right) \right| \\ &= \left| \frac{1}{2} \cdot \left(\Pr \left[H_0^{\mathcal{A}} = 1 \right] - \Pr \left[H_{n_p^2 n_{ss}}^{\mathcal{A}} = 1 \right] \right) \right| \\ &\leq \frac{1}{2} \cdot \left(\left| \Pr \left[H_0^{\mathcal{A}} = 1 \right] - \Pr \left[H_1^{\mathcal{A}} = 1 \right] \right| + \dots \right. \\ &\quad \left. + \left| \Pr \left[H_{n_p^2 n_{ss}-1}^{\mathcal{A}} = 1 \right] - \Pr \left[H_{n_p^2 n_{ss}}^{\mathcal{A}} = 1 \right] \right| \right) \\ &= \frac{1}{2} \cdot \sum_{i=1}^{n_p^2 n_{ss}} \left| \Pr \left[H_{i-1}^{\mathcal{A}} = 1 \right] - \Pr \left[H_i^{\mathcal{A}} = 1 \right] \right| \\ &= \sum_{i=1}^{n_p^2 n_{ss}} \frac{1}{2} \cdot \left| \Pr[1 \leftarrow_{\$} \mathcal{B} | b_{\text{srchid}} = 1] - \Pr[1 \leftarrow_{\$} \mathcal{B} | b_{\text{srchid}} = 0] \right| \\ &= \sum_{i=1}^{n_p^2 n_{ss}} \left| \Pr[1 \leftarrow_{\$} \mathcal{B} | b_{\text{srchid}} = 1] \cdot \Pr[b_{\text{srchid}} = 1] \right. \\ &\quad \left. + \Pr[0 \leftarrow_{\$} \mathcal{B} | b_{\text{srchid}} = 0] \cdot \Pr[b_{\text{srchid}} = 0] - \frac{1}{2} \right| \\ &= \sum_{i=1}^{n_p^2 n_{ss}} \left| \Pr \left[\mathcal{G}_{\text{DVS}}^{\text{srchid}}(\mathcal{B}) = 1 - \frac{1}{2} \right] \right| \\ &= \sum_{i=1}^{n_p^2 n_{ss}} \text{Adv}_{\text{DVS}}^{\text{srchid}}(\mathcal{B}) = n_p^2 n_{ss} \cdot \text{Adv}_{\text{DVS}}^{\text{srchid}}(\mathcal{B}). \end{aligned}$$

Hence by this analysis, it follows that \mathcal{A} 's probability of winning the deniability game is bounded by $\epsilon' \leq n_p^2 n_{ss} \cdot \epsilon_{\text{srchid}}$. \square

9 DISCUSSION AND LIMITATIONS

Our protocols demonstrate that designated verifier signatures are helpful for constructing practical authenticated key exchange protocols with constraints on the message flow (asynchronicity) and with specialized security properties (deniability).

The key ingredient in our approach for achieving post-quantum asynchronous DAKE is a post-quantum designated verifier signature scheme. While there are several lattice-based DVS schemes in the literature as described in Section 3.1.1, we believe that their security merits further scrutiny before adoption. In the meantime, we propose instantiations via 2-user ring signatures, for which we discussed post-quantum candidates in Section 3.2.3.

We believe SPQR is a good start as a PQ replacement for the Signal X3DH handshake, but in any real-world protocol deployment there are many subtleties, some of which we now highlight.

The way Signal is used in practice has the semi-static keys signed under the long-term key. In SPQR the long-term key is not suitable for this purpose, so an additional long-term signing key might have to be introduced solely for the purposes of signing the other keys; note this could be done without undermining deniability. This characteristic was likewise not considered in the provable security analysis of Signal of [22].

SPQR is solely a replacement for the initial handshake (X3DH). A fully post-quantum Signal would require quantum-resistance in the ratcheting and message encryption; fortunately there are several generic treatments of ratcheting [1, 6, 74].

As Signal does not use certificates or a PKI, long-term public keys must be manually authenticated out-of-band, and that remains the case with SPQR.

Our analysis of SPQR considers disclosure of randomness, but not use of malicious randomness. This has been considered for ratcheting [1], but not yet in the initial handshake. Our security analysis shows that SPQR, as an authenticated key exchange protocol, has offline deniability. One should be careful with deniability as a cryptographic property. How cryptographers understand deniability may be different from how a judge in the legal system would understand it [82]. Additionally there are stronger notions of deniability [34] that SPQR (and the Signal handshake) does not achieve, such as if one party maliciously generates messages or colludes in real-time with the judge. One should also confirm deniability at all levels of the protocol, and that deniability of individual components composes appropriately. Despite all these subtleties, steps toward deniability are helpful, as Unger and Goldberg write [82]: “we should strive to design deniable protocols to avoid unintentionally incriminating users.”

Acknowledgements

We thank anonymous reviewers who pointed out a flaw in our DVS constructions in an earlier version of this paper.

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. F.G. was supported in part by German Research Foundation (DFG) Research Fellowship grant GU 1859/1-1. C.J. was (partially) funded by the Deutsche Forschungsgemeinschaft (DFG) – SFB 1119 – 236615297. D.S. was supported by Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery grant RGPIN-2016-05146.

