

Revisiting PQ WireGuard: A Comprehensive Security Analysis With a New Design Using Reinforced KEMs Full Version

Keitaro Hashimoto*, Shuichi Katsumata*[†], Guilhem Niot[†], Thom Wiggers[†]

*National Institute of Advanced Industrial Science and Technology (AIST), Tokyo, Japan

[†]PQShield, <https://pqshield.com>

Abstract—WireGuard is a VPN based on the Noise protocol, known for its high performance, small code base, and unique security features. Recently, Hülsing et al. (IEEE S&P’21) presented post-quantum (PQ) WireGuard, replacing the Diffie-Hellman (DH) key exchange underlying the Noise protocol with key-encapsulation mechanisms (KEMs). Since WireGuard requires the handshake message to fit in one UDP packet of size roughly 1200 B, they combined Classic McEliece and a modified variant of Saber. However, as Classic McEliece public keys are notoriously large, this comes at the cost of severely increasing the server’s memory requirement. This hinders deployment, especially in environments with constraints on memory (allocation), such as a kernel-level implementations.

In this work, we revisit PQ WireGuard and improve it on three fronts: design, (computational) security, and efficiency. As KEMs are *semantically* but not *syntactically* the same as DH key exchange, there are many (in hindsight) ad-hoc design choices being made, further amplified by the recent finding on the binding issues with PQ KEMs (Cremers et al., CCS’24). We redesign PQ WireGuard addressing these issues, and prove it secure in a new computational model by fixing and capturing new security features that were not modeled by Hülsing et al. We further propose *reinforced KEM (RKEM)* as a natural building block for key exchange protocols, enabling a PQ WireGuard construction where the server no longer needs to store Classical McEliece keys, reducing public key memory by 190 to 390×. In essence, we construct a **RKEM** named **Rebar** to compress two ML-KEM-like ciphertexts which may be of an independent interest.

1. Introduction

WireGuard [3] is a VPN protocol available for all major operating systems. It offers an efficient one-roundtrip, mutually authenticated handshake protocol following the “IK” pattern of the Noise protocol [4], based on the Diffie-Hellman (DH) key exchange. Compared to other VPNs like OpenVPN [5] or IPsec [6], it is well-known for its high performance and small code base. It implements the

entire protocol in the Linux kernel and is “cryptographically opinionated” [3] by only allowing a small set of modern cryptographic primitives.

In addition, WireGuard offers several unique security properties, those not typically considered by traditional handshake protocols and other VPNs, such as session-key uniqueness, Denial of Service (DoS) mitigation, identity hiding, and so on. These security properties, informally laid out in the original work by Donenfeld [3], has been formally (and partially) analyzed in the computational model by Dowling and Paterson [7] and Lipp et al. [8], and in the symbolic model by Donenfeld and Milner [9].

With the continuing demand for transitioning to post-quantum cryptography, Hülsing et al. [1] recently proposed a *post-quantum* variant of WireGuard (PQ WireGuard). The main technical contribution of PQ WireGuard is in the successful replacement of the DH key exchange with key-encapsulation mechanisms (KEMs). At a theoretical level, it is well-understood that one can emulate (almost) any DH key exchange by KEMs, however, this is no longer the case when considering real-world constraints.

WireGuard’s handshake is designed to have as little state as possible. Additionally, it operates over UDP for performance, which does not transparently handle too-large packets. This means it requires its handshake messages to fit into a single UDP packet, which max out at 1232 B for IPv6. This makes the creation of a post-quantum version of WireGuard challenging, because the main post-quantum KEM, ML-KEM [10], is too large to replace DH within this limit. Hülsing et al. [1] overcome this by combining two different KEMs: Classic McEliece [11] and a modified, IND-CPA version of Saber [12]. Capitalizing on Classic McEliece’s small ciphertext sizes and the fact that its large (long-term) public keys do not need to be transmitted during the protocol, Hülsing et al. fit the handshake in a single UDP packet. Lastly, they analyzed their PQ WireGuard’s security both in the computational and symbolic models previously used by WireGuard [7], [9].

1.1. Our Contribution

In this paper, we revisit the PQ WireGuard protocol by Hülising et al., improving it on three fronts: design, computational security, and efficiency. Concretely, we provide a leaner generic construction of PQ WireGuard based on a novel primitive, coined as *reinforcing* KEMs (RKEM); a special type of KEMs explained shortly. Our PQ WireGuard, instantiated with a novel lattice-based RKEM, replaces the client-side long-term public key that was a Classic McEliece key with a much shorter ML-KEM-like key. This reduces the server’s public key memory requirements by 190 to 390×, making PQ WireGuard easier to deploy and operate, especially in environments with constraints on memory (allocation), such as WireGuard’s kernel-level implementations. We show that PQ WireGuard and RKEM are practical by implementing them and measuring their performance; our implementation and measurement results are [available online](#).

We explain our three contributions in more detail below.

Design. At a high level, WireGuard is a simple DH key exchange protocol; however, under the hood, it is quite involved. Technically, it follows the “IK” pattern of the Noise protocol [4], where the exchanged session key is systematically derived by a key derivation function (KDF), taking as input the cryptographic materials generated within the protocol. In addition, a hash chain is locally maintained, authenticating the views of the client and server. WireGuard further supports an optional pre-shared (symmetric) key (PSK) to be used, enhancing the security in case all other secrets are compromised. An interested reader may refer to the detailed description in [Fig. 1](#) for a better intuition.

While Hülising et al. [1] replaced the DH messages by KEM public keys and ciphertexts, (in hindsight) these were handled in an ad-hoc manner. The issue is that while a KEM may *semantically* replace DH key exchange, it is *syntactically* quite different, raising the question of which values should be absorbed, and in what order, into the KDF and hash chain. This issue is further amplified by how Hülising et al. used PSKs. They now had to use it much earlier in the protocol since, unlike a DH key exchange, a KEM lacks client-side authentication in the first round; a property used by WireGuard to attain DoS protection. This lead Hülising et al. to absorb the PSK *twice* into the KDF (cf. [Fig. 1](#)). While, subsequently, there was work on *post-quantum* (PQ) Noise [13], it did not incorporate the later-discovered concerns surrounding KEM binding [2] nor discuss how to use PSKs. Consequently, while the “IK” pattern of Noise readily yields WireGuard, no such analog exists in the post-quantum setting.

In our work, we revisit Hülising et al.’s PQ WireGuard and systemically handle the KEM public keys and ciphertexts, leading to a more streamlined design. As an independent interest, we also show how to translate our findings to PQ Noise, naturally bypassing the KEM binding issue [2].

Computational security. As mentioned above, PQ WireGuard aims to be as secure as WireGuard, offering several unique security notions. While Hülising et al. analyzed most

of these notions in their symbolic model, only a basic security notion (i.e., key indistinguishability) was analyzed in their computational model. As a correct computational proof gives stronger security guarantees by making fewer idealizing assumptions, – a point also highlighted by Hülising et al. – our aim is to make the computational model as (or even more) expressive as the symbolic model.

In our work, we revisit the computational model by Hülising et al. [1], which they based on those by Dowling and Paterson [7], and formalize the remaining security notions that were covered by the symbolic model. For some notions (e.g., identity hiding), our computational model provides better security guarantee since these were only symbolically analyzed on a sub-protocol of PQ WireGuard [1].

In fact, we notice a subtle bug in the previous computational models [1], [7] due to the handling of the PSKs in the *post-specified peer model*, rendering their security proof incomplete. Unlike standard mutually-authenticated key exchange, in PQ WireGuard, the server identifies and authenticates the identity of the client *during* the protocol. For completeness, this issue is detailed in [Sec. A.1](#). In our work, we fix this by explicitly defining the server as a two-step algorithm, consisting of an *identification phase* followed by a *key-derivation phase*. The former is used by the server to identify and authenticate the client, and the latter is used to complete the key exchange with the identified client. This not only fixes the subtle issue in the formalization but also turns out to be essential for defining DoS security.

Lastly, we explicitly define *entity authentication* (e.g., security against *key-compromise impersonation* attacks) as a stand-alone security requirement following recent works [14], [15], as opposed to implicitly handling them through key indistinguishability. This allows capturing a broader authentication guarantee; even those which may not necessarily contradict key indistinguishability such as unknown key share (UKS) attacks. Indeed, due to the aforementioned KEM binding issue, we observe an explicit attack on Hülising et al.’s PQ WireGuard when instantiating it with general KEMs (see [Sec. B](#) for the details).¹ This is akin to the KEM re-encapsulation attack on Signal’s PQXDH protocol [16].

Efficiency. In PQ WireGuard, the server uses the client’s long-term and ephemeral KEM keys pk_C and pk_e , respectively, to create the second handshake message (i.e., ciphertexts ct_C and ct_e). Since $|ct_C| + |ct_e|$ must fit in a single UDP packet to comply with WireGuard’s design choice, maxing out at 1232 B for IPv6, Hülising et al. could not rely on ML-KEM which would result in size 2×768 B. Instead, they relied on Classic McEliece [11] for a very small ct_C , and a modified, IND-CPA-secure version of Saber [12] for a modest sized ct_e . As Classic McEliece has very large public key size, the downside is its high server memory footprint, which may be an issue when operating in memory constrained environments.

1. We note that this does not contradict their symbolic proof, because they analyze a slightly different PQ WireGuard computing the key chain value C_5 from ct_C instead of ct_e . By their modeling of KEM ciphertexts this binds public key pk_C [2, Sec. 5.1], thwarting the UKS attack.

As our final contribution, we introduce *reinforcing* KEMs (RKEM). This is a concise abstraction of how the server uses the client public keys. Informally, this is a special KEM whose security of the long-term public key pk_C can be *reinforced* by an ephemeral public key pk_e . While running two independent KEMs in parallel is one naive construction of an RKEM, as used in PQ WireGuard, it is clearly wasteful since the ephemeral pk_e does not need to be as secure as pk_C . For instance, RKEM does not need to be secure when pk_C is maliciously generated since this does not relate to any PQ WireGuard security. However, the naive RKEM construction would be secure under such unnecessary scenarios.

In our work, we construct a novel ML-KEM inspired RKEM called Rebar whose ciphertext is roughly 40% smaller than running two ML-KEM, allowing to fit comfortably in a single UDP packet. As the long-term public key is reduced to roughly 1.4kB, this leads to a 190 to 390 $\times$ saving in server’s public key memory requirement. Technically, we give a generic construction of a CCA secure RKEM from a one-way secure reinforced *double-message* PKE (RPKE₂) by applying the Fujisaki-Okamoto (FO) transform [17] (cf. Fig. 7) in the *quantum* random oracle model (QROM). We then construct a lattice-based RPKE₂, resembling the multi-recipient adaptation [18] of ML-KEM [19], with a novel twist. We believe our RKEM to be useful in other handshake protocols that have mutual KEM-based authentication, e.g., FSXY [20], [21], SKEME/IKEv2 [6], [22], and KEMTLS-PDK [23].

1.2. Related Work

We have already discussed prior work on (PQ) WireGuard [1], [3], [7]–[9]. As described above WireGuard builds on Noise protocols [4]. The first attempt at PQ Noise was in [24], which was later developed further by Angel et al. into PQ Noise [13]. Both proposals do not consider Noise’s PSK extension and pre-date KEM misbinding attacks [2]. Hülsing et al. [1]’s PQ WireGuard was further developed into the Rosenpass application [25], which instead of integrating PQ KEMs into WireGuard itself runs a “side-car” key exchange protocol, and then injects keys as PSKs into an otherwise still classically-secure WireGuard session. The authors took a more conservative approach to its key schedule, and include all public keys and ciphertexts, cutting off the flaws we found in [1]. However, it still uses Classic McEliece, and its side-car design results in additional overhead. An alternative approach to PQ WireGuard was put forward by Raynal [26], who proposed using now-broken signature scheme Rainbow [27] for authentication. As Rainbow has very large public keys like McEliece, this presents the same implementation hurdles. Some commercial VPN services now offer “PQ” WireGuard-based VPNs [28], [29], though few public details are available. This is seemingly done through the PSK mechanism in WireGuard with an out-of-band PQ key exchange. It is not clear what exact security is provided or what protocol is used.

Other work on PQ VPN protocols includes several forks of OpenVPN [30]–[32]. As OpenVPN is based on

TLS 1.3 [33], this amounts to integrating PQ TLS [34]. PQ extensions to the IKEv2 AKE [35] underlying IPsec [6] have been studied in the symbolic model [36], and its performance measured [37]. There are further proposals for PQ mechanisms at the IETF [38], [39]. A standard was also published for integrating optional PSKs (like in WireGuard) [40].

Related to our reinforced KEM, Xue et al. [41] proposed a similar primitive called *double-key* KEM (2KEM). While it builds on similar motivations, 2KEM still views the two KEM keys symmetrically, and their proposed lattice-based 2KEM essentially runs two ML-KEMs in parallel.

1.3. Concurrent and Independent Work

At USENIX Security 2025, Lafourcade et al. [42] put forward a new formal, symbolic analysis of PQ WireGuard using SAPIC+ [43], as well as showing how PQ WireGuard can be hybridized. In their analysis they independently found the UKS attack and PQ WireGuard’s reliance on additional assumptions on KEMs. They propose a fix, concatenating static public keys to the PSK value in the computation of transcript value H_4 . Similarly to the original PQ WireGuard by Hülsing et al. [1], they still include the PSK twice in the key computations. In contrast, in our streamlined design of PQ WireGuard, we only have one inclusion. Additionally, they argue that WireGuard’s design assumption that identity public keys are *usually* unknown to adversaries in the real world is unjust; without this assumption, the MAC values break identity hiding. Lafourcade et al. insert the PSK into the computation of these anti-DoS MACs to prove identity hiding without this assumption, a countermeasure that was suggested in the original paper by Donenfeld [3]. The downside is that this seems to go against the purpose of the anti-DoS feature as now the recipient needs to perform public-key operations to identify the sender to look up the appropriate PSK. In this work we model and prove identity hiding preserving the assumption that public keys are hidden from the adversary, and additionally prove *unlinkability* of independent sessions between users.

We were recently made aware that Beguinet et al. [44] constructed an efficient lattice-based 2KEM named Maul. Unlike the 2KEM proposed by Xue et al. [41], they exploit the asymmetry in the security strengths of the two KEMs being combined. While the terminologies are different, the targeted security properties of their 2KEM is identical to those of RKEM. The design of Maul and Rebar are similar at a high level. In more detail, Rebar additionally performs bit-dropping on the reinforcing (i.e., ephemeral) public key, and also relies on standard Gaussians while Maul relies on sum of uniforms (cf. Sec. 8.1). The way security is proven is also different: while Rebar is constructed following the standard FO transform, breaking it down into the two T and U_m^χ transforms [45], Maul is constructed from a tailored FO transform. Lastly, while we show the usefulness of RKEM using PQ WireGuard as a leading concrete example, they use it to build general bilateral and unilateral AKE protocols

in two- and three-message variants (corresponding to weak and full forward secrecy).

2. Preliminaries

2.1. Notation

For a set $\mathcal{M}$, we denote $m \xleftarrow{\$} \mathcal{M}$ as the process of sampling uniformly random over $\mathcal{M}$. We may make the uniform distribution explicit by denoting $\mathcal{U}(\mathcal{M})$. For a randomized algorithm $f : \mathcal{X} \times \mathcal{R} \rightarrow \mathcal{Y}$, $y \xleftarrow{\$} f(x)$ denotes the process of executing f on x with randomness $r \xleftarrow{\$} \mathcal{R}$. We write $y = f(x; r)$ when fixing the randomness r , and write $y \in f(x)$ to denote that $\exists r \in \mathcal{R}$ such that $y = f(x; r)$. We further write $f(\mathcal{X})$ for the set $\{f(x; r) \mid (x, r) \in \mathcal{X} \times \mathcal{R}\}$, where “ $\mid$ ” is interpreted as “such that”. $\Pr[X : Y]$ denotes the probability that event Y occurs on the randomness of variable X . For $a, x \in \mathbb{N}$, $a \stackrel{\pm}{\leftarrow} x$ is a shorthand for $a \leftarrow a + x$. We use this for sets S, X such that $S \stackrel{\pm}{\leftarrow} X$ denotes $S \leftarrow S \cup X$. $\oplus$ denotes the XOR operation over bit strings. Lastly, a random variable y following distribution D_y over support $\mathcal{Y}$ has *high min-entropy* if $\max_{y \in \mathcal{Y}} \Pr_{y' \leftarrow D_y}[y = y']$ is negligible.

2.2. Lattices

Cyclotomic Rings. Let d be a power-of-two integer and q a prime. Let $R = \mathbb{Z}[x]/(x^d + 1)$ be the cyclotomic ring of degree d and denote $R_q = R \bmod q$. For a real matrix $\mathbf{M} \in \mathbb{R}^{k \times \ell}$, we note $s_1(\mathbf{M})$ and call spectral norm of $\mathbf{M}$ the value $\max_{\mathbf{x} \neq 0} \frac{\|\mathbf{M} \cdot \mathbf{x}\|}{\|\mathbf{x}\|}$, where $\|\cdot\|$ denotes the L_2 -norm. The spectral norm of $\mathbf{M}$ is also the (unique non-negative) square root of the largest eigenvalue of $\mathbf{M} \cdot \mathbf{M}^\top$. We recall that if $\mathbf{M}$ is symmetric, then its singular values are the square roots of its eigenvalues. If $\mathbf{B} \in \mathcal{R}^{k \times \ell}$ has its entries in $\mathcal{R}$, we identify $\mathbf{B}$ with its associated anti-circulant matrix $\mathbf{M} \in \mathbb{Z}^{nk \times n\ell}$ and abusively say that the spectral norm of $\mathbf{B}$ is the spectral norm of $\mathbf{M}$.

Hardness Assumption. In this work, we rely on the standard module learning with errors (MLWE) problem along with (a generalization of) the recent *hint* MLWE problem by Kim et al. [46], stating that MLWE remains hard even if some leakage of the secret is provided.

Definition 1 (MLWE). Let n, k, q be integers and χ, χ' be probability distributions over R_q . The advantage $\text{Adv}_{\mathcal{A}}^{\text{MLWE}}(1^\lambda)$ of an adversary $\mathcal{A}$ against the *Module Learning with Errors* (MLWE $_{q,n,k,\chi,\chi'}$) problem is defined as:

$$\left| \Pr[\mathcal{A}(\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{x}) = 1] - \Pr[\mathcal{A}(\mathbf{A}, \mathbf{b}) = 1] \right|,$$

where $(\mathbf{A}, \mathbf{b}, \mathbf{s}, \mathbf{x}) \xleftarrow{\$} R_q^{n \times k} \times R_q^k \times \chi^k \times \chi'^n$. The MLWE $_{q,n,k,\chi,\chi'}$ assumption states that any efficient adversary $\mathcal{A}$ has negligible advantage.

Above, we may assume the noise $\mathbf{x}$ has entries sampled from different distributions, i.e., the i -th entry of $\mathbf{x}$ is sampled from χ'_i . In which case, we may write MLWE $_{q,n,k,\chi,(\chi_i)_{i \in [n]}}$.

The following is a slight generalization of the original hint-MLWE problem [46], formally introduced by [47].

Definition 2 (hint-MLWE). Let n, k, ℓ, q be integers, χ and $\tilde{\chi}$ be probability distributions over R_q and $\mathcal{F}$ be a probability distribution over $R_q^{\ell \times (n+k)}$. The advantage $\text{Adv}_{\mathcal{A}}^{\text{hint-MLWE}}(1^\lambda)$ of an adversary $\mathcal{A}$ against the *Hint Module Learning with Errors* (hint-MLWE $_{q,n,k,\ell,\chi,\tilde{\chi},\mathcal{F}}$) problem is defined as:

$$\left| \Pr[\mathcal{A}(\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{x}, \mathbf{M}, \mathbf{h}) = 1] - \Pr[\mathcal{A}(\mathbf{A}, \mathbf{b}, \mathbf{M}, \mathbf{h}) = 1] \right|,$$

where $(\mathbf{A}, \mathbf{b}, \mathbf{s}, \mathbf{x}, \mathbf{M}) \xleftarrow{\$} R_q^{n \times k} \times R_q^k \times \chi^k \times \chi^n \times \mathcal{F}$.

Moreover, the *hint* is defined as $\mathbf{h} = \mathbf{M} \begin{bmatrix} \mathbf{s} \\ \mathbf{x} \end{bmatrix} + \mathbf{z}$ where $\mathbf{z} \xleftarrow{\$} \tilde{\chi}^\ell$. The hint-MLWE $_{q,n,k,\ell,\chi,\tilde{\chi},\mathcal{F}}$ assumption states that any efficient adversary $\mathcal{A}$ has negligible advantage.

Lastly, we recall the following result which establishes the hardness of the hint-MLWE problem based on the MLWE problem. This is a simple adaptation of the original proof [46], formally appearing in [47]. Below, we denote $\mathcal{D}_\sigma$ as a discrete Gaussian distribution with standard deviation σ .

Theorem 1 (Hardness of hint-MLWE). *For any integers n, k, ℓ, q, d , let $\mathcal{F}$ be a probability distribution over $R_q^{\ell \times (n+k)}$, χ and $\tilde{\chi}$ be discrete Gaussian distributions $\mathcal{D}_{\sigma_1}$ and $\mathcal{D}_{\sigma_2}$, respectively, and B, σ a positive real such that*

$$\Pr[s_1(\mathbf{M} \cdot \mathbf{M}^\top) < B : \mathbf{M} \xleftarrow{\$} \mathcal{F}] \geq 1 - \text{negl}(\lambda),$$

and $\sigma = \omega(\sqrt{\log d})$ and $\frac{1}{\sigma^2} = 2 \cdot \left(\frac{1}{\sigma_1^2} + \frac{B}{\sigma_2^2} \right)$. Under these conditions, the hint-MLWE $_{q,n,k,\ell,\mathcal{D}_{\sigma_1},\mathcal{D}_{\sigma_2},\mathcal{F}}$ problem is as hard as the MLWE $_{q,n,k,\mathcal{D}_\sigma}$ problem.

2.3. Rounding

In our work, we use the rounding definition used by Kyber [19]. Below, we briefly recall their definition.

For an even (resp. odd) positive integer q , we define $x' = x \bmod^\pm q$ to be the unique element x' in the range $\frac{-q}{2} < x' \leq \frac{q}{2}$ (resp. $-\frac{q-1}{2} < x' \leq \frac{q-1}{2}$) such that $x' = x \bmod q$. For any positive integer q , we define $x' = x \bmod^+ q$ to be the unique element x' in the range $0 \leq x' < q$ such that $x' = x \bmod q$. We simply write $x \bmod q$ when the representation is not important. Also, for an element in $x \in \mathbb{Q}$, $\lfloor x \rfloor$ denotes the rounding to the nearest integer, where in case of a tie, we take the larger integer.

Compression and Decompression. We define the following compression and decompression algorithms for positive integers d and q such that $d < \lceil \log_2(q) \rceil$:

$$\text{Compress}_q : \mathbb{Z}_q \longrightarrow \mathbb{Z}_{2^d}$$

$$x \longmapsto \left\lfloor \frac{2^d}{q} \cdot x \right\rfloor \bmod^+ 2^d.$$

$$\text{Decompress}_q : \mathbb{Z}_{2^d} \longrightarrow \mathbb{Z}_q$$

$$y \longmapsto \left\lfloor \frac{q}{2^d} \cdot y \right\rfloor.$$

For these functions, we have the following:

Lemma 1. Let d and q be positive integers such that $d < \lceil \log_2(q) \rceil$. Then, for any $x \in \mathbb{Z}_q$, we have

$$|x' - x \bmod^\pm q| \leq \left\lfloor \frac{q}{2^{d+1}} \right\rfloor,$$

where $x' = \text{Decompress}_q(\text{Compress}_q(x, d), d)$.

When, Compress_q or Decompress_q is used with $x \in R_q$ or $\mathbf{x} \in R_q^k$, the procedure is applied to each coefficient individually.

2.4. Quantum Tools

The following is taken from [48], slightly adapted as in [49].

Lemma 2 (Generic Distinguishing Problem with Bounded Probabilities). Let $\mathcal{S}$ be a finite set. Let $F : \mathcal{X} \rightarrow \{0, 1\}$ be the following function: for each $x \in \mathcal{X}$, $F(x) = 1$ with probability $\delta_x \leq \delta$ and $F(x) = 0$ otherwise. Let $Z : \mathcal{X} \rightarrow \{0, 1\}$ be the zero function, that is $Z(x) = 0$ for all x . For any unbounded-time quantum adversary $\mathcal{A}$ making at most Q quantum queries to F or Z , we have

$$\left| \Pr[\mathcal{A}^{F(\cdot)}() = 1] - \Pr[\mathcal{A}^{Z(\cdot)}() = 1] \right| \leq 8 \cdot (Q + 1)^2 \cdot \delta.$$

Following [50], we model a quantum random oracle $\mathcal{O} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ as a mapping $|x\rangle|y\rangle \rightarrow |x\rangle|y \oplus \mathcal{O}(x)\rangle$, where $\mathcal{O}$ is a uniformly sampled random function. In the quantum random oracle model (QROM), all parties can evaluate $\mathcal{O}$ in superposition. Below are some useful lemmas in the QROM used in prior works [51]–[53].

Lemma 3. Let ℓ be a positive integer. Let $\mathcal{X}$ and $\mathcal{Y}$ be finite sets. Let $H_{\text{prf}} : \{0, 1\}^\ell \times \mathcal{X} \rightarrow \mathcal{Y}$ and $H_q : \mathcal{X} \rightarrow \mathcal{Y}$ be two independent quantum random oracles. Then for any unbounded-time quantum adversary $\mathcal{A}$ making at most Q quantum queries to the random oracles, we have

$$\left| \Pr[s \xleftarrow{\$} \{0, 1\}^\ell : \mathcal{A}^{H_{\text{prf}}(\cdot, \cdot), H_q(\cdot)}() = 1] - \Pr[\mathcal{A}^{H_{\text{prf}}(\cdot, \cdot), H_q(\cdot)}() = 1] \right| \leq Q \cdot 2^{-\frac{\ell}{2} + 1}.$$

Lemma 4 (Oneway to Hiding, simplified). Let $\mathcal{X}$ and $\mathcal{Y}$ be finite sets and $\mathcal{S} \subseteq \mathcal{X}$ a random set. Let $G, H : \mathcal{X} \rightarrow \mathcal{Y}$ be random functions satisfying $G(x) = H(x)$ for all $x \notin \mathcal{S}$. Let z be an arbitrary bit string. ($\mathcal{S}, G, H, z$ may have an arbitrary joint distribution).

Let $\mathcal{A}$ be a binary-output algorithm making at most Q oracle queries. Let $\mathcal{B}^H$ be a quantum algorithm that on input z does the following: sample $i \xleftarrow{\$} [Q]$, run $\mathcal{A}^{H(\cdot)}(z)$ until (just before) the i -th query, measure all query input registers in the computational basis, output the set $\mathcal{T}$ of measurement outcomes.

Let

$$P_{\text{left}} := \Pr_{H, z}[\mathcal{A}^{H(\cdot)}(z) = 1], \quad P_{\text{right}} := \Pr_{G, z}[\mathcal{A}^{G(\cdot)}(z) = 1],$$

$$P_{\text{guess}} := \Pr_{H, G, \mathcal{S}, z}[\mathcal{T} \xleftarrow{\$} \mathcal{B}^{G(\cdot)}(z) : \mathcal{S} \cap \mathcal{T} \neq \emptyset].$$

Then,

$$|P_{\text{left}} - P_{\text{right}}| \leq 2 \cdot Q \cdot \sqrt{P_{\text{guess}}}.$$

2.5. General Cryptographic Primitives

2.5.1. Key Derivation Functions. In the protocol construction we use key derivation functions, which map arbitrary-length keys and labels to an output $\text{KDF} : \mathcal{K} \times \mathcal{L} \rightarrow \mathcal{X}$.

We require a KDF to be *dual-PRF-secure* [54], [55] and *collision resistant*.

Definition 3 (Dual-PRF-Security). A KDF is *PRF-secure* (in the first input) if for any efficient adversary $\mathcal{A}$, the following advantage is negligible:

$$\left| \Pr[k \xleftarrow{\$} \mathcal{K} : \mathcal{A}^{\text{KDF}(k, \cdot)}() = 1] - \Pr[\mathcal{A}^{R(\cdot)}() = 1] \right|,$$

where $R(\cdot)$ is sampled uniformly over all functions mapping $\mathcal{L}$ to $\mathcal{X}$. Moreover, a KDF is *dual-PRF-secure* if we further have the above for the second input: $\text{KDF}(\cdot, \ell)$ for a randomly sampled $\ell \xleftarrow{\$} \mathcal{L}$ is indistinguishable from a truly random function mapping $\mathcal{K}$ to $\mathcal{X}$.

Definition 4 (Collision Resistance). A KDF is *collision-resistant* if for any efficient adversary $\mathcal{A}$, the following advantage is negligible:

$$\Pr \left[\begin{array}{l} (s, s', m, m') \xleftarrow{\$} \mathcal{A}(1^\lambda) : \\ \text{KDF}(s, m) = \text{KDF}(s', m') \wedge (s, m) \neq (s', m') \end{array} \right].$$

2.5.2. Authenticated Encryption with Associated Data.

We recall the definition of an AEAD.

Definition 5 (AEAD). An authenticated encryption with associated data (AEAD) with key, nonce, header, and message spaces $\mathcal{K}$, $\mathcal{N} = \{0, 1\}^\ell$, $\mathcal{H} \subseteq \{0, 1\}^*$, and $\mathcal{M} \subseteq \{0, 1\}^*$, respectively, consists of the following algorithms.

AEAD.Enc(k, N, H, M) $\rightarrow C$: On input a key $k \in \mathcal{K}$, a nonce $N \in \mathcal{N}$, a header $H \in \mathcal{H}$ and a message $M \in \mathcal{M}$, it outputs a ciphertext $C \in \{0, 1\}^*$.

AEAD.Dec(k, N, H, C) $\rightarrow M$: On input a key $k \in \mathcal{K}$, a nonce $N \in \mathcal{N}$, a header $H \in \mathcal{H}$, and a ciphertext $C \in \{0, 1\}^*$, it outputs a message $M \in \mathcal{M}$ or $\perp \notin \mathcal{M}$ indicating decryption failure.

We require an AEAD to be correct, IND-CPA secure, and INT-CTXT secure [56].

Definition 6 (Correctness). An AEAD is correct if $\text{AEAD.Dec}(k, N, H, \text{AEAD.Enc}(k, N, H, M)) = M$ holds 1 for all k, N, H, M in the appropriate spaces.

Definition 7 (IND-CPA Security). An AEAD is IND-CPA secure if for any efficient adversary $\mathcal{A}$, the following advantage is negligible:

$$\left| \Pr[k \xleftarrow{\$} \mathcal{K} : \mathcal{A}^{\mathcal{O}_{\text{Enc}}(k, \cdot, \cdot, \cdot)}(1^\lambda) = 1] - \Pr[\mathcal{A}^{\mathcal{S}(\cdot, \cdot, \cdot)}(1^\lambda) = 1] \right|,$$

where on query (N, H, M) , $\mathcal{O}_{\text{Enc}}(k, \cdot, \cdot, \cdot)$ outputs $\text{AEAD.Enc}(k, N, H, M)$ and $\mathcal{S}(\cdot, \cdot, \cdot)$ outputs a random bit string with length $|M|$. When queried on the same input twice, the oracles are assumed to output to the previous value.

Definition 8 (INT-CTXT Security). An AEAD is INT-CTXT secure if for any efficient adversary $\mathcal{A}$, the following advantage is negligible:

$$\Pr \left[\begin{array}{l} k \xleftarrow{\$} \mathcal{K}, (N^*, H^*, C^*) \xleftarrow{\$} \mathcal{A}^{\mathcal{O}_{\text{Enc}}(k, \cdot, \cdot, \cdot)}(1^\lambda) : \\ \text{AEAD.Dec}(k, N^*, H^*, C^*) \neq \perp \wedge (N^*, H^*, C^*) \notin \mathcal{Q}, \end{array} \right]$$

where on input (N, H, M) , $\mathcal{O}_{\text{Enc}}(k, \cdot, \cdot, \cdot)$ outputs $\text{AEAD.Enc}(k, N, H, M)$, and $\mathcal{Q}$ is the list of (N, H, M) that $\mathcal{A}$ submitted to oracle $\mathcal{O}_{\text{Enc}}$.

2.5.3. Key Encapsulation Mechanisms. We recall the definition of a KEM.

Definition 9 (KEM). A key encapsulation mechanism (KEM) an efficiently sampleable key space $\mathcal{K}$ consists of the following algorithms.

KEM.KeyGen $(1^\lambda) \rightarrow (\text{pk}, \text{sk})$: On input the security parameter 1^λ , it outputs a public and secret key pair (pk, sk) .

KEM.Enc $(\text{pk}) \rightarrow (\text{ct}, \text{ss})$: On input a public key pk , it outputs a ciphertext ct and a key $\text{ss} \in \mathcal{K}$.

KEM.Dec $(\text{sk}, \text{ct}) \rightarrow \text{ss}$: On input a secret key sk and a ciphertext ct , it outputs a key $\text{ss} \in \mathcal{K}$.

We require a KEM to be correct with all but probability δ . This is set to be negligible in the instantiation.

Definition 10 (δ -Correctness). A KEM is δ -correct if for any $\lambda \in \mathbb{N}$, we have $\Pr[\text{ss} = \text{Dec}(\text{sk}, \text{ct})] \geq 1 - \delta$, where the probability is taken over $(\text{pk}, \text{sk}) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$ and $(\text{ct}, \text{ss}) \xleftarrow{\$} \text{Enc}(\text{pk})$.

We further require the SPR-CCA security for KEMs [49].