REFERENCES

- [1] Joël Alwen, Sandro Coretti, and Yevgeniy Dodis. 2019. The Double Ratchet: Security Notions, Proofs, and Modularization for the Signal Protocol. In *EUROCRYPT 2019, Part I (LNCS)*, Yuval Ishai and Vincent Rijmen (Eds.), Vol. 11476. Springer, Heidelberg, 129–158. https://doi.org/10.1007/978-3-030-17653-2_5
- [2] Reza Azarderakhsh, David Jao, and Christopher Leonardi. 2017. Post-Quantum Static-Static Key Agreement Using Multiple Protocol Instances. In *SAC 2017 (LNCS)*, Carlisle Adams and Jan Camenisch (Eds.), Vol. 10719. Springer, Heidelberg, 45–63. https://doi.org/10.1007/978-3-319-72565-9_3
- [3] Mihir Bellare and Phillip Rogaway. 1994. Entity Authentication and Key Distribution. In *CRYPTO'93 (LNCS)*, Douglas R. Stinson (Ed.), Vol. 773. Springer, Heidelberg, 232–249. https://doi.org/10.1007/3-540-48329-2_21
- [4] Mihir Bellare and Phillip Rogaway. 1995. Optimal Asymmetric Encryption. In *EUROCRYPT'94 (LNCS)*, Alfredo De Santis (Ed.), Vol. 950. Springer, Heidelberg, 92–111. <https://doi.org/10.1007/BFb0053428>
- [5] Mihir Bellare and Phillip Rogaway. 1996. The Exact Security of Digital Signatures: How to Sign with RSA and Rabin. In *EUROCRYPT'96 (LNCS)*, Ueli M. Maurer (Ed.), Vol. 1070. Springer, Heidelberg, 399–416. https://doi.org/10.1007/3-540-68339-9_34
- [6] Mihir Bellare, Asha Camper Singh, Joseph Jaeger, Maya Nyayapati, and Igor Stepanovs. 2017. Ratcheted Encryption and Key Exchange: The Security of Messaging. In *CRYPTO 2017, Part III (LNCS)*, Jonathan Katz and Hovav Shacham (Eds.), Vol. 10403. Springer, Heidelberg, 619–650. https://doi.org/10.1007/978-3-319-63697-9_21
- [7] Adam Bender, Jonathan Katz, and Ruggero Morselli. 2009. Ring Signatures: Stronger Definitions, and Constructions without Random Oracles. *Journal of Cryptology* 22, 1 (Jan. 2009), 114–138. <https://doi.org/10.1007/s00145-007-9011-9>
- [8] Ward Beullens, Shuichi Katsumata, and Federico Pintore. 2020. Calamari and Falafel: Logarithmic (Linkable) Ring Signatures from Isogenies and Lattices. In *ASIACRYPT 2020, Part II (LNCS)*, Shiho Moriai and Huaxiong Wang (Eds.), Vol. 12492. Springer, Heidelberg, 464–492. https://doi.org/10.1007/978-3-030-64834-3_16
- [9] Simon Blake-Wilson, Don Johnson, and Alfred Menezes. 1997. Key Agreement Protocols and Their Security Analysis. In *6th IMA International Conference on Cryptography and Coding (LNCS)*, Michael Darnell (Ed.), Vol. 1355. Springer, Heidelberg, 30–45.
- [10] Dan Boneh, Darren Glass, Daniel Krashen, Kristin Lauter, Shahed Sharif, Alice Silverberg, Mehdi Tibouchi, and Mark Zhandry. 2020. Multiparty Non-Interactive Key Exchange and More From Isogenies on Elliptic Curves. *Journal of Mathematical Cryptology* 14, 1 (2020), 5–14.
- [11] Xavier Bonnetain and André Schrottenloher. 2020. Quantum Security Analysis of CSIDH. In *EUROCRYPT 2020, Part II (LNCS)*, Anne Canteaut and Yuval Ishai (Eds.), Vol. 12106. Springer, Heidelberg, 493–522. https://doi.org/10.1007/978-3-030-45724-2_17
- [12] Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, and Damien Stehlé. 2018. CRYSTALS – Kyber: a CCA-secure module-lattice-based KEM. In *2018 IEEE European Symposium on Security and Privacy, EuroS&P 2018*. IEEE, 353–367. <https://cryptojedi.org/papers/#kyber>.
- [13] Colin Boyd, Yvonne Cliff, Juan Manuel Gonzalez Nieto, and Kenneth G. Paterson. 2009. One-round key exchange in the standard model. *IJACT* 1 (2009), 181–199. Issue 3.
- [14] Colin Boyd and Kai Gellert. 2020. A Modern View on Forward Security. *Comput. J.* 64, 4 (08 2020), 639–652. <https://doi.org/10.1093/comjnl/bxaa104> arXiv:<https://academic.oup.com/comjnl/article-pdf/64/4/639/37161647/bxaa104.pdf>
- [15] Jacqueline Brendel, Marc Fischlin, Felix Günther, Christian Janson, and Douglas Stebila. 2020. Towards Post-Quantum Security for Signal’s X3DH Handshake. In *27th Conference on Selected Areas in Cryptography (SAC)*. Springer.
- [16] Jie Cai, Han Jiang, Pingyuan Zhang, Zhihua Zheng, Hao Wang, Guangshi Lü, and Qiuliang Xu. 2019. ID-Based Strong Designated Verifier Signature over \mathcal{R} -SIS Assumption. *Secur. Commun. Networks* 2019 (2019), 9678095:1–9678095:8. <https://doi.org/10.1155/2019/9678095>
- [17] Ran Canetti and Hugo Krawczyk. 2001. Analysis of Key-Exchange Protocols and Their Use for Building Secure Channels. In *EUROCRYPT 2001 (LNCS)*, Birgit Pfitzmann (Ed.), Vol. 2045. Springer, Heidelberg, 453–474. https://doi.org/10.1007/3-540-44987-6_28
- [18] Ran Canetti and Hugo Krawczyk. 2002. Security Analysis of IKE’s Signature-based Key-Exchange Protocol. In *CRYPTO 2002 (LNCS)*, Moti Yung (Ed.), Vol. 2442. Springer, Heidelberg, 143–161. https://doi.org/10.1007/3-540-45708-9_10 <https://eprint.iacr.org/2002/120/>.
- [19] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. 2012. Bonsai Trees, or How to Delegate a Lattice Basis. *Journal of Cryptology* 25, 4 (Oct. 2012), 601–639. <https://doi.org/10.1007/s00145-011-9105-2>
- [20] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. 2018. CSIDH: An Efficient Post-Quantum Commutative Group Action. In *ASIACRYPT 2018, Part III (LNCS)*, Thomas Peyrin and Steven Galbraith (Eds.), Vol. 11274. Springer, Heidelberg, 395–427. https://doi.org/10.1007/978-3-030-03332-3_15
- [21] Wouter Castryck, Jana Sotáková, and Frederik Vercauteren. 2020. Breaking the Decisional Diffie-Hellman Problem for Class Group Actions Using Genus Theory. In *CRYPTO 2020, Part II (LNCS)*, Daniele Micciancio and Thomas Ristenpart (Eds.), Vol. 12171. Springer, Heidelberg, 92–120. https://doi.org/10.1007/978-3-030-56880-1_4
- [22] Katriel Cohn-Gordon, Cas J. F. Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. 2017. A Formal Security Analysis of the Signal Messaging Protocol. In *IEEE European Symposium on Security and Privacy, EuroS&P 2017*. 451–466. <https://doi.org/10.1109/EuroSP.2017.27>
- [23] Katriel Cohn-Gordon, Cas J. F. Cremers, and Luke Garratt. 2016. On Post-compromise Security. In *CSF 2016 Computer Security Foundations Symposium*, Michael Hicks and Boris Köpf (Eds.). IEEE Computer Society Press, 164–178. <https://doi.org/10.1109/CSF.2016.19>
- [24] Cas Cremers and Michele Feltz. 2011. One-round Strongly Secure Key Exchange with Perfect Forward Secrecy and Deniability. *Cryptology ePrint Archive, Report 2011/300*. (2011). <https://eprint.iacr.org/2011/300>.
- [25] Özgür Dagdelen, Marc Fischlin, Tommaso Gagliardoni, Giorgia Azzurra Marson, Arno Mittelbach, and Cristina Onete. 2013. A Cryptographic Analysis of OPACITY - (Extended Abstract). In *ESORICS 2013 (LNCS)*, Jason Crampton, Sushil Jajodia, and Keith Mayes (Eds.), Vol. 8134. Springer, Heidelberg, 345–362. https://doi.org/10.1007/978-3-642-40203-6_20
- [26] Ivan Damgård, Helene Haagh, Rebekah Mercer, Anca Nitulescu, Claudio Orlandi, and Sophia Yakoubov. 2020. Stronger Security and Constructions of Multi-designated Verifier Signatures. In *TCC 2020, Part II (LNCS)*, Rafael Pass and Krzysztof Pietrzak (Eds.), Vol. 12551. Springer, Heidelberg, 229–260. https://doi.org/10.1007/978-3-030-64378-2_9