Definition 11 (SPR-CCA Security). For a simulator $\mathcal{S}$, we define the advantage $\text{Adv}_{\text{KEM}, \mathcal{A}}^{\text{SPR-CCA}}(\lambda)$ of $\mathcal{A}$ against the SPR-CCA security game as follows:

$$\Pr \left[\begin{array}{l} b \xleftarrow{\$} \{0, 1\}, \\ (\text{pk}, \text{sk}) \xleftarrow{\$} \text{KEM.KeyGen}(1^\lambda), \\ (\text{ct}_0, \text{ss}_0) \xleftarrow{\$} \text{KEM.Encaps}(\text{pk}); \quad : b = b' \\ (\text{ct}_1, \text{ss}_1) \xleftarrow{\$} \mathcal{S}(1^\lambda) \times \mathcal{K}, \\ b' \xleftarrow{\$} \mathcal{A}^{\mathcal{O}_{\text{Dec}}(\text{sk}, \cdot)}(\text{ek}, \text{ct}_b, \text{ss}_b) \end{array} \right] - \frac{1}{2},$$

where on input ct , $\mathcal{O}_{\text{Dec}}(\text{sk}, \cdot)$ outputs $\text{KEM.Dec}(\text{sk}, \text{ct})$, and $\mathcal{A}$ is not allowed to query the challenge ciphertexts ct_0 and ct_1 to $\mathcal{O}_{\text{Dec}}$. A KEM is SPR-CCA secure if there exists an efficient simulator $\mathcal{S}$ such that $\text{Adv}_{\text{KEM}, \mathcal{A}}^{\text{SPR-CCA}}(\lambda) = \text{negl}(\lambda)$ for any efficient adversary $\mathcal{A}$.

3. Reviewing PQ WireGuard by Hülsing et al.

We briefly review the design and security of PQ WireGuard [1], a post-quantum adaptation of WireGuard [3].

3.1. Construction Overview

Recall that WireGuard is based on the Noise [4] “IKpsk2” pattern. This is a one-roundtrip DH key exchange between a

client and server,² where the client knows the server’s long-term public key, while the server learns the client’s identity after the initial handshake message. A pre-shared key (PSK) can be mixed into the key exchange, making the protocol more secure in case all other secrets are compromised. Following the design of Noise (with the pre-shared key extension [4, Sec. 9]), the transmitted ephemeral DH keys are hashed with the shared secrets in a *key chain* (C_i ’s), and a separate *hash chain* (H_i ’s) of all messages is computed and confirmed through AEAD ciphertexts in both messages. (Note that although it will appear that some processing steps could be merged, the systematic, piecewise processing relates to the token-processing-based nature of Noise protocols.)

At a high level, PQ WireGuard by Hülsing et al. [1] replaces the DH components with KEMs, using Classic McEliece 460896 for the long-term keys and a modified, IND-CPA version of Saber for the ephemeral keys. A detailed description of PQ WireGuard, and on how the key and hash chains evolve is provide in Fig. 1. The timestamp prevents replay attacks; the server keeps track of the greatest timestamp received per client and discards messages containing timestamps less than or equal to it [3, Sec. 5.1]. This lets the server avoid DoS from half-open connections, as each client message can only open one connection. MAC tags m1 , m2 are further used to prevent DoS attacks from drive-by attackers and spoofed initiators, letting WireGuard avoid having to do public-key operations on such message. Tag m1 , keyed with the client’s public key, proves that the server is familiar with the client’s identity (assuming that the public key is not widely disseminated, as is often the case in practice). Tag m2 is keyed with a *cookie* value only used when the server is under load. If cookies are to be used but not provided, the server will first reply with a randomized value derived from its public key and a short-term secret. It is sent encrypted using a key derived from the server’s public key and the source IP; otherwise it is just 16 zero bytes [3, Sec. 5.3]. A response using a correct, non-zero cookie proves that the client actually receives messages at its specified source IP address. Lastly, values $\text{lb1}_1, \text{lb1}_2, \text{lb1}_3$ are distinct fixed strings stemming from the Noise protocol [3, Sec. 5.4]. Here, we note that if the server’s public key is unknown to the adversary, then the cookie value has high min-entropy (i.e., it cannot be guessed by the adversary).

Most relevant to the context of our paper is the key schedule of Hülsing et al. [1]. Since KEMs come with both a public key pk and a ciphertext ct , they are syntactically quite different from DH key exchanges, and as such, the key and hash chains evolve quite differently. Their protocol mixes only pk_e and ct_e in the key chain, and only pk_s , pk_e , ct_e , and (an AEAD encryption of) pk_C and ct_C in the hash chain. They further add another use of the PSK to the computation of K_{time} to authenticate the *time* message (in lieu of static-static DH key exchange, which KEMs do not support), resulting in mixing in PSK *twice* into the key chain.

2. While “client” and “server” is used throughout this section for readability, “initiator” and “responder” would be the more general term. This is terminology we use in Sec. 4.




Figure 1. Description of PQ WireGuard by Hülsing et al. [1] and ours in the left column. Variant-specific functionality is boxed. Our protocol based on standard KEMs is identical to Hülsing et al.'s, except for the header information in the second to last AEAD ciphertext, marked in green. The middle and right columns show the key schedule (i.e., evolution of the key and hash chains and generated cryptographic keys K); KDF_i -labeled transitions indicate domain-separated calls to KDF. The green insets show the computations for our protocol when using KEM instead of RKEM.

We note that while [1] uses a NAXOS-like construction and protect against bad randomness, we simplify our analysis and presentation by assuming secure random generation.

Looking ahead, since pk_C (or ct_C) is not hashed into the key chain, PQ WireGuard of Hülsing et al. [1] is susceptible against unknown key share (UKS) attacks. In their symbolic model, they prove it secure against it by hashing ct_C . While the proof requires the KEM to satisfy some binding properties [2], this holds in their way of modeling KEM. See Sec. B for more details.

3.2. Security Properties

PQ WireGuard aims to achieve the same level of security as WireGuard, even against quantum adversaries. While a formal definition is given in Sec. 4, we provide a brief overview here. The baseline security property is key indistinguishability, requiring that the established session key is pseudorandom. This should hold even if the adversary compromises any (non-trivial) combination of the secret information held by the parties.

It further aims to achieve several unique properties: *session key uniqueness*, *entity authentication*, *identity hiding*, and *DoS mitigation* [1, Sec. II-B]. The first notion requires the established keys to be unique, and the second requires that the parties communicate with their intended peer, implying that no adversary can impersonate another party through *key-compromise impersonation* (KCI) or *unknown key share* (UKS) attacks. The third notion requires that the transcript alone does not leak the communicating parties (i.e., the long-term keys) to the adversary.

Lastly, while the DoS mitigation entails many meanings, in [1], it is the requirement that the first client handshake message is authenticated. This allows a server to reject forged messages without having to explicitly respond to an unauthenticated client. Concretely, looking at PQ WireGuard in Fig. 1, ltk is an AEAD encryption of (the hash of the) client public key pk_C . The server, first decrypts $H(pk_C)$, and checks if any of the client public key in memory hashes to this value. If so it identifies its peer as pk_C , and otherwise, it halts. As PSK is used to derive the AEAD key, an adversary without knowledge of PSK cannot impersonate the client.

Hülsing et al. [1] analyze all the security properties above in their symbolic model using the Tamarin prover, building on the symbolic model by [9]. In contrast, only key-indistinguishability (and, implicitly, some authentication notion) is analyzed in their computational model, building on the model by [7].

4. Computational Model for PQ WireGuard

We provide a formal computational (i.e., game-based) security model for PQ WireGuard.

4.1. Syntax and Execution Environment

We formally define the key exchange protocol used by PQ WireGuard as a WGKE protocol: a one-roundtrip protocol where an *initiator* (i.e., client) starts the conversation, and a *responder* (i.e., server) talks back once. As explained in Sec. 1.1, the responder algorithm is split into two phases *ResIdentify* and *ResDer*, making the *post-specified peer* model explicit. Below, we assume ϵ is a special symbol denoting an empty string. This is used for instance to make it clear that pre-shared keys psk are optional inputs.

Definition 12. A WGKE protocol consists of the following algorithms, where $\mathcal{K}$ and $\mathcal{K}_{\text{psk}}$ are the session key and pre-shared key spaces, respectively.

KeyGen(1^λ) $\rightarrow$ (**lpk**, **lsk**): On input the security parameter 1^λ , it outputs a public and secret key pair (**lpk**, **lsk**).

Init(**lsk_i**, **lpk_r**, **psk**) $\rightarrow$ (**m₁**, **st**): The initiator on input its secret key **lsk_i**, the responder's public key **lpk_r**, and a pre-shared key $\text{psk} \in \mathcal{K}_{\text{psk}} \cup \{\epsilon\}$, outputs a first message **m₁** and a session state **st**.

ResIdentify(**lsk_r**, **m₁**) $\rightarrow$ **lpk_i** or $\perp$: The responder on input its secret key **lsk_r** and a first message **m₁**, outputs either the intended initiator's public key **lpk_i**, or $\perp$ indicating that the initiator could not be identified.

ResDer(**lsk_r**, **lpk_i**, **m₁**, **psk**) $\rightarrow$ (**K**, **m₂**): The responder on input its secret key **lsk_r**, intended initiator's public key **lpk_i**, a first message **m₁**, and a pre-shared key $\text{psk} \in \mathcal{K}_{\text{psk}} \cup \{\epsilon\}$, outputs a session key **K** and a second message **m₂**.

InitDer(**lsk_i**, **m₂**, **st**, **psk**) $\rightarrow$ **K**: The initiator on input secret key **lsk_i**, a second message **m₂**, a session state **st**, and a pre-shared key $\text{psk} \in \mathcal{K}_{\text{psk}} \cup \{\epsilon\}$, outputs a session key **K**.

We consider a unified computational (i.e., game-based) security model, formally described in Fig. 3, capturing the various types of security guarantees. Each party $u \in \mathbb{N}$ possess a public and secret key pair (**lpk_u**, **lsk_u**), and as in [3, Sec. 2], u is identified by its public key **lpk_u** and parties are unable to register an existing public key. For convenience, we assume a function π_{userid} mapping **lpk_u** to a party $u \in \mathbb{N}$ maintained by the game.

The game allows an initiator and a responder to establish a key exchange session. However, as an initiator and responder may have diverging views on the session, the game creates an *instance* $\text{iID} \in \mathbb{N}$ for each of them and keeps track of the following associated information.

role[**iID**] $\in \{\text{init}, \text{resp}\}$: instance iID 's role, i.e., whether the instance acts as the initiator or the responder.

(**Init**[**iID**], **Resp**[**iID**]) $\in \mathbb{N} \cup \{\epsilon\} \times \mathbb{N}$: the identities of the initiator and the responder relative to the instance iID . The initiator may be ϵ , modeling the fact that WireGuard allows the server to learn the client's identity during the protocol.

psk[**iID**]: the pre-shared key psk used by the instance iID , where $\text{psk} = \epsilon$ if it is unused.

sent[**iID**]: the message sent by the instance iID .

recv[**iID**]: the message received by the instance iID .

key[**iID**]: the session key computed by the instance iID . This is set to $\perp$ if iID fails to establish a key or ϵ if the session has not either started or finished yet.

While there are other information kept tracked by the game (i.e., corruption states), they are explained when needed. We also denote $\mathcal{S}_{\text{iID}} \subset \mathbb{N}$ as all the instances generated in the game (cf. Fig. 3). Following prior AKE definitions [57]–[61], we use the notion of *partners* to define the set of trivial attacks; since there are two instances iID, iID' that may share the same session key, we need a mechanism to identify them. While there are various ways to define the predicate *Partner*, we use the following simplistic form as they are equivalent for natural protocols (cf. [61]).

Definition 13 (*Partner*). The predicate *Partner*(iID, iID') holds true if and only if $\forall \text{iID}, \text{iID}' \in \mathcal{S}_{\text{iID}}$, we have $\llbracket \text{role}[\text{iID}] \neq \text{role}[\text{iID}'] \rrbracket \wedge \llbracket \text{key}[\text{iID}] = \text{key}[\text{iID}'] \notin \{\epsilon, \perp\} \rrbracket$.

Lastly, we require WGKE to be *correct*. This requires that when all the users in the system honestly execute the WGKE protocol without the adversary tampering the protocol messages, then they derive an identical session key all but with a negligible probability. Formally, we model this through a game between a *passive* adversary $\mathcal{P}$ and a challenger. Here, a passive adversary $\mathcal{P}$ can arbitrary interact with the users under the restriction that it must honestly deliver the protocol messages. For instance, if an initiator i outputs a first message **m₁** to a receiver r , then $\mathcal{P}$ cannot invoke the receiver r on anything other than **m₁**.

Definition 14 (*Correctness*). We define the correctness game in Fig. 2 and define the advantage of a *passive* adversary $\mathcal{P}$ as

$$\text{Adv}_{\text{WGKE}, \mathcal{P}}^{\text{CORR}}(1^\lambda) := \Pr \left[\text{Game}_{\text{WGKE}, \mathcal{P}}^{\text{CORR}}(1^\lambda) = 1 \right].$$

We say a WGKE protocol is *correct* if $\text{Adv}_{\text{WGKE}, \mathcal{P}}^{\text{CORR}}(1^\lambda) = \text{negl}(\lambda)$ for any efficient passive adversary $\mathcal{P}$.

4.2. Security Guarantees

We now explain the four types of security guarantees satisfied by PQ WireGuard: authentication, DoS mitigation, identity hiding, and key indistinguishability.

4.2.1. Entity Authentication. In [1], [7], authentication was not defined explicitly and was merely a means to break key indistinguishability. This is why attacks such as UKS attacks

```

1: function GameWGKE, PCORR(1λ)
2:   SIID := ∅ ▷ Set of existing instances
3:   NumIID := 0 ▷ Number of instances
4:   UHon := [μ] ▷ μ ∈ ℕ honest parties
5:   for u ∈ UHon do
6:     (lpku, lsku) ←$ KeyGen(1λ)
7:   1 ← PO* ((lpku)u ∈ UHon) ▷ P always terminates with 1
8:   for (iID, iID') ∈ SIID × SIID do
9:     cond := [[role[iID] ≠ role[iID']]
10:      ∧ [[Init[iID] = Init[iID']]]
11:      ∧ [[Resp[iID] = Resp[iID']]]
12:      ∧ [[psk[iID] = psk[iID']]]
13:      ∧ [[sent[iID] = rcv[iID']]] ∧ [[rcv[iID] = sent[iID']]]
14:      ∧ [[key[iID] ≠ key[iID']]]]
15:     if cond then
16:       return 1
17:   return 0

```

Figure 2. Correctness game where $\mathcal{O}^*$ denotes the set $\{\mathcal{O}_{\text{Init}}, \mathcal{O}_{\text{InitDer}}, \mathcal{O}_{\text{ResDer}}\}$ defined in Fig. 4.

```

1: function GameWGKE, Atype(1λ)
2:   b ←$ {0, 1}
3:   SIID := ∅ ▷ Set of instances
4:   NumIID := 0 ▷ Number of instances
5:   iID* := ε ▷ Challenge instance for Key ind. game
6:   (UHon, Uchall, UMal) := ([μ], ∅, ∅) ▷ μ honest parties
7:   numu := μ ▷ Number of registered parties
8:   for u ∈ UHon do
9:     (lpku, lsku) ←$ KeyGen(1λ)
10:  if [type ≠ IDHIDE] then
11:    b' ←$ AO(1λ, (lpku)u ∈ UHon)
12:    if [type = MATCH] then
13:      return [[Match(SIID) = false] ▷ Def. 15
14:    else if [type = DoS] then
15:      return [[DoSSec(SIID) = false] ▷ Def. 16
16:    else ▷ type = KIND
17:      if [[safeWGKE(iID*) = false]] then b' ←$ {0, 1} ▷ Def. 26
18:      return [b = b']
19:  else ▷ type = IDHIDE
20:    b' ←$ AO*(1λ, UHon) ▷ No public keys of honest parties
21:    if [[∃u ∈ Uchall : RevLSK[u] = true]] then b' ←$ {0, 1}
22:    return [b = b']

```

Figure 3. Games for match soundness, DoSSec security, identity hiding, and key indistinguishability, where type ∈ {MATCH, DoSSec, IDHIDE, KIND}. $\mathcal{O}$ denotes the set of all oracles in Fig. 4 excluding $\mathcal{O}_{\text{Test}}^{\text{IDHIDE}}$. $\mathcal{O}^*$ denotes the set of oracles $\mathcal{O}_{\text{Test}}^{\text{IDHIDE}}$, $\mathcal{O}_{\text{RevSessKey}}$, $\mathcal{O}_{\text{RevLSK}}$, $\mathcal{O}_{\text{RevState}}$, and $\mathcal{O}_{\text{RevPSK}}$.

were not captured in their computational models as they do not necessarily break key indistinguishability.

Following recent AKE definitions [14], [15], we model this explicitly via a game called *match soundness*. At a high level, this defines the expected behavior of partnered instances (i.e., instances with the same session key). Citing from [14], *implicit authentication* requires “should there be an instance iID that possesses the same session key as instance iID', then the owner of instance iID' must be the identity designated as the peer of instance iID.” Indeed, this covers UKS and KCI attacks. Moreover, our model captures session key uniqueness as sessions have at most one partner.

Definition 15 (Match Soundness). We define the predicate Match such that Match(S_{IID}) holds true if and only if

∀iID, iID', iID'' ∈ S_{IID},

- 1) If [[Partner(iID, iID') = true]], then [[Init[iID] = Init[iID']] ∧ [[Resp[iID] = Resp[iID']]].
- 2) If [[Partner(iID, iID') = Partner(iID, iID'') = true]], then [[iID' = iID'']].

Formally, WGKE is *match sound* if Adv_{WGKE, A}^{MATCH}(1^λ) := Pr[Game_{WGKE, A}^{MATCH}(1^λ) = 1] (cf. Fig. 3) is negligible for all efficient A.

4.2.2. DoS Mitigation. We aim to address DoS attacks in which an adversary attempts to create load and half-open connections on the server by transmitting forged client messages and thereby eliciting responses that require expensive operations to craft. Following the symbolic model of PQ-WireGuard by [1], we require that an adversary cannot pass ResIdentify (which only identifies the sender) without knowledge of the PSK. To formalize this requirement, we introduce the DoSSec security game defined in Fig. 3 with the predicate in Def. 16. In this game, the responder is guaranteed to execute ResDer (i.e., send a message and derive a session key, possibly ⊥, after running ResIdentify) at most L times, provided that an initiator with the correct PSK has been executed L times. This definition relies on the following observation: the adversary may obtain legitimate client messages via the $\mathcal{O}_{\text{Init}}$ oracle and forward them to the server, thereby eliciting valid responses. Such behavior is not regarded as a DoS attack, as the messages are legitimately generated by the client for session establishment. Instead, the adversary is deemed successful if it can cause the server to produce L + 1 responses, namely, at least one more than the number of legitimate initiations by the client. The DoSSec security therefore captures the requirement that the responder maintains no half-open connections beyond those explicitly requested by the client.

Definition 16 (DoS Security). We define the predicate DoSSec such that DoSSec(S_{IID}) holds true if and only if ∀i, j ∈ U_{Hon} such that RevPSK[{i, j}] = false, |RESP_{i, j}| ≤ |INIT_{i, j}| holds, where

$$\text{INIT}_{i, j} := \left\{ iID \mid \begin{array}{l} iID \in S_{IID} \wedge \text{role}[iID] = \text{init} \\ \wedge \text{Init}[iID] = i \wedge \text{Resp}[iID] = j \end{array} \right\}, \text{ and}$$

$$\text{RESP}_{i, j} := \left\{ iID \mid \begin{array}{l} iID \in S_{IID} \wedge \text{role}[iID] = \text{resp} \\ \wedge \text{Init}[iID] = i \wedge \text{Resp}[iID] = j \wedge \text{key}[iID] \neq \epsilon \end{array} \right\}.$$

Formally, WGKE is *DoS secure* if Adv_{WGKE, A}^{DoSSec}(1^λ) := Pr[Game_{WGKE, A}^{DoSSec}(1^λ) = 1] (cf. Fig. 3) is negligible for all efficient A.

4.2.3. Identity Hiding. Identity hiding ensures that a passive adversary cannot identify the communicating parties as long as their public keys are unrevealed; this is a required assumption since MAC tags, keyed with the public keys, are included in the transcript. We model this by only giving the number of parties |U_{Hon}| to the adversary. In fact, our model captures a stronger property *unlinkability* that has not been formalized even in the symbolic models [1], [9], further requiring that an adversary cannot tell if two session executions were from the same parties. We note that this

```

1: function  $\mathcal{O}_{\text{RegisterPK}}(\text{lpk})$ 
2:   require  $\llbracket \forall u \in \mathcal{U}_{\text{Hon}}, \text{lpk} \neq \text{lpk}_u \rrbracket$ 
3:    $\text{num}_u \stackrel{\$}{\leftarrow} 1$ 
4:    $\mathcal{U}_{\text{Mal}} \stackrel{\$}{\leftarrow} \{\text{num}_u\}$ 
5:    $\text{RevLSK}[\text{num}_u] \leftarrow \text{true}$ 
6: function  $\mathcal{O}_{\text{CreatePSK}}(i, j)$ 
7:   require  $\llbracket i \neq j \rrbracket \wedge \llbracket \text{PSK}[\{i, j\}] = \epsilon \rrbracket$ 
8:    $\text{PSK}[\{i, j\}] \stackrel{\$}{\leftarrow} \mathcal{K}_{\text{psk}}$ 
9:    $\text{RevPSK}[\{i, j\}] \leftarrow \text{false}$ 
10: function  $\mathcal{O}_{\text{Init}}(i, r)$ 
11:   require  $\llbracket i \neq r \rrbracket \wedge \llbracket i \in \mathcal{U}_{\text{Hon}} \rrbracket$ 
12:   require  $\llbracket r \in \mathcal{U}_{\text{Hon}} \cup \mathcal{U}_{\text{Mal}} \rrbracket$ 
13:    $\text{iID} \leftarrow \text{iID} + 1 \triangleright \text{Create new instance}$ 
14:    $\mathcal{S}_{\text{iID}} \leftarrow \mathcal{S}_{\text{iID}} \cup \{\text{iID}\}$ 
15:    $\text{psk}_{i,r} \leftarrow \text{PSK}[\{i, r\}]$ 
16:    $(\text{m}_1, \text{st}) \stackrel{\$}{\leftarrow} \text{Init}(\text{lsk}_i, \text{lpk}_r, \text{psk}_{i,r})$ 
17:    $(\text{role}[\text{iID}], \text{Init}[\text{iID}], \text{Resp}[\text{iID}]) \leftarrow (\text{init}, i, r)$ 
18:    $(\text{sent}[\text{iID}], \text{state}[\text{iID}], \text{psk}[\text{iID}]) \leftarrow (\text{m}_1, \text{st}, \text{psk}_{i,r})$ 
19:   return  $(\text{iID}, \text{m}_1)$ 
20: function  $\mathcal{O}_{\text{Res}}(r, \text{m}_1)$ 
21:   require  $\llbracket r \in \mathcal{U}_{\text{Hon}} \rrbracket$ 
22:    $\text{iID} \leftarrow \text{iID} + 1 \triangleright \text{Create new instance}$ 
23:    $\mathcal{S}_{\text{iID}} \leftarrow \mathcal{S}_{\text{iID}} \cup \{\text{iID}\}$ 
24:    $\triangleright \text{Identify initiator lpk}$ 
25:    $\text{lpk} \leftarrow \text{ResIdentify}(\text{lsk}_r, \text{m}_1)$ 
26:    $\triangleright \text{lpk is not registered}$ 
27:   if  $\llbracket \text{lpk} \notin \{\text{lpk}_u\}_{u \in \mathcal{U}_{\text{Hon}} \cup \mathcal{U}_{\text{Mal}}} \rrbracket$ 
28:     then
29:        $(\text{role}[\text{iID}], \text{Init}[\text{iID}], \text{Resp}[\text{iID}]) \leftarrow (\text{resp}, \perp, r)$ 
30:        $(\text{sent}[\text{iID}], \text{recv}[\text{iID}], \text{psk}[\text{iID}]) \leftarrow (\perp, \text{m}_1, \perp)$ 
31:       key  $[\text{iID}] \leftarrow \epsilon \triangleright \text{Key not set}$ 
32:     else  $\triangleright \text{Run rest of responder}$ 
33:        $\triangleright i \in \mathcal{U}_{\text{Hon}} \cup \mathcal{U}_{\text{Mal}}$ 
34:        $i := \pi_{\text{userid}}(\text{lpk})$ 
35:        $\text{psk}_{i,r} \leftarrow \text{PSK}[\{i, r\}]$ 
36:        $(\text{m}_2, K) \stackrel{\$}{\leftarrow} \text{ResDer}(\text{lsk}_r, \text{lpk}, \text{m}_1, \text{psk}_{i,r})$ 
37:        $(\text{role}[\text{iID}], \text{Init}[\text{iID}], \text{Resp}[\text{iID}]) \leftarrow (\text{resp}, i, r)$ 
38:        $(\text{sent}[\text{iID}], \text{recv}[\text{iID}], \text{psk}[\text{iID}]) \leftarrow (\text{m}_2, \text{m}_1, \text{psk}_{i,r})$ 
39:       key  $[\text{iID}] \leftarrow K$ 
40:       return  $(\text{iID}, \text{m}_2)$ 
41: function  $\mathcal{O}_{\text{InitDer}}(\text{iID}, \text{m}_2)$ 
42:   if  $\llbracket \text{role}[\text{iID}] \neq \text{init} \rrbracket \vee \llbracket \text{key}[\text{iID}] \neq \epsilon \rrbracket$  then
43:     return  $\perp \triangleright \text{Require initiator's session is open}$ 
44:    $(i, r) := (\text{Init}[\text{iID}], \text{Resp}[\text{iID}])$ 
45:    $K \stackrel{\$}{\leftarrow} \text{InitDer}(\text{lsk}_i, \text{m}_2, \text{st}[\text{iID}], \text{psk}[\text{iID}])$ 
46:    $\text{recv}[\text{iID}] \leftarrow \text{m}_2$ 
47:   key  $[\text{iID}] \leftarrow K$ 
48: function  $\mathcal{O}_{\text{RevSessKey}}(\text{iID})$ 
49:   require  $\llbracket \text{role}[\text{iID}] = \text{init} \rrbracket$ 
50:   return  $\text{key}[\text{iID}]$ 
51: function  $\mathcal{O}_{\text{RevLSK}}(u)$ 
52:   require  $\llbracket u \in \mathcal{U}_{\text{Hon}} \rrbracket$ 
53:    $\text{RevLSK}[u] \leftarrow \text{true}$ 
54:   return  $\text{lsk}_u$ 
55: function  $\mathcal{O}_{\text{RevState}}(\text{iID})$ 
56:   require  $\llbracket \text{role}[\text{iID}] = \text{init} \rrbracket \triangleright \text{Responder has no state to reveal}$ 
57:    $\text{RevState}[\text{iID}] \leftarrow \text{true}$ 
58:   return  $\text{state}[\text{iID}]$ 
59: function  $\mathcal{O}_{\text{RevPSK}}(i, j)$ 
60:   require  $\llbracket \text{PSK}[\{i, j\}] \neq \epsilon \rrbracket$ 
61:    $\text{RevPSK}[\{i, j\}] \leftarrow \text{true}$ 
62:   return  $\text{PSK}[\{i, j\}]$ 
63: function  $\mathcal{O}_{\text{IDHIDE}}^{\text{Test}}(i_0, r_0, i_1, r_1)$ 
64:   require  $\llbracket i_0, r_0, i_1, r_1 \in \mathcal{U}_{\text{Hon}} \rrbracket \wedge \llbracket i_0 \neq r_0 \rrbracket \wedge \llbracket i_1 \neq r_1 \rrbracket$ 
65:    $\triangleright \text{Hon. exec. of the protocol}$ 
66:    $(\text{iID}_i, \text{m}_1) \leftarrow \mathcal{O}_{\text{Init}}(i_b, r_b)$ 
67:    $(\text{iID}_r, \text{m}_2) \leftarrow \mathcal{O}_{\text{ResDer}}(r_b, \text{m}_1)$ 
68:    $\mathcal{O}_{\text{InitDer}}(\text{iID}_i, \text{m}_2)$ 
69:    $\mathcal{U}_{\text{challed}} \stackrel{\$}{\leftarrow} \{i_0, r_0, i_1, r_1\}$ 
70:   return  $(\text{iID}_i, \text{iID}_r, \text{m}_1, \text{m}_2)$ 
71: function  $\mathcal{O}_{\text{IDHIDE}}^{\text{KIND}}(\text{iID})$ 
72:   require  $\llbracket \text{iID}^* = \epsilon \rrbracket \wedge \llbracket \text{key}[\text{iID}] \notin \{\epsilon, \perp\} \rrbracket$ 
73:    $\text{iID}^* \leftarrow \text{iID}$ 
74:    $K_0 := \text{key}[\text{iID}]; K_1 \stackrel{\$}{\leftarrow} \mathcal{K}$ 
75:   return  $K_b$ 

```

Figure 4. Oracles used in the security games. Above, if any condition in **require** does not hold, then it immediately returns $\perp$.

holds even if the associated session keys and session states leak, modeling the fact that identity hiding is a property that solely relies on the security of the identity key.

Formally, WGKE is *identity hiding* if $\text{Adv}_{\text{WGKE}, \mathcal{A}}^{\text{IDHIDE}}(1^\lambda) := \left| 2 \cdot \Pr[\text{Game}_{\text{WGKE}, \mathcal{A}}^{\text{IDHIDE}}(1^\lambda) = 1] - 1 \right|$ (cf. Fig. 3) is negligible for all efficient $\mathcal{A}$.

4.2.4. Key Indistinguishability. We define it similarly to prior works [1], [7] which use eCK [62]-style modeling with ephemeral key reveals, with the following differences:

- (i) We fix a subtle bug in the security game by accurately modeling the post-specified peer model via `ResIdentify`.
- (ii) We allow malicious public keys to be registered.
- (iii) We model *weak* forward secrecy, as opposed to *full*.

Weak forward secrecy (wFS) guarantees that the session key remains secure even if the public keys are compromised at a later point *as long as* the adversary is passive during session establishment. In contrast, full forward secrecy (fFS) allows the adversary to actively interfere with the session. [1] proved PQ WireGuard satisfies fFS by adding an explicit key-confirmation message from the initiator to the responder, making the protocol three-rounds in total.

In our work, we keep our analysis consistent with WireGuard [3], without running such explicit key-confirmation (though WireGuard defers responder traffic until it has received an initiator message serving as key-confirmation). However, we highlight that we can modularly achieve the same fFS result as [1], [7] since an wFS protocol with implicit authentication can be upgraded to satisfy fFS by adding explicit key-confirmation [14], [63], [64]. The formal definition of our key indistinguishability appears in Sec. A.2.

5. PQ WireGuard from Reinforced KEM

We introduce a new primitive called *reinforced* KEM (RKEM) as a natural building block for many key exchange protocols. We then propose a leaner design of PQ WireGuard compared to Hulsing et al. [1]'s design and prove it secure under our new computational model.

5.1. Introducing Reinforced KEM

An RKEM allows to *reinforce* the security of the public key pk with a *reinforcing* public key rpki. To give more context, one can think of rpki as an ephemeral public key epk used in a key exchange protocol, where recall epk provides additional security when the (long-term) public key is compromised. It is informative to keep this correspondence in mind when reading the requirements of RKEM below.

Definition 17. A *reinforced* KEM (RKEM) with an efficiently sampleable key space $\mathcal{K}$ consists of the following algorithms.

KeyGen $(1^\lambda) \rightarrow (\text{pk}, \text{sk})$: On input the security parameter 1^λ , it outputs a (standard) public and secret key pair (pk, sk) .

RKeyGen $(\text{pk}) \rightarrow (\text{rpki}, \text{rsk})$: On input a public key pk, it outputs a *reinforcing* public and secret key pair $(\text{rpki}, \text{rsk})$. We assume pk defines the set of reinforced public keys $\mathcal{RPK}(\text{pk})$, whose set membership can be efficiently checked.

REnc $(\text{pk}, \text{rpki}) \rightarrow (\text{ct}, \text{ss})$: On input a standard and reinforcing public key pair (pk, rpki) , it outputs a ciphertext ct and a key $\text{ss} \in \mathcal{K}$.

RDec $(\text{sk}, \text{rsk}, \text{ct}) \rightarrow \text{ss}$: On input a standard and reinforcing secret key pair (sk, rsk) and a ciphertext ct, it outputs a key $\text{ss} \in \mathcal{K}$.

Definition 18 (δ -Correctness). We say that RKEM is δ -correct if for any $\lambda \in \mathbb{N}$, we have

$\Pr[ss = \text{RDec}(sk, rsk, ct)] \geq 1 - \delta$, where the probability is taken over $(pk, sk) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$, $(rpk, rsk) \xleftarrow{\$} \text{RKeyGen}(pk)$, and $(ct, ss) \xleftarrow{\$} \text{REnc}(pk, rpk)$.

Also, we require KeyGen (resp. RKeyGen) and REnc to output (resp. reinforcing) public keys and session keys that have high min-entropy under honest execution, respectively.

We further require an RKEM to satisfy two types of CCA-like security, each associated to the standard and reinforced public keys. The first requires a strong level of security for the standard public key, reflecting the fact that reinforcing keys cannot harm the security. In short, even a *maliciously generated* reinforcing public key rpk' should not weaken the security of the standard public key. That is, a ciphertext generated with (pk^*, rpk') remains secure as long as pk^* is not compromised. This holds even giving access to a decryption oracle using sk^* . Moreover, the ciphertext must be simulatable without knowing pk^* ; a property necessary for identity hiding. This notion is akin to SPR-CCA security of standard KEMs [49].

```

1: function  $\text{Game}_{\text{RKEM}, \mathcal{A}}^{\text{SPR-CCA}}(1^\lambda)$ 
2:    $b \xleftarrow{\$} \{0, 1\}$ ;  $L := \emptyset$ ;  $rpk' \leftarrow \perp$ 
3:    $(pk^*, sk^*) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$ 
4:    $(rpk', \text{state}) \xleftarrow{\$} \mathcal{A}^{\text{ORKeyGen}(\cdot), \text{OMalDec}, \perp(\cdot, \cdot)}(pk^*)$ 
5:    $(ct_0, ss_0) \xleftarrow{\$} \text{REnc}(pk^*, rpk')$ 
6:    $(ct_1, ss_1) \xleftarrow{\$} \mathcal{S}(1^\lambda, rpk') \times \mathcal{K}$ 
7:    $b' \xleftarrow{\$} \mathcal{A}^{\text{ORKeyGen}(\cdot), \text{OMalDec}, ct_b(\cdot, \cdot)}(ct_b, ss_b, \text{state})$ 
8:   return  $\llbracket b' = b \rrbracket$ 

9: function  $\mathcal{O}_{\text{RKeyGen}}()$ 
10:   $(rpk, rsk) \leftarrow \text{RKeyGen}(pk^*)$ 
11:   $L \leftarrow L \cup \{(rpk, rsk)\}$ 
12:  return  $(rpk, rsk)$ 

13: function  $\mathcal{O}_{\text{MalDec}, \alpha}(rpk, ct)$ 
14:  if  $\llbracket ct = a \rrbracket \wedge \llbracket rpk = rpk' \rrbracket$  then return  $\perp \triangleright$  trivial attack
15:   $\mathcal{V} := \{rsk \mid (rpk, rsk) \in L\}$ 
16:  if  $\llbracket |\mathcal{V}| \neq 1 \rrbracket$  then return  $\perp \triangleright$  no unique rsk exists
17:   $rsk \leftarrow \mathcal{V} \triangleright \mathcal{V} := \{rsk\}$ 
18:  return  $ss := \text{RDec}(sk^*, rsk, ct)$ 

19: function  $\text{Game}_{\text{RKEM}, \mathcal{A}}^{\text{IND-RCCA}}(1^\lambda)$ 
20:   $b \xleftarrow{\$} \{0, 1\}$ 
21:   $(pk^*, sk^*) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$ 
22:   $(rpk^*, rsk^*) \xleftarrow{\$} \text{RKeyGen}(pk^*)$ 
23:   $(ct^*, ss_0) \xleftarrow{\$} \text{REnc}(pk^*, rpk^*)$ 
24:   $ss_1 \xleftarrow{\$} \mathcal{K}$ 
25:   $b' \xleftarrow{\$} \mathcal{A}^{\text{OHonDec}, ct^*(\cdot)}(pk^*, sk^*, rpk^*, ct^*, ss_b)$ 
26:  return  $\llbracket b' = b \rrbracket$ 

27: function  $\mathcal{O}_{\text{HonDec}, ct^*}(ct)$ 
28:  if  $\llbracket ct = ct^* \rrbracket$  then return  $\perp \triangleright$  trivial attack
29:  return  $ss := \text{RDec}(sk^*, rsk^*, ct)$ 

```

Figure 5. Security games for an RKEM. The core differences between the SPR-CCA security and IND-RCCA security are highlighted in blue.

Definition 19 (SPR-CCA Security). For a simulator $\mathcal{S}$, we define the advantage of $\mathcal{A}$ against the *strong pseudorandom-*

ness under chosen-ciphertext attack (SPR-CCA) security game as follows:

$$\text{Adv}_{\text{RKEM}, \mathcal{A}}^{\text{SPR-CCA}}(1^\lambda) := \left| 2 \cdot \Pr \left[\text{Game}_{\text{RKEM}, \mathcal{A}}^{\text{SPR-CCA}}(1^\lambda) = 1 \right] - 1 \right|,$$

where $\text{Game}_{\text{RKEM}, \mathcal{A}}^{\text{SPR-CCA}}(1^\lambda)$ is defined in Fig. 5. We say RKEM is SPR-CCA *secure* if there exists an efficient simulator $\mathcal{S}$ such that $\text{Adv}_{\text{RKEM}, \mathcal{A}}^{\text{SPR-CCA}}(\lambda) = \text{negl}(\lambda)$ for any efficient adversary $\mathcal{A}$.

The second requires that the reinforcing public key provides additional security in case the *honestly generated* standard public key is compromised. This reflects the reinforcing nature of rpk . We note that we do not require strong pseudorandomness as this is not required by PQ WireGuard, i.e., rpk is not tied to the initiator's identity.

Definition 20 (IND-RCCA Security). We define the advantage of $\mathcal{A}$ against the *indistinguishable against reinforced chosen-ciphertext attack* (IND-RCCA) security game as follows:

$$\text{Adv}_{\text{RKEM}, \mathcal{A}}^{\text{IND-RCCA}}(1^\lambda) := \left| 2 \cdot \Pr \left[\text{Game}_{\text{RKEM}, \mathcal{A}}^{\text{IND-RCCA}}(1^\lambda) = 1 \right] - 1 \right|,$$

where $\text{Game}_{\text{RKEM}, \mathcal{A}}^{\text{IND-RCCA}}(1^\lambda)$ is defined in Fig. 5. We say RKEM is IND-RCCA *secure* if $\text{Adv}_{\text{RKEM}, \mathcal{A}}^{\text{IND-RCCA}}(\lambda) = \text{negl}(\lambda)$ for any efficient adversary $\mathcal{A}$.

Notice that running two standard KEMs in parallel is one naive way to construct an RKEM. That is, RKeyGen outputs a standard KEM key, and REnc generates two KEM ciphertexts with session keys ss_1, ss_2 and sets $ss := H_1(ss_1, H_2(ss_2))$ for some hash functions H_1, H_2 .³ Indeed, this is the implicit RKEM used in a typical key exchange protocol, including PQ WireGuard, where recall rpk act as an ephemeral public key. However, this RKEM is *strictly* more secure than what RKEM requires. Since the two KEMs are symmetrical, it does not exploit the asymmetry in the security strength required by pk and rpk ; that is, no security is required when mixing rpk with a *maliciously generated* pk . Looking ahead, in Sec. 6, we show how to construct a more efficient RKEM that takes advantage of this asymmetry.

5.2. Building PQ WireGuard from RKEM

The protocol description and key schedule of our PQ WireGuard is depicted in the left and right columns of Fig. 1, respectively. We explain our design in a bottom-up fashion.

Similarly to how the “IKpsk2” Noise pattern [4] yields WireGuard, we hope to explain PQ WireGuard as a particular PQ Noise pattern. However, as explained in Sec. 1.1, PQ Noise [13] does not incorporate the later-discovered concerns surrounding KEM binding [2] nor discusses how to use PSKs. To this end, we streamline the design of PQ Noise by slightly changing the key schedule to ensure randomized keys without relying on KEM binding [2]. In short, we require the inclusion of all public keys in the key schedule, where

3. We intentionally compute ss in a cascaded manner to be consistent with the key schedule in (PQ) WireGuard (cf. Fig. 1).

more details are given in Sec. F. This also then matches the behavior for ephemeral public keys in classic Noise extended with PSK, which hashes them to ensure randomized keys; the key computations for PSK itself also follow classic Noise.

At this point, our PQ WireGuard based on KEMs simply follows the “IKpsk1” PQ Noise pattern. The protocol description (i.e., left column of Fig. 1) is almost identical to Hülising et al. [1]’s. The only difference is the key schedule depicted in the right column with green insets, derived naturally from the (streamlined) PQ Noise pattern. Note that it is not the “IKpsk2” pattern as in WireGuard. This is because to authenticate the time message, as identified by [1], we must process the PSK earlier.

Lastly, as the way we chain the two KEM ciphertexts is exactly the naive instantiation of RKEM explained in the previous section, we can simply bundle them together to arrive at our desired PQ WireGuard based on RKEMs. For completeness, we explain how to extend PQ Noise with RKEM in Sec. F.

5.3. Security of Our RKEM-based PQ WireGuard

We prove the security of our PQ WireGuard. The details on key indistinguishability are provided in Sec. C; the proof follows closely to [1], [7]. Below, we first show the authentication guarantee.

Theorem 2. *Our PQ WireGuard satisfies match soundness if KDF is collision-resistant and RKEM outputs reinforcing public and session keys with high min-entropy.*

Proof. Since we assume a collision-resistant KDF, the keys pk_S, rp_k, pk_C input to the KDF by two instances iID and iID' such that $key[iID] = key[iID'] \notin \{\epsilon, \perp\}$ must be identical. The first property of match soundness then follows from the restriction $role[iID] \neq role[iID']$ for partners. Moreover, the second property in case $role[iID] = resp$ (resp. $role[iID] = init$) immediately follows from the high min-entropy of the RKEM reinforcing public (resp. session) key. $\square$

We next show DoS security, where the proof assumes a feature of WireGuard [3, Sec. 5.1], preventing replay attacks by only accepting monotonically increasing timestamps.

Theorem 3. *Our PQ WireGuard is DoS secure if KDF is PRF secure and AEAD is INT-CTXT secure.*

Proof. Assume an initiator is executed L times and generated $time_i := AEAD.Enc(K_{6,i}, H_{6,i}, t_{now,i})$ for $i \in [L]$. As psk is unknown and KDF is PRF secure, each $K_{6,i}$ is distributed uniformly at random. Then, relying on the INT-CTXT security of the AEAD, no responder can accept more than L (distinct) AEAD outputs. $\square$

Lastly, we show identity hiding. For this, the MAC is assumed to output pseudorandom tags. Otherwise, different sessions between the same parties may become linkable by observing the tags $m1$, keyed by the public keys. As Hülising et al. [1] symbolically analyzed identity hiding on a subprotocol of their PQ WireGuard, this assumption was not

documented. In our work, we rely on the fact that WireGuard uses HMAC [3, Sec. 5.4] which can be viewed as a PRF. Moreover, we assume the hash function H is a random oracle in the proof. While the proof can be completed in the standard model assuming that H is a seeded randomness extractor, we choose to analyze it otherwise as it more accurately models the practical instantiation. Indeed, in WireGuard [3, Sec. 5.4], H is an *unkeyed* BLAKE2s [65].

Theorem 4. *Our PQ WireGuard is identity hiding if the KEM and RKEM are SPR-CCA secure and their (standard) public keys have high min-entropy, AEAD is IND-CPA secure, KDF is dual-PRF secure, MAC is instantiated by an HMAC, and the hash function H is modeled as a random oracle.*

Proof. To prove the theorem, we will show that the responses of Test-queries are independent of the initiator’s and responder’s long-term keys, i.e., the challenge bit b . In the identity hiding game, the adversary can make multiple Test-queries. Let q_T and q_H be the maximum number of Test-queries and random oracle queries made by the adversary. To prove all responses of Test-queries are independent of the challenge bit, we consider the following sequence of games. In the following, let $N = |\mathcal{U}|$ be the number of honest users in the system.

Game₀: This game is identical to the original IDHIDE security game. We thus have

$$\Pr[\text{Game}_0 = 1] = \Pr[\text{Game}_{\text{WGKE}, \mathcal{A}}^{\text{IDHIDE}} = 1].$$

Game_{1,0}: In this game, we change how the KEM ciphertexts and session key (ct_S, ss_S) are generated. When a Test-query is made, if the challenge bit is $b = 0$, a random ss_S and simulated ct_S generated with the simulator $\mathcal{S}(1^\lambda)$ are used to answer the query. We can prove Game₀ and Game_{1,0} are indistinguishable assuming the SPR-CCA security of KEM. To this end, we consider the following sub-hybrids between Game₀ and Game_{1,0}.

Game_{0,t,0}: For the 1st through t -th Test-queries, if the challenge bit b is 0, a random ss_S and simulated ct_S are used to answer the Test-queries. Note that $\text{Game}_{0,0,0} = \text{Game}_0$ and $\text{Game}_{0,q_T,0} = \text{Game}_{1,0}$.

We will construct a reduction $\mathcal{B}_{1,0}$ that breaks the SPR-CCA security of KEM as follows.

$\mathcal{B}_{1,0}$ receives a tuple (pk^*, ct^*, ss^*) from its challenger. Then, $\mathcal{B}_{1,0}$ samples a random $\hat{r} \xleftarrow{\$} [N]$, which is the guess of the responder to the t -th Test-query, and $t \xleftarrow{\$} [q_T]$. Then, it samples the challenge bit b and generates the long-term key pairs as follows. For the user $\hat{r}$, $\mathcal{B}_{1,0}$ uses pk^* as its long-term public key, and for all the other parties, $\mathcal{B}_{1,0}$ computes the long-term key pairs correctly. Finally, $\mathcal{B}_{1,0}$ invokes $\mathcal{A}$ and answers the queries made by $\mathcal{A}$ as follows:

- $\mathcal{O}_{\text{Test}}^{\text{IDHIDE}}(i_0, r_0, i_1, r_1)$: If this is the t -th Test-query and $r_0 = \hat{r}$, $\mathcal{B}_{1,0}$ responds as in the previous game except that (ct^*, ss^*) is used as (ct_S, ss_S) if $b = 0$. Else if this is the t -th Test-query and $r_0 \neq \hat{r}$, $\mathcal{B}_{1,0}$ aborts. Else, $\mathcal{B}_{1,0}$ responds as in the previous game.
- $\mathcal{O}_{\text{RevLSK}}(u)$: If $u = \hat{r}$ is queried, $\mathcal{B}_{1,0}$ aborts. Otherwise, $\mathcal{B}_{1,0}$ responds as in the previous game.

- $\mathcal{O}_{\text{CreatePSK}}(u)$, $\mathcal{O}_{\text{RevPSK}}(u)$, $\mathcal{O}_{\text{RevState}}(\text{ilD})$, $\mathcal{O}_{\text{RevSessKey}}(\text{ilD})$: $\mathcal{B}_{1,0}$ responds as in the previous game.

Finally, $\mathcal{A}$ outputs a guess b' . If $b = b'$, $\mathcal{B}_{1,0}$ outputs 0, else outputs 1 as its guess. If $\mathcal{B}_{1,0}$ does not abort, $\mathcal{B}_{1,0}$ perfectly simulates game $\text{Game}_{0,t-1,0}$ (resp. $\text{Game}_{0,t,0}$) to $\mathcal{A}$ when the challenge $(\text{ct}^*, \text{ss}^*)$ is the ciphertext and real key (resp. a simulated ciphertext and a random key). Also, the probability $\mathcal{B}_{1,0}$ does not abort is $1/N$ because there are N candidates of the tested responder. Thus, we have

$$|\Pr[\text{Game}_0 = 1] - \Pr[\text{Game}_{1,0} = 1]| = Nq_T \cdot \text{Adv}_{\text{KEM}, \mathcal{B}_{1,0}}^{\text{SPR-CCA}}(1^\lambda).$$

Game_{1,1}: In this game, we change how the KEM ciphertexts and session key $(\text{ct}_S, \text{ss}_S)$ are generated. When a Test-query is made, if the challenge bit is $b = 1$, a random ss_S and simulated ct_S generated with the simulator $\mathcal{S}(1^\lambda)$ are used to answer the query. Due to the same argument as in $\text{Game}_{1,0}$, we can prove $\text{Game}_{1,0}$ and $\text{Game}_{1,1}$ are indistinguishable assuming the SPR-CCA security of KEM, and we have

$$|\Pr[\text{Game}_{1,0} = 1] - \Pr[\text{Game}_{1,1} = 1]| \leq Nq_T \cdot \text{Adv}_{\text{KEM}, \mathcal{B}_{1,1}}^{\text{SPR-CCA}}(\lambda).$$

Game_{2,0}: In this game, we change how the RKEM ciphertexts and session key (ct, ss) are generated. When a Test-query is made, if the challenge bit is $b = 0$, a random ss and simulated ct generated with the simulator $\mathcal{S}(1^\lambda, \text{rpk}')$ are used, where rpk' is the reinforcing public key generated during this Test-query. We can prove $\text{Game}_{1,1}$ and $\text{Game}_{2,0}$ are indistinguishable assuming the SPR-CCA security of RKEM. To this end, we consider the following sub-hybrid between $\text{Game}_{1,1}$ and $\text{Game}_{2,0}$.

Game_{1,t,0}: For the 1st through t -th Test-queries, a random ss and simulated ct are used if $b = 0$. Note that $\text{Game}_{1,0,0} = \text{Game}_{1,1}$ and $\text{Game}_{1,q_T,0} = \text{Game}_{2,0}$.

We will construct a reduction $\mathcal{B}_{2,0}$ that breaks the SPR-CCA security of RKEM as follows.

$\mathcal{B}_{2,0}$ receives a challenge public key pk^* from its challenger. $\mathcal{B}_{2,0}$ then samples the challenge bit b , a random $\hat{i} \xleftarrow{\$} [N]$ and $t \xleftarrow{\$} [q_T]$ and generates the long-term key pairs as follows. For the user $\hat{i}$, $\mathcal{B}_{2,0}$ uses pk^* as its long-term public key, and for all the other parties, $\mathcal{B}_{2,0}$ computes the long-term key pairs correctly. Finally, $\mathcal{B}_{2,0}$ invokes $\mathcal{A}$ and answers the queries made by $\mathcal{A}$ as follows:

- $\mathcal{O}_{\text{Test}}^{\text{IDHIDE}}(i_0, r_0, i_1, r_1)$: If this is the t -th Test-query and $i_0 = \hat{i}$, $\mathcal{B}_{2,0}$ responds as in the previous game except that, if the challenge bit is $b = 0$, it generates rpk and submits it to its challenger. It then receives $(\text{ct}^*, \text{ss}^*)$, and sets $(\text{ct}, \text{ss}) := (\text{ct}^*, \text{ss}^*)$ instead of generating them on its own. If this is the t -th Test-query but $i_0 \neq \hat{i}$, $\mathcal{B}_{2,0}$ aborts. Otherwise, $\mathcal{B}_{2,0}$ responds as in the previous game.
- $\mathcal{O}_{\text{RevLSK}}(u)$: If $u = \hat{i}$ is queried, $\mathcal{B}_{2,0}$ aborts. Otherwise, $\mathcal{B}_{2,0}$ responds as in the previous game.
- $\mathcal{O}_{\text{CreatePSK}}(u)$, $\mathcal{O}_{\text{RevPSK}}(u)$, $\mathcal{O}_{\text{RevState}}(\text{ilD})$, $\mathcal{O}_{\text{RevSessKey}}(\text{ilD})$: $\mathcal{B}_{2,0}$ responds as in the previous game.

If $\mathcal{A}$ outputs a guess b' , $\mathcal{B}_{2,0}$ outputs 0 if $b = b'$ and 1 otherwise as its guess. If $\mathcal{B}_{2,0}$ does not abort, $\mathcal{B}_{2,0}$ perfectly simulates game $\text{Game}_{1,t-1,0}$ (resp. $\text{Game}_{1,t,0}$) to $\mathcal{A}$ when the challenge $(\text{ct}^*, \text{ss}^*)$ is the ciphertext and real key (resp. a simulated ciphertext and a random key). Also, the probability

$\mathcal{B}_{2,0}$ does not abort is $1/N$ because there are N candidates of the tested initiator. Thus, we have

$$\begin{aligned} |\Pr[\text{Game}_{1,1} = 1] - \Pr[\text{Game}_{2,0} = 1]| \\ = Nq_T \cdot \text{Adv}_{\text{RKEM}, \mathcal{B}_{2,0}}^{\text{SPR-CCA}}(1^\lambda). \end{aligned}$$

Game_{2,1}: In this game, we change how the RKEM ciphertexts and session key $(\text{ct}_S, \text{ss}_S)$ are generated with the simulator $\mathcal{S}(1^\lambda, \text{rpk}')$. When a Test-query is made, if the challenge bit is $b = 1$, a random ss and simulated ct are used to answer the query. Due to the same argument that prove the indistinguishability of $\text{Game}_{1,1}$ and $\text{Game}_{2,0}$, we can prove $\text{Game}_{2,0}$ and $\text{Game}_{2,1}$ are indistinguishable assuming the SPR-CCA security of RKEM, and we have

$$|\Pr[\text{Game}_{2,0} = 1] - \Pr[\text{Game}_{2,1} = 1]| \leq Nq_T \cdot \text{Adv}_{\text{KEM}, \mathcal{B}_{2,1}}^{\text{SPR-CCA}}(\lambda).$$

Game_{3,0}: In this game, we change how AEAD keys are generated. When a Test-query is made, the game samples random K_{ltk} and C_4 instead of computing them with the KDF. Since ss_S is chosen uniformly at random in the previous game, the dual-PRF security guarantees that $\text{Game}_{2,1}$ and $\text{Game}_{3,0}$ are indistinguishable. Since q_T Test-queries are made, we have

$$|\Pr[\text{Game}_{2,1} = 1] - \Pr[\text{Game}_3 = 1]| \leq q_T \cdot \text{Adv}_{\text{KDF}, \mathcal{B}_{3,0}}^{\text{PRF}}(\lambda).$$

Game_{3,1}: In this game, we change how AEAD keys are generated. When a Test-query is made, the game samples a random C_5 instead of computing it with the KDF. Since C_4 is chosen uniformly at random in the previous game, the dual-PRF security guarantees that $\text{Game}_{3,0}$ and $\text{Game}_{3,1}$ are indistinguishable. Since q_T Test-queries are made, we have

$$|\Pr[\text{Game}_{3,0} = 1] - \Pr[\text{Game}_{3,1} = 1]| \leq q_T \cdot \text{Adv}_{\text{KDF}, \mathcal{B}_{3,1}}^{\text{PRF}}(\lambda).$$

Game_{3,2}: In this game, we change how AEAD keys are generated. When a Test-query is made, the game samples a random key K_{time} instead of computing it with the KDF. Since C_5 is chosen uniformly at random in the previous game, the dual-PRF security guarantees that $\text{Game}_{3,1}$ and $\text{Game}_{3,2}$ are indistinguishable. Since q_T Test-queries are made, we have

$$|\Pr[\text{Game}_{3,1} = 1] - \Pr[\text{Game}_{3,2} = 1]| \leq q_T \cdot \text{Adv}_{\text{KDF}, \mathcal{B}_{3,2}}^{\text{PRF}}(\lambda).$$

Game_{3,3}: In this game, we change how AEAD keys are generated. When a Test-query is made, the game samples a random key K_{empty} instead of computing it with the KDF. Since ss is chosen uniformly at random in this game, the dual-PRF security guarantees that $\text{Game}_{3,2}$ and $\text{Game}_{3,3}$ are indistinguishable. Since q_T Test-queries are made, we have

$$|\Pr[\text{Game}_{3,2} = 1] - \Pr[\text{Game}_{3,3} = 1]| \leq q_T \cdot \text{Adv}_{\text{KDF}, \mathcal{B}_{3,3}}^{\text{PRF}}(\lambda).$$

Game₄: In this game, the AEAD ciphertexts ltk , time , rkem and empty are replaced with a random bit-string of length $|\mathcal{H}(\text{pk}_G)|$, $|t_{\text{now}}|$ and $|\text{ct}|$, respectively. Since the AEAD keys K_{ltk} , K_{time} and K_{empty} are chosen uniformly at random in the previous game, the IND-CPA security guarantees that $\text{Game}_{3,3}$ and Game_4 are indistinguishable. Since there are $3 \cdot q_T$ AEAD ciphertexts, we have

$$|\Pr[\text{Game}_{3,3} = 1] - \Pr[\text{Game}_4 = 1]| \leq 3q_T \cdot \text{Adv}_{\text{AEAD}, \mathcal{B}_4}^{\text{IND-CPA}}(\lambda).$$

Game₅ : In this game, the MAC key $H(1b1_3 || pk_S)$ is replaced with a uniformly random key for all tested responders. Here, note that $1b1_3$ is a fixed value across the system [3, Sec. 5.4]. In the classical ROM, the probability that an adversary notices this change is easily bounded by $q_H \cdot p$, where p is the maximum probability of any outcome of the server's long-term public key pk_S . In the QROM, we can instead use the *semi-classical* One-Way to Hiding [53] to bound this by $4 \cdot q_H \cdot \sqrt{p}$; we note that the semi-classical variant gives tighter bounds compared to Lemma 4 which would give us an additional $\sqrt{q_H}$. Since we assume that the server's long-term public key pk_S has high min-entropy (i.e., $p = \text{negl}(\lambda)$), we have

$$|\Pr[\text{Game}_4 = 1] - \Pr[\text{Game}_5 = 1]| = \text{negl}(\lambda).$$

Game₆: In this game, the MAC key $H(1b1_3 || pk_C)$ is replaced with a uniformly random key for all tested initiators. Following the same argument as in Game₅, we have the following assuming the high min-entropy of the client's long-term public key pk_C :

$$|\Pr[\text{Game}_5 = 1] - \Pr[\text{Game}_6 = 1]| = \text{negl}(\lambda).$$

Game₇: In this game, the MAC tag $m1$ is replaced with a random value. Since we assume MAC is instantiated with HMAC and HMAC is a PRF, the output of MAC is indistinguishable from a random value. Therefore, the two games are indistinguishable. Since at most $2q_T$ MAC tags $m1$ are generated through Test-queries in total, we have

$$|\Pr[\text{Game}_6 = 1] - \Pr[\text{Game}_7 = 1]| \leq 2q_T \cdot \text{Adv}_{\text{HMAC}, B_5}^{\text{PRF}}(\lambda).$$

Game₈: In this game, if the challenged users use a non-zero *cookie* value, the MAC key $H(\text{cookie})$ is replaced with a random key. Recall that when *cookie* is a non-zero value, it has high min-entropy conditioned that the server's long-term public key pk_S has high min-entropy (cf. Sec. 3.1). Therefore, following the same argument as in Game₅, the two games are indistinguishable, and we have

$$|\Pr[\text{Game}_7 = 1] - \Pr[\text{Game}_8 = 1]| = \text{negl}(\lambda).$$

Game₉: In this game, if the challenged users use a non-zero *cookie* value, the MAC tag $m2$ is replaced with a random value. Since we assume MAC is instantiated with HMAC and HMAC is a PRF, the output of MAC is indistinguishable from a random value. Therefore, the two games are indistinguishable. Since at most $2q_T$ MAC tags $m2$ are generated through Test-queries in total, we have

$$|\Pr[\text{Game}_8 = 1] - \Pr[\text{Game}_9 = 1]| \leq 2q_T \cdot \text{Adv}_{\text{HMAC}, B_6}^{\text{PRF}}(\lambda).$$

In Game₉, the responses of Test-queries are independent of the public keys of the initiator and responder, i.e., the challenge bit b . Therefore, $\Pr[\text{Game}_9 = 1] = 1/2$. Combining all arguments together, we obtain

$$\begin{aligned} & \text{Adv}_{\text{WGKE}, A}^{\text{IDHIDE}}(\lambda) \\ & \leq 2Nq_T \cdot (\text{Adv}_{\text{KEM}, B_1}^{\text{SPR-CCA}}(\lambda) + \text{Adv}_{\text{RKEM}, B_1}^{\text{SPR-CCA}}(\lambda)) \\ & + q_T \cdot (4\text{Adv}_{\text{KDF}, B_3}^{\text{PRF}}(\lambda) + 3\text{Adv}_{\text{AEAD}, B_4}^{\text{IND-CPA}}(\lambda)) \\ & + 2q_T \cdot (\text{Adv}_{\text{HMAC}, B_5}^{\text{PRF}}(\lambda) + \text{Adv}_{\text{HMAC}, B_6}^{\text{PRF}}(\lambda)) + \text{negl}(\lambda). \end{aligned}$$

□

We note in Game₈ and Game₉ above, we implicitly restrict the adversary to query the test oracle $\mathcal{O}_{\text{Test}}^{\text{IDHIDE}}$ so that *cookie* is either a fixed string or not for *both* challenged initiators, as otherwise, there would be a trivial distinguishing attack observing $m2$.

6. Generic Construction of Reinforced KEM

We show how to construct the desired CCA secure RKEM by applying the Fujisaki-Okamoto (FO) transform [17] to a specific one-way (OW) secure PKE called a *reinforced double-message* PKE (RPKE₂).⁴ Once this “right” definition of RPKE₂ is set in place, we can naturally apply the FO transform, with a security proof following closely to those used by standard KEMs. Notably, the security proof is given in the QROM, building on the recent progress in FO-based KEMs, e.g., [49], [51], [66], [67].

6.1. Reinforced Double-Message PKE

Similarly to an RKEM, an RPKE₂ comes with a reinforcing public key rpk to strengthen the security of the standard public key pk . The main difference is that RPKE₂ is capable of encrypting and decrypting *two* messages, each corresponding to pk and rpk , as opposed to a single session key as in RKEM. We further require an RPKE₂ to come with a *half* decryption algorithm RHalfDec_2 , allowing to decrypt the message for pk . More formally, we have the following.

Definition 21. A *reinforced double-message* PKE (RPKE₂) with message space $\mathcal{M}$ consists of the following algorithms. **KeyGen₂**(1^λ) $\rightarrow$ (pk, sk): On input the security parameter 1^λ , it outputs a (standard) public and secret key pair (pk, sk).

RKeyGen₂(pk) $\rightarrow$ (rpk, rsk): On input a public key pk , it outputs a *reinforcing* public and secret key pair (rpk, rsk). We assume pk defines the set of reinforced public keys $\mathcal{RPK}(pk)$, whose set membership can be efficiently checked.

REnc₂(pk, rpk, μ, μ_r) $\rightarrow$ ct : On input a standard and reinforcing public key pair (pk, rpk) and two messages $(\mu, \mu_r) \in \mathcal{M} \times \mathcal{M}$, it outputs a ciphertext ct .

RHalfDec₂(sk, ct) $\rightarrow$ μ : On input a standard secret key sk and a ciphertext ct , it outputs a message $\mu \in \mathcal{M}$.

RDec₂(sk, rsk, ct) $\rightarrow$ (μ, μ_r): On input a standard and reinforcing secret key pair (sk, rsk) and a ciphertext ct , it outputs two messages $(\mu, \mu_r) \in \mathcal{M} \times \mathcal{M}$.

We require a RPKE₂ to be *correct* with overwhelming probability over the random choice of the keys and encryption algorithm.

Definition 22 (δ_c -Correctness). We say that RPKE₂ is δ_c -correct if for any $\lambda \in \mathbb{N}$, the following is greater than equal to $1 - \delta$:

$$\text{Exp} \left[\max_{(\mu, \mu_r) \in \mathcal{M} \times \mathcal{M}} \Pr[(\mu, \mu_r) = \text{RDec}_2(sk, rsk, \text{REnc}_2(pk, rpk, \mu, \mu_r))] \right],$$

4. While the RPKE₂ requires some level of indistinguishability flavor, we call it one-way secure to be consistent with prior works on FO transforms.

where the expectation is taken over the randomness of sampling $(pk, sk) \xleftarrow{\$} \text{KeyGen}_2(1^\lambda)$ and $(rp_k, rsk) \xleftarrow{\$} \text{RKeyGen}_2(pk)$.

We also require it to be *half-correct*, that is, even for a possibly *maliciously generated* reinforcing public key rp_k , we have $\mu = \text{RHalfDec}_2(sk, \text{REnc}_2(pk, rp_k, \mu, \mu_r))$ with overwhelming probability. Looking ahead, this property is crucial for proving the security of the FO transform. More formally, we have the following.

Definition 23 (δ_{hc} -Half Correctness). We say that RPKE_2 is δ_{hc} -half-correct if for any $\lambda \in \mathbb{N}$, the following is greater than or equal to $1 - \delta_{hc}$:

$$\mathbb{E} \left[\max_{\substack{(\mu, \mu_r) \in \mathcal{M} \times \mathcal{M} \\ rp_k \in \mathcal{RPK}(\text{pk})}} \Pr [\mu = \text{RHalfDec}_2(sk, \text{REnc}_2(pk, rp_k, \mu, \mu_r))] \right],$$

where the expectation is taken over the randomness of sampling $(pk, sk) \xleftarrow{\$} \text{KeyGen}_2(1^\lambda)$.

We further require RPKE_2 to be OW secure. An adversary given the secret key of either the standard or reinforcing public keys should not be able to find the message associated to the public key it does not know.