- [27] Bor de Kock, Kristian Gjøsteen, and Mattia Veroni. 2020. Practical Isogeny-Based Key-exchange with Optimal Tightness. In *27th Conference on Selected Areas in Cryptography (SAC)*. Springer.
- [28] Cyprien Delpèch de Saint Guilhem, Nigel P. Smart, and Bogdan Warinschi. 2017. Generic Forward-Secure Key Agreement Without Signatures. In *ISC 2017 (LNCS)*, Phong Q. Nguyen and Jianying Zhou (Eds.), Vol. 10599. Springer, Heidelberg, 114–133.
- [29] David Derler, Sebastian Ramacher, and Daniel Slamanig. 2018. Post-Quantum Zero-Knowledge Proofs for Accumulators with Applications to Ring Signatures from Symmetric-Key Primitives. In *Post-Quantum Cryptography - 9th International Conference, PQCrypto 2018*, Tanja Lange and Rainer Steinwandt (Eds.), Springer, Heidelberg, 419–440. https://doi.org/10.1007/978-3-319-79063-3_20
- [30] Mario Di Raimondo and Rosario Gennaro. 2005. New Approaches for Deniable Authentication. In *ACM CCS 2005*, Vijayalakshmi Atluri, Catherine Meadows, and Ari Juels (Eds.), ACM Press, 112–121. <https://doi.org/10.1145/1102120.1102137>
- [31] Mario Di Raimondo, Rosario Gennaro, and Hugo Krawczyk. 2006. Deniable authentication and key exchange. In *ACM CCS 2006*, Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati (Eds.), ACM Press, 400–409. <https://doi.org/10.1145/1180405.1180454>
- [32] Samuel Dobson, Steven D. Galbraith, Jason T. LeGrow, Yan Bo Ti, and Lukas Zobernig. 2020. An adaptive attack on 2-SIDH. *Int. J. Comput. Math. Comput. Syst. Theory* 5, 4 (2020), 282–299. <https://doi.org/10.1080/23799927.2020.1822446>
- [33] Samuel Dobson, Trey Li, and Lukas Zobernig. 2019. A Note on a Static SIDH Protocol. Cryptology ePrint Archive, Report 2019/1244. (2019). <https://eprint.iacr.org/2019/1244>.
- [34] Yevgeniy Dodis, Jonathan Katz, Adam Smith, and Shabsi Walfish. 2009. Compositability and On-Line Deniability of Authentication. In *TCC 2009 (LNCS)*, Omer Reingold (Ed.), Vol. 5444. Springer, Heidelberg, 146–162. https://doi.org/10.1007/978-3-642-00457-5_10
- [35] Léo Ducas and Daniele Micciancio. 2014. Improved Short Lattice Signatures in the Standard Model. In *CRYPTO 2014, Part I (LNCS)*, Juan A. Garay and Rosario Gennaro (Eds.), Vol. 8616. Springer, Heidelberg, 335–352. https://doi.org/10.1007/978-3-662-44371-2_19
- [36] Ines Duits. 2019. *The Post-Quantum Signal Protocol: Secure Chat in a Quantum World*. Master's thesis. University of Twente. <http://essay.utwente.nl/77239/>
- [37] Cynthia Dwork, Moni Naor, and Amit Sahai. 1998. Concurrent Zero-Knowledge. In *30th ACM STOC*. ACM Press, 409–418. <https://doi.org/10.1145/276698.276853>
- [38] Muhammed F. Esgin, Ron Steinfeld, Joseph K. Liu, and Dongxi Liu. 2019. Lattice-Based Zero-Knowledge Proofs: New Techniques for Shorter and Faster Constructions and Applications. In *CRYPTO 2019, Part I (LNCS)*, Alexandra Boldyreva and Daniele Micciancio (Eds.), Vol. 11692. Springer, Heidelberg, 115–146. https://doi.org/10.1007/978-3-030-26948-7_5
- [39] Muhammed F. Esgin, Ron Steinfeld, Amin Sakzad, Joseph K. Liu, and Dongxi Liu. 2019. Short Lattice-Based One-out-of-Many Proofs and Applications to Ring Signatures. In *ACNS 19 (LNCS)*, Robert H. Deng, Valérie Gauthier-Umaña, Martín Ochoa, and Moti Yung (Eds.), Vol. 11464. Springer, Heidelberg, 67–88. https://doi.org/10.1007/978-3-030-21568-2_4
- [40] Muhammed F. Esgin, Raymond K. Zhao, Ron Steinfeld, Joseph K. Liu, and Dongxi Liu. 2019. MatRiCT: Efficient, Scalable and Post-Quantum Blockchain Confidential Transactions Protocol. In *ACM CCS 2019*, Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz (Eds.), ACM Press, 567–584. <https://doi.org/10.1145/3319535.3354200>
- [41] Amos Fiat and Adi Shamir. 1987. How to Prove Yourself: Practical Solutions to Identification and Signature Problems. In *CRYPTO '86 (LNCS)*, Andrew M. Odlyzko (Ed.), Vol. 263. Springer, Heidelberg, 186–194. https://doi.org/10.1007/3-540-47721-7_12
- [42] Marc Fischlin and Sogol Mazaheri. 2015. Notions of Deniable Message Authentication. In *Proceedings of the 14th ACM Workshop on Privacy in the Electronic Society, WPES 2015, Denver, Colorado, USA, October 12, 2015*, Indrajit Ray, Nicholas Hopper, and Rob Jansen (Eds.), ACM, 55–64. <https://doi.org/10.1145/2808138.2808143>
- [43] Eduarda S. V. Freire, Dennis Hofheinz, Eike Kiltz, and Kenneth G. Paterson. 2013. Non-Interactive Key Exchange. In *PKC 2013 (LNCS)*, Kaoru Kurosawa and Goichiro Hanaoka (Eds.), Vol. 7778. Springer, Heidelberg, 254–271. https://doi.org/10.1007/978-3-642-36362-7_17
- [44] Atsushi Fujioka, Koutarou Suzuki, Keita Xagawa, and Kazuki Yoneyama. 2012. Strongly Secure Authenticated Key Exchange from Factoring, Codes, and Lattices. In *PKC 2012 (LNCS)*, Marc Fischlin, Johannes Buchmann, and Mark Manulis (Eds.), Vol. 7293. Springer, Heidelberg, 467–484. https://doi.org/10.1007/978-3-642-30057-8_28
- [45] Eiichiro Fujisaki and Tatsuaki Okamoto. 1999. Secure Integration of Asymmetric and Symmetric Encryption Schemes. In *CRYPTO '99 (LNCS)*, Michael J. Wiener (Ed.), Vol. 1666. Springer, Heidelberg, 537–554. https://doi.org/10.1007/3-540-48405-1_34
- [46] Steven D. Galbraith, Christophe Petit, Barak Shani, and Yan Bo Ti. 2016. On the Security of Supersingular Isogeny Cryptosystems. In *ASIACRYPT 2016, Part I (LNCS)*, Jung Hee Cheon and Tsuyoshi Takagi (Eds.), Vol. 10031. Springer, Heidelberg, 63–91. https://doi.org/10.1007/978-3-662-53887-6_3