<pre> 1: function Adv^{OW-CPA}_{RPKE₂, A}(1^λ) 2: (pk*, sk*) $\xleftarrow{\\$}$ KeyGen₂(1^λ) 3: (rp_k', μ_r', state) $\xleftarrow{\\$}$ A(pk*) 4: μ* $\xleftarrow{\\$}$ U(M) 5: ct $\xleftarrow{\\$}$ REnc₂(pk*, rp_k', μ*, μ_r') 6: μ' $\xleftarrow{\\$}$ A(ct, state) 7: return [μ' = μ*] </pre>	<pre> 8: function Adv^{OW-RCPA}_{RPKE₂, A}(1^λ) 9: (pk*, sk*) $\xleftarrow{\\$}$ KeyGen₂(1^λ) 10: (rp_k*, rsk*) $\xleftarrow{\\$}$ RKeyGen₂(pk*) 11: (μ', state) $\xleftarrow{\\$}$ A(pk*, sk*, rp_k*) 12: μ_r' $\xleftarrow{\\$}$ U(M) 13: ct $\xleftarrow{\\$}$ REnc₂(pk*, rp_k*, μ', μ_r') 14: μ_r' $\xleftarrow{\\$}$ A(ct, state) 15: return [μ_r' = μ_r*] </pre>
<pre> 16: function Adv^{C-IND}_{RPKE₂, A}(1^λ) 17: b $\xleftarrow{\\$}$ {0, 1} 18: (pk*, sk*) $\xleftarrow{\\$}$ KeyGen₂(1^λ) 19: (rp_k', μ_r', state) $\xleftarrow{\\$}$ A(pk*) 20: μ* $\xleftarrow{\\$}$ U(M) 21: ct₀ $\xleftarrow{\\$}$ REnc₂(pk*, rp_k', μ*, μ_r') 22: ct₁ $\xleftarrow{\\$}$ S(1^λ, rp_k', μ_r') 23: b' $\xleftarrow{\\$}$ A(ct_b, state) 24: return [b' = b] </pre>	<pre> 25: function Adv^{CR-IND}_{RPKE₂, A}(1^λ) 26: b $\xleftarrow{\\$}$ {0, 1} 27: (pk*, sk*) $\xleftarrow{\\$}$ KeyGen₂(1^λ) 28: (rp_k*, rsk*) $\xleftarrow{\\$}$ RKeyGen₂(pk*) 29: (μ', state) $\xleftarrow{\\$}$ A(pk*, sk*, rp_k*) 30: μ_r' $\xleftarrow{\\$}$ U(M) 31: ct₀ $\xleftarrow{\\$}$ REnc₂(pk*, rp_k*, μ', μ_r') 32: ct₁ $\xleftarrow{\\$}$ S_r(1^λ, pk*, sk*, rp_k*, μ_r') 33: b' $\xleftarrow{\\$}$ A(ct_b, state) 34: return [b' = b] </pre>

Figure 6. Security games for an RPKE_2 . The core differences between the standard and reinforced security notions are highlighted in blue.

Definition 24 (OW-XCPA Security). For $X \in \{\emptyset, R\}$, we define the advantage of $\mathcal{A}$ against the *onewayness under (reinforced) chosen-ciphertext attack* (OW-XCPA) security game as follows:

$$\text{Adv}_{\text{RPKE}_2, \mathcal{A}}^{\text{OW-XCPA}}(1^\lambda) := \Pr [\text{Game}_{\text{RPKE}_2, \mathcal{A}}^{\text{OW-XCPA}}(1^\lambda) = 1],$$

where $\text{Game}_{\text{RPKE}_2, \mathcal{A}}^{\text{OW-XCPA}}(1^\lambda)$ is defined in Fig. 6. We say RPKE_2 is OW-XCPA secure if $\text{Adv}_{\text{RPKE}_2, \mathcal{A}}^{\text{OW-XCPA}}(\lambda) = \text{negl}(\lambda)$ for any efficient adversary $\mathcal{A}$. In the following, we denote OW-XCPA as OW-CPA when $X = \emptyset$.

Lastly, we require an appropriate extension of the *disjoint simulatability* for a standard PKE [51], a useful security property enabling a proof in the QROM, e.g., [48], [49], [66], [67]. In the standard PKE setting, this assumes there exists a simulator on input only a public key, generates a *fake* ciphertext that is indistinguishable from a real ciphertext of a random message. Moreover, a fake ciphertext falls in a valid ciphertext space with negligible probability. For RPKE_2 , we instead require two simulators to capture the fact that the ciphertexts remain indistinguishable against an adversary knowing either the standard or reinforcing secret key.

Definition 25 (Disjoint Simulatability). A RPKE_2 with message space $\mathcal{M}$ is *disjoint simulatable* if there exists efficient simulators $\mathcal{S}$ and $\mathcal{S}_r$ that satisfy the followings:

Statistical disjointness: For any $\lambda \in \mathbb{N}$, $\text{Disj}_{\text{RPKE}_2}(1^\lambda)$ defined below is negligible:

$$\max_{\substack{(pk, sk) \in \text{KeyGen}_2(1^\lambda) \\ (rp_k, \mu_r) \in \mathcal{RPK}(pk) \times \mathcal{M}}} \Pr \left[\begin{array}{l} ct \xleftarrow{\$} \mathcal{S}(1^\lambda, rp_k, \mu_r) : \\ ct \in \text{REnc}_2(pk, rp_k, \mathcal{M}, \mu_r) \end{array} \right].$$

Ciphertext X-indistinguishability: For $X \in \{\emptyset, R\}$, we define the advantage of $\mathcal{A}$ against the *ciphertext (reinforced)-indistinguishable* (CX-IND) security game as follows:

$$\text{Adv}_{\text{RPKE}_2, \mathcal{A}}^{\text{CX-IND}}(1^\lambda) := \left| 2 \cdot \Pr [\text{Game}_{\text{RPKE}_2, \mathcal{A}}^{\text{CX-IND}}(1^\lambda) = 1] - 1 \right|,$$

where $\text{Game}_{\text{RPKE}_2, \mathcal{A}}^{\text{CX-IND}}(1^\lambda)$ is defined in Fig. 6. We say RPKE_2 is CX-IND if $\text{Adv}_{\text{RPKE}_2, \mathcal{A}}^{\text{CX-IND}}(1^\lambda) = \text{negl}(\lambda)$ for any efficient adversary $\mathcal{A}$. In the following, we denote CX-IND as C-IND when $X = \emptyset$.

6.2. OW-Secure RPKE_2 to CCA Secure RKEM

Following [45], [51], we break down the FO transform into two transforms T and U_m^λ . A summary of the transform is given in Fig. 7. Other than the fact that the underlying RPKE_2 comes with two message spaces, the transforms are essentially identical to the standard PKE setting. This enables us to base the security proof of our transform on prior works. Below, we explain in more detail.




Figure 7. Transforming a (probabilistic) OW secure RPKE_2 into a CCA secure RKEM in the QROM. DS stands for disjoint simulatability, which consists of three properties: C-IND and CR-IND security and statistical disjointness. In Sec. 7, we build RPKE_2 from (Hint)-MLWE. The T and U_m^λ transforms are given in Fig. 8 and 9, respectively.

```

1: function  $\overline{\text{KeyGen}}_2(1^\lambda)$ 
2:    $(pk, sk) \xleftarrow{\$} \text{KeyGen}_2(1^\lambda)$ 
3:   return  $pk$ 

4: function  $\overline{\text{RKeyGen}}_2(pk)$ 
5:    $(rpk, rsk) \xleftarrow{\$} \text{RKeyGen}_2(pk)$ 
6:   return  $rpk$ 

7: function  $\overline{\text{REnc}}_2(pk, rpk, \mu, \mu_r)$ 
8:    $r := F(rpk, \mu, \mu_r)$ 
9:    $ct := \text{REnc}_2(pk, rpk, \mu, \mu_r; r)$ 
10:  return  $ct$ 

11: function  $\overline{\text{RDec}}_2(sk, rsk, ct)$ 
12:    $(\mu', \mu'_r) \xleftarrow{\$} \text{RDec}_2(sk, rsk, ct)$ 
13:   return  $(\mu', \mu'_r)$ 

```

Figure 8. The T transform: a deterministic $\overline{\text{RPKE}}_2 := T[\text{RPKE}_2, F]$ from a probabilistic RPKE_2 and a hash function F .

6.2.1. T Transform: Probabilistic $\text{RPKE}_2 \Rightarrow$ Deterministic $\overline{\text{RPKE}}_2$. This transforms a probabilistic RPKE_2 into a deterministic $\overline{\text{RPKE}}_2$, formally depicted in Fig. 8. Derandomization is done by deterministically deriving the encryption randomness via a hash of the reinforcing public key and two messages (cf. Ln. 8).

As in the standard PKE setting, the T transform preserves the CX-IND security of disjoint simulatability (cf. Def. 25). Looking ahead, we directly establish the statistical disjointness of our lattice-based RPKE_2 from the MLWE assumption (cf. Fig. 7). The following theorem establishes the security of the T transform.

Theorem 5. *Let RPKE_2 be a probabilistic RPKE_2 that is CX-IND and OW-XCPA secure for $X \in \{\emptyset, R\}$. Then, $\overline{\text{RPKE}}_2 = T[\text{RPKE}_2, F]$ in Fig. 8 is CX-IND secure for $X \in \{\emptyset, R\}$ in the QROM.*

Formally, for all $X \in \{\emptyset, R\}$, any quantum adversary $\mathcal{A}$ against CX-IND security of $\overline{\text{RPKE}}_2$ making at most Q_F queries to the random oracle F , there exists quantum adversaries $\mathcal{B}_1$ against CX-IND security and $\mathcal{B}_2$ against OW-XCPA security of RPKE_2 such that

$$\text{Adv}_{\overline{\text{RPKE}}_2, \mathcal{A}}^{\text{CX-IND}}(1^\lambda) \leq 2 \cdot Q_F \cdot \sqrt{\text{Adv}_{\text{RPKE}_2, \mathcal{B}_2}^{\text{OW-XCPA}}(1^\lambda)} + \text{Adv}_{\text{RPKE}_2, \mathcal{B}_1}^{\text{CX-IND}}(1^\lambda).$$

Proof Sketch. The proof is almost identical to those in the standard PKE setting [49], [51], [66], [67]. In particular, the oneway to hiding lemma [53] allows us to argue that sampling an encryption randomness r uniformly random is indistinguishable from setting it as $r := F(rpk, \mu, \mu_r)$, assuming the hardness of the OW security of the probabilistic RPKE_2 . The full proof is in Sec. D.1. $\square$

6.2.2. U_m^λ Transform: Deterministic $\overline{\text{RPKE}}_2 \Rightarrow \text{RKEM}$. This transforms a OW secure deterministic $\overline{\text{RPKE}}_2$ into a CCA secure RKEM with *implicit* rejection, formally depicted in Fig. 9. Namely, when the underlying $\overline{\text{RPKE}}_2$ fails the re-encryption check (i.e., decryption fails), we output a random value derived by seeds included in the secret keys, as opposed to outputting a special $\perp$ symbol indicating an explicit decryption failure.

As in the standard PKE setting, the resulting RKEM is SPR-CCA and IND-RCCA secure if the underlying deterministic $\overline{\text{RPKE}}_2$ is disjoint-simulatable.

Theorem 6. *Let RPKE_2 be a probabilistic RPKE_2 . Let $\overline{\text{RPKE}}_2$ be a derandomized $T[\text{RPKE}_2, F]$ in Fig. 8 such that it is δ_{hc} -half and δ_c -correct and disjoint-simulatable. Then, $\text{RKEM} := U_m^\lambda[\overline{\text{RPKE}}_2, H, H_{\text{prf}}]$ in Fig. 9 is δ_c -correct, SPR-CCA, and IND-RCCA secure in the QROM.*

```

1: function  $\text{KeyGen}(1^\lambda)$ 
2:    $(pk, sk') \xleftarrow{\$} \text{KeyGen}_2(1^\lambda)$ 
3:    $s \xleftarrow{\$} \{0, 1\}^\ell$ 
4:    $sk := (sk', s)$ 
5:   return  $(pk, sk)$ 

6: function  $\text{RKeyGen}(pk)$ 
7:    $(rpk, rsk') \xleftarrow{\$} \overline{\text{RKeyGen}}_2(pk)$ 
8:    $s_r \xleftarrow{\$} \{0, 1\}^\ell$ 
9:    $rsk := (rsk', s_r)$ 
10:  return  $(rpk, rsk)$ 

11: function  $\text{REnc}(pk, rpk)$ 
12:    $(\mu, \mu_r) \xleftarrow{\$} \mathcal{M} \times \mathcal{M}$ 
13:    $ct := \overline{\text{REnc}}_2(pk, rpk, \mu, \mu_r)$ 
14:    $ss := H(rpk, \mu, \mu_r)$ 
15:  return  $(ct, ss)$ 

16: function  $\text{RDec}(sk, rsk, ct)$ 
17:  parse  $(sk', s) \leftarrow sk$ 
18:  parse  $(rsk', s_r) \leftarrow rsk$ 
19:   $(\mu', \mu'_r) \xleftarrow{\$}$ 
20:    $\text{RDec}_2(sk', rsk', ct)$ 
21:  if  $\llbracket \mu' = \perp \rrbracket \vee \llbracket \mu'_r = \perp \rrbracket \vee$ 
22:    $\llbracket ct \neq \overline{\text{REnc}}_2(pk, rpk, \mu', \mu'_r) \rrbracket$ 
23:   then
24:     return  $ss :=$ 
25:        $H_{\text{prf}}(s, s_r, rpk, ct)$ 
26:   return  $ss := H(rpk, \mu', \mu'_r)$ 

```

Figure 9. The U_m^λ transform: an $\text{RKEM} := U_m^\lambda[\overline{\text{RPKE}}_2, H, H_{\text{prf}}]$ from a deterministic $\overline{\text{RPKE}}_2 := T[\text{RPKE}_2, F]$ (cf. Fig. 8).

Formally, for any quantum adversary $\mathcal{A}$ against SPR-CCA security of RKEM making at most Q_{dec} , Q_{rpk} queries to the decapsulation and reinforced key generation oracles, respectively, and Q_F , Q_H , $Q_{H_{\text{prf}}}$ queries to the random oracles F , H , H_{prf} , respectively, there exists a quantum adversary $\mathcal{B}$ against CR-IND security of $\overline{\text{RPKE}}_2$ such that

$$\begin{aligned} \text{Adv}_{\text{RKEM}, \mathcal{A}}^{\text{SPR-CCA}}(1^\lambda) &\leq \text{Adv}_{\overline{\text{RPKE}}_2, \mathcal{B}}^{\text{CR-IND}}(1^\lambda) + \text{Disj}_{\overline{\text{RPKE}}_2}(1^\lambda) \\ &\quad + 4 \cdot (\delta_{hc} + Q_{\text{rpk}} \cdot \delta_c) + (Q_{\text{dec}} + Q_{H_{\text{prf}}}) \cdot 2^{-\frac{\ell}{2}+2} \\ &\quad + 16 \cdot (Q_F + Q_{\text{dec}} + 2)^2 \cdot (\delta_{hc} + Q_{\text{rpk}} \cdot \delta_c) \\ &\quad + 16 \cdot (Q_F + Q_H + 2)^2 \cdot (\delta_{hc} + Q_{\text{rpk}} \cdot \delta_c). \end{aligned}$$

Moreover, for any adversary $\mathcal{A}$ against IND-RCCA security of RKEM making at most Q_{dec} queries to the decapsulation oracle and Q_F , Q_H , $Q_{H_{\text{prf}}}$ queries to F , H , H_{prf} , respectively, there exists a quantum adversary $\mathcal{B}$ against C-IND security of $\overline{\text{RPKE}}_2$ such that

$$\begin{aligned} \text{Adv}_{\text{RKEM}, \mathcal{A}}^{\text{IND-RCCA}}(1^\lambda) &\leq \text{Adv}_{\overline{\text{RPKE}}_2, \mathcal{B}}^{\text{C-IND}}(1^\lambda) + \text{Disj}_{\overline{\text{RPKE}}_2}(1^\lambda) \\ &\quad + 4 \cdot \delta_c + (Q_{\text{dec}} + Q_{H_{\text{prf}}}) \cdot 2^{-\frac{\ell}{2}+2} \\ &\quad + 16 \cdot (Q_F + Q_{\text{dec}} + 2)^2 \cdot \delta_c + 16 \cdot (Q_F + Q_H + 2)^2 \cdot \delta_c. \end{aligned}$$

Proof Sketch. We provide the full proof in Sec. D.2 and provide a proof sketch below. Correctness is directly inherited from the underlying $\overline{\text{RPKE}}_2$. Below, we only focus on SPR-CCA security since the proof of IND-RCCA security is analogous. Our high level proof strategy follows those of [49, Theorem 4.1], which is also an extension of those used in [51], [66], [67]. The main difference is that our challenge ciphertext may be generated with respect to a *malicious* reinforcing public key.

We first modify $\overline{\text{RPKE}}_2$ to be *perfectly* correct by tweaking the random oracle F (cf. Fig. 8) to only output randomness that leads to correct (half) decryption. This is possible relying on (half) correctness. Now, let us consider the following two functions:

$$\begin{aligned} g_L(rpk, \mu, \mu_r) &:= (rpk, \overline{\text{REnc}}_2(pk, rpk, \mu, \mu_r)) \\ g_{-L}(rpk, \mu, \mu_r) &:= (rpk, \overline{\text{REnc}}_2(pk, rpk, \mu, \mu_r), \mu_r). \end{aligned}$$

If $rpk \in L$, g_L is injective due to perfect correctness, where recall L is the list of honestly generated rpks (cf. Fig. 5).

TABLE 1. PARAMETERS AND NOTATIONS WHERE ALL THE HASH FUNCTIONS ARE MODELED AS RANDOM ORACLES. SEE SEC. 2.3 FOR THE DEFINITIONS OF Compress_q AND Decompress_q .

Notation	Explanation
(R_q, d)	Polynomial ring with prime modulus q and degree d
(n, k)	Dimension of public matrix $\mathbf{A} \in R_q^{n \times k}$
d_r	Amount of bits after compressing rpk such that $d_r < \lceil \log_2(q) \rceil$
(d_u, d_{v_0}, d_{v_1})	Amount of bits after compressing ct such that $d_a < \lceil \log_2(q) \rceil$ with $a \in \{u, v_0, v_1\}$
(χ, χ_z, χ_r)	Noise distributions over R_q
Compress_q	Rounding operations from Kyber [19]
Decompress_q	
$\lfloor \cdot \rfloor_{q/2}$	Compression function $\lfloor a \rfloor_{q/2} := \text{Compress}_q(a, 1)$
G_{Decomp}	Hash function for randomized decompression
G	Hash function $\text{G} : \{0, 1\}^{256} \rightarrow R_q^{n \times k}$
G_r	Hash function $\text{G}_r : \mathcal{M} \rightarrow \mathbb{Z}_{2^{d_{v_1}}}^d$

However, for a maliciously generated $\text{rpk} \notin L$, g_L may not be injective as correctness is no longer guaranteed. This is why we introduce g_{-L} , a function guaranteed to be injective relying only on perfect *half*-correctness. We can then simulate the quantum random oracle $H(\text{rpk}, \mu, \mu_r)$ as $H_L \circ g_L(\text{rpk}, \mu, \mu_r)$ if $\text{rpk} \in L$ and as $H_{-L} \circ g_{-L}(\text{rpk}, \mu, \mu_r)$ otherwise, where H_L and H_{-L} are two independent random oracles that are only accessible to the adversary via querying H .

Now, assume an adversary queries a challenge ciphertext ct^* on a maliciously generated $\text{rpk}' \in \neg L$; we ignore the case $\text{rpk}' \in L$ as it is trivial. Notice the key is generated as $\text{ss}^* := H_{-L} \circ g_{-L}(\text{rpk}', \mu^*, \mu_r^*) = H_{-L}(\text{rpk}', \text{REnc}_2(\text{pk}, \text{rpk}, \mu^*, \mu_r^*), \mu_r^*) = H_{-L}(\text{rpk}', \text{ct}^*, \mu_r^*)$. Relying on CR-IND security of RPKE_2 , a ct^* is indistinguishable from the simulator's output $\text{ct}^* \xleftarrow{\$} \bar{S}(1^\lambda, \text{rpk}', \mu_r^*)$. Moreover, the simulated ct^* is an invalid ciphertext due to statistical disjointness. Since the adversary can only learn the outputs of H_{-L} on valid ciphertexts, ss^* is information theoretically hidden from the adversary. This completes the proof. $\square$

7. Our Lattice-based Reinforced KEM: Rebar

In this section, we construct a lattice-based OW secure RPKE_2 . The construction is given in Fig. 10, where the notations are summarized in Tab. 1. Our desired CCA secure RKEM, Rebar, is then obtained by applying the T and U_m^x transforms from Sec. 6.2.

At a high-level, our construction resembles the multi-recipient adaptation [18] of ML-KEM [19]. It uses MLWE samples $\mathbf{A} \cdot \mathbf{s} + \mathbf{x}$ for both the public key $\mathbf{b}$ and reinforcing key $\mathbf{b}_r$. Then, it encrypts a pair of messages (μ, μ_r) to the two keys by sampling a fresh MLWE sample $\mathbf{u} = \mathbf{r}^\top \cdot \mathbf{A} + \mathbf{z}^\top$ and two polynomials $v_0 = \mathbf{r}^\top \cdot \mathbf{b} + z + \mu \cdot \lfloor q/2 \rfloor$, $v_1 = \mathbf{r}^\top \cdot \mathbf{b}_r + z_r + \mu_r \cdot \lfloor q/2 \rfloor$. Our security model considers the possibility that the reinforcing key is chosen maliciously, which could undermine security. We resolve this challenge by deriving a uniform vector Δ from μ and masking v_1 with it to prevent extracting information from v_1 without knowing the secret key of $\mathbf{b}$ as well. On top of that, we introduce rounding optimizations, with the novel rounding of the reinforcing public key and reliance on a randomized decompression G_{Decomp} to preserve security.

```

1: function KeyGen2(1λ)
2:   seedA ← {0, 1}256
3:   A := G(seedA) ∈ Rqn × k
4:   (s, x) ← χk × χn
5:   b := A · s + x
6:   pk := (seedA, b)
7:   sk := (seedA, s)
8:   return (pk, sk)

9: function RKeyGen2(pk)
10:  parse (seedA, b) ← pk
11:  A := G(seedA)
12:  (sr, xr) ← χk × χn
13:  br := A · sr + xr
14:  br := Compressq(br, dbr)
15:  return (rpk, rsk) := (br, sr)

16: function REnc2(pk, rpk, μ, μr)
17:  parse ((seedA, b), br) ← (pk, rpk)
18:  A := G(seedA)
19:  br* := GDecomp(br, dbr)
20:  (r, z, zr) ← χn × χk × χz × χr

21:  ▷ Remaining is deterministic ◁
22:  Δ := Gr(μ) ∈ Z2dv1d
23:  u := r⊤ · A + z⊤
24:  v0 := r⊤ · b + z + μ · ⌊q/2⌋
25:  v1 := r⊤ · br*
26:   + zr + μr · ⌊q/2⌋
27:  uu := Compressq(u, du)
28:  v0 := Compressq(v0, dv0)
29:  v1 := Compressq(v1, dv1)
30:  return ct := (uu, v0, v1 ⊕ Δ)

31: function RDec2(sk, rsk, ct)
32:  parse ((seedA, s), sr) ← (sk, rsk)
33:  parse (uu, v0, v1) ← ct
34:  Δ := Gr(μ')
35:  uu := Decompressq(uu, du)
36:  v0 := Decompressq(v0, dv0)
37:  v1 := Decompressq(v1 ⊕ Δ, dv1)
38:  μ' := ⌊v0 - uu⊤ · s⌋q/2
39:  μ'r := ⌊v1 - uu⊤ · sr⌋q/2
40:  return (μ', μ'r)

```

Figure 10. Our lattice-based (probabilistic) RPKE_2 with hash functions G and G_r . RHalfDec_2 is the same as RDec_2 except that it terminates on Ln. 39 and returns μ' .

7.1. OW-(R)CPA Security

OW-CPA and OW-RCPA security of RPKE_2 follows immediately from the C-IND and CR-IND security, respectively, which we show in Sec. 7.2. Below, note that hash functions $\text{G}, \text{G}_r, \text{G}_{\text{Decomp}}$ are all modeled as a (quantum) random oracle.

OW-CPA security.

Theorem 7 (OW-CPA Security). *The (probabilistic) RPKE_2 in Fig. 10 is OW-CPA in the QROM assuming the hardness of the $\text{MLWE}_{q,n,k,\chi,\chi}$ and $\text{MLWE}_{q,k+1,n,\chi_r,(\chi_i)_{i \in [k+1]}}$ assumptions, where $\chi_i := \chi_r$ for all $i \in [k]$ and $\chi_{k+1} := \chi_z$.*

Formally, there exists an efficient simulator S such that for any quantum adversary $\mathcal{A}$ against C-IND security of RPKE_2 making at most Q_G queries to the random oracle $\bar{G}$, there exists quantum adversaries $\mathcal{B}_1$ against the $\text{MLWE}_{q,n,k,\chi,\chi}$ problem and $\mathcal{B}_2, \mathcal{B}_3$ against the $\text{MLWE}_{q,k+1,n,\chi,(\chi_i)_{i \in [k+1]}}$ problem such that

$$\begin{aligned} \text{Adv}_{\text{RPKE}_2, \mathcal{A}}^{\text{C-IND}}(1^\lambda) &\leq \text{Adv}_{\mathcal{B}_1}^{\text{MLWE}}(1^\lambda) + 2 \cdot Q_G \cdot \sqrt{\text{Adv}_{\mathcal{B}_2}^{\text{MLWE}}(1^\lambda) + \frac{1}{|\mathcal{M}|}} \\ &\quad + \text{Adv}_{\mathcal{B}_3}^{\text{MLWE}}(1^\lambda) + \frac{1}{|\mathcal{M}|}. \end{aligned}$$

Proof. The proof follows almost identically to the proof of C-IND security, which we later show in Thm. 9. Concretely, we go through the same game transitions from Game_0 to Game_7 in Thm. 9, where the only modification is that the adversary $\mathcal{A}$ outputs a message μ' , as opposed to a bit b' at the end of the game. In Game_7 , the challenge message μ^* is information theoretically hidden from $\mathcal{A}$, and the advantage of $\mathcal{A}$ is at most $\frac{1}{|\mathcal{M}|}$. Collecting all the bounds, we arrive at the theorem statement. $\square$

OW-CPA security. Below, we rely on the hint-MLWE $_{q,k+1,n,\chi,\chi_z,\mathcal{F}}$ problem, where $\mathcal{F}$ is a distribution over $R_q^{1 \times (k+n+1)}$ that outputs $[\mathbf{x}^\top \mid -\mathbf{s}^\top \mid 0]$ with $(\mathbf{s}, \mathbf{x}) \xleftarrow{\$} \chi^k \times \chi^n$.

Theorem 8 (OW-RCPA Security). *The (probabilistic) RPKE $_2$ in Fig. 10 is OW-RCPA in the QROM assuming the hardness of the hint-MLWE $_{q,k+1,n,\chi,\chi_z,\mathcal{F}}$ assumption.*

Formally, there exists an efficient simulator $\mathcal{S}_r$ such that for any quantum adversary $\mathcal{A}$ against OW-CPA security of RPKE $_2$, there exist quantum adversaries $\mathcal{B}_1$ against the MLWE $_{q,k,n,\chi,\chi}$ problem and $\mathcal{B}_2$ against the hint-MLWE $_{q,k+1,n,\chi_r,\chi_z,\mathcal{F}}$ problem such that

$$\text{Adv}_{\text{RPKE}_2,\mathcal{A}}^{\text{OW-CPA}}(1^\lambda) \leq \text{Adv}_{\mathcal{B}_1}^{\text{MLWE}}(1^\lambda) + \text{Adv}_{\mathcal{B}_2}^{\text{hint-MLWE}}(1^\lambda) + \frac{1}{|\mathcal{M}|}.$$

Proof. The proof follows almost identically to the proof of CR-IND security, which we later show in Thm. 7. Concretely, we go through the same game transitions from Game $_0$ to Game $_5$ in Thm. 7, where the only modification is that the adversary $\mathcal{A}$ outputs a message μ'_r , as opposed to a bit b' at the end of the game. In Game $_5$, the challenge message μ'_r is information theoretically hidden from $\mathcal{A}$, and the advantage of $\mathcal{A}$ is at most $\frac{1}{|\mathcal{M}|}$. Collecting all the bounds, we arrive at the theorem statement. $\square$

To apply the T transform, our lattice-based RPKE $_2$ must satisfy OW-(R)CPA and C(R)-IND security (cf. Def. 24 and 25); see Fig. 7 for the overview. As the proof of OW security follows naturally from the proof of IND security, we refer the formal statement to Sec. E.

7.2. Ciphertext (R)-Indistinguishability

The following two theorems establish the C-IND and CR-IND security, respectively. Below, we rely on the hint-MLWE $_{q,k+1,n,\chi,\chi_z,\mathcal{F}}$ problem as defined in the previous section.

Theorem 9 (C-IND Security). *The (probabilistic) RPKE $_2$ in Fig. 10 is C-IND in the QROM assuming the hardness of the MLWE $_{q,n,k,\chi,\chi}$ and MLWE $_{q,k+1,n,\chi_r,(\chi_i)_{i \in [k+1]}}$ assumptions, where $\chi_i := \chi_r$ for all $i \in [k]$ and $\chi_{k+1} := \chi_z$.*

Formally, there exists an efficient simulator $\mathcal{S}$ such that for any quantum adversary $\mathcal{A}$ against C-IND security of RPKE $_2$ making at most Q_G queries to the random oracle G_r , there exist quantum adversaries $\mathcal{B}_1$, $\mathcal{B}_4$ against the MLWE $_{q,n,k,\chi,\chi}$ problem and $\mathcal{B}_2$, $\mathcal{B}_3$ against the MLWE $_{q,k+1,n,\chi,(\chi_i)_{i \in [k+1]}}$ problem such that

$$\begin{aligned} \text{Adv}_{\text{RPKE}_2,\mathcal{A}}^{\text{C-IND}}(1^\lambda) &\leq \text{Adv}_{\mathcal{B}_1}^{\text{MLWE}}(1^\lambda) + 2 \cdot Q_{G_r} \cdot \sqrt{\text{Adv}_{\mathcal{B}_2}^{\text{MLWE}}(1^\lambda) + \frac{1}{|\mathcal{M}|}} \\ &\quad + \text{Adv}_{\mathcal{B}_3}^{\text{MLWE}}(1^\lambda) + \text{Adv}_{\mathcal{B}_4}^{\text{MLWE}}(1^\lambda). \end{aligned}$$

Theorem 10 (CR-IND Security). *The (probabilistic) RPKE $_2$ in Fig. 10 is CR-IND in the QROM assuming the hardness of the hint-MLWE $_{q,k+1,n,\chi_r,\chi_z,\mathcal{F}}$ assumption.*

Formally, for any quantum adversary $\mathcal{A}$ against CR-IND security of RPKE $_2$, there exist quantum adversaries $\mathcal{B}_1$ and

$\mathcal{B}_3$ against the MLWE $_{q,k,n,\chi,\chi}$ problem, and $\mathcal{B}_2$ against the hint-MLWE $_{q,k+1,n,\chi_r,\chi_z,\mathcal{F}}$ problem such that

$$\begin{aligned} \text{Adv}_{\text{RPKE}_2,\mathcal{A}}^{\text{CR-IND}}(1^\lambda) &\leq \text{Adv}_{\mathcal{B}_1}^{\text{MLWE}}(1^\lambda) \\ &\quad + \text{Adv}_{\mathcal{B}_2}^{\text{hint-MLWE}}(1^\lambda) + \text{Adv}_{\mathcal{B}_3}^{\text{MLWE}}(1^\lambda) + \text{negl}(\lambda). \end{aligned}$$

Proof Sketch. We provide the full proof in Secs. E.1 and E.2 and provide a proof sketch below. Let us first consider C-IND security. We define the simulator $\mathcal{S}$ as follows: on input $(1^\lambda, \text{rpk}, \mu_r)$, it samples random $(\mathbf{u}, v_0, v_1) \xleftarrow{\$} R_q^k \times R_q \times R_q$, and outputs a simulated ciphertext ct as $(\bar{\mathbf{u}}, \bar{v}_0, \bar{v}_1) := (\text{Compress}_q(\mathbf{u}, d_u), \text{Compress}_q(v_0, d_{v_0}), \text{Compress}_q(v_1, d_{v_1}))$. Proving that this is indistinguishable from the real game follows from the key observation that $\bar{v}_1$ is one-time padded with $\Delta = G_r(\mu)$, where μ is encrypted under the standard public key. Notably, while the adversary possesses a reinforcing secret key, $\bar{v}_1$ looks pseudorandom since $(\bar{\mathbf{u}}, \bar{v}_0)$ is pseudorandom under the MLWE assumption.

CR-IND security is quite different since the adversary has the standard secret key $\text{sk}^* = (\text{seed}_A, \mathbf{s})$. Namely, we must prove that μ'_r remains hidden even if it can half decrypt the ciphertext. The key observation is that we can rewrite v_0 (i.e., the component associated to $\mathbf{s}$) as $v_0 - \mu'_r \cdot [q/2] = \mathbf{r}^\top \cdot \mathbf{b} + z = \mathbf{r}^\top \cdot (\mathbf{A} \cdot \mathbf{s} + \mathbf{x}) + z = \mathbf{u} \cdot \mathbf{s} - \mathbf{z}^\top \cdot \mathbf{s} + \mathbf{r}^\top \cdot \mathbf{x} + z$, where μ'_r is the adversarially chosen message. Let us denote the underlined component as the *hint* h . Then the above shows that given $\mathbf{u}$ and h , we can simulate v_0 from $\mathbf{s}$, and this hint h is exactly what the hint-MLWE provides. Notably, given an hint-MLWE instance $(\mathbf{u}, v_1, h)$, the simulator simulates the challenge ciphertext component v_0 using $\mathbf{s}$. As hint-MLWE stipulates that $(\mathbf{u}, v_1)$ is pseudorandom, this completes the proof. Note that G_{Decomp} is used to argue $\mathbf{b}_r^*$ in v_1 is distributed uniformly random, required to correctly invoke hint-MLWE. $\square$

7.3. Remaining Properties

To complete the transform (cf. Fig. 7), it remains to prove that our deterministic $\bar{\text{RPKE}}_2 = \text{T}[\text{RPKE}_2, \text{F}]$ satisfies statistical disjointness and (half)-correctness. As these are simple statistical properties, we refer the details to Secs. E.3 and E.4. By construction, we note that RKEM outputs standard and reinforcing public keys and session keys that have high min-entropy.

8. Implementation and Experiments

We first discuss how we instantiated and implemented our RKEM, Rebar, and our experiments with PQ WireGuard, followed by a discussion of benchmark results. Our implementations of PQ WireGuard, Rebar, and collected measurement data are available at github.com/pqwg/experiment.