- [47] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. 2008. Trapdoors for hard lattices and new cryptographic constructions. In *40th ACM STOC*, Richard E. Ladner and Cynthia Dwork (Eds.), ACM Press, 197–206. <https://doi.org/10.1145/1374376.1374407>
- [48] Keitaro Hashimoto, Shuichi Katsumata, Kris Kwiatkowski, and Thomas Prest. 2021. An Efficient and Generic Construction for Signal's Handshake (X3DH): Post-Quantum, State Leakage Secure, and Deniable. In *PKC 2021, Part II (LNCS)*, Juan Garay (Ed.), Vol. 12711. Springer, Heidelberg, 410–440. https://doi.org/10.1007/978-3-030-75248-4_15
- [49] Javier Herranz. 2005. *Some digital signature schemes with collective signers*. Ph.D. Dissertation. Universitat Politècnica de Catalunya, Barcelona. <https://upcommons.upc.edu/bitstream/handle/2117/94334/01Jhs01de01.pdf>
- [50] Andreas Hülsing and Florian Weber. 2021. Epochal Signatures for Deniable Group Chats. In *2021 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, 997–1015. <https://doi.org/10.1109/SP40001.2021.00058>
- [51] Markus Jakobsson, Kazuo Sako, and Russell Impagliazzo. 1996. Designated Verifier Proofs and Their Applications. In *EUROCRYPT '96 (LNCS)*, Ueli M. Maurer (Ed.), Vol. 1070. Springer, Heidelberg, 143–154. https://doi.org/10.1007/3-540-68339-9_13
- [52] David Jao and Luca De Feo. 2011. Towards Quantum-Resistant Cryptosystems from Supersingular Elliptic Curve Isogenies. In *Post-Quantum Cryptography - 4th International Workshop, PQCrypto 2011*, Bo-Yin Yang (Ed.), Springer, Heidelberg, 19–34. https://doi.org/10.1007/978-3-642-25405-5_2
- [53] Jonathan Katz, Vladimir Kolesnikov, and Xiao Wang. 2018. Improved Non-Interactive Zero Knowledge with Applications to Post-Quantum Signatures. In *ACM CCS 2018*, David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang (Eds.), ACM Press, 525–537. <https://doi.org/10.1145/3243734.3243805>
- [54] Tomoki Kawashima, Katsuyuki Takashima, Yusuke Aikawa, and Tsuyoshi Takagi. 2020. An Efficient Authenticated Key Exchange from Random Self-reducibility on CSIDH. In *ICISC 20 (LNCS)*, Deukjo Hong (Ed.), Vol. 12593. Springer, Heidelberg, 58–84. https://doi.org/10.1007/978-3-030-68890-5_4
- [55] Hugo Krawczyk. 1996. SKEME: A versatile secure key exchange for Internet. In *Network and Distributed Systems Security Symposium (NDSS)*. Internet Society and IEEE, 114–127.
- [56] Hugo Krawczyk. 2005. HMQV: A High-Performance Secure Diffie-Hellman Protocol. In *CRYPTO 2005 (LNCS)*, Victor Shoup (Ed.), Vol. 3621. Springer, Heidelberg, 546–566. https://doi.org/10.1007/11535218_33
- [57] Hugo Krawczyk and Tal Rabin. 2000. Chameleon Signatures. In *NDSS 2000*. The Internet Society.
- [58] Kaoru Kurosawa and Jun Furukawa. 2014. 2-Pass Key Exchange Protocols from CPA-Secure KEM. In *CT-RSA 2014 (LNCS)*, Josh Benaloh (Ed.), Vol. 8366. Springer, Heidelberg, 385–401. https://doi.org/10.1007/978-3-319-04852-9_20
- [59] Fabien Laguillaumie and Damien Vergnaud. 2005. Designated Verifier Signatures: Anonymity and Efficient Construction from Any Bilinear Map. In *SCN 04 (LNCS)*, Carlo Blundo and Stelvio Cimato (Eds.), Vol. 3352. Springer, Heidelberg, 105–119. https://doi.org/10.1007/978-3-540-30598-9_8
- [60] Brian A. LaMacchia, Kristin Lauter, and Anton Mityagin. 2007. Stronger Security of Authenticated Key Exchange. In *ProvSec 2007 (LNCS)*, Willy Susilo, Joseph K. Liu, and Yi Mu (Eds.), Vol. 4784. Springer, Heidelberg, 1–16.
- [61] BaoHong Li, YanZhi Liu, and Sai Yang. 2018. Lattice-based universal designated verifier signatures. In *2018 IEEE 15th International Conference on e-Business Engineering (ICEBE)*. IEEE, 329–334. <https://doi.org/10.1109/ICEBE.2018.00062>
- [62] Yong Li, Willy Susilo, Yi Mu, and Dingyi Pei. 2007. Designated Verifier Signature: Definition, Framework and New Constructions. In *Ubiquitous Intelligence and Computing, 4th International Conference, UIC 2007, Hong Kong, China, July 11-13, 2007, Proceedings (Lecture Notes in Computer Science)*, Jadwiga Indulska, Jianhua Ma, Laurence Tianruo Yang, Theo Ungerer, and Jiannong Cao (Eds.), Vol. 4611. Springer, 1191–1200. https://doi.org/10.1007/978-3-540-73549-6_116
- [63] Benoît Libert, San Ling, Khoa Nguyen, and Huaxiong Wang. 2016. Zero-Knowledge Arguments for Lattice-Based Accumulators: Logarithmic-Size Ring Signatures and Group Signatures Without Trapdoors. In *EUROCRYPT 2016, Part II (LNCS)*, Marc Fischlin and Jean-Sébastien Coron (Eds.), Vol. 9666. Springer, Heidelberg, 1–31. https://doi.org/10.1007/978-3-662-49896-5_1
- [64] Xingye Lu, Man Ho Au, and Zhenfei Zhang. 2019. Raptor: A Practical Lattice-Based (Linkable) Ring Signature. In *ACNS 19 (LNCS)*, Robert H. Deng, Valérie Gauthier-Umaña, Martín Ochoa, and Moti Yung (Eds.), Vol. 11464. Springer, Heidelberg, 110–130. https://doi.org/10.1007/978-3-030-21568-2_6
- [65] Vadim Lyubashevsky. 2009. Fiat-Shamir with Aborts: Applications to Lattice and Factoring-Based Signatures. In *ASIACRYPT 2009 (LNCS)*, Mitsuru Matsui (Ed.), Vol. 5912. Springer, Heidelberg, 598–616. https://doi.org/10.1007/978-3-642-10366-7_35
- [66] Vadim Lyubashevsky. 2012. Lattice Signatures without Trapdoors. In *EUROCRYPT 2012 (LNCS)*, David Pointcheval and Thomas Johansson (Eds.), Vol. 7237. Springer, Heidelberg, 738–755. https://doi.org/10.1007/978-3-642-29011-4_43
- [67] Vadim Lyubashevsky and Gregory Neven. 2017. One-Shot Verifiable Encryption from Lattices. In *EUROCRYPT 2017, Part I (LNCS)*, Jean-Sébastien Coron and Jesper Buus Nielsen (Eds.), Vol. 10210. Springer, Heidelberg, 293–323. https://doi.org/10.1007/978-3-030-56544-4_16