8.1. Parameters and Instantiation of Our RKEM

Parameters of Rebar are chosen targeting NIST level I security with a 256-bit message space, overwhelming correctness and statistical disjointness. The hardness of MLWE is evaluated using the *Lattice Estimator* tool [68]. For

TABLE 2. INSTANTIATIONS OF PQ-WIREGUARD. HANDSHAKE TIME IS FOR AN INITIATOR. MESSAGE SIZES EXCLUDE OVERHEAD.

Name	Identity			Per-peer identity storage		Message size		Handshake time
	Client	Server	Ephemeral	Client	Server	Initial	Response	
Classic	X25519	X25519	X25519	32 B	32 B	64 B	64 B	33.602 ms
PQWG [1]	McEliece 460896	McEliece 460896	Dagger [1]	524 160 B	524 160 B	1084 B	1148 B	36.292 ms
PQWG (Ours)	Rebar	McEliece 460896	Rebar	524 160 B	1376 B	1052 B	1088 B	36.887 ms

hint-MLWE, we follow prior works [69]–[71], and use the MLWE parameters provided by its proven reduction for Gaussian secrets [47] with the estimator. Concretely, we bound overwhelmingly the spectral norm of the hint matrix $\mathbf{M}$ in the hint $\mathbf{h} = \mathbf{M}[\mathbf{s}^\top \ \mathbf{x}^\top]^\top + \mathbf{z}$ by a constant B , and compute a reduced MLWE standard deviation $\sigma_{\text{red}}^{-2} = \frac{1}{\sigma^2} + \frac{B^2}{\sigma_r^2}$. In our case, $\mathbf{M}$ is composed of Gaussian sampled polynomials, previously studied in [72]. Note that we can improve sizes and efficiency relying on *sums of uniforms* instead of Gaussians as in [71] while relying on heuristic security of hint-MLWE. As we aimed for more conservative and robust parameters, we chose Gaussians.

We base the implementation of **Rebar** on the reference implementation of ML-KEM from PQClean [73]. Notably we change the NTT algorithm to tackle our different modulus q , and use discrete Gaussian sampling based on reverse cumulative distribution tables (RCDT) as in [74] for χ, χ_r, χ_z . Its implementation remains simpler and easier to verify than NTRU-based KEMs [75], [76]. All secret-dependent operations are constant-time.

We compute the RCDT constants with 72-bit precision as prescribed in [77]. This analysis relies on Rényi divergence which applies to the search games OW-CPA and OW-RCPA. For C-IND and CR-IND, this slightly modifies the distributions used in hint-MLWE and MLWE. As the concrete security of these assumptions mainly relies on the entropy of the input distributions, we argue that this precision does not substantially affect our concrete security.

We compare the performance against reference implementation of ML-KEM-512 used for forward secrecy and ephemeral key exchange in Tab. 3. Note that as RKEM encapsulations target two public keys, KEM.Encaps needs to be called twice for equivalent use.

TABLE 3. PERFORMANCE AND SIZES OF ML-KEM VERSUS REBAR. BENCHMARKS WERE OBTAINED ON AN INTEL CORE-I7 12700H CPU.

ML-KEM-512			Rebar		
Operation	Cycles	Size	Operation	Cycles	Size
KEM.KeyGen	60 364	800 B	KeyGen	239 557	1376 B
			RKeyGen	239 717	864 B
2 × KEM.Encaps	136 660	1536 B	REnc	658 294	1088 B
2 × KEM.Decaps	175 800		RDec	708 649	

8.2. Integration in WireGuard and Experiments

Our implementation of PQ WireGuard is based on the most up-to-date version of the WireGuard kernel module

(Linux 6.14). We modified the protocol messages and cryptographic computations according to our description in Sec. 5.2. We use this implementation for all instantiations (including Classic); the difference in symmetric operations (ChaCha20-Poly1305 for AEAD and Blake2 for hashes) should be negligible compared to asymmetric operations and network overhead. For the KEMs, we use the kernel-ready implementations of Classic McEliece and Dagger from [1] and the implementation of **Rebar** described above. In the implementation of **Rebar** we were careful to limit the amount of stack memory used to fit in kernel restrictions.

For benchmarking, we follow the approach from [23], [34], [78], [79]. We create a virtualized network environment based on Linux network namespaces, set the bandwidth to 1000 Mbps and use the `netem` network emulator to add 30.9 ms of latency as a representation of a reasonably high-quality network link. Measurements are obtained by logging the time it takes to complete the handshake. We changed the rekey interval, so that WireGuard performs a handshake every 5 seconds. We measured 1500 handshakes on an AWS `c5d.4xlarge` instance with 16 cores and 32 GiB of memory, running Amazon Linux 2023 with kernel 6.12.25-32.101. Tab. 2 shows the message sizes and the initiator’s handshake times for different instantiations of WireGuard. The results in the table show that the handshake times for the classically instantiated handshake are slightly faster, but that PQ WireGuard’s performance is similar while requiring much less server memory. RKEM performs slightly slower than Dagger, but it should be noted that the implementations of Dagger and Classic McEliece are heavily optimized using AVX2, unlike the reference implementation of **Rebar**. Further optimizing RKEM remains for future work. Finally, note that using RKEM gives the necessary space to use PQ-X25519 hybrids. This uses an additional 64 B in both messages, bringing them up to 1232 B and 1212 B respectively. The limit for UDP is 1232 B; with McEliece-Dagger, there are less than 64 bytes remaining for hybridization.

Acknowledgements

We thank the anonymous reviewers and shepherd for their comments. We also thank Thomas Prest for sharing his idea for randomizing rounded MLWE samples used as public keys during encryption to improve the parameter selection. This paper is partially based on results obtained from a project, JPNP24003, commissioned by the New Energy and Industrial Technology Development Organization (NEDO). Our implementations are archived at [10.5281/zenodo.17207294](https://doi.org/10.5281/zenodo.17207294).

References

- [1] A. Hülsing, K.-C. Ning, P. Schwabe, F. J. Weber, and P. R. Zimmermann, “Post-quantum WireGuard,” in *2021 IEEE Symposium on Security and Privacy*, IEEE Computer Society Press, May 2021, pp. 304–321. DOI: [10.1109/SP40001.2021.00030](https://doi.org/10.1109/SP40001.2021.00030)
- [2] C. Cremers, A. Dax, and N. Medinger, “Keeping up with the KEMs: Stronger security notions for KEMs and automated analysis of KEM-based protocols,” in *ACM CCS 2024*, ACM, Oct. 2024, pp. 1046–1060. DOI: [10.1145/3658644.3670283](https://doi.org/10.1145/3658644.3670283)
- [3] J. A. Donenfeld, “WireGuard: Next generation kernel network tunnel,” in *NDSS 2017*, Updated version: <https://www.wireguard.com/papers/wireguard.pdf>, The Internet Society, 2017. DOI: [10.14722/ndss.2017.23160](https://doi.org/10.14722/ndss.2017.23160)
- [4] T. Perrin, “Noise protocol framework,” Tech. Rep., Jul. 11, 2018. Accessed: Apr. 2, 2025. [Online]. Available: <https://noiseprotocol.org/noise.html>
- [5] “OpenVPN,” Accessed: May 30, 2025. [Online]. Available: <https://openvpn.net/>
- [6] K. Seo and S. Kent, *Security Architecture for the Internet Protocol*, RFC 4301, Dec. 2005. DOI: [10.17487/RFC4301](https://doi.org/10.17487/RFC4301) [Online]. Available: <https://www.rfc-editor.org/info/rfc4301>
- [7] B. Dowling and K. G. Paterson, “A cryptographic analysis of the WireGuard protocol,” in *ACNS 2018*, ser. LNCS, vol. 10892, Springer, Jul. 2018, pp. 3–21. DOI: [10.1007/978-3-319-93387-0_1](https://doi.org/10.1007/978-3-319-93387-0_1)
- [8] B. Lipp, B. Blanchet, and K. Bhargavan, “A Mechanised Cryptographic Proof of the WireGuard Virtual Private Network Protocol,” Inria Paris, Research Report RR-9269, Apr. 2019, p. 50. [Online]. Available: <https://inria.hal.science/hal-02100345>
- [9] J. A. Donenfeld and K. Milner, “Formal verification of the wireguard protocol,” Jun. 7, 2017. Accessed: Apr. 2, 2025. [Online]. Available: <https://www.wireguard.com/papers/wireguard-formal-verification.pdf>
- [10] NIST, “Module-lattice-based key-encapsulation mechanism standard,” NIST, Tech. Rep., Aug. 13, 2024. DOI: [10.6028/NIST.FIPS.203](https://doi.org/10.6028/NIST.FIPS.203)
- [11] M. R. Albrecht, D. J. Bernstein, T. Chou, C. Cid, J. Gilcher, T. Lange, V. Maram, I. von Maurich, R. Misoczki, R. Niederhagen, K. G. Paterson, E. Persichetti, C. Peters, P. Schwabe, N. Sendrier, J. Szefer, C. J. Tjhai, M. Tomlinson, and W. Wang, “Classic McEliece,” National Institute of Standards and Technology, Tech. Rep., 2022, available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-4-submissions>.
- [12] J.-P. D’Anvers, A. Karmakar, S. S. Roy, F. Vercauteren, J. M. B. Mera, M. Van Beirendonck, and A. Basso, “SABER,” National Institute of Standards and Technology, Tech. Rep., 2020, available at <https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions>.
- [13] Y. Angel, B. Dowling, A. Hülsing, P. Schwabe, and F. J. Weber, “Post quantum noise,” in *ACM CCS 2022*, ACM, Nov. 2022, pp. 97–109. DOI: [10.1145/3548606.3560577](https://doi.org/10.1145/3548606.3560577)
- [14] C. Delpech de Saint Guilhem, M. Fischlin, and B. Warinschi, “Authentication in key-exchange: Definitions, relations and composition,” in *CSF 2020 Computer Security Foundations Symposium*, IEEE Computer Society Press, 2020, pp. 288–303. DOI: [10.1109/CSF49147.2020.00028](https://doi.org/10.1109/CSF49147.2020.00028)
- [15] K. Gellert, K. Gjøsteen, H. Jacobsen, and T. Jager, “On optimal tightness for key exchange with full forward secrecy via key confirmation,” in *CRYPTO 2023*, ser. LNCS, vol. 14084, Springer, Aug. 2023, pp. 297–329. DOI: [10.1007/978-3-031-38551-3_10](https://doi.org/10.1007/978-3-031-38551-3_10)
- [16] E. Kret and R. Schmidt, “The PQXDH key agreement protocol,” Tech. Rep., Oct. 18, 2023. [Online]. Available: <https://signal.org/docs/specifications/pqxdh/>
- [17] E. Fujisaki and T. Okamoto, “Secure integration of asymmetric and symmetric encryption schemes,” in *CRYPTO’99*, ser. LNCS, vol. 1666, Springer, Aug. 1999, pp. 537–554. DOI: [10.1007/3-540-48405-1_34](https://doi.org/10.1007/3-540-48405-1_34)
- [18] S. Katsumata, K. Kwiatkowski, F. Pintore, and T. Prest, “Scalable ciphertext compression techniques for post-quantum KEMs and their applications,” in *ASIACRYPT 2020*, ser. LNCS, vol. 12491, Springer, Dec. 2020, pp. 289–320. DOI: [10.1007/978-3-030-64837-4_10](https://doi.org/10.1007/978-3-030-64837-4_10)
- [19] P. Schwabe, R. Avanzi, J. Bos, L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, J. M. Schanck, G. Seiler, D. Stehlé, and J. Ding, “CRYSTALS-KYBER,” National Institute of Standards and Technology, Tech. Rep., 2022, available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- [20] A. Fujioka, K. Suzuki, K. Xagawa, and K. Yoneyama, “Practical and post-quantum authenticated key exchange from one-way secure key encapsulation mechanism,” in *ASIACCS 13*, ACM Press, May 2013, pp. 83–94. DOI: [10.1145/2484313.2484323](https://doi.org/10.1145/2484313.2484323)
- [21] A. Fujioka, K. Suzuki, K. Xagawa, and K. Yoneyama, “Strongly secure authenticated key exchange from factoring, codes, and lattices,” *DCC*, vol. 76, no. 3, pp. 469–504, 2015. DOI: [10.1007/s10623-014-9972-2](https://doi.org/10.1007/s10623-014-9972-2)
- [22] H. Krawczyk, “SKEME: A versatile secure key exchange mechanism for internet,” in *NDSS’96*, IEEE Computer Society, Feb. 1996, pp. 114–127. DOI: [10.1109/NDSS.1996.492418](https://doi.org/10.1109/NDSS.1996.492418)
- [23] P. Schwabe, D. Stebila, and T. Wiggers, “More efficient post-quantum KEMTLS with pre-distributed public keys,” in *ESORICS 2021*, ser. LNCS, vol. 12972, Springer, Oct. 2021, pp. 3–22. DOI: [10.1007/978-3-030-88418-5_1](https://doi.org/10.1007/978-3-030-88418-5_1)
- [24] T. Perrin, *KEM-based hybrid forward secrecy for noise*, revision 1, unofficial/unstable, Nov. 17, 2018. [Online]. Available: https://github.com/noiseprotocol/noise_hfs_spec/blob/025f0f60cb3b94ad75b68e3a4158b9aac234f8cb/output/noise_hfs.pdf

- [25] K. Varner, B. Lipp, W. Zaeske, L. Schmidt, and P. Dua, “Rosenpass,” 2024-08-28. [Online]. Available: <https://rosenpass.eu/whitepaper.pdf>
- [26] M. Raynal, *PQ-WireGuard: We did it again!* Presentation at NIST PQC Standardization Conference, 2021. Accessed: Apr. 2, 2025. [Online]. Available: <https://csrc.nist.gov/Presentations/2021/pq-wireguard-we-did-it-again>
- [27] J. Ding, M.-S. Chen, A. Petzoldt, D. Schmidt, B.-Y. Yang, M. J. Kannwischer, and J. Patarin, “Rainbow,” National Institute of Standards and Technology, Tech. Rep., 2020, available at <https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions>.
- [28] Mullvad VPN. “Quantum-resistant tunnels with WireGuard,” Accessed: Apr. 2, 2025. [Online]. Available: <https://mullvad.net/en/help/quantum-resistant-tunnels-with-wireguard>
- [29] NordVPN. “Post-quantum encryption explained.” The NordLynx protocol is based on WireGuard, Accessed: Apr. 2, 2025. [Online]. Available: <https://support.nordvpn.com/hc/en-us/articles/30046321712529-NordVPN-Post-quantum-encryption-explained>
- [30] D. Stebila and M. Mosca, “Post-quantum key exchange for the internet and the open quantum safe project,” in *SAC 2016*, ser. LNCS, vol. 10532, Springer, Aug. 2016, pp. 14–37. DOI: [10.1007/978-3-319-69453-5_2](https://doi.org/10.1007/978-3-319-69453-5_2)
- [31] M. van Heesch, N. van Adrichem, T. Attema, and T. Veugen, *Towards quantum-safe VPNs and internet*, Cryptology ePrint Archive, Report 2019/1277, 2019. [Online]. Available: <https://eprint.iacr.org/2019/1277>
- [32] K. Eastbrook, K. Kane, B. LaMacchia, D. Shumow, G. Zaverucha, and C. Paquin. “Post-quantum cryptography VPN,” Accessed: Jun. 1, 2025. [Online]. Available: <https://www.microsoft.com/en-us/research/project/post-quantum-crypto-vpn/>
- [33] E. Rescorla, *The Transport Layer Security (TLS) Protocol Version 1.3*, RFC 8446, Aug. 2018. DOI: [10.17487/RFC8446](https://doi.org/10.17487/RFC8446)
- [34] T. Wiggers, “Post-quantum TLS,” Ph.D. dissertation, Radboud University, NL, Jan. 9, 2024. DOI: [2066/300702](https://doi.org/10.26661/300702)
- [35] C. Kaufman, P. E. Hoffman, Y. Nir, P. Eronen, and T. Kivinen, *Internet Key Exchange Protocol Version 2 (IKEv2)*, RFC 7296, Oct. 2014. DOI: [10.17487/RFC7296](https://doi.org/10.17487/RFC7296) [Online]. Available: <https://www.rfc-editor.org/info/rfc7296>
- [36] S.-L. Gazdag, S. Grundner-Culemann, T. Guggemos, T. Heider, and D. Loebenberger, “A formal analysis of IKEv2’s post-quantum extension,” in *ACM CSAC*, ser. ACSAC ’21, Virtual Event, USA: ACM, 2021, pp. 91–105. DOI: [10.1145/3485832.3485885](https://doi.org/10.1145/3485832.3485885)
- [37] D. Herzinger, S.-L. Gazdag, and D. Loebenberger, “Real-world quantum-resistant ipsec,” in *SIN*, vol. 1, 2021, pp. 1–8. DOI: [10.1109/SIN54109.2021.9699255](https://doi.org/10.1109/SIN54109.2021.9699255)
- [38] P. Kampanakis and G. Ravago, “Post-quantum Hybrid Key Exchange with ML-KEM in the Internet Key Exchange Protocol Version 2 (IKEv2),” IETF, I-D draft-ietf-ipsecme-ikev2-mlkem-00, May 2025. [Online]. Available: <https://datatracker.ietf.org/doc/draft-ietf-ipsecme-ikev2-mlkem/00/>
- [39] T. Reddy.K, V. Smyslov, and S. Fluhrer, “Signature Authentication in the Internet Key Exchange Version 2 (IKEv2) using PQC,” IETF, I-D draft-ietf-ipsecme-ikev2-pqc-auth-02, Apr. 2025. [Online]. Available: <https://datatracker.ietf.org/doc/draft-ietf-ipsecme-ikev2-pqc-auth/02/>
- [40] S. Fluhrer, P. Kampanakis, D. McGrew, and V. Smyslov, *Mixing Preshared Keys in the Internet Key Exchange Protocol Version 2 (IKEv2) for Post-quantum Security*, RFC 8784, Jun. 2020. DOI: [10.17487/RFC8784](https://doi.org/10.17487/RFC8784)
- [41] H. Xue, X. Lu, B. Li, B. Liang, and J. He, “Understanding and constructing AKE via double-key key encapsulation mechanism,” in *ASIACRYPT 2018*, ser. LNCS, vol. 11273, Springer, Dec. 2018, pp. 158–189. DOI: [10.1007/978-3-030-03329-3_6](https://doi.org/10.1007/978-3-030-03329-3_6)
- [42] P. Lafourcade, D. Mahmoud, S. Ruhault, and A. R. Taleb, “A tale of two worlds, a formal story of WireGuard hybridization,” in *USENIX Security 2025*, USENIX Association, Aug. 2025, pp. 4937–4956. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity25/presentation/lafourcade>
- [43] V. Cheval, C. Jacomme, S. Kremer, and R. Künnemann, “SAPIC+: Protocol verifiers of the world, unite!” In *USENIX Security 2022*, USENIX Association, Aug. 2022, pp. 3935–3952. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity22/presentation/cheval>
- [44] H. Beguin, C. Chevalier, G. Lebrun, T. Legavre, T. Ricosset, M. Roméas, and E. Sageloli, “DAKE: Bandwidth-Efficient (U)AKE from Double-KEM,” Preprint, 2025. [Online]. Available: <https://eprint.iacr.org/2025/1755>
- [45] D. Hofheinz, K. Hövelmanns, and E. Kiltz, “A modular analysis of the Fujisaki-Okamoto transformation,” in *TCC 2017*, ser. LNCS, vol. 10677, Springer, Nov. 2017, pp. 341–371. DOI: [10.1007/978-3-319-70500-2_12](https://doi.org/10.1007/978-3-319-70500-2_12)
- [46] D. Kim, D. Lee, J. Seo, and Y. Song, “Toward practical lattice-based proof of knowledge from hint-MLWE,” in *CRYPTO 2023*, ser. LNCS, vol. 14085, Springer, Aug. 2023, pp. 549–580. DOI: [10.1007/978-3-031-38554-4_18](https://doi.org/10.1007/978-3-031-38554-4_18)
- [47] T. Espitau, G. Niot, and T. Prest, “Flood and submerge: Distributed key generation and robust threshold signature from lattices,” in *CRYPTO 2024*, ser. LNCS, vol. 14926, Springer, Aug. 2024, pp. 425–458. DOI: [10.1007/978-3-031-68394-7_14](https://doi.org/10.1007/978-3-031-68394-7_14)
- [48] K. Hövelmanns, E. Kiltz, S. Schäge, and D. Unruh, “Generic authenticated key exchange in the quantum random oracle model,” in *PKC 2020*, ser. LNCS, vol. 12111, Springer, May 2020, pp. 389–422. DOI: [10.1007/978-3-030-45388-6_14](https://doi.org/10.1007/978-3-030-45388-6_14)
- [49] K. Xagawa, “Anonymity of NIST PQC round 3 KEMs,” in *EUROCRYPT 2022*, ser. LNCS, vol. 13277,

- Springer, 2022, pp. 551–581. DOI: [10.1007/978-3-031-07082-2_20](https://doi.org/10.1007/978-3-031-07082-2_20)
- [50] D. Boneh, Ö. Dagdelen, M. Fischlin, A. Lehmann, C. Schaffner, and M. Zhandry, “Random oracles in a quantum world,” in *ASIACRYPT 2011*, ser. LNCS, vol. 7073, Springer, Dec. 2011, pp. 41–69. DOI: [10.1007/978-3-642-25385-0_3](https://doi.org/10.1007/978-3-642-25385-0_3)
 - [51] T. Saito, K. Xagawa, and T. Yamakawa, “Tightly-secure key-encapsulation mechanism in the quantum random oracle model,” in *EUROCRYPT 2018*, ser. LNCS, vol. 10822, Springer, 2018, pp. 520–551. DOI: [10.1007/978-3-319-78372-7_17](https://doi.org/10.1007/978-3-319-78372-7_17)
 - [52] H. Jiang, Z. Zhang, L. Chen, H. Wang, and Z. Ma, “IND-CCA-secure key encapsulation mechanism in the quantum random oracle model, revisited,” in *CRYPTO 2018*, ser. LNCS, vol. 10993, Springer, Aug. 2018, pp. 96–125. DOI: [10.1007/978-3-319-96878-0_4](https://doi.org/10.1007/978-3-319-96878-0_4)
 - [53] A. Ambainis, M. Hamburg, and D. Unruh, “Quantum security proofs using semi-classical oracles,” in *CRYPTO 2019*, ser. LNCS, vol. 11693, Springer, Aug. 2019, pp. 269–295. DOI: [10.1007/978-3-030-26951-7_10](https://doi.org/10.1007/978-3-030-26951-7_10)
 - [54] M. Bellare, “New proofs for NMAC and HMAC: Security without collision-resistance,” in *CRYPTO 2006*, ser. LNCS, vol. 4117, Springer, Aug. 2006, pp. 602–619. DOI: [10.1007/11818175_36](https://doi.org/10.1007/11818175_36)
 - [55] M. Bellare and A. Lysyanskaya, “Symmetric and dual PRFs from standard assumptions: A generic validation of a prevailing assumption,” *Journal of Cryptology*, vol. 37, no. 4, p. 33, Oct. 2024. DOI: [10.1007/s00145-024-09513-6](https://doi.org/10.1007/s00145-024-09513-6)
 - [56] P. Rogaway, “Authenticated-encryption with associated-data,” in *ACM CCS 2002*, ACM, Nov. 2002, pp. 98–107. DOI: [10.1145/586110.586125](https://doi.org/10.1145/586110.586125)
 - [57] A. Fujioka, K. Suzuki, K. Xagawa, and K. Yoneyama, “Strongly secure authenticated key exchange from factoring, codes, and lattices,” in *PKC 2012*, ser. LNCS, vol. 7293, Springer, May 2012, pp. 467–484. DOI: [10.1007/978-3-642-30057-8_28](https://doi.org/10.1007/978-3-642-30057-8_28)
 - [58] T. Jager, E. Kiltz, D. Riepel, and S. Schäge, “Tightly-secure authenticated key exchange, revisited,” in *EUROCRYPT 2021*, ser. LNCS, vol. 12696, Springer, Oct. 2021, pp. 117–146. DOI: [10.1007/978-3-030-77870-5_5](https://doi.org/10.1007/978-3-030-77870-5_5)
 - [59] K. Hashimoto, S. Katsumata, K. Kwiatkowski, and T. Prest, “An efficient and generic construction for Signal’s handshake (X3DH): Post-quantum, state leakage secure, and deniable,” *Journal of Cryptology*, vol. 35, no. 3, p. 17, Jul. 2022. DOI: [10.1007/s00145-022-09427-1](https://doi.org/10.1007/s00145-022-09427-1)
 - [60] J. Brendel, R. Fiedler, F. Günther, C. Janson, and D. Stebila, “Post-quantum asynchronous deniable key exchange and the Signal handshake,” in *PKC 2022*, ser. LNCS, vol. 13178, Springer, Mar. 2022, pp. 3–34. DOI: [10.1007/978-3-030-97131-1_1](https://doi.org/10.1007/978-3-030-97131-1_1)
 - [61] C. Brzuska, C. Cremers, H. Jacobsen, D. Stebila, and B. Warinschi, *Falsifiability, composability, and comparability of game-based security models for key exchange protocols*, Cryptology ePrint Archive, Report 2024/1215, 2024. [Online]. Available: <https://eprint.iacr.org/2024/1215>
 - [62] B. A. LaMacchia, K. Lauter, and A. Mityagin, “Stronger security of authenticated key exchange,” in *ProvSec 2007*, ser. LNCS, vol. 4784, Springer, Nov. 2007, pp. 1–16. DOI: [10.1007/978-3-540-75670-5_1](https://doi.org/10.1007/978-3-540-75670-5_1)
 - [63] M. Bellare, D. Pointcheval, and P. Rogaway, “Authenticated key exchange secure against dictionary attacks,” in *EUROCRYPT 2000*, ser. LNCS, vol. 1807, Springer, May 2000, pp. 139–155. DOI: [10.1007/3-540-45539-6_11](https://doi.org/10.1007/3-540-45539-6_11)
 - [64] H. Krawczyk, “HMQV: A high-performance secure Diffie-Hellman protocol,” in *CRYPTO 2005*, ser. LNCS, vol. 3621, Springer, Aug. 2005, pp. 546–566. DOI: [10.1007/11535218_33](https://doi.org/10.1007/11535218_33)
 - [65] J.-P. Aumasson, S. Neves, Z. Wilcox-O’Hearn, and C. Winnerlein, “BLAKE2: Simpler, smaller, fast as MD5,” in *ACNS 2013*, ser. LNCS, vol. 7954, Springer, Jun. 2013, pp. 119–135. DOI: [10.1007/978-3-642-38980-1_8](https://doi.org/10.1007/978-3-642-38980-1_8)
 - [66] K. Xagawa and T. Yamakawa, “(Tightly) QCCA-secure key-encapsulation mechanism in the quantum random oracle model,” in *PQCrypto 2019*, Springer, 2019, pp. 249–268. DOI: [10.1007/978-3-030-25510-7_14](https://doi.org/10.1007/978-3-030-25510-7_14)
 - [67] X. Liu and M. Wang, “QCCA-secure generic key encapsulation mechanism with tighter security in the quantum random oracle model,” in *PKC 2021*, ser. LNCS, vol. 12710, Springer, May 2021, pp. 3–26. DOI: [10.1007/978-3-030-75245-3_1](https://doi.org/10.1007/978-3-030-75245-3_1)
 - [68] M. R. Albrecht, R. Player, and S. Scott, *On the concrete hardness of learning with errors*, Cryptology ePrint Archive, Report 2015/046, 2015. [Online]. Available: <https://eprint.iacr.org/2015/046>
 - [69] M. F. Esgin, T. Espitau, G. Niot, T. Prest, A. Sakzad, and R. Steinfeld, “Plover: Masking-friendly hash-and-sign lattice signatures,” in *EUROCRYPT 2024*, ser. LNCS, vol. 14657, Springer, May 2024, pp. 316–345. DOI: [10.1007/978-3-031-58754-2_12](https://doi.org/10.1007/978-3-031-58754-2_12)
 - [70] R. Del Pino, S. Katsumata, M. Maller, F. Mouhartem, T. Prest, and M.-J. O. Saarinen, “Threshold raccoon: Practical threshold signatures from standard lattice assumptions,” in *EUROCRYPT 2024*, ser. LNCS, vol. 14652, Springer, May 2024, pp. 219–248. DOI: [10.1007/978-3-031-58723-8_8](https://doi.org/10.1007/978-3-031-58723-8_8)
 - [71] Y. Dodis, D. Jost, S. Katsumata, T. Prest, and R. Schmidt, *Triple ratchet: A bandwidth efficient hybrid-secure signal protocol*, Cryptology ePrint Archive, Paper 2025/078, 2025. [Online]. Available: <https://eprint.iacr.org/2025/078>
 - [72] G. Niot, *Practical deniable post-quantum X3DH: A lightweight split-KEM for K-Waay*, Cryptology ePrint Archive, Paper 2025/853, 2025. [Online]. Available: <https://eprint.iacr.org/2025/853>
 - [73] M. J. Kannwischer, P. Schwabe, D. Stebila, and T. Wiggers, “Improving software quality in cryptography

- standardization projects,” in *IEEE EuroS&P 2022 Workshops*, USA: IEEE, 2022, pp. 19–30. DOI: [10.1109/EuroSPW55150.2022.00010](https://doi.org/10.1109/EuroSPW55150.2022.00010) [Online]. Available: <https://eprint.iacr.org/2022/337>
- [74] T. Prest, P.-A. Fouque, J. Hoffstein, P. Kirchner, V. Lyubashevsky, T. Pornin, T. Ricosset, G. Seiler, W. Whyte, and Z. Zhang, “FALCON,” National Institute of Standards and Technology, Tech. Rep., 2022, available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>.
- [75] P.-A. Fouque, P. Kirchner, T. Pornin, and Y. Yu, “BAT: Small and fast KEM over NTRU lattices,” *IACR TCHES*, vol. 2022, no. 2, pp. 240–265, 2022. DOI: [10.46586/tches.v2022.i2.240-265](https://doi.org/10.46586/tches.v2022.i2.240-265)
- [76] J. Zhang, D. Feng, and D. Yan, “NEV: Faster and smaller NTRU encryption using vector decoding,” in *ASIACRYPT 2023*, ser. LNCS, vol. 14444, Springer, Dec. 2023, pp. 157–189. DOI: [10.1007/978-981-99-8739-9_6](https://doi.org/10.1007/978-981-99-8739-9_6)
- [77] T. Prest, “Sharper bounds in lattice-based cryptography using the Rényi divergence,” in *ASIACRYPT 2017*, ser. LNCS, vol. 10624, Springer, Dec. 2017, pp. 347–374. DOI: [10.1007/978-3-319-70694-8_13](https://doi.org/10.1007/978-3-319-70694-8_13)
- [78] C. Paquin, D. Stebila, and G. Tamvada, “Benchmarking post-quantum cryptography in TLS,” in *PQCrypto 2020*, Springer, 2020, pp. 72–91. DOI: [10.1007/978-3-030-44223-1_5](https://doi.org/10.1007/978-3-030-44223-1_5)
- [79] P. Schwabe, D. Stebila, and T. Wiggers, “Post-quantum TLS without handshake signatures,” in *ACM CCS 2020*, ACM, Nov. 2020, pp. 1461–1480. DOI: [10.1145/3372297.3423350](https://doi.org/10.1145/3372297.3423350)
- [80] C. J. F. Cremers and M. Feltz, “Beyond eCK: Perfect forward secrecy under actor compromise and ephemeral-key reveal,” in *ESORICS 2012*, ser. LNCS, vol. 7459, Springer, Sep. 2012, pp. 734–751. DOI: [10.1007/978-3-642-33167-1_42](https://doi.org/10.1007/978-3-642-33167-1_42)
- [81] M. Zhandry, “Secure identity-based encryption in the quantum random oracle model,” in *CRYPTO 2012*, ser. LNCS, vol. 7417, Springer, Aug. 2012, pp. 758–775. DOI: [10.1007/978-3-642-32009-5_44](https://doi.org/10.1007/978-3-642-32009-5_44)

Appendix A. Omitted Details of Security Model

A.1. Subtle Issue in Prior Computational Model

We provide a self-contained explanation on the subtle issue in how Dowling and Paterson [7] and Hülsing et al. [1] handle the pre-shared keys (PSKs) in the post-specified peer model. This renders their computational model ill-defined, and consequently, making their proofs incomplete.

They base their computational model on Cremers and Feltz [80], allowing to capture post-specified peers, and extends it to handle pre-shared keys. Using their notations, an adversary can query the oracle $\text{Create}(i, j, \text{role})$ on input $\text{role} = \text{resp}$ to set up responder i ’s s -th session⁵ π_i^s , where

5. In our work, we refer to a session as an *instance*.

i ’s intended communication partner $\pi_i^s.\text{pid}$ is set to j . To model post-specified peer, $j = \perp$ is allowed. Moreover, the PSK used by the session π_i^s is set as $\pi_i^s.\text{psk} \leftarrow \text{PSK}_i[j]$, where the content of $\text{PSK}_i[j]$ is set when the adversary queries the oracle $\text{CreatePSK}(i, j)$.

The issue with this is that the oracle $\text{CreatePSK}(i, j)$ implicitly assumes an input $i, j \neq \perp$ since it is supposed to generate a PSK between parties i and j . However, this means that $\text{PSK}_i[j]$ with $j = \perp$ is undefined. As a result, $\text{Create}(i, j, \text{role})$ cannot define $\pi_i^s.\text{psk}$ when $j = \perp$. In words, the responder’s PSK cannot be set when the initiator is unknown, and there is no simple mechanism for the game to update it when the initiator becomes known.

In our work, we fix this subtle issue by the ResIdentify algorithm. The game explicitly runs ResIdentify to identify the initiator, and only then defines the PSK of the session (cf. Fig. 4, Ln. 34). Importantly, the responder’s PSK remains undefined either before executing ResIdentify or ResIdentify fails to identify the initiator.

A.2. Key Indistinguishability

To formalize key indistinguishability, the security game needs to know what secret information are revealed. To do this, the game keeps track of the adversary’s knowledge by managing the following lists. Note that every list except for RevPSK is initialized with false, and set to true when the associated secret information is revealed. This is because pre-shared secrets are optional; RevPSK is initialized with true and only set to false once $\mathcal{O}_{\text{CreatePSK}}$ is queried.

RevSessKey[iID] is true if the session key of instance iID is revealed.

RevIK[u] is true if the long-term secret key of u is revealed.

RevState[iID] is true if the session state of iID is revealed.

RevPSK[{u, v}] becomes false if the pre-shared key for u and v is generated, and is set to true when revealed.

Security covered by our definition. We define key indistinguishability similarly to [1], [7] that uses eCK [62]-style modeling with ephemeral key reveals. The following set of common security guarantees are covered:

Security against Maximal exposure attack: Session keys remains secure against an adversary compromising any non-trivial combinations of the long-term secret key, session state, and pre-shared key. As the responder has no session state, we assume the adversary always compromises the responder’s long-term secret key as long as it does not trivialize the game. Note that this implicitly captures security against *key-compromise impersonation* (KCI) attacks, where an adversary compromising a party’s long-term key and tries to impersonate another party against him.

(Weak) forward security: The session key of a party remains secure against an adversary that does not compromise the long-term secret key of the peer (and the session state when the role of the tested instance is initiator). Note that *full* forward security (ffs) considers the scenario where the secret information of the peer can be compromised *after* the

session key has been computed. As explained in [Sec. 4.2](#), we can achieve fFS by adding explicit key-confirmation.

Formal definition. Key indistinguishability is defined with respect to a predicate $\text{safe}_{\text{WGKE}}$ that checks whether an adversary performed an attack that trivializes the game, where note that the aforementioned attacks are allowed. We require that the session key of the tested instance iID^* is indistinguishable from random if $\text{safe}_{\text{WGKE}}(\text{iID}^*) = \text{true}$.

Definition 26. Let $i^* := \text{Init}[\text{iID}^*]$ and $r^* := \text{Resp}[\text{iID}^*]$ and let $\mathfrak{P}(\text{iID}^*) := \{\text{iID} \in \mathcal{S}_{\text{iID}} \mid \text{Partner}(\text{iID}^*, \text{iID}) = \text{true}\}$ be the partner instances of the tested instance iID^* .

We define the predicate $\text{safe}_{\text{WGKE}}$ such that $\text{safe}_{\text{WGKE}}(\text{iID}^*)$ holds true if and only if $\text{RevSessKey}[\text{iID}^*] = \text{false}$, $\forall \text{iID} \in \mathfrak{P}(\text{iID}^*) : \text{RevSessKey}[\text{iID}] = \text{false}$, and one of the following conditions is satisfied:

Case (1-1): $|\mathfrak{P}(\text{iID}^*)| = 1$ and $\text{RevPSK}[\{i^*, r^*\}] = \text{false}$ (i.e., iID^* has a partner and the used pre-shared key is not compromised);

Case (1-2): $|\mathfrak{P}(\text{iID}^*)| = 1$ and $\text{RevLSK}[i^*] = \text{false}$ (i.e., iID^* has a partner and the initiator's long-term secret key is not compromised);

Case (1-3): $|\mathfrak{P}(\text{iID}^*)| = 1$ and $\text{RevState}[\text{iID}_{i^*}] = \text{false}$ (i.e., iID^* has a partner and the initiator's session state is not compromised);

Case (2-1): $\text{role}[\text{iID}^*] = \text{init}$, $|\mathfrak{P}(\text{iID}^*)| = 0$, $\text{RevLSK}[r^*] = \text{false}$ and $\text{RevState}[\text{iID}_{i^*}] = \text{false}$ (i.e., iID^* has no partner and the responder's long-term secret key and the initiator's session state are not compromised)⁶;

Case (2-2): $\text{role}[\text{iID}^*] = \text{resp}$, $|\mathfrak{P}(\text{iID}^*)| = 0$, $\text{RevLSK}[i^*] = \text{false}$ (i.e., iID^* has no partner and the initiator's long-term secret key is not compromised).

Then, WGKE is *key indistinguishable* for the predicate $\text{safe}_{\text{WGKE}}$ if $\text{Adv}_{\text{WGKE}, \mathcal{A}}^{\text{KIND}}(1^\lambda) := |2 \cdot \Pr[\text{Game}_{\text{WGKE}, \mathcal{A}}^{\text{KIND}}(1^\lambda) = 1] - 1|$ (cf. [Fig. 3](#)) is negligible for all efficient $\mathcal{A}$.

Appendix B.

UKS Attack Without Proper Key Schedule

We provide a simple UKS attack on the PQ WireGuard by Hulsing et al. [1], as depicted in [Fig. 1](#). More formally, we explain an attack that breaks match soundness (cf. [Def. 15](#)). The attack relies on a KEM such that for a (possibly malicious) $(\text{pk}^*, \text{sk}^*)$ and ct generated as $(\text{ct}, \text{ss}) \xleftarrow{\$} \text{KEM}.\text{Enc}(\text{pk}^*)$, it is easy to construct ct' such that $\text{ss} = \text{KEM}.\text{Dec}(\text{sk}', \text{ct}')$ for an honestly generated (pk', sk') . That is, a KEM that is not *LEAK/MAL-BIND-K-PK* [2].

The adversary invokes the initiator and compromises its long-term and ephemeral public keys pk_C and pk_e . It also compromises the long-term public key of the responder pk_S , and PSK if it is used. It then decrypts ct_S , derives the key K_{ltk} , and generates a modified ltk' injecting another long-term public key pk_A , and computes the associated time' . When the server replies, it decrypts ct_A and ct_e , and performs the above “re-encapsulation” attack on KEM to prepare ct_C ,

6. This is the only case where we consider the responder's long-term secret key to be uncompromised.

encapsulating the same secret as ct_A . It then replaces the server message with ct_C and prepares skem' and empty' . As the key chain (C_i) does not take as input pk_C nor ct_C , the client and server derive the same session key K_{out} . However, since the server believes it's talking to $\mathcal{A}$, this breaks the first requirement of match soundness.

We note that in their symbolic model [1], the key chain value C_5 is computed from ct_C instead of ct_e . As the session key is now tied to ct_C , the above attack, injecting ct_A no longer works. Nevertheless, the KEM is still required to be *LEAK/MAL-BIND-K, CT-PK* [2] secure (which is satisfied by how they model KEMs). In our work, we simply include pk_C in the key chain to rely on no extra assumption.

Appendix C.

Key-Indistinguishability of Our PQ WireGuard

In this section, we provide the formal proof of key indistinguishability of our PQ WireGuard based on RKEM.

Theorem 11. *Our PQ WireGuard is key indistinguishable if it is match-sound (cf. [Thm. 2](#)), KEM is SPR-CCA secure, RKEM is SPR-CCA and IND-RCCA secure, and KDF is dual-PRF secure.*

Proof. Let $\mathcal{A}$ be an adversary that plays the security game $\text{Game}_{\mathcal{A}}^{\text{KIND}}(1^\lambda)$ and let ϵ be its advantage. In order to prove the theorem, we show that our PQ WireGuard is secure against all cases listed in [Def. 26](#). For each case, we construct an algorithm that breaks one of the underlying assumptions by using $\mathcal{A}$ as a subroutine. Let $N = |\mathcal{U}|$ be the number of honest users in the system and M be the upper bound of NumiID (i.e., the maximum number of iID generated by $\mathcal{A}$).

We present a security proof structured as a sequence of games. Without loss of generality, we assume that $\mathcal{A}$ always issues $\mathcal{O}_{\text{Test}}$ -query. Regardless of the cases caused by $\mathcal{A}$, all proofs share a common game sequence Game_0 and Game_1 as described below.

Game₀: This game is identical to the original KIND security game. We thus have

$$\Pr[\text{Game}_0 = 1] = \Pr[\text{Game}_{\text{WGKE}, \mathcal{A}}^{\text{KIND}} = 1].$$

Game₁: This game is identical to Game_0 , except that it aborts if the predicate *Match* returns false. Since we proved that WGKE is match-sound, we have

$$|\Pr[\text{Game}_0 = 1] - \Pr[\text{Game}_1 = 1]| \leq \text{Adv}_{\text{WGKE}, \mathcal{A}}^{\text{MATCH}}(\lambda).$$

In the following, we can assume that the tested instance iID^* has a unique partner if it exists.

We now divide the game sequence depending on the case made by the adversary $\mathcal{A}$. Regardless of the cases, we prove that ϵ_1 is negligible, which in particular implies that ϵ_0 is also negligible. Formally, this is shown in [Lemmas 5 to 9](#) provided below. We first complete the proof of the theorem. Specifically, by combining all the lemmas together,

we obtain the following bound, where we take the maximum value inside $\{\dots\}$:

$$\text{Adv}_{\text{WGKE}, \mathcal{A}}^{\text{KIND}}(\lambda) \leq \text{negl}(\lambda) + \text{Adv}_{\text{WGKE}, \mathcal{A}}^{\text{MATCH}}(\lambda) + \left\{ \begin{array}{l} 3N^2 \cdot \text{Adv}_{\mathcal{B}_1}^{\text{PRF}}(\lambda) \\ NM \cdot (\text{Adv}_{\text{RKEM}, \mathcal{B}_2}^{\text{SPR-CCA}}(\lambda) + 2\text{Adv}_{\text{KDF}, \mathcal{B}_2'}^{\text{PRF}}(\lambda)) \\ NM^2 \cdot (\text{Adv}_{\text{RKEM}, \mathcal{B}_3}^{\text{IND-RCCA}}(\lambda) + 2\text{Adv}_{\text{KDF}, \mathcal{B}_3'}^{\text{PRF}}(\lambda)) \\ NM \cdot (\text{Adv}_{\text{RKEM}, \mathcal{B}_4}^{\text{SPR-CCA}}(\lambda) + 4\text{Adv}_{\text{KDF}, \mathcal{B}_4'}^{\text{PRF}}(\lambda)) \\ NM \cdot (\text{Adv}_{\text{RKEM}, \mathcal{B}_5}^{\text{SPR-CCA}}(\lambda) + 2\text{Adv}_{\text{KDF}, \mathcal{B}_5'}^{\text{PRF}}(\lambda)) \end{array} \right\}.$$

□

It remains to prove [Lemmas 5 to 8](#).

Proof of Case (1-1): PSK is not leaked.

Lemma 5. *For any efficient adversary $\mathcal{A}$, there exists a reduction $\mathcal{B}_1$ that breaks the PRF security of KDF such that*

$$\left| \Pr[\text{Game}_1 = 1] - \frac{1}{2} \right| \leq 3M^2 \cdot \text{Adv}_{\text{KDF}, \mathcal{B}_1}^{\text{PRF}}(\lambda).$$

Proof. We present the rest of the sequence of games from game Game_2 .

Game₂: This game guesses two instances that will be tested and be partnered to the tested instance. At the beginning of the game, it chose iID_i and iID_r at random from $[M]$. Let E_{Tested} be the event that the tested instance iID^* satisfies either $\text{iID}^* = \text{iID}_i$ or $\text{iID}^* = \text{iID}_r$ and iID_i and iID_r are partner (W.l.o.g, we can assume iID_i is an initiator and iID_r is a responder). Since the E_{Tested} event can be efficiently checked, the game outputs a random bit as soon as it detects that E_{Tested} does not occur. The probability is $\Pr[E_{\text{Tested}}] = 1/M^2$, so we have

$$\left| \Pr[\text{Game}_2 = 1] - \frac{1}{2} \right| = \frac{1}{M^2} \left| \Pr[\text{Game}_1 = 1] - \frac{1}{2} \right|.$$

Game₃: In this game, when the users $i^* = \text{Init}[\text{iID}^*]$ and $r^* = \text{Resp}[\text{iID}^*]$ compute their session keys, C_6 is replaced with a uniformly random value. If E_{Tested} occurs, since we consider case (1-1), the PSK psk_{i^*, j^*} is not leaked. Further, PSK is chosen uniformly at random by definition. Therefore, due to the PRF security of KDF, there exists a reduction $\mathcal{B}_1$ such that

$$|\Pr[\text{Game}_2 = 1] - \Pr[\text{Game}_3 = 1]| \leq \text{Adv}_{\text{KDF}, \mathcal{B}_1}^{\text{PRF}}(\lambda).$$

Game₄: In this game, when iID_i and iID_r compute their session keys, C_7 is replaced with a uniformly random value. Since C_6 is chosen uniformly at random now, due to the dual-PRF security of KDF, there exists a reduction $\mathcal{B}_1'$ such that

$$|\Pr[\text{Game}_3 = 1] - \Pr[\text{Game}_4 = 1]| \leq \text{Adv}_{\text{KDF}, \mathcal{B}_1'}^{\text{PRF}}(\lambda).$$

Game₅: In this game, iID_i and iID_r replace K_{out} with a uniformly random value. Since C_7 is chosen uniformly at random, due to the dual-PRF security of KDF, there exists a reduction $\mathcal{B}_2'$ such that

$$|\Pr[\text{Game}_4 = 1] - \Pr[\text{Game}_5 = 1]| \leq \text{Adv}_{\text{KDF}, \mathcal{B}_2'}^{\text{PRF}}(\lambda).$$

In game Game_5 , the session key K_{out} is an independent and uniformly random value. Since our PQ WireGuard satisfies match soundness (cf. [Thm. 2](#)), the only instance that share the same K_{out} must be the tested instance and its partner (i.e., iID_i and iID_r conditioning on E_{Tested}). Since the adversary cannot reveal the tested session key and its partner by definition, the tested session key looks random to the adversary regardless of the challenge bit b . Therefore, we conclude that $\Pr[\text{Game}_5 = 1] = 1/2$. Combining all arguments together, we obtain

$$\left| \Pr[\text{Game}_1 = 1] - \frac{1}{2} \right| \leq 3M^2 \cdot 3\text{Adv}_{\text{KDF}, \mathcal{B}_1}^{\text{PRF}}(\lambda).$$

□

Proof of Cases (1-2): initiator's long-term key is not leaked

Lemma 6. *For any efficient adversary $\mathcal{A}$, there exists a reduction $\mathcal{B}_2$ that breaks the SPR-CCA security of RKEM and $\mathcal{B}_2'$ that breaks the PRF security of KDF such that*

$$\left| \Pr[\text{Game}_1 = 1] - \frac{1}{2} \right| \leq NM \cdot (\text{Adv}_{\text{RKEM}, \mathcal{B}_2}^{\text{SPR-CCA}}(\lambda) + 2\text{Adv}_{\text{KDF}, \mathcal{B}_2'}^{\text{PRF}}(\lambda)).$$

Proof. We present the rest of the sequence of games from game Game_2 .

Game₂: This game guesses a responder instance and its initiator, which will be the tested instance and its partner. At the beginning of the game, it chose iID_r at random from $[M]$ and $\hat{i}$ at random from $\mathcal{U}$. Let E_{Tested} be the event that $\text{Init}[\text{iID}^*] = \hat{i}$ and either $\text{iID}_r = \text{iID}^*$ (when $\text{role}[\text{iID}^*] = \text{resp}$) or iID_r is a partner of iID^* (when $\text{role}[\text{iID}^*] = \text{init}$). Since E_{Tested} is an efficiently checkable event, the game outputs a random bit as soon as it detects that event E_{Tested} does not occur. The probability the choice made by $\mathcal{A}$ is correctly guessed with probability $1/NM$, so we have

$$\left| \Pr[\text{Game}_2 = 1] - \frac{1}{2} \right| = \frac{1}{NM} \left| \Pr[\text{Game}_1 = 1] - \frac{1}{2} \right|.$$

Game₃: In this game, we modify how the user $\hat{i}$ computes the session key. Let $(\text{ct}_C, \text{ss}_C) \xleftarrow{\$} \text{REnc}(\text{pk}_{\hat{i}}, \text{rpk}_{\hat{i}})$ be the RKEM key-ciphertext pair generated by the instance iID_r with the user $\hat{i}$'s long-term key $\text{pk}_{\hat{i}}$ and some $\text{rpk}_{\hat{i}}$ generate by the $\hat{i}$'s initiator instance. Then, when iID' such that $\text{Init}[\text{iID}'] = \hat{i}$ and its session state has a secret key of $\text{rpk}_{\hat{i}}$ is invoked with ct_C , it uses the key ss_C that was generated by iID_r instead of decrypting ct_C . If decryption errors do not occur in RKEM, the two games are identical. Hence,

$$|\Pr[\text{Game}_2 = 1] - \Pr[\text{Game}_3 = 1]| = \text{negl}(\lambda).$$

Game₄: In this game, we modify the way the responder instance iID_r computes the RKEM session key and ciphertext. Let rpk be the RKEM reinforcing public key received by iID_r . iID_r samples a random RKEM secret ss_C and simulates ct_C with the simulator $\mathcal{S}(1^\lambda, \text{rpk})$ instead of computing $(\text{ct}_C, \text{ss}_C) \xleftarrow{\$} \text{REnc}(\text{pk}_{\hat{i}}, \text{rpk})$. Note that due to the modification we made in the previous game, when $\hat{i}$ receives ct_C

and decrypts it with rpk 's secret key, it also uses the random key ss_C generated by iID_r . Since we consider the case (1-2), the long-term secret key of the user $\hat{i}$, who is the initiator of the tested instance, is never leaked if E_{Tested} occurs. Thus, due to the SPR-CCA security of RKEM, Game_3 and Game_4 are indistinguishable. More precisely, we can construct an algorithm $\mathcal{B}_2$ breaking the SPR-CCA security of RKEM as follows.

$\mathcal{B}_2$ receives a challenge public key pk^* from its challenger. $\mathcal{B}_2$ then samples a random $(\hat{i}, \text{iID}_r) \xleftarrow{\$} [N] \times [M]$ and generates the long-term key pairs as follows. For the user $\hat{i}$, $\mathcal{B}_2$ uses pk^* as its long-term public key, and for all the other parties, $\mathcal{B}_2$ computes the long-term key pairs correctly. Finally, $\mathcal{B}_2$ invokes $\mathcal{A}$ and answers the queries made by $\mathcal{A}$ as follows:

- $\mathcal{O}_{\text{Init}}(i, r)$: $\mathcal{B}_2$ responds as in the previous game except that it generates a reinforcing public key rpk through $\mathcal{O}_{\text{RKeyGen}}$ oracle.
- $\mathcal{O}_{\text{Res}}(r, m_1)$: Let iID be the assigned iID and rpk be the received reinforcing public key in m_1 .
 - If $\text{Init}[\text{iID}] = \hat{i}$ and $\text{iID} = \text{iID}_r$, then $\mathcal{B}_2$ responds as in the previous game except that it outputs rpk to its challenger, receives $(\text{ct}^*, \text{ss}^*)$, and sets $(\text{ct}_C, \text{ss}_C) := (\text{ct}^*, \text{ss}^*)$ instead of generating them on its own.
 - Otherwise, $\mathcal{B}_2$ responds as in the previous game.
- $\mathcal{O}_{\text{InitDer}}(\text{iID}, m_2)$: Depending on the value of $i = \text{Init}[\text{iID}]$, it performs the following:
 - If $i = \hat{i}$, then $\mathcal{B}_2$ checks if $\text{ct} = \text{ct}^*$. If so, it responds as in the previous game (i.e., it uses $\text{ss} := \text{ss}^*$). Otherwise, if $\text{ct} \neq \text{ct}^*$, then it sends ct and its reinforcing public key rpk to $\mathcal{O}_{\text{MalDec}}$ oracle and receives back ss' . $\mathcal{B}_2$ then responds with ss' .
 - If $i \neq \hat{i}$, then $\mathcal{B}_2$ responds as in the previous game since it knows both the long-term secret key and reinforcing public secret key.
- $\mathcal{O}_{\text{CreatePSK}}(u)$, $\mathcal{O}_{\text{RevLSK}}(u)$, $\mathcal{O}_{\text{RevPSK}}(u)$, $\mathcal{O}_{\text{RegisterPK}}(u)$, $\text{RevState}(\text{iID})$, $\text{RevSessKey}(\text{iID})$, $\text{Test}(\text{iID})$: $\mathcal{B}_2$ responds as in the previous game. Here, note that since we consider case (1-2), $\mathcal{B}_2$ can answer all the $\mathcal{O}_{\text{RevLSK}}$ -query since $\mathcal{A}$ never queries $\text{RevLSK}(\hat{i})$ conditioning on E_{Tested} not occurring, which is the only query that $\mathcal{B}_2$ cannot answer.

If $\mathcal{A}$ outputs a guess b' , $\mathcal{B}_2$ outputs 1 if $b = b'$ and 0 otherwise as its guess. It can be checked that $\mathcal{B}_2$ perfectly simulates game Game_3 (resp. Game_4) to $\mathcal{A}$ when the answer in the challenge phase is the real key and ciphertext (resp. a random key and a simulated ciphertext). Thus, we have

$$|\Pr[\text{Game}_3 = 1] - \Pr[\text{Game}_4 = 1]| \leq \text{Adv}_{\text{RKEM}, \mathcal{B}_2}^{\text{SPR-CCA}}(\lambda).$$

Game₅: We change how iID_r (and its partner instance) computes C_7 . In this game, a uniformly random value is used as C_7 instead of computing it as $\text{KDF}(\text{ss}_C, C_6)$. Due to the modification we made in the previous game, ss_C is chosen uniformly at random. Therefore, by the PRF security of KDF, the two games are indistinguishable. Thus, we have

$$|\Pr[\text{Game}_4 = 1] - \Pr[\text{Game}_5 = 1]| \leq \text{Adv}_{\text{KDF}, \mathcal{B}_2}^{\text{PRF}}(\lambda).$$

Game₆: In this game, a uniformly random value is used as K_{out} . Due to the modification we made in the previous game, C_7 is chosen uniformly at random. Also, since our PQ WireGuard satisfies Match soundness, only iID_i and iID_r have the same C_7 . Therefore, by the dual-PRF security of KDF, the two games are indistinguishable. Thus, we have

$$|\Pr[\text{Game}_5 = 1] - \Pr[\text{Game}_6 = 1]| \leq \text{Adv}_{\text{KDF}, \mathcal{B}_2'}^{\text{PRF}}(\lambda).$$

In game Game_6 , the tested session key K_{out} is an independent and uniformly random value. Following a similar argument for the case (1-1), K_{out} looks random to the adversary regardless of the challenge bit b . Therefore, we conclude that $\Pr[\text{Game}_6 = 1] = 1/2$. Combining all arguments together, we obtain

$$\begin{aligned} & \left| \Pr[\text{Game}_1 = 1] - \frac{1}{2} \right| \\ & \leq NM \cdot (\text{Adv}_{\text{RKEM}, \mathcal{B}_2}^{\text{SPR-CCA}}(\lambda) + 2\text{Adv}_{\text{KDF}, \mathcal{B}_2'}^{\text{PRF}}(\lambda)). \end{aligned}$$

□

Proof of Cases (1-3): initiator's session state is not leaked.

Lemma 7. For any efficient adversary $\mathcal{A}$, there exists a reduction $\mathcal{B}_3$ that breaks the IND-RCCA security of RKEM and $\mathcal{B}_3'$ that breaks the PRF security of KDF such that

$$\begin{aligned} & \left| \Pr[\text{Game}_1 = 1] - \frac{1}{2} \right| \\ & \leq NM^2 \cdot (\text{Adv}_{\text{RKEM}, \mathcal{B}_3}^{\text{IND-RCCA}}(\lambda) + 2\text{Adv}_{\text{KDF}, \mathcal{B}_3'}^{\text{PRF}}(\lambda)). \end{aligned}$$

Proof. We present the rest of the sequence of games from game Game_2 .

Game₂: This game guesses the tested instance and its partner instance. At the beginning of the game, it chooses i at random from $[N]$ and iID_i and iID_r at random from $[M]$. Let E_{Tested} be the event that iID_i and iID_r are partner, $\text{Init}[\text{iID}^*] = i$ and either $\text{iID}_i = \text{iID}^*$ or $\text{iID}_r = \text{iID}^*$. (W.l.o.g., we assume iID_i is initiator and iID_r is a responder) Since E_{Tested} is an efficiently checkable event, the game outputs a random bit as soon as it detects that event E_{Tested} does not occur. The probability the choice made by $\mathcal{A}$ is correctly guessed with probability $1/NM^2$, so we have

$$\left| \Pr[\text{Game}_2 = 1] - \frac{1}{2} \right| = \frac{1}{NM^2} \left| \Pr[\text{Game}_1 = 1] - \frac{1}{2} \right|.$$

Game₃: In this game, we modify how iID_i and iID_r computes the session key. Let $(\text{ct}_C, \text{ss}_C) \xleftarrow{\$} \text{REnc}(\text{pk}_{\hat{i}}, \text{rpk}_{\hat{i}})$ be the RKEM key-ciphertext pair generated by the instance iID_r with the user $\hat{i}$'s long-term key $\text{pk}_{\hat{i}}$ and the reinforcing public key $\text{rpk}_{\hat{i}}$ generated by iID_i . Then, when iID_i is invoked with ct_C , it uses the key ss_C that was generated by iID_r instead of decrypting ct_C . If decryption errors do not occur in RKEM, the two games are identical. Hence,

$$|\Pr[\text{Game}_2 = 1] - \Pr[\text{Game}_3 = 1]| = \text{negl}(\lambda).$$

Game₄: In this game, we modify the way the responder instance iID_r computes the RKEM session key and ciphertext.

iID_r samples a random RKEM secret ss_C instead of ss_C computed as $(ct_C, ss_C) \xleftarrow{\$} \text{REnc}(pk_i, rpki)$, where $rpki$ is the reinforcing public key generated by iID_i . Note that due to the modification we made in the previous game, when iID_i receives ct_C , it also uses the random key ss_C generated by iID_r . Since we consider the case (1-3), the reinforced secret key of $rpki$, which is stored on the session state of iID_i is not leaked if E_{Tested} occurs. Thus, due to the IND-RCCA security of RKEM, Game_3 and Game_4 are indistinguishable. More precisely, we can construct an algorithm $\mathcal{B}_3$ breaking the IND-RCCA security of RKEM as follows.

$\mathcal{B}_3$ receives $(pk^*, sk^*, rpki^*, ct^*, ss^*)$ from its challenger. $\mathcal{B}_3$ then samples a random $(\hat{i}, iID_i, iID_r) \xleftarrow{\$} [N] \times [M]^2$ and generates the long-term key pairs as follows. For the user $\hat{i}$, $\mathcal{B}_3$ uses (pk^*, sk^*) as its long-term public key, and for all the other parties, $\mathcal{B}_3$ computes the long-term key pairs correctly. Finally, $\mathcal{B}_3$ invokes $\mathcal{A}$ and answers the queries made by $\mathcal{A}$ as follows:

- $\mathcal{O}_{\text{Init}}(i, r)$: Let iID be the assigned iID . $\mathcal{B}_3$ responds as in the previous game except that, if $iID = iID_i$, it uses $rpki^*$ as the reinforcing public key.
- $\mathcal{O}_{\text{Res}}(r, m_1)$: Let iID be the assigned iID and $rpki$ be the received reinforcing public key in m_1 .
 - If $\text{Init}[iID] = \hat{i}$ and $iID = iID_r$, then $\mathcal{B}_3$ responds as in the previous game except that it sets $(ct_C, ss_C) := (ct^*, ss^*)$ instead of generating them on its own.
 - Otherwise, $\mathcal{B}_3$ responds as in the previous game.
- $\mathcal{O}_{\text{InitDer}}(iID, m_2)$: Depending on the value of $i = \text{Init}[iID]$, it performs the following:
 - If $iID = iID_i$, then $\mathcal{B}_3$ checks if $ct = ct^*$. If so, it responds as in the previous game (i.e., it uses $ss := ss^*$). Otherwise, if $ct \neq ct^*$, then it sends ct to $\mathcal{O}_{\text{HonDec}}$ oracle and receives back ss' . $\mathcal{B}_3$ then responds with ss' .
 - Otherwise, then $\mathcal{B}_3$ responds as in the previous game since it knows both long-term and reinforcing public secret key.
- $\mathcal{O}_{\text{CreatePSK}}(u)$, $\mathcal{O}_{\text{RevLSK}}(u)$, $\mathcal{O}_{\text{RevPSK}}(u)$, $\mathcal{O}_{\text{RegisterPK}}(u)$, $\text{RevState}(iID)$, $\text{RevSessKey}(iID)$, $\text{Test}(iID)$: $\mathcal{B}_3$ responds as in the previous game. Here, note that since we consider case (1-3), $\mathcal{B}_3$ can answer all the $\mathcal{O}_{\text{RevState}}$ -query since $\mathcal{A}$ never queries $\text{RevState}(iID_i)$ conditioning on E_{Tested} not occurring, which is the only query that $\mathcal{B}_3$ cannot answer.

If $\mathcal{A}$ outputs a guess b' , $\mathcal{B}_3$ outputs 1 if $b = b'$ and 0 otherwise as its guess. It can be checked that $\mathcal{B}_3$ perfectly simulates game Game_3 (resp. Game_4) to $\mathcal{A}$ when the challenged RKEM session key is the real key (resp. a random key). Thus, we have

$$|\Pr[\text{Game}_3 = 1] - \Pr[\text{Game}_4 = 1]| \leq \text{Adv}_{\text{RKEM}, \mathcal{B}_3}^{\text{IND-RCCA}}(\lambda).$$

Game₅: We change how iID_i and iID_r compute C_7 . In this game, we sample a random C_7 instead of computing it as $\text{KDF}(ss_C, C_6)$. Due to the modification we made in the previous game, ss_C is chosen uniformly at random. Therefore, by the PRF security of KDF, the two games

are indistinguishable. Thus, we have

$$|\Pr[\text{Game}_4 = 1] - \Pr[\text{Game}_5 = 1]| \leq \text{Adv}_{\text{KDF}, \mathcal{B}_3'}^{\text{PRF}}(\lambda).$$

Game₆: We change how iID_i and iID_r compute the session key K_{out} . In this game, we sample a random K_{out} instead of computing it as $\text{KDF}(\emptyset, C_7)$. Due to the modification we made in the previous game, C_7 is chosen uniformly at random. Therefore, by the dual-PRF security of KDF, the two games are indistinguishable. Thus, we have

$$|\Pr[\text{Game}_5 = 1] - \Pr[\text{Game}_6 = 1]| \leq \text{Adv}_{\text{KDF}, \mathcal{B}_3''}^{\text{PRF}}(\lambda).$$

In game Game_6 , the tested session key K_{out} is an independent and uniformly random value. Following a similar argument for the case (1-1), K_{out} looks random to the adversary regardless of the challenge bit b . Therefore, we conclude that $\Pr[\text{Game}_6 = 1] = 1/2$. Combining all arguments together, we obtain

$$\begin{aligned} & \left| \Pr[\text{Game}_1 = 1] - \frac{1}{2} \right| \\ & \leq NM^2 \cdot (\text{Adv}_{\text{RKEM}, \mathcal{B}_3}^{\text{IND-RCCA}}(\lambda) + 2\text{Adv}_{\text{KDF}, \mathcal{B}_3'}^{\text{PRF}}(\lambda)). \end{aligned}$$

□

Proof of Case (2-1): responder's long-term key and initiator's session state are not compromised.

Lemma 8. For any efficient adversary $\mathcal{A}$, there exists a reduction $\mathcal{B}_4$ that breaks the SPR-CCA security of KEM and $\mathcal{B}_4'$ that breaks the PRF security of KDF such that

$$\begin{aligned} & \left| \Pr[\text{Game}_1 = 1] - \frac{1}{2} \right| \\ & \leq NM \cdot (\text{Adv}_{\text{KEM}, \mathcal{B}_4}^{\text{SPR-CCA}}(\lambda) + 4\text{Adv}_{\text{KDF}, \mathcal{B}_4'}^{\text{PRF}}(\lambda)). \end{aligned}$$

Proof. We present the rest of the sequence of games from game Game_2 .

Game₂: This game guesses the tested instance and its peer responder. At the beginning of the game, it choses iID_i at random from $[M]$ and $\hat{r}$ at random from $\mathcal{U}$. Let E_{Tested} be the event that $iID_i = iID^*$ and $\text{Resp}[iID^*] = \hat{r}$. Since E_{Tested} is an efficiently checkable event, the game outputs a random bit as soon as it detects that event E_{Tested} does not occur. The probability the choice made by $\mathcal{A}$ is correctly guessed with probability $1/NM$, so we have

$$\left| \Pr[\text{Game}_2 = 1] - \frac{1}{2} \right| = \frac{1}{NM} \left| \Pr[\text{Game}_1 = 1] - \frac{1}{2} \right|.$$

In the following, we can assume iID_i is an initiator since the tested instance is initiator in this case and its peer is $\hat{r}$.

Game₃: In this game, we modify how the user $\hat{r}$ decrypts the KEM ciphertext. Let $(ct_S, ss_S) \xleftarrow{\$} \text{Enc}(pk_{\hat{r}})$ be the KEM key-ciphertext pair generated by some instance iID with the user $\hat{r}$'s long-term key $pk_{\hat{r}}$. Then, when the user $\hat{r}$ receives such a ciphertext ct_S , it uses the corresponding key ss_S that was generated by iID instead of decrypting ct_S on the fly. If decryption errors do not occur in KEM, the two games are identical. Hence,

$$|\Pr[\text{Game}_2 = 1] - \Pr[\text{Game}_3 = 1]| = \text{negl}(\lambda).$$

Game₄: In this game, we modify the way the initiator instance iD_i computes the KEM session key and ciphertext. iD_i samples a random KEM secret ss_S and simulates ct_S with the simulator $\mathcal{S}(1^\lambda)$ instead of computing $(ct_S, ss_S) \xleftarrow{\$} \text{Enc}(pk_{\hat{r}})$. Note that due to the modification we made in the previous game, when $\hat{r}$ receives ct_S and decrypts it, it also uses the random key ss_S generated by iD_i . Since we consider the case (2-1), the long-term secret key of the user $\hat{j}$ is not compromised if E_{Tested} occurs. Thus, due to the SPR-CCA security of KEM, Game₃ and Game₄ are indistinguishable. More precisely, we can construct an algorithm $\mathcal{B}_4$ breaking the SPR-CCA security of KEM as follows.