- //doi.org/10.1007/978-3-319-56620-7_11
- [68] Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. 2021. SMILE: Set Membership from Ideal Lattices with Applications to Ring Signatures and Confidential Transactions. *IACR Cryptol. ePrint Arch.* 2021 (2021), 564. <https://eprint.iacr.org/2021/564>
- [69] Moxie Marlinspike and Trevor Perrin. November 2016. The double ratchet algorithm. (November 2016). <https://www.signal.org/docs/specifications/doublerratchet/>
- [70] Moxie Marlinspike and Trevor Perrin. November 2016. The X3DH key agreement protocol. (November 2016). <https://signal.org/docs/specifications/x3dh/>
- [71] Geontae Noh and Ik Rae Jeong. 2017. Strong designated verifier signature scheme from lattices in the standard model. *Security Comm. Networks* 9 (Feb. 2017), 6202–6214. <https://doi.org/10.1002/sec.1766>
- [72] Rafael Pass. 2003. On Deniability in the Common Reference String and Random Oracle Model. In *CRYPTO 2003 (LNCS)*, Dan Boneh (Ed.), Vol. 2729. Springer, Heidelberg, 316–337. https://doi.org/10.1007/978-3-540-45146-4_19
- [73] Chris Peikert. 2020. He Gives C-Sieves on the CSIDH. In *EUROCRYPT 2020, Part II (LNCS)*, Anne Canteaut and Yuval Ishai (Eds.), Vol. 12106. Springer, Heidelberg, 463–492. https://doi.org/10.1007/978-3-030-45724-2_16
- [74] Bertram Poettering and Paul Rösler. 2018. Towards Bidirectional Ratcheted Key Exchange. In *CRYPTO 2018, Part I (LNCS)*, Hovav Shacham and Alexandra Boldyreva (Eds.), Vol. 10991. Springer, Heidelberg, 3–32. https://doi.org/10.1007/978-3-319-96884-1_1
- [75] Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. 2020. *FALCON*. Technical Report. National Institute of Standards and Technology. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions>.
- [76] Oded Regev. 2005. On lattices, learning with errors, random linear codes, and cryptography. In *37th ACM STOC*, Harold N. Gabow and Ronald Fagin (Eds.). ACM Press, 84–93. <https://doi.org/10.1145/1060590.1060603>
- [77] Ronald L. Rivest, Adi Shamir, and Yael Tauman. 2001. How to Leak a Secret. In *ASIACRYPT 2001 (LNCS)*, Colin Boyd (Ed.), Vol. 2248. Springer, Heidelberg, 552–565. https://doi.org/10.1007/3-540-45682-1_32
- [78] Shahrokh Saeednia, Steve Kremer, and Olivier Markowitch. 2004. An Efficient Strong Designated Verifier Signature Scheme. In *ICISC 03 (LNCS)*, Jong In Lim and Dong Hoon Lee (Eds.), Vol. 2971. Springer, Heidelberg, 40–54.
- [79] Peter Schwabe, Douglas Stebila, and Thom Wiggers. 2020. Post-Quantum TLS Without Handshake Signatures. In *ACM CCS 2020*, Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna (Eds.). ACM Press, 1461–1480. <https://doi.org/10.1145/3372297.3423350>
- [80] Ron Steinfeld, Laurence Bull, Huaxiong Wang, and Josef Pieprzyk. 2003. Universal Designated-Verifier Signatures. In *ASIACRYPT 2003 (LNCS)*, Chi-Sung Laih (Ed.), Vol. 2894. Springer, Heidelberg, 523–542. https://doi.org/10.1007/978-3-540-40061-5_33
- [81] Xi Sun, Haibo Tian, and Yumin Wang. 2012. Toward Quantum-Resistant Strong Designated Verifier Signature from Isogenies. In *4th International Conference on Intelligent Networking and Collaborative Systems*. IEEE, 292–296. <https://doi.org/10.1109/iNCoS.2012.70>
- [82] Nik Unger and Ian Goldberg. 2015. Deniable Key Exchanges for Secure Messaging. In *ACM CCS 2015*, Indrajit Ray, Ninghui Li, and Christopher Kruegel (Eds.). ACM Press, 1211–1223. <https://doi.org/10.1145/2810103.2813616>
- [83] Nik Unger and Ian Goldberg. 2018. Improved Strongly Deniable Authenticated Key Exchanges for Secure Messaging. *PoPETs 2018*, 1 (Jan. 2018), 21–66. <https://doi.org/10.1515/popets-2018-0003>
- [84] Nihal Vatandas, Rosario Gennaro, Bertrand Ithurburn, and Hugo Krawczyk. 2020. On the Cryptographic Deniability of the Signal Protocol. In *ACNS 20, Part II (LNCS)*, Mauro Conti, Jianying Zhou, Emiliano Casalicchio, and Angelo Spognardi (Eds.), Vol. 12147. Springer, Heidelberg, 188–209. https://doi.org/10.1007/978-3-030-57878-7_10
- [85] Fenghe Wang, Yupu Hu, and Baocang Wang. 2012. Lattice-based strong designate verifier signature and its applications. *Malaysian Journal of Computer Science* 25 (2012), 11–22. Issue 1.
- [86] Fenghe Wang, Yupu Hu, and Baocang Wang. 2014. Identity-based strong designate verifier signature over lattices. *The Journal of China Universities of Post and Telecommunications* 21 (2014), 52–60. Issue 6. [https://doi.org/10.1016/S1005-8885\(14\)60345-9](https://doi.org/10.1016/S1005-8885(14)60345-9)
- [87] Bo Yang, Yong Yu, and Ying Sun. 2013. A novel construction of SDVS with secure disavowability. *Clust. Comput.* 16, 4 (2013), 807–815. <https://doi.org/10.1007/s10586-013-0254-y>
- [88] Greg Zaverucha, Melissa Chase, David Derler, Steven Goldfeder, Claudio Orlandi, Sebastian Ramacher, Christian Rechberger, Daniel Slamanig, Jonathan Katz, Xiao Wang, Vladimir Kolesnikov, and Daniel Kales. 2020. *Picnic*. Technical Report. National Institute of Standards and Technology. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions>.
- [89] Yongqiang Zhang, Qiang Liu, Chengpei Tang, and Haibo Tian. 2015. A lattice-based designated verifier signature for cloud computing. *International Journal*

of High Performance Computing and Networking 8 (June 2015), 135–143. Issue 2. <https://doi.org/10.1504/IJHPCN.2015.070013>

A RELATED WORK ON DENIABILITY

Deniability allows a party to deny having interacted with a peer. In particular, the peer cannot convince a judge of the first party having interacted with itself. *Online deniability* is concerned with the scenario of the judge interacting with the peer during the protocol execution with the first party. This notion is not achievable in the asynchronous setting. [82] Hence, we address *offline deniability*, where the peer presents data to the judge *after* the protocol execution has taken place.

Prior work defined several notions of offline deniability for authenticated key exchange [24, 25, 31, 82, 83]. Based on the work of Dwork, Naor, and Sahai [37] on deniable authentication, Di Raimondo, Gennaro, and Krawczyk defined *concurrently deniable (or fully deniable) authenticated key exchange* using the simulation paradigm in [31]. Given the list of all public keys and some auxiliary information (e.g., some legal transcripts), the adversary may freely interact with honest parties as either initiator or as responder, interleaving between executions at will. The view of the adversary then consists of the transcripts, session keys, and random coins of the protocol executions it took part in. The session key is included in the view as it may be used as part of another protocol for which deniability is desirable. This view needs to be indistinguishable from the output of a simulator running on the same inputs as the adversary.