$\mathcal{B}_4$ receives a tuple of challenge (pk^*, ct^*, ss^*) from its challenger. $\mathcal{B}_4$ then samples a random $(\hat{r}, iD_i) \xleftarrow{\$} [N] \times [M]$ and generates the long-term key pairs as follows. For the user $\hat{r}$, $\mathcal{B}_4$ uses pk^* as its long-term public key, and for all the other parties, $\mathcal{B}_4$ computes the long-term key pairs correctly. Finally, $\mathcal{B}_4$ invokes $\mathcal{A}$ and answers the queries made by $\mathcal{A}$ as follows:

- $\mathcal{O}_{\text{Init}}(i, r)$: Let iD be the assigned iD .
 - If $iD = iD_i$, then $\mathcal{B}_4$ responds as in the previous game except that it sets $(ct_S, ss_S) := (ct^*, ss^*)$ instead of generating them on its own.
 - Otherwise, $\mathcal{B}_4$ responds as in the previous game.
- $\mathcal{O}_{\text{Res}}(r, m_1)$: Let iD be the assigned iD and ct be the KEM ciphertext in m_1 .
 - If $r = \hat{r}$, then $\mathcal{B}_4$ checks if $ct = ct^*$. If so, it responds as in the previous game (i.e., it uses $ss := ss^*$). Otherwise, if $ct \neq ct^*$, then it sends ct to $\mathcal{O}_{\text{Dec}}$ oracle and receives back ss' . $\mathcal{B}_4$ then responds with ss' .
 - Otherwise, $\mathcal{B}_4$ responds as in the previous game.
- $\mathcal{O}_{\text{InitDer}}(iD, m_2)$: $\mathcal{B}_4$ responds as in the previous game.
- $\mathcal{O}_{\text{CreatePSK}}(u)$, $\mathcal{O}_{\text{RevLSK}}(u)$, $\mathcal{O}_{\text{RevPSK}}(u)$, $\mathcal{O}_{\text{RegisterPK}}(u)$, $\text{RevState}(iD)$, $\text{RevSessKey}(iD)$, $\text{Test}(iD)$: $\mathcal{B}_4$ responds as in the previous game. Here, note that since we consider case (2-1), $\mathcal{B}_4$ can answer all the $\mathcal{O}_{\text{RevLSK}}$ -query and $\mathcal{O}_{\text{RevState}}$ -query since $\mathcal{A}$ never queries $\text{RevLSK}(\hat{r})$ and $\text{RevState}(iD_i)$ conditioning on E_{Tested} not occurring, which is the only query that $\mathcal{B}_4$ cannot answer.

If $\mathcal{A}$ outputs a guess b' , $\mathcal{B}_4$ outputs 1 if $b = b'$ and 0 otherwise as its guess. It can be checked that $\mathcal{B}_4$ perfectly simulates game Game₃ (resp. Game₄) to $\mathcal{A}$ when the challenge is the real key and ciphertext (resp. a random key and a simulated ciphertext). Thus, we have

$$|\Pr[\text{Game}_3 = 1] - \Pr[\text{Game}_4 = 1]| \leq \text{Adv}_{\text{KEM}, \mathcal{B}_4}^{\text{SPR-CCA}}(\lambda).$$

Game₅: We change how iD_i computes C_4 . In this game, we replace C_4 with a uniformly random value instead of computing it as $\text{KDF}(ss_S, C_3)$. Due to the modification we made in the previous game, ss_S is chosen uniformly at random. Therefore, by the PRF security of KDF, the two games are indistinguishable. Thus, we have

$$|\Pr[\text{Game}_4 = 1] - \Pr[\text{Game}_5 = 1]| \leq \text{Adv}_{\text{KDF}, \mathcal{B}_4'}^{\text{PRF}}(\lambda).$$

Game₆: We change how iD_i computes C_5 . In this game, we replace C_5 with a uniformly random value instead of

computing it as $\text{KDF}(pk_C, C_4)$. Due to the modification we made in the previous game, C_4 is chosen uniformly at random. Therefore, by the dual-PRF security of KDF, the two games are indistinguishable. Thus, we have

$$|\Pr[\text{Game}_5 = 1] - \Pr[\text{Game}_6 = 1]| \leq \text{Adv}_{\text{KDF}, \mathcal{B}_4''}^{\text{PRF}}(\lambda).$$

Game₇: We change how iD_i computes C_6 . In this game, we replace C_6 with a uniformly random value instead of computing it as $\text{KDF}(psk, C_5)$, where psk is the pre-shared key used in iD_i . Due to the modification we made in the previous game, C_5 is chosen uniformly at random. Therefore, by the dual-PRF security of KDF, the two games are indistinguishable. Thus, we have

$$|\Pr[\text{Game}_6 = 1] - \Pr[\text{Game}_7 = 1]| \leq \text{Adv}_{\text{KDF}, \mathcal{B}_4'''}^{\text{PRF}}(\lambda).$$

Game₇: We change how iD_i computes C_7 . In this game, we replace C_7 with a uniformly random value instead of computing it as $\text{KDF}(ss, C_6)$, where ss is the RKEM session key decrypted by iD_i . Due to the modification we made in the previous game, C_6 is chosen uniformly at random. Also, in case (2-1), the adversary does not know the seed C_6 stored in the session state of $iD_i = iD^*$ since it never queries iD^* to $\mathcal{O}_{\text{RevState}}$. Therefore, by the dual-PRF security of KDF, the two games are indistinguishable. Thus, we have

$$|\Pr[\text{Game}_7 = 1] - \Pr[\text{Game}_8 = 1]| \leq \text{Adv}_{\text{KDF}, \mathcal{B}_4'''}^{\text{PRF}}(\lambda).$$

Game₉: We change how iD_r computes the session key K_{out} . In this game, K_{out} is replaced with a uniformly random value. Due to the modification we made in the previous game, C_7 is chosen uniformly at random. Therefore, by the dual-PRF security of KDF, the two games are indistinguishable. Thus, we have

$$|\Pr[\text{Game}_8 = 1] - \Pr[\text{Game}_9 = 1]| \leq \text{Adv}_{\text{KDF}, \mathcal{B}_4'''}^{\text{PRF}}(\lambda).$$

In game Game₉, the tested session key K_{out} is an independent and uniformly random value. Following a similar argument for the case (1-1), K_{out} looks random to the adversary regardless of the challenge bit b . Therefore, we conclude that $\Pr[\text{Game}_9 = 1] = 1/2$. Combining all arguments together, we obtain

$$\begin{aligned} & \left| \Pr[\text{Game}_1 = 1] - \frac{1}{2} \right| \\ & \leq NM \cdot (\text{Adv}_{\text{KEM}, \mathcal{B}_4}^{\text{SPR-CCA}}(\lambda) + 4\text{Adv}_{\text{KDF}, \mathcal{B}_4'}^{\text{PRF}}(\lambda)). \end{aligned}$$

□

Proof of Cases (2-2): initiator's long-term key is not leaked

Lemma 9. For any efficient adversary $\mathcal{A}$, there exists a reduction $\mathcal{B}_5$ that breaks the SPR-CCA security of RKEM and $\mathcal{B}_5'$ that breaks the PRF security of KDF such that

$$\begin{aligned} & \left| \Pr[\text{Game}_1 = 1] - \frac{1}{2} \right| \\ & \leq NM \cdot (\text{Adv}_{\text{RKEM}, \mathcal{B}_5}^{\text{SPR-CCA}}(\lambda) + 2\text{Adv}_{\text{KDF}, \mathcal{B}_5'}^{\text{PRF}}(\lambda)). \end{aligned}$$

Proof. We present the rest of the sequence of games from game Game_2 .

Game₂: This game guesses a responder instance (which will be tested) and its initiator. At the beginning of the game, it choses iID_r at random from $[M]$ and $\hat{i}$ at random from $\mathcal{U}$. Let E_{Tested} be the event that $\text{Init}[\text{iID}^*] = \hat{i}$ and $\text{iID}_r = \text{iID}^*$. Since E_{Tested} is an efficiently checkable event, the game outputs a random bit as soon as it detects that event E_{Tested} does not occur. The probability the choice made by $\mathcal{A}$ is correctly guessed with probability $1/NM$, so we have

$$\left| \Pr[\text{Game}_2 = 1] - \frac{1}{2} \right| = \frac{1}{NM} \left| \Pr[\text{Game}_1 = 1] - \frac{1}{2} \right|.$$

Game₃: In this game, we modify how the user $\hat{i}$ computes the session key. Let $(\text{ct}_C, \text{ss}_C) \xleftarrow{\$} \text{REnc}(\text{pk}_{\hat{i}}, \text{rpk}_{\hat{i}})$ be the RKEM key-ciphertext pair generated by the instance iID_r with the user $\hat{i}$'s long-term key $\text{pk}_{\hat{i}}$ and some $\text{rpk}_{\hat{i}}$ generate by the $\hat{i}$'s initiator instance. Then, when some iID' such that $\text{Init}[\text{iID}'] = \hat{i}$ and its session state has a secret key of $\text{rpk}_{\hat{i}}$ is invoked with ct_C , it uses the key ss_C that was generated by iID_r instead of decrypting ct_C . If decryption errors do not occur in RKEM, the two games are identical. Hence,

$$|\Pr[\text{Game}_2 = 1] - \Pr[\text{Game}_3 = 1]| = \text{negl}(\lambda).$$

Game₄: In this game, we modify the way the responder instance iID_r computes the RKEM session key and ciphertext. Let rpk be the RKEM reinforcing public key received by iID_r . iID_r samples a random RKEM secret ss_C and simulates ct_C with the simulator $\mathcal{S}(1^\lambda, \text{rpk})$ instead of computing $(\text{ct}_C, \text{ss}_C) \xleftarrow{\$} \text{REnc}(\text{pk}_{\hat{i}}, \text{rpk})$. Note that due to the modification we made in the previous game, when $\hat{i}$ receives ct_C and decrypts it with rpk 's secret key, it also uses the random key ss_C generated by iID_r . Since we consider the case (2-2), the long-term secret key of the user $\hat{i}$, who is the initiator of the tested instance, is never leaked if E_{Tested} occurs. Thus, due to the SPR-CCA security of RKEM, Game_3 and Game_4 are indistinguishable. More precisely, we can construct an algorithm $\mathcal{B}_5$ breaking the SPR-CCA security of RKEM as follows.

$\mathcal{B}_5$ receives a challenge public key pk^* from its challenger. $\mathcal{B}_5$ then samples a random $(\hat{i}, \text{iID}_r) \xleftarrow{\$} [N] \times [M]$ and generates the long-term key pairs as follows. For the user $\hat{i}$, $\mathcal{B}_5$ uses pk^* as its long-term public key, and for all the other parties, $\mathcal{B}_5$ computes the long-term key pairs correctly. Finally, $\mathcal{B}_5$ invokes $\mathcal{A}$ and answers the queries made by $\mathcal{A}$ as follows:

- $\mathcal{O}_{\text{Init}}(i, r)$: $\mathcal{B}_5$ responds as in the previous game except that it generates a reinforcing public key rpk through $\mathcal{O}_{\text{RKeyGen}}$ oracle.
- $\mathcal{O}_{\text{Res}}(r, m_1)$: Let iID be the assigned iID and rpk be the received reinforcing public key in m_1 .
 - If $\text{Init}[\text{iID}] = \hat{i}$ and $\text{iID} = \text{iID}_r$, then $\mathcal{B}_5$ responds as in the previous game except that it outputs rpk to its challenger, receives $(\text{ct}^*, \text{ss}^*)$, and sets $(\text{ct}_C, \text{ss}_C) := (\text{ct}^*, \text{ss}^*)$ instead of generating them on its own.
 - Otherwise, $\mathcal{B}_5$ responds as in the previous game.

- $\mathcal{O}_{\text{InitDer}}(\text{iID}, m_2)$: Depending on the value of $i = \text{Init}[\text{iID}]$, it performs the following:
 - If $i = \hat{i}$, then $\mathcal{B}_5$ checks if $\text{ct} = \text{ct}^*$. If so, it responds as in the previous game (i.e., it uses $\text{ss} := \text{ss}^*$). Otherwise, if $\text{ct} \neq \text{ct}^*$, then it sends ct and its reinforcing public key rpk to $\mathcal{O}_{\text{MalDec}}$ oracle and receives back ss' . $\mathcal{B}_5$ then responds with ss' .
 - If $i \neq \hat{i}$, then $\mathcal{B}_5$ responds as in the previous game since it knows both the long-term secret key and reinforcing public secret key.
- $\mathcal{O}_{\text{CreatePSK}}(u)$, $\mathcal{O}_{\text{RevLSK}}(u)$, $\mathcal{O}_{\text{RevPSK}}(u)$, $\mathcal{O}_{\text{RegisterPK}}(u)$, $\text{RevState}(\text{iID})$, $\text{RevSessKey}(\text{iID})$, $\text{Test}(\text{iID})$: $\mathcal{B}_5$ responds as in the previous game. Here, note that since we consider case (2-2), $\mathcal{B}_5$ can answer all the $\mathcal{O}_{\text{RevLSK}}$ -query since $\mathcal{A}$ never queries $\text{RevLSK}(\hat{i})$ conditioning on E_{Tested} not occurring, which is the only query that $\mathcal{B}_5$ cannot answer.

If $\mathcal{A}$ outputs a guess b' , $\mathcal{B}_5$ outputs 1 if $b = b'$ and 0 otherwise as its guess. It can be checked that $\mathcal{B}_5$ perfectly simulates game Game_3 (resp. Game_4) to $\mathcal{A}$ when the answer in the challenge phase is the real key and ciphertext (resp. a random key and a simulated ciphertext). Thus, we have

$$|\Pr[\text{Game}_3 = 1] - \Pr[\text{Game}_4 = 1]| \leq \text{Adv}_{\text{RKEM}, \mathcal{B}_5}^{\text{SPR-CCA}}(\lambda).$$

Game₅: We change how iID_r computes C_7 . In this game, we replace C_7 with a uniformly random value instead of computing it as $\text{KDF}(\text{ss}_C, C_6)$. Due to the modification we made in the previous game, ss_C is chosen uniformly at random. Therefore, by the PRF security of KDF, the two games are indistinguishable. Thus, we have

$$|\Pr[\text{Game}_4 = 1] - \Pr[\text{Game}_5 = 1]| \leq \text{Adv}_{\text{KDF}, \mathcal{B}_5'}^{\text{PRF}}(\lambda).$$

Game₆: In this game, K_{out} is replaced with a uniformly random value. Due to the modification we made in the previous game, C_7 is chosen uniformly at random. Therefore, by the dual-PRF security of KDF, the two games are indistinguishable. Thus, we have

$$|\Pr[\text{Game}_5 = 1] - \Pr[\text{Game}_6 = 1]| \leq \text{Adv}_{\text{KDF}, \mathcal{B}_5''}^{\text{PRF}}(\lambda).$$

In game Game_6 , the tested session key K_{out} is an independent and uniformly random value. Following a similar argument for the case (1-1), K_{out} looks random to the adversary regardless of the challenge bit b . Therefore, we conclude that $\Pr[\text{Game}_6 = 1] = 1/2$. Combining all arguments together, we obtain

$$\begin{aligned} & \left| \Pr[\text{Game}_1 = 1] - \frac{1}{2} \right| \\ & \leq NM \cdot (\text{Adv}_{\text{RKEM}, \mathcal{B}_5}^{\text{SPR-CCA}}(\lambda) + 2\text{Adv}_{\text{KDF}, \mathcal{B}_5'}^{\text{PRF}}(\lambda)). \end{aligned}$$

□

Appendix D.

Omitted Details of Generic Construction of RKEM

D.1. Proof of T Transform

Proof. As the proof for C-IND and CR-IND security proceed almost identically, we only prove C-IND security below. We consider a sequence of hybrid games, where we denote E_i as the event that $\mathcal{A}$ outputs $b' = 1$ in Game_i . The games are formally depicted in Fig. 11.

Game₀: In this game, the adversary is given an honestly generated ciphertext. This corresponds to the C-IND security game where $b = 0$ is sampled.

Game₁: This is the same as Game_0 except that it samples μ^* at the beginning of the game. This is only a conceptual change and we have $\Pr[E_1] = \Pr[E_0]$.

Game₂: This is the same as Game_1 except that the adversary $\mathcal{A}$ is given access to a different oracle G . Let $\mathcal{RPK}(\text{pk}^*)$ denote the space of reinforcing public keys, possibly including rpks that are not an output of $\text{RKeyGen}(\text{pk}^*)$. Then, G is sampled uniformly random over all the functions mapping $\mathcal{RPK}(\text{pk}^*) \times \mathcal{M} \times \mathcal{M}$ to $\mathcal{R}$ conditioned on that $F(\text{rp}_k, \mu, \mu_r) = G(\text{rp}_k, \mu, \mu_r)$ for all $(\text{rp}_k, \mu, \mu_r) \in \mathcal{RPK}(\text{pk}^*) \times \mathcal{M} \times \mathcal{M} \setminus S_{\mu^*}$, where $S_{\mu^*} := \{(\text{rp}_k, \mu^*, \mu_r) : (\text{rp}_k, \mu_r) \in \mathcal{RPK}(\text{pk}^*) \times \mathcal{M}\}$. We later show in Lemma 10 that there exists an adversary $\mathcal{B}_2$ against OW-CPA security such that

$$|\Pr[E_1] - \Pr[E_2]| \leq 2 \cdot Q_F \cdot \sqrt{\text{Adv}_{\text{RPKE}_2, \mathcal{B}_2}^{\text{OW-CPA}}(1^\lambda)}.$$

Game₃: This is the same as Game_2 except that the randomness rand is generated uniformly at random. Notice that the random oracle F is no longer accessible by the adversary $\mathcal{A}$, and therefore, $F(\text{rp}_k', \mu^*, \mu_r)$ for all $(\text{rp}_k', \mu^*, \mu_r) \in \mathcal{RPK}(\text{pk}^*) \times S_{\mu^*}$ are information theoretically hidden from $\mathcal{A}$. Therefore, this modification cannot be noticed by $\mathcal{A}$. Hence, we have $\Pr[E_3] = \Pr[E_2]$.

Game₄: In this game the adversary is given a simulated ciphertext. This corresponds to the C-IND security game where $b = 1$ is sampled. It is clear that we can construct an adversary $\mathcal{B}_1$ against C-IND security such that

$$|\Pr[E_3] - \Pr[E_4]| \leq \text{Adv}_{\text{RPKE}_2, \mathcal{B}_1}^{\text{C-IND}}(1^\lambda).$$

Now notice that we have

$$\begin{aligned} \text{Adv}_{\text{RPKE}_2, \mathcal{A}}^{\text{C-IND}}(1^\lambda) &:= \left| 2 \cdot \Pr \left[\text{Game}_{\text{RPKE}_2, \mathcal{A}}^{\text{C-IND}}(1^\lambda) = 1 \right] - 1 \right| \\ &= \left| (1 - \Pr[E_0]) + \Pr[E_4] - 1 \right| \\ &= |\Pr[E_0] - \Pr[E_4]|. \end{aligned}$$

Combining all the bounds, we arrive at the bound in the theorem statement. It remains to prove Lemma 10 below.

Lemma 10. *There exists an adversary $\mathcal{B}_2$ against OW-CPA security such that*

$$|\Pr[E_1] - \Pr[E_2]| \leq 2 \cdot Q_F \cdot \sqrt{\text{Adv}_{\text{RPKE}_2, \mathcal{B}_2}^{\text{OW-CPA}}(1^\lambda)}.$$

Proof. We invoke the O2H lemma (cf. Lemma 4) with oracles $F, G : \mathcal{RPK}(\text{pk}^*) \times \mathcal{M} \times \mathcal{M} \rightarrow \mathcal{R}$ as defined in Game_2 . We define the random variable z as (pk^*, μ^*, L) , where L is a list such that for all $(\text{rp}_k, \mu_r) \in \mathcal{RPK}(\text{pk}^*) \times \mathcal{M}$, $L[\text{rp}_k, \mu_r] := F(\text{rp}_k, \mu^*, \mu_r)$. This is without loss of generality as z can be any bit string that can depend arbitrary on F and S_{μ^*} . We then define the oracle-calling algorithms $\mathcal{A}_{\text{O2H}}$ and $\mathcal{B}_{\text{O2H}}$ as in Fig. 11. If $H := F$ (resp. $H := G$), then $\mathcal{A}_{\text{O2H}}^H(z)$ is identical to Game_1 (resp. Game_2), assuming $\mu^* \xleftarrow{\$} \mathcal{D}_{\mathcal{M}}$ and $(\text{pk}^*, \text{sk}^*) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$. Then, using Lemma 4, we have

$$\begin{aligned} &|\Pr[E_1] - \Pr[E_2]| \\ &\leq 2 \cdot Q_F \cdot \sqrt{\Pr_{F, G, S_{\mu^*}, z} \left[\mathcal{T} \xleftarrow{\$} \mathcal{B}_{\text{O2H}}^{G(\cdot)}(z) : S_{\mu^*} \cap \mathcal{T} \neq \emptyset \right]}. \end{aligned}$$

It remains to upper bound probability inside the square root, which we denote by p . To this end, we construct an adversary $\mathcal{B}_2$ against the OW-CPA security that simulates the behavior of $\mathcal{B}_{\text{O2H}}$. Namely, $\mathcal{B}_2$ is given pk^* and access to a random oracle G as part of the OW-CPA security game. It first samples $i \xleftarrow{\$} [Q_F]$ and runs $\mathcal{A}$ on input pk^* , relaying any oracle queries $\mathcal{A}$ makes to its own oracle G . Once $\mathcal{A}$ outputs $(\text{rp}_k', \mu_r', \text{state})$, it outputs this as its own output. Once it receives ct from the game, it further invokes $\mathcal{A}$ on input $(\text{ct}, \text{state})$. In the above, when $\mathcal{A}$ makes its i -th query to G , which is assumed to exist as $\mathcal{A}$ makes Q_F queries, $\mathcal{B}_2$ measures it and outputs the first register μ^* as the solution to its OW-CPA security game.

It can be checked that $\mathcal{B}_2$ simulates $\mathcal{B}_{\text{O2H}}^{G(\cdot)}(z)$ perfectly, where $z = (\text{pk}^*, \mu^*, L)$ and μ^* is the challenge message sampled by the OW-CPA security game. Here, importantly, while $\mathcal{B}_2$ cannot prepare (the exponentially large unknown string) z , this is not an issue as z is not explicitly used by $\mathcal{B}_{\text{O2H}}^{G(\cdot)}(z)$ (and hence by $\mathcal{A}_{\text{O2H}}^{G(\cdot)}(z)$). Thus, if $\mathcal{B}_{\text{O2H}}^{G(\cdot)}(z)$ outputs some $\mathcal{T} = (\text{rp}_k, \mu^*, \mu_r)$ as its output, then $\mathcal{B}_2$ wins the OW-CPA security game. In particular, we have the desired bound $p \leq \text{Adv}_{\text{RPKE}_2, \mathcal{B}_2}^{\text{OW-CPA}}(1^\lambda)$. This completes the proof $\square$

D.2. Proof of $U_m^{\mathcal{Y}}$ Transform

Proof. It is immediate that δ_c -correctness is inherited from RPKE_2 . As the proof for IND-RCCA security is a simplified version of the proof of SPR-CCA security, we focus on SPR-CCA security below.

We consider a sequence of hybrid games, where we denote E_i as the event that $\mathcal{A}$ outputs $b' = 1$ in Game_i . The core changes between the games are depicted in Fig. 6. While there are many subtle differences in how we manipulate the random oracles, the high level game transition follows the proof structure of [49, Theorem 4.1], which is also an extension of techniques used in [51], [66], [67]. Below, for simplicity and without loss of generality, we assume $\mathcal{A}$ only asks for a challenge ciphertext with respect to a reinforcing key rp_k' that it has not obtained from $\mathcal{O}_{\text{RKeyGen}}$.

Game₀ 1 : $(pk^*, sk^*) \xleftarrow{\$} \text{KeyGen}_2(1^\lambda)$ 2 : $(rp_k', \mu_r', \text{state}) \xleftarrow{\$} \mathcal{A}^{F(\cdot)}(pk^*)$ 3 : $\mu^* \xleftarrow{\$} \mathcal{D}_{\mathcal{M}}$ 4 : $\text{rand} := F(rp_k', \mu^*, \mu_r')$ 5 : $\text{ct} := \text{REnc}_2(pk^*, rp_k', \mu^*, \mu_r'; \text{rand})$ 6 : $b' \xleftarrow{\$} \mathcal{A}^{F(\cdot)}(\text{ct}, \text{state})$ 7 : return b'	Game₃ and Game₄ 1 : $\mu^* \xleftarrow{\$} \mathcal{D}_{\mathcal{M}}$ 2 : $(pk^*, sk^*) \xleftarrow{\$} \text{KeyGen}_2(1^\lambda)$ 3 : $(rp_k', \mu_r', \text{state}) \xleftarrow{\$} \mathcal{A}^{G(\cdot)}(pk^*)$ 4 : $\text{rand} \xleftarrow{\$} \mathcal{R}$ 5 : $\text{ct} := \text{REnc}_2(pk^*, rp_k', \mu^*, \mu_r'; \text{rand})$ / If Game ₁ 6 : $\text{ct} \xleftarrow{\$} \mathcal{S}(1^\lambda, rp_k', \mu_r')$ / If Game ₂ 7 : $b' \xleftarrow{\$} \mathcal{A}^{G(\cdot)}(\text{ct}, \text{state})$ 8 : return b'
Game₁ 1 : $\mu^* \xleftarrow{\$} \mathcal{D}_{\mathcal{M}}$ 2 : $(pk^*, sk^*) \xleftarrow{\$} \text{KeyGen}_2(1^\lambda)$ 3 : $(rp_k', \mu_r', \text{state}) \xleftarrow{\$} \mathcal{A}^{F(\cdot)}(pk^*)$ 4 : $\text{rand} := F(rp_k', \mu^*, \mu_r')$ 5 : $\text{ct} := \text{REnc}_2(pk^*, rp_k', \mu^*, \mu_r'; \text{rand})$ 6 : $b' \xleftarrow{\$} \mathcal{A}^{F(\cdot)}(\text{ct}, \text{state})$ 7 : return b'	$\mathcal{A}_{\text{O2H}}^{H(\cdot)}(z)$ 1 : parse $(rp_k^*, \mu^*, L) \leftarrow z$ 2 : $(rp_k', \mu_r', \text{state}) \xleftarrow{\$} \mathcal{A}^{H(\cdot)}(pk^*)$ 3 : $\text{rand} := L[rp_k', \mu_r']$ 4 : $\text{ct} := \text{REnc}_2(pk^*, rp_k', \mu^*, \mu_r'; \text{rand})$ 5 : return $\mathcal{A}^{H(\cdot)}(\text{ct}, \text{state})$
Game₂ 1 : $\mu^* \xleftarrow{\$} \mathcal{D}_{\mathcal{M}}$ 2 : $(pk^*, sk^*) \xleftarrow{\$} \text{KeyGen}_2(1^\lambda)$ 3 : $(rp_k', \mu_r', \text{state}) \xleftarrow{\$} \mathcal{A}^{G(\cdot)}(pk^*)$ 4 : $\text{rand} := F(rp_k', \mu^*, \mu_r')$ 5 : $\text{ct} := \text{REnc}_2(pk^*, rp_k', \mu^*, \mu_r'; \text{rand})$ 6 : $b' \xleftarrow{\$} \mathcal{A}^{G(\cdot)}(\text{ct}, \text{state})$ 7 : return b'	$\mathcal{B}_{\text{O2H}}^{H(\cdot)}(z)$ 1 : $i \xleftarrow{\$} [Q_F]$ 2 : Run $\mathcal{A}_{\text{O2H}}^{H(\cdot)}(z)$ until i -th query $ (rp_k, \mu, \mu_r)\rangle$ to H 3 : return Measurement of $ (rp_k, \mu, \mu_r)\rangle$

Figure 11. Hybrid games for [Thm. 5](#) and tools to invoke the O2H lemma (cf. [Lemma 4](#)). The blue highlights the differences in the games.

Game₀: This corresponds to the SPR-CCA game where $b = 0$ is sampled. We thus have

$$\Pr[E_0] = 1 - \Pr \left[\text{Game}_{\text{RKEM}, \mathcal{A}}^{\text{SPR-CCA}}(1^\lambda) = 1 \mid b = 0 \right].$$

Game₁: This is the same as Game₀ except that $H(rp_k, \mu, \mu_r)$ in the decapsulation oracle is replaced by $H_L(rp_k, \mu, \mu_r)$, where $H_L : \mathcal{RPK}(pk^*) \times \mathcal{M} \times \mathcal{M} \rightarrow \mathcal{R}$ is another random oracle. From [Lemma 3](#), we have

$$|\Pr[E_0] - \Pr[E_1]| \leq (Q_{\text{dec}} + Q_{H_{\text{prf}}}) \cdot 2^{-\frac{\ell}{2} + 1}.$$

Game₂: This is the same as Game₁ except that the game samples the Q_{rp_k} reinforcing keys $(rp_{k_i}, \text{rsk}_i)_{i \in [Q_{\text{rp}_k}]}$ after it samples (pk^*, sk^*) . On the i -th query to $\mathcal{O}_{\text{RKeyGen}}$ from the adversary, it simply returns (rp_{k_i}, rsk_i) , as opposed to sampling them as in the previous game. This is merely a conceptual change. In this game, the game further samples the random oracle $F : \mathcal{RPK}(pk^*) \times \mathcal{M} \times \mathcal{M} \rightarrow \mathcal{R}$ from $\mathcal{F}_{pk^*, (rp_{k_i}, \mu_i, \mu_r, i) \in [Q_{\text{rp}_k}]}^{\text{good}}$ defined below, where $\mathcal{R}$ denotes the encryption randomness space of the underlying RPKE₂.

For any $N \in \mathbb{N}$, $(pk^*, sk^*) \in \text{KeyGen}(1^\lambda)$, $(rp_{k_i}, \text{rsk}_i) \in \text{RKeyGen}(pk^*)$, and $(\mu_i, \mu_{r,i}) \in \mathcal{M} \times \mathcal{M}$ for $i \in [Q_{\text{rp}_k}]$, we define two sets of *good* randomness as

$$\begin{aligned} \mathcal{R}_{pk^*, \mu}^{\text{good}} &:= \{r \in \mathcal{R} : \forall rp_k \in \mathcal{RPK}(pk^*), \\ &\quad \mu = \text{RHalfDec}_2(sk, \text{REnc}_2(pk, rp_k, \mu, \mu_r; r))\}, \text{ and} \\ \mathcal{R}_{pk^*, (rp_{k_i}, \mu_i, \mu_r, i) \in [Q_{\text{rp}_k}]}^{\text{good}} &:= \{(r_i)_{i \in [Q_{\text{rp}_k}]} \in \mathcal{R}^{Q_{\text{rp}_k}} : \forall i \in [Q_{\text{rp}_k}], \\ &\quad \text{RDec}_2(sk^*, \text{rsk}_i, \text{REnc}_2(pk^*, rp_{k_i}, \mu_i, \mu_{r,i}; r_i)) = (\mu_i, \mu_{r,i})\}. \end{aligned}$$

The first (resp. second) set is the set of good randomness such that half (resp. standard) decryption succeeds. Here, while we omit them for better readability, we note that $\mathcal{R}_{pk^*}^{\text{good}}$ and $\mathcal{R}_{pk^*, (rp_{k_i}, \mu_i, \mu_r, i) \in [Q_{\text{rp}_k}]}^{\text{good}}$ also depend on the secret keys sk^* and $(\text{rsk}_i)_{i \in [Q_{\text{rp}_k}]}$. Then, let $\mathcal{F}_{pk^*, (rp_{k_i}, \mu_i, \mu_r, i) \in [Q_{\text{rp}_k}]}^{\text{good}}$ be the set of functions $F : \mathcal{RPK}(pk^*) \times \mathcal{M} \times \mathcal{M} \rightarrow \mathcal{R}$ satisfying $F(rp_k, \mu, \mu_r) \in \mathcal{R}_{pk^*}^{\text{good}}$ for all $(rp_k, \mu, \mu_r) \in \mathcal{RPK}(pk^*) \times \mathcal{M} \times \mathcal{M}$ and $(F(rp_{k_i}, \mu_i, \mu_{r,i}))_{i \in [Q_{\text{rp}_k}]} \in \mathcal{R}_{pk^*, (rp_{k_i}, \mu_i, \mu_r, i) \in [Q_{\text{rp}_k}]}^{\text{good}}$ for all $(\mu_i, \mu_{r,i}) \in \mathcal{M} \times \mathcal{M}$. In Game₂, we sample F from this $\mathcal{F}_{pk^*, (rp_{k_i}, \mu_i, \mu_r, i) \in [Q_{\text{rp}_k}]}^{\text{good}}$. Note that

TABLE 4. SUMMARY OF GAMES FOR THE PROOF OF SPR-CCA IN THM. 6.