Di Raimondo, Gennaro, and Krawczyk [31] also proposed a weaker notion called *partial deniability*, which formalizes the intuition that it is indistinguishable whether an (honest) user interacted with party A or party B. Based on the definition of partial deniability, Cremers and Feltz [24] proposed *peer deniability* and *peer-and-time deniability*. For either notion the simulator does not have to output the session key and gets access to the secret key of corrupted parties. Peer-deniability intuitively allows a user to deny its communication peer, while peer-and-time deniability allows a party to deny that it was alive during a certain time frame.

Dagdelen, Fischlin, Gagliardini, Marson, Mittelbach, and Onete [25] proposed a game-based definition called *outsider deniability*. Here, the adversary has access to Init, Exec, Send, Reveal, Corrupt, and Register oracles (identical to the key secrecy game) and a modified challenge oracle. Depending on the secret bit, the challenge oracle returns either a real transcript and session key or a transcript and session key simulated based on public data. Intuitively, this allows parties to deny having engaged in a protocol run against an eavesdropper that frames a party.

In [82, 83], Unger and Goldberg have given deniability definitions in the UC model. For this they define an ideal functionality called post-specified peer key exchange with incrementing abort that unifies the model of contributiveness, deniability with abort, and their model [82] of post-specified peers.

In [84], Vatandas, Gennaro, Ithurburn, and Krawczyk provided an analysis that Signal’s X3DH is deniable (wrt. full deniability of [31]) under a general extractability assumption. The authors emphasize the observation of Pass in [72] that a simulator for deniability must

be a real algorithm (unlike a Zero-Knowledge simulator, which can be thought of as a thought experiment allowing, e.g., re-winding).

Recently, Hülsing and Weber have defined deniability for group chats (and not just key exchange) in [50]. They formalize a stronger notion than ours that allows an unbounded judge to choose all long-term key pairs and learn all short-term keys, and the simulator does not get access to any secret key. Furthermore, the judge chooses the instructions (i.e., messages and group actions) to be executed. However, the group setting requires a restriction: Informally, they need one message that authentically reaches all other group members.

We propose a game-based definition, Definition 7.2, where the adversary interacts with a real-or-random challenge oracle. Intuitively, the “real” part relates to the view of the adversary and the “random” part to the simulated view. However this simulated view cannot make use of features like re-winding and is a plain probabilistic classical algorithm. At the same time, the distinguisher (or judge) of simulation-based definitions relates to the adversary in the game-based definition.

We further take into account the informal requirement on deniability for asynchronous DAKE in [70, §4.4]: *Neither party has a proof of the fact that both parties communicated and a purported transcript of the execution can be produced by any party that has access to one of the party’s secret keys.* This informal description implies a relaxation compared to previous definitions: The simulated transcript may make use of the secret key of either party.

Please observe that we only consider asynchronous DAKE, which consist of only one message with a specified peer. Hence, we do not need to take any special precautions to achieve deniability for concurrent executions. Furthermore, we give the adversary (i.e., the distinguisher) access to all secret keys. This models the scenario where a party is framed in court and the judge (in a legal sense) learns the secret keys of all involved parties through a subpoena. Hence, our distinguisher is significantly stronger than previous distinguishers. As the distinguisher has access to all secret keys, the challenge oracle does not return the random coins used. Otherwise, the distinguisher could compute both the real and simulated execution of the protocol and check which result is identical to the return value of the oracle. One could prevent this by requiring identical outputs instead of indistinguishably distributed outputs. We deem this impractical, though.

B DESIGNATED VERIFIER SIGNATURES FROM CHAMELEON HASH FUNCTIONS

As mentioned in Section 3.1, in an earlier version we attempted to build post-quantum designated verifier signatures in a direct manner, following full-domain-hash and Fiat-Shamir-type approaches and involving chameleon hash functions. For transparency and educational purposes, we comment on these attempts a bit further in the following.

We first recall the definition of a chameleon hash function, more concretely, the formalization of Cash, Hofheinz, Kiltz, and Peikert [19], where the trapdoor enables preimage sampling (unlike [57], where the trapdoor enables collision sampling).

Definition B.1. A *chameleon hash function* (CHF) is a tuple of algorithms $\text{CHF} = (\text{KGen}, \text{Hash}, \text{Inv})$ with public key space \mathcal{P} ,

message space \mathcal{M} , digest space \mathcal{D} , randomness space \mathcal{R} , and a (not necessarily uniform) distribution $\mathcal{R}_{\text{dist}}$ over \mathcal{R} :

- $\text{KGen}() \xrightarrow{s} (pk, sk)$: A probabilistic key generation algorithm.
- $\text{Hash}(pk, m; r) \rightarrow h$: A hashing algorithm that takes as input a public key pk and a message $m \in \mathcal{M}$ along with randomness $r \in \mathcal{R}$, and outputs a digest $h \in \mathcal{D}$.
- $\text{Inv}(sk, h, m) \xrightarrow{s} r$: A probabilistic hash inversion algorithm that takes as input a secret key sk , digest $h \in \mathcal{D}$, and message $m \in \mathcal{M}$, and outputs randomness $r \in \mathcal{R}$.

Chameleon hash functions provide the same security properties as standard hash functions, with the addition of the *chameleon properties* introduced by the trapdoor:

Definition B.2. A CHF is (t, ϵ) -secure if it satisfies:

- Uniformity** For $(pk, sk) \leftarrow \text{KGen}()$, $m \in \mathcal{M}$, and $r \leftarrow \mathcal{R}_{\text{dist}}$, we have that $(pk, \text{Hash}(pk, m; r))$ is ϵ -close to uniform over $\mathcal{P} \times \mathcal{D}$.
- Chameleon** For $(pk, sk) \leftarrow \text{KGen}()$, $h \in \mathcal{D}$, $m \in \mathcal{M}$, we have that $h = \text{Hash}(pk, m; \text{Inv}(sk, h, m))$.
- Collision resistance** Given $pk \in \mathcal{P}$, no time- t -bounded adversary can find distinct $(m, r), (m', r')$ with $\text{Hash}(pk, m; r) = \text{Hash}(pk, m'; r')$ with probability greater than ϵ .
- Chameleon indistinguishability** For all $(pk, sk) \leftarrow \text{KGen}()$, $m \in \mathcal{M}$, and $h \in \mathcal{D}$, $\text{Inv}(sk, h, m)$ is ϵ -close to the distribution of $r \leftarrow \mathcal{R}_{\text{dist}}$ conditioned on $\text{Hash}(pk, m; r) = h$.

In the initial version of this paper, we used such chameleon hash functions within full-domain-hash and Fiat-Shamir-style signatures to build designated verifier signatures. This required the modeling of the chameleon hash function as a random oracle. However, as was kindly pointed out to us, chameleon hash functions have strong algebraic properties, which makes them inappropriate for instantiating random oracles.

Since random oracles are crucial in the proofs of the DVS constructions in question, we tried to salvage the situation by additionally introducing a standard hash function that could either be applied before or after chameleon hashing. This newly-introduced hash function can be modeled as a random oracle RO, while the chameleon hash function CHF is assumed to only provide the above-defined security properties. Unfortunately, neither order of application (*CHF-then-RO* or *RO-then-CHF*) lead to a secure construction. We will briefly outline the reasons next, using Fiat-Shamir-style signatures as illustrative example; the obstacles are analogous for our attempted full-domain-hash based GPV [47] construction.

Recall that Fiat-Shamir signatures are built on top of a passively secure canonical identification scheme $\text{CID} = (\text{KGen}, \text{P} = (\text{P}_1, \text{P}_2), \text{V} = (\text{V}_1, \text{V}_2))$. Signing in our failed DVS construction then was defined as in Figure 12, with the chameleon hash function replacing the regular hash function in the Fiat-Shamir transform:

B.1 CHF-then-RO

As CHF cannot be modeled as a random oracle, one loses the required programmability for the security proofs of the FSDVS construction in Figure 12. One could try to salvage this construction by first applying the chameleon hash function, and then a hash function modeled as a random oracle within the signature. Figure 13 shows this approach, with the changes marked in boxed code.

```

FSDVS.Sign( $sk_S, pk_D, m$ ):
1 ( $com, st$ )  $\leftarrow$  CID.P1( $sk_S$ )
2  $r \leftarrow$  CHF. $\mathcal{R}_{dist}$ 
3  $ch \leftarrow$  CHF.Hash( $pk_D, com || m; r$ )
4  $rsp \leftarrow$  CID.P2( $ch, st$ )
5  $\sigma \leftarrow (r, (com, ch, rsp))$ 
6 return  $\sigma$ 

```

Figure 12: Fiat-Shamir-style signature using a chameleon hash function.

```

FSDVS'.Sign( $sk_S, pk_D, m$ ):
1 ( $com, st$ )  $\leftarrow$  CID.P1( $sk_S$ )
2  $r \leftarrow$  CHF. $\mathcal{R}_{dist}$ 
3  $h \leftarrow$  CHF.Hash( $pk_D, m; r$ )
4  $ch \leftarrow$  RO( $com || h$ )
5  $rsp \leftarrow$  CID.P2( $ch, st$ )
6  $\sigma \leftarrow (r, (com, ch, rsp))$ 
7 return  $\sigma$ 

```

Figure 13: Fiat-Shamir-style signature with RO applied to a CHF digest; changes from Figure 12 in boxes.

Unfortunately, this construction cannot provide source-hiding anymore. Recall that source-hiding requires the existence of an efficient simulator Sim which, given the secret key of the designated verifier and a signer’s public key, can output signatures that are indistinguishable from signatures by the sender for the designated verifier.

The original construction given in Figure 12 achieved source hiding with a simulator that first retrieved an accepting conversation (com, ch, rsp) , and then utilized the trapdoor sk_D of the chameleon hash function to find randomness r such that $ch = \text{Hash}(pk_D, com || m; r)$.

However, such a simulator does not exist anymore for signatures generated as in Figure 13. Starting from an accepting conversation (com, ch, rsp) , the simulator cannot utilize the trapdoor in the chameleon hash function, as it does not know the value h that led to ch , i.e., finding this value h would imply an efficient algorithm that can find preimages of RO. The same obstacle arises in full-domain-hash style signatures based on GPV.

B.2 RO-then-CHF

The above issue clearly stemmed from the fact that the random oracle hid the relevant information for the simulator. But what if we first used the random oracle to generate a uniformly random value and then applied the chameleon hash function to that? Figure 14 depicts this approach, again marking changes in boxed code:

Here, unforgeability fails to hold. Recall that the Fiat-Shamir transform crucially relies on the programming of the random oracle that outputs the challenge ch for the prover P_2 to achieve unforgeability of the resulting signature scheme. This is not possible here, as the object that outputs the challenge for the prover is the chameleon hash function and not the random oracle.

For our GPV construction this approach also fails. In the original proof, unforgeability hinges on a reduction from a successful forger to the collision-resistance of the employed preimage-sampleable

```

FSDVS''.Sign( $sk_S, pk_D, m$ ):
1 ( $com, st$ )  $\leftarrow$  CID.P1( $sk_S$ )
2  $h \leftarrow$  RO( $com || m$ )
3  $r \leftarrow$  CHF. $\mathcal{R}_{dist}$ 
4  $ch \leftarrow$  CHF.Hash( $pk_D, h; r$ )
5  $rsp \leftarrow$  CID.P2( $ch, st$ )
6  $\sigma \leftarrow (r, (com, ch, rsp))$ 
7 return  $\sigma$ 

```

Figure 14: Fiat-Shamir-style signatures with CHF applied to a RO digest; changes from Figure 12 in boxes.

function. However, since chameleon hashing is a probabilistic process (as opposed to regular hashing) and the forger controls the randomness leading to the eventual signature forgery, no efficient collision-finder can be constructed.

C SUMMARY OF MAJOR CHANGES

- **version 1.0 and 1.0.1 - June 2021:** Initial release and minor editorial changes
- **version 1.1 - August 2021:**
 - removed direct DVS constructions using full-domain-hash GPV and Fiat-Shamir-type approaches leveraging chameleon hash functions due to flawed security proofs; added a discussion of the obstacles in Appendix B
 - added RingDVS construction in Section 3.2