Game	H	F	ct*	ss*	Decapsulation		Justification
					valid ct	invalid ct	
Game ₀	H	$\mathcal{F}_{pk^*}$	$\text{REnc}(pk^*, rp_k')$	$H(rp_k', \mu^*, \mu_r^*)$	$H(rp_k, \mu, \mu_r)$	$H_{\text{prf}}(s, s_r, rp_k, ct)$	
Game ₁	H	$\mathcal{F}_{pk^*}$	$\text{REnc}(pk^*, rp_k')$	$H(rp_k', \mu^*, \mu_r^*)$	$H(rp_k, \mu, \mu_r)$	$H_L(rp_k, ct)$	Lemma 3
Game ₂	H	$\mathcal{F}_{pk^*}^{\text{good}}$	$\text{REnc}(pk^*, rp_k')$	$H(rp_k', \mu^*, \mu_r^*)$	$H(rp_k, \mu, \mu_r)$	$H_L(rp_k, ct)$	Lemma 2 + δ_{hc} , δ_c -correctness
Game ₃	$H'_L \circ g_L$ $H'_{-L} \circ g_{-L}$	$\mathcal{F}_{pk^*}^{\text{good}}$	$\text{REnc}(pk^*, rp_k')$	$H(rp_k', \mu^*, \mu_r^*)$	$H(rp_k, \mu, \mu_r)$	$H_L(rp_k, ct)$	by ensuring perfect correctness
Game ₄	$H_L \circ g_L$ $H_{-L} \circ g_{-L}$	$\mathcal{F}_{pk^*}^{\text{good}}$	$\text{REnc}(pk^*, rp_k')$	$H(rp_k', \mu^*, \mu_r^*)$	$H(rp_k, \mu, \mu_r)$	$H_L(rp_k, ct)$	by ensuring perfect correctness
Game ₅	$H_L \circ g_L$ $H_{-L} \circ g_{-L}$	$\mathcal{F}_{pk^*}^{\text{good}}$	$\text{REnc}(pk^*, rp_k')$	$H_{-L}(rp_k', ct^*, \mu_r^*)$	$H_L(rp_k, ct)$	$H_L(rp_k, ct)$	conceptual
Game ₆	$H_L \circ g_L$ $H_{-L} \circ g_{-L}$	$\mathcal{F}_{pk^*}$	$\text{REnc}(pk^*, rp_k')$	$H_{-L}(rp_k', ct^*, \mu_r^*)$	$H_L(rp_k, ct)$	$H_L(rp_k, ct)$	Lemma 2 + δ_{hc} -half and δ_c -correctness
Game ₇	$H_L \circ g_L$ $H_{-L} \circ g_{-L}$	$\mathcal{F}_{pk^*}$	$\overline{S}(1^\lambda, rp_k', \mu_r^*)$	$H_{-L}(rp_k', ct^*, \mu_r^*)$	$H_L(rp_k, ct)$	$H_L(rp_k, ct)$	CR-IND security
Game ₈	$H_L \circ g_L$ $H_{-L} \circ g_{-L}$	$\mathcal{F}_{pk^*}$	$\overline{S}(1^\lambda, rp_k', \mu_r^*)$	$\mathcal{U}(\mathcal{K})$	$H_L(rp_k, ct)$	$H_L(rp_k, ct)$	statistical disjointness
Game ₉	$H_L \circ g_L$ $H_{-L} \circ g_{-L}$	$\mathcal{F}_{pk^*}^{\text{good}}$	$\overline{S}(1^\lambda, rp_k', \mu_r^*)$	$\mathcal{U}(\mathcal{K})$	$H_L(rp_k, ct)$	$H_L(rp_k, ct)$	Lemma 2 + δ_{hc} -half, δ_c -correctness
Game ₁₀	$H_L \circ g_L$ $H_{-L} \circ g_{-L}$	$\mathcal{F}_{pk^*}^{\text{good}}$	$\overline{S}(1^\lambda, rp_k', \mu_r^*)$	$\mathcal{U}(\mathcal{K})$	$H(rp_k, \mu, \mu_r)$	$H_L(rp_k, ct)$	conceptual
Game ₁₁	$H'_L \circ g_L$ $H'_{-L} \circ g_{-L}$	$\mathcal{F}_{pk^*}^{\text{good}}$	$\overline{S}(1^\lambda, rp_k', \mu_r^*)$	$\mathcal{U}(\mathcal{K})$	$H(rp_k, \mu, \mu_r)$	$H_L(rp_k, ct)$	by ensuring perfect correctness
Game ₁₂	H	$\mathcal{F}_{pk^*}^{\text{good}}$	$\overline{S}(1^\lambda, rp_k', \mu_r^*)$	$\mathcal{U}(\mathcal{K})$	$H(rp_k, \mu, \mu_r)$	$H_L(rp_k, ct)$	by ensuring perfect correctness
Game ₁₃	H	$\mathcal{F}_{pk^*}$	$\overline{S}(1^\lambda, rp_k', \mu_r^*)$	$\mathcal{U}(\mathcal{K})$	$H(rp_k, \mu, \mu_r)$	$H_L(rp_k, ct)$	Lemma 2 + δ_{hc} -half, δ_c -correctness
Game ₁₄	H	$\mathcal{F}_{pk^*}$	$\overline{S}(1^\lambda, rp_k', \mu_r^*)$	$\mathcal{U}(\mathcal{K})$	$H(rp_k, \mu, \mu_r)$	$H_{\text{prf}}(s, s_r, rp_k, ct)$	Lemma 3

in the previous game, we sampled F from $\mathcal{F}_{pk^*}$, the set of all functions mapping $\mathcal{RPK}(pk^*) \times \mathcal{M} \times \mathcal{M}$ to $\mathcal{R}$.

Let us analyze the difference in the adversary's advantage. For any $N \in \mathbb{N}$, let us define

$$\delta_{pk^*, (rp_k, \mu, \mu_r)_{i \in [Q_{\text{rpk}}]}} := \left| \mathcal{R} \setminus \left(\mathcal{R}_{pk^*}^{\text{good}} \cap \mathcal{R}_{pk^*, (rp_k, \mu, \mu_r)_{i \in [Q_{\text{rpk}}]}}^{\text{good}} \right) \right| / |\mathcal{R}|,$$

the fraction of the bad randomness. That is the fraction of randomness for which half or standard decryption fails. Further define

$$\delta_{pk^*, (rp_k)_{i \in [Q_{\text{rpk}}]}} := \max_{(\mu_i, \mu_r, i)_{i \in [Q_{\text{rpk}}]} \in (\mathcal{M} \times \mathcal{M})^{Q_{\text{rpk}}}} \delta_{pk^*, (rp_k, \mu_i, \mu_r, i)_{i \in [Q_{\text{rpk}}]}}.$$

Relying on Lemma 2 and following a similar analysis in the proof of [48, Theorem 3.2], we have

$$\begin{aligned} & \left| \Pr[E_1 \mid (pk^*, sk^*, (rp_k, rsk_i)_{i \in [Q_{\text{rpk}}]})] \right. \\ & \quad \left. - \Pr[E_2 \mid (pk^*, sk^*, (rp_k, rsk_i)_{i \in [Q_{\text{rpk}}]})] \right| \\ & \leq 8 \cdot (Q_F + Q_{\text{dec}} + 2)^2 \cdot \delta_{pk^*, (rp_k)_{i \in [Q_{\text{rpk}}]}}. \end{aligned}$$

By definition of δ_{hc} -half and δ_c -correctness of RPKE_2 (cf. Def. 22 and 23), and using the union bound, we have $\text{Exp}[\delta_{pk^*, (rp_k)_{i \in [Q_{\text{rpk}}]}}] \leq \delta_{\text{hc}} + Q_{\text{rpk}} \cdot \delta_c$, where the expectation is taken over the randomness of sampling $(pk^*, sk^*) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$ and $(rp_k, rsk_i) \xleftarrow{\$} \text{RKeyGen}(pk^*)$ for $i \in [Q_{\text{rpk}}]$. Taking the average, we have

$$|\Pr[E_1] - \Pr[E_2]| \leq 8 \cdot (Q_F + Q_{\text{dec}} + 2)^2 \cdot (\delta_{\text{hc}} + Q_{\text{rpk}} \cdot \delta_c).$$

Below, let BAD denote the event for which either of the following conditions hold: there exists some $(rp_k, \mu, \mu_r) \in \mathcal{RPK}(pk^*) \times \mathcal{M} \times \mathcal{M}$ such that no $r \in \mathcal{R}$ satisfies $\mu = \text{RHalfDec}_2(sk, \text{REnc}_2(pk, rp_k, \mu, \mu_r; r))$; and there exists some $i \in [Q_{\text{rpk}}]$ and $(\mu_i, \mu_r, i)_{i \in [Q_{\text{rpk}}]} \in (\mathcal{M} \times \mathcal{M})^{Q_{\text{rpk}}}$ such that no $r_i \in \mathcal{R}$ satisfies $(\mu_i, \mu_r, i) = \text{RDec}_2(sk^*, rsk_i, \text{REnc}_2(pk^*, rp_k, \mu_i, \mu_r, i; r_i))$. Put differently, BAD denotes the event that there are no randomness for which either the half or standard decryption succeeds. Following a similar argument to [67, Claim 3] and using the union bound, we have $\Pr[\text{BAD}] \leq Q_{\text{rpk}} \delta_{\text{hc}} + Q_{\text{rpk}} \cdot \delta_c$. Lastly, we denote GOOD as the event that BAD does not occur.

Game₃: This is the same as Game₂ except that the random oracle H is simulated as follows, where L denotes the set of

Q_{rpk} reinforcing public keys $(\text{rpki})_{i \in [Q_{\text{rpk}}]}$ sampled by the game at the beginning of the game:

$$H(\text{rpki}, \mu, \mu_r) := \begin{cases} H'_L \circ g_L(\text{rpki}, \mu, \mu_r) & \text{if } \text{rpki} \in L \\ H'_{-L} \circ g_{-L}(\text{rpki}, \mu, \mu_r) & \text{if } \text{rpki} \notin L, \end{cases}$$

where $H'_L : \mathcal{RPK}(\text{pk}^*) \times \mathcal{CT} \rightarrow \mathcal{K}$ and $H'_{-L} : \mathcal{RPK}(\text{pk}^*) \times \mathcal{CT} \times \mathcal{M} \rightarrow \mathcal{K}$ are randomly sampled random oracles and

$$\begin{aligned} g_L(\text{rpki}, \mu, \mu_r) &:= (\text{rpki}, \overline{\text{REnc}_2}(\text{pk}, \text{rpki}, \mu, \mu_r)) \\ g_{-L}(\text{rpki}, \mu, \mu_r) &:= (\text{rpki}, \overline{\text{REnc}_2}(\text{pk}, \text{rpki}, \mu, \mu_r), \mu_r). \end{aligned}$$

Here, $\mathcal{CT}$ denotes the ciphertext space of the underlying RPKE_2 .

Notice that conditioned on GOOD, we have perfect half and standard decryption. Therefore, g_L and g_{-L} are injective functions. Note that g_{-L} outputs μ_r because in case $\text{rpki} \notin L$, we cannot have no guarantee that $\text{ct} = \overline{\text{REnc}_2}(\text{pk}, \text{rpki}, \mu, \mu_r)$ correctly decrypts the message μ_r . Thus, if GOOD occurs, then H'_L and H'_{-L} are both random functions and this game is equivalent to the previous game. We have,

$$\Pr[E_2 \wedge \text{GOOD}] = \Pr[E_3 \wedge \text{GOOD}]$$

Moreover, using a simple statistical lemma [49, Lemma A.1], we have

$$\begin{aligned} |\Pr[E_1] - \Pr[E_2]| &\leq |\Pr[E_1] - \Pr[E_2 \wedge \text{GOOD}]| + \Pr[\text{BAD}] \\ &\leq |\Pr[E_1] - \Pr[E_2 \wedge \text{GOOD}]| + \delta_{\text{hc}} + Q_{\text{rpki}} \cdot \delta_c. \end{aligned}$$

Game₄: This is the same as Game₃ except that H'_L is simulated using H_L . For consistency, we also make a conceptual modification by rewriting H'_{-L} as H_{-L} .

A ciphertext ct (with respect to $\text{rpki} \in L$) is said to be *valid* if we have $\text{ct} = \overline{\text{REnc}_2}(\text{pk}^*, \text{rpki}, \text{RDec}_2(\text{sk}^*, \text{rsk}_i, \text{ct}))$, and *invalid* otherwise. Notice that in the previous game, H_L is only used for invalid ciphertexts, and an adversary cannot access H_L on a valid ciphertext. On the other hand, conditioned on GOOD, H'_L is only used for valid ciphertexts owing to g_L , and an adversary cannot access H'_L on an invalid ciphertext. Therefore, conditioned on GOOD, the adversary cannot detect if we replace H'_L by H_L , and we have

$$\Pr[E_3 \wedge \text{GOOD}] = \Pr[E_4 \wedge \text{GOOD}]$$

Game₅: This game is the same as Game₄ except that the challenge session key ss^* is set as $H_{-L}(\text{rpki}', \text{ct}^*, \mu_r^*)$ and the decryption oracle always returns $H_L(\text{rpki}, \text{ct})$ as long as $\text{ct} \neq \text{ct}^*$. Conditioned on GOOD and the assumption that $\text{rpki}' \notin L$, this modification is merely conceptual due to the definition of g_L and g_{-L} . We thus have

$$\Pr[E_4 \wedge \text{GOOD}] = \Pr[E_5 \wedge \text{GOOD}]$$

Using the same statistical argument, we have

$$\begin{aligned} |\Pr[E_5] - \Pr[E_5]| &\leq |\Pr[E_5] - \Pr[E_5 \wedge \text{GOOD}]| + \Pr[\text{BAD}] \\ &\leq |\Pr[E_5] - \Pr[E_5 \wedge \text{GOOD}]| + \delta_{\text{hc}} + Q_{\text{rpki}} \cdot \delta_c. \end{aligned}$$

Game₆: This game is the same as Game₅ except that we undo the change we made in Game₂. Namely, we now sample F from $\mathcal{F}_{\text{pk}^*}$, the set of all functions mapping $\mathcal{RPK}(\text{pk}^*) \times$

$\mathcal{M} \times \mathcal{M}$ to $\mathcal{R}$. Following the same argument as before, we have

$$|\Pr[E_5] - \Pr[E_6]| \leq 8 \cdot (Q_F + Q_H + 2)^2 \cdot (\delta_{\text{hc}} + Q_{\text{rpki}} \cdot \delta_c).$$

Here note that Q_{dec} is replaced by Q_H as H internally calls F when invoking g_L and g_{-L} .

Game₇: This game is the same as Game₆ except that the challenge ciphertext ct^* is generated using the simulator $\overline{S}(1^\lambda, \text{rpki}', \mu_r^*)$ defined by the CR-IND security game of the underlying RPKE_2 .

Notice that in Game₇, the decapsulation oracle no longer needs any secret key as the response is $H_L(\text{rpki}, \text{ct})$ regardless of the decryption result. Moreover, the challenge session key ss^* no longer needs the challenge message μ^* as it can use ct^* instead. Therefore, it is easy to check that we can construct an adversary $\mathcal{B}$ against the CR-IND security of RPKE_2 such that

$$|\Pr[E_6] - \Pr[E_7]| \leq \text{Adv}_{\text{RPKE}_2, \mathcal{B}}^{\text{CR-IND}}(1^\lambda).$$

Game₈: This game is the same as Game₇ except that we sample $\text{ss}^* \xleftarrow{\$} \mathcal{K}$, as opposed to $\text{ss}^* := H_{-L}(\text{rpki}', \text{ct}^*, \mu_r^*)$.

Notice that if $\text{ct}^* \xleftarrow{\$} \overline{S}(1^\lambda, \text{rpki}', \mu_r^*)$ is not in $\overline{\text{REnc}_2}(\text{pk}, \text{rpki}, \mathcal{M}, \mathcal{M})$, then the adversary has no information on $H_{-L}(\text{rpki}', \text{ct}^*, \mu_r^*)$. This is because the only way the adversary can access H_{-L} is through H but it internally runs $\overline{\text{REnc}_2}$ thanks to g_{-L} . Hence, relying on the statistical disjointness of the underlying RPKE_2 , we have

$$|\Pr[E_7] - \Pr[E_8]| \leq \text{Disj}_{\overline{\text{RPKE}_2}}(1^\lambda).$$

Game₉ to Game₁₄: The remaining games consist of undoing all the modifications we made on the random oracles and decapsulation oracle from Game₀ to Game₅. The exact game transition is given in Tab. 4. As it is just a rewrite of Game₀ to Game₅ with the same bound, we omit the details. We note that Game₁₄ corresponds to the SPR-CCA game where $b = 1$ is sampled. We thus have

$$\Pr[E_{14}] = \Pr \left[\text{Game}_{\text{RKEM}, \mathcal{A}}^{\text{SPR-CCA}}(1^\lambda) = 1 \mid b = 1 \right].$$

Combining all the bounds together, we arrive at the bound in the theorem statement.

Lastly, we point out that the proof of IND-RCCA is a simplification of the above proof for SPR-CCA. This is because in the SPR-CCA security game, the challenge ciphertext was created using a malicious reinforcing public key rpki' and the decapsulation oracle could be queried on any honest $\text{rpki}_i \in L$ generated by the challenger. In contrast, in the IND-RCCA security game, the challenge ciphertext and decapsulation oracle are defined by a unique honestly generated rpki^* . The proof for IND-RCCA consists of almost the same game hybrids for SPR-CCA, but removes the discussion on δ_{hc} -half correctness and sets $Q_{\text{rpki}} = 1$. As the proof is almost identical, we omit the details. $\square$

Appendix E.

Omitted Details of Lattice-based RPKE₂

E.1. C-IND Security of RPKE₂

Proof. We consider a sequence of hybrid games, where we denote E_i as the event that $\mathcal{A}$ outputs $b' = 1$ in Game_i . The simulator $\mathcal{S}$ takes as input $(1^\lambda, \text{rpk}, \mu_r)$, samples random $(\mathbf{u}, v_0, v_1) \xleftarrow{\$} R_q^k \times R_q \times R_q$, and outputs a ciphertext ct as $(\bar{\mathbf{u}}, \bar{v}_0, \bar{v}_1) := (\text{Compress}_q(\mathbf{u}, d_u), \text{Compress}_q(v_0, d_{v_0}), \text{Compress}_q(v_1, d_{v_1}))$. Some of the critical game transitions are depicted in Fig. 12.

Game₀: This corresponds to the C-IND game where $b = 0$ is sampled. We thus have

$$\Pr[E_0] = 1 - \Pr[\text{Game}_{\text{RPKE}_2, \mathcal{A}}^{\text{C-IND}}(1^\lambda) = 1 \mid b = 0].$$

Game₁: This is the same as Game_0 except that the game modifies how the random oracle G is simulated. It first samples $(\text{seed}_A, \mathbf{A}) \xleftarrow{\$} \{0, 1\}^{256} \times R_q^{n \times k}$ and a $2Q_G$ -wise independent function $G : \{0, 1\}^{256} \times R_q^{n \times k}$ conditioned on $G(\text{seed}_A) = \mathbf{A}$, where Q_G is the number of oracle queries to G . One way to sample such function would be to first sample a random G' (which is possible since q is a prime, cf. [81, Theorem 6.1]). Let $\mathbf{M} = G'(\text{seed}_A)$. Then simply define G as computing G' and adding $\mathbf{A} - \mathbf{M}$. Using [81, Theorem 6.1], a $2Q_G$ -wise independent function is perfectly indistinguishable from a random oracle to an adversary making at most Q_G oracle queries. Moreover, for oracle G_{Decomp} , we can simply use a $2Q_{\text{Decomp}}$ -wise independent function to sample the randomness to run G_{Decomp} , where Q_{Decomp} is the number of oracle queries to G_{Decomp} . We thus have $\Pr[E_0] = \Pr[E_1]$.

Game₂: This is the same game as Game_1 except that the public key $\mathbf{b}$ is sampled uniformly over R_q^n as opposed to $\mathbf{b} := \mathbf{A} \cdot \mathbf{s} + \mathbf{x}$ as in the previous game. As we modified $\mathbf{A}$ to be directly sampled by the game, it is easy to check that this game is indistinguishable from the previous game relying on MLWE. Formally, there exists an adversary $\mathcal{B}_1$ against the $\text{MLWE}_{q, n, k, \chi, \chi}$ problem such that

$$|\Pr[E_1] - \Pr[E_2]| \leq \text{Adv}_{\mathcal{B}_1}^{\text{MLWE}}(1^\lambda).$$

Game₃: This is the same as Game_2 except that it samples the challenge message μ^* and computes $\Delta := G_r(\mu^*)$ at the outset of the game. See Fig. 12 for the exact description. Clearly this produces the same view to the adversary as in the previous game. We thus have $\Pr[E_2] = \Pr[E_3]$.

Game₄: This is the same as Game_3 except that the adversary $\mathcal{A}$ is given access to a different oracle F as opposed to G_r . In particular, the game samples a random $F : \mathcal{M} \rightarrow R_q$ such that $F(\mu) = G_r(\mu)$ for all $\mu \in \mathcal{M} \setminus \{\mu^*\}$. See Fig. 12 for the formal description of the game. We later show in Lemma 11 that there exists an adversary $\mathcal{B}_2$ against the $\text{MLWE}_{q, k+1, n, \chi_r, (\chi_i)_{i \in [k+1]}}$ problem such that

$$|\Pr[E_3] - \Pr[E_4]| \leq 2 \cdot Q_{G_r} \cdot \sqrt{\text{Adv}_{\mathcal{B}_2}^{\text{MLWE}}(1^\lambda) + \frac{1}{|\mathcal{M}|}}.$$

Game₅: This is the same as Game_4 except that it samples $\Delta \xleftarrow{\$} \mathbb{Z}_{2^{d_{v_1}}}^d$ as opposed to setting it as $\Delta := G_r(\mu^*)$. See Fig. 12 for the formal description of the game. Notice that the random oracle G_r is no longer accessible by the adversary $\mathcal{A}$, and therefore, Δ is information theoretically hidden from $\mathcal{A}$. Therefore, this modification cannot be noticed by $\mathcal{A}$. Hence, we have $\Pr[E_4] = \Pr[E_5]$.

Game₆: This is the same game as Game_5 except that the challenger ciphertext $\text{ct}^* = (\bar{\mathbf{u}}^*, \bar{v}_0^*, \bar{v}_1^*)$ is sampled slightly differently. It generates $\bar{\mathbf{u}}^*$ and $\bar{v}_0^*$ as in the previous game but samples $\bar{v}_1^* \xleftarrow{\$} \mathbb{Z}_{2^{d_{v_1}}}^d$. Due to the modification we made in the previous game, Δ is information theoretically hidden to the adversary. Therefore, $\bar{v}_1^*$ is distributed uniformly random over $\mathbb{Z}_{2^{d_{v_1}}}^d$. We therefore have $\Pr[E_5] = \Pr[E_6]$.

Game₇: This is the same game as Game_6 except that the remaining part of the challenger ciphertext ct^* is sampled uniformly at random. It samples $(\mathbf{u}^*, v_0^*) \xleftarrow{\$} R_q^k \times R_q$ and sets $(\bar{\mathbf{u}}^*, \bar{v}_0^*) := (\text{Compress}_q(\mathbf{u}^*, d_u), \text{Compress}_q(v_0^*, d_{v_0}))$. See Fig. 12 for the formal description of the game. It is easy to check that this game is indistinguishable from the previous game relying on MLWE. Formally, there exists an adversary $\mathcal{B}_3$ against the $\text{MLWE}_{q, k+1, n, \chi_r, (\chi_i)_{i \in [k+1]}}$ problem such that

$$|\Pr[E_6] - \Pr[E_7]| \leq \text{Adv}_{\mathcal{B}_3}^{\text{MLWE}}(1^\lambda).$$

At this point, observe that the challenge ciphertext ct^* is the same as an output of the simulator $\mathcal{S}(1^\lambda, \text{rpk}', \mu_r')$.

Game₈: This is the same as Game_7 except that it no longer samples Δ and the adversary $\mathcal{A}$ is given access to the oracle G_r as opposed to F in the prior game. See Fig. 12 for the formal description of the game. As Δ is no longer used by the game, this is a conceptual change. Moreover, as G_r and F have the same functionality, this is also only conceptual. Therefore, we have $\Pr[E_7] = \Pr[E_8]$.

Game₉: The remaining game is to undo the changes from Game_0 to Game_2 . Namely, the public key is sampled as a valid MLWE sample and the oracle G is sampled randomly. Following the exact same argument to move from Game_0 to Game_2 , we can show that there exists an adversary $\mathcal{B}_4$ against the $\text{MLWE}_{q, n, k, \chi, \chi}$ problem such that

$$|\Pr[E_8] - \Pr[E_9]| \leq \text{Adv}_{\mathcal{B}_4}^{\text{MLWE}}(1^\lambda).$$

This corresponds to the C-IND game where $b = 1$ is sampled. We thus have

$$\Pr[E_9] = \Pr[\text{Game}_{\text{RPKE}_2, \mathcal{A}}^{\text{C-IND}}(1^\lambda) = 1 \mid b = 1].$$

Combining all the bounds together, we arrive at the bound in the theorem statement. It remains to prove the following lemma to complete the proof.

Lemma 11. *There exists an adversary $\mathcal{B}_2$ against the $\text{MLWE}_{q, k+1, n, \chi_r, (\chi_i)_{i \in [k+1]}}$ problem such that*

$$|\Pr[E_3] - \Pr[E_4]| \leq 2 \cdot Q_{G_r} \cdot \sqrt{\text{Adv}_{\mathcal{B}_2}^{\text{MLWE}}(1^\lambda) + \frac{1}{|\mathcal{M}|}}.$$

Proof. We invoke the O2H lemma (cf. Lemma 4) with oracles $G_r, F : \mathcal{M} \rightarrow R_q$ as defined in Game_4 , where

Game₃, Game₄, Game₅ <pre> 1 : $\mu^* \xleftarrow{\\$} \mathcal{U}(\mathcal{M})$ 2 : $\Delta := G_r(\mu^*) \in \mathbb{Z}_{2^{d_{v_1}}}^d$ / If Game₃, Game₄ 3 : $\Delta \xleftarrow{\\$} \mathbb{Z}_{2^{d_{v_1}}}^d$ / If Game₅ 4 : $(\text{seed}_A, A) \xleftarrow{\\$} \{0, 1\}^{256} \times R_q^{n \times k}$ 5 : Sample $2Q_G$-wise independent function G s.t. $G(\text{seed}_A) = A$ 6 : $b \xleftarrow{\\$} R_q^n$ 7 : $\text{pk}^* := (\text{seed}_A, b)$ 8 : $s \xleftarrow{\\$} \{0, 1\}^\ell$ 9 : $\text{sk}^* := (\text{seed}_A, s)$ 10 : $(\text{rpk}', \mu_r', \text{state}) \xleftarrow{\\$} \mathcal{A}^{G(\cdot), G_{\text{Decomp}}(\cdot), G_r(\cdot)}(\text{pk}^*)$ / If Game₃ 11 : $(\text{rpk}', \mu_r', \text{state}) \xleftarrow{\\$} \mathcal{A}^{G(\cdot), G_{\text{Decomp}}(\cdot), F(\cdot)}(\text{pk}^*)$ / If Game₄, Game₅ 12 : $\text{parse } \overline{b}_r \leftarrow \text{rpk}'$ 13 : $b_r^* \xleftarrow{\\$} \text{RandDecomp}_{q, d_{b_r}}(\overline{b}_r, d_{b_r})$ 14 : $(r, z, z_r) \xleftarrow{\\$} \chi_r^n \times \chi_r^k \times \chi_z \times \chi_r$ 15 : $u^* := r^\top \cdot A + z^\top$ 16 : $v_0^* := r^\top \cdot b + z + \mu^* \cdot [q/2]$ 17 : $v_1^* := r^\top \cdot b_r^* + z_r + \mu_r' \cdot [q/2]$ 18 : $\overline{u}^* := \text{Compress}_q(u, d_u)$ 19 : $\overline{v}_0^* := \text{Compress}_q(v_0, d_{v_0})$ 20 : $\overline{v}_1^* := \text{Compress}_q(v_1, d_{v_1}) \oplus \Delta$ 21 : $\text{ct}^* := (\overline{u}^*, \overline{v}_0^*, \overline{v}_1^*)$ 22 : $b' \xleftarrow{\\$} \mathcal{A}^{G(\cdot), G_{\text{Decomp}}(\cdot), G_r(\cdot)}(\text{ct}^*, \text{state})$ / If Game₃ 23 : $b' \xleftarrow{\\$} \mathcal{A}^{G(\cdot), G_{\text{Decomp}}(\cdot), F(\cdot)}(\text{ct}^*, \text{state})$ / If Game₄, Game₅ 24 : return b' </pre>	Game₇, Game₈ <pre> 1 : $\mu^* \xleftarrow{\\$} \mathcal{U}(\mathcal{M})$ 2 : $\Delta \xleftarrow{\\$} \mathbb{Z}_{2^{d_{v_1}}}^d$ / If Game₇ 3 : $(\text{seed}_A, A) \xleftarrow{\\$} \{0, 1\}^{256} \times R_q^{n \times k}$ 4 : Sample $2Q_G$-wise independent function G s.t. $G(\text{seed}_A) = A$ 5 : $b \xleftarrow{\\$} R_q^n$ 6 : $\text{pk}^* := (\text{seed}_A, b)$ 7 : $s \xleftarrow{\\$} \{0, 1\}^\ell$ 8 : $\text{sk}^* := (\text{seed}_A, s)$ 9 : $(\text{rpk}', \mu_r', \text{state}) \xleftarrow{\\$} \mathcal{A}^{G(\cdot), G_{\text{Decomp}}(\cdot), G_r(\cdot)}(\text{pk}^*)$ / If Game₈ 10 : $(\text{rpk}', \mu_r', \text{state}) \xleftarrow{\\$} \mathcal{A}^{G(\cdot), G_{\text{Decomp}}(\cdot), F(\cdot)}(\text{pk}^*)$ / If Game₇ 11 : $\text{ct}^* \xleftarrow{\\$} \mathcal{S}(1^\lambda, \text{rpk}', \mu_r')$ / From Game₇ 12 : $b' \xleftarrow{\\$} \mathcal{A}^{G(\cdot), G_{\text{Decomp}}(\cdot), G_r(\cdot)}(\text{ct}^*, \text{state})$ / If Game₈ 13 : $b' \xleftarrow{\\$} \mathcal{A}^{G(\cdot), G_{\text{Decomp}}(\cdot), F(\cdot)}(\text{ct}^*, \text{state})$ / If Game₇ 14 : return b' </pre> $\mathcal{A}_{\text{O2H}}^{H(\cdot)}(w)$ <pre> 1 : $\text{parse } (\text{pk}^*, G, \Delta) \leftarrow w$ 2 : $(\text{rpk}', \mu_r', \text{state}) \xleftarrow{\\$} \mathcal{A}^{G(\cdot), G_{\text{Decomp}}(\cdot), H(\cdot)}(\text{pk}^*)$ 3 : $\text{parse } \overline{b}_r \leftarrow \text{rpk}'$ 4 : $A := G(\text{seed}_A)$ 5 : $b_r^* \xleftarrow{\\$} \text{RandDecomp}_{q, d_{b_r}}(\overline{b}_r, d_{b_r})$ 6 : $(r, z, z_r) \xleftarrow{\\$} \chi_r^n \times \chi_r^k \times \chi_z \times \chi_r$ 7 : $u^* := r^\top \cdot A + z^\top$ 8 : $v_0^* := r^\top \cdot b + z + \mu^* \cdot [q/2]$ 9 : $v_1^* := r^\top \cdot b_r^* + z_r + \mu_r' \cdot [q/2] + \Delta$ 10 : $\overline{u}^* := \text{Compress}_q(u, d_u)$ 11 : $\overline{v}_0^* := \text{Compress}_q(v_0, d_{v_0})$ 12 : $\overline{v}_1^* := \text{Compress}_q(v_1, d_{v_1})$ 13 : $\text{ct}^* := (\overline{u}^*, \overline{v}_0^*, \overline{v}_1^*)$ 14 : $b' \xleftarrow{\\$} \mathcal{A}^{G(\cdot), G_{\text{Decomp}}(\cdot), H(\cdot)}(\text{ct}^*, \text{state})$ 15 : return b' </pre> $\mathcal{B}_{\text{O2H}}^{H(\cdot)}(w)$ <pre> 1 : $i \xleftarrow{\\$} [Q_F]$ 2 : Run $\mathcal{A}_{\text{O2H}}^{H(\cdot)}(w)$ until i-th query $\mu\rangle$ to H 3 : return Measurement of $\mu\rangle$ </pre>
---	---

Figure 12. Game₃ to Game₅, and Game₇ and Game₈ for the proof of C-IND security, and tools to invoke O2H lemma.

$\mathcal{S} := \{\mu^*\} \subset \mathcal{M}$ for a randomly chosen μ^* . We define the random variable w as (pk^*, G, Δ) , where $\Delta := G_r(\mu^*)$. This is without loss of generality as z can be any bit string that can depend arbitrary on G_r and $\mathcal{S}$. We then define the oracle-calling algorithms $\mathcal{A}_{\text{O2H}}$ and $\mathcal{B}_{\text{O2H}}$ as in Fig. 12. If $H := G_r$ (resp. $H := F$), then $\mathcal{A}_{\text{O2H}}^{H(\cdot)}(w)$ is identical to Game₃ (resp. Game₄). Then, using Lemma 4, we have

$$\begin{aligned}
& |\Pr[E_3] - \Pr[E_4]| \\
& \leq 2 \cdot Q_{G_r} \cdot \sqrt{\Pr_{G_r, F, \mathcal{S}, w} \left[\mathcal{T} \xleftarrow{\$} \mathcal{B}_{\text{O2H}}^{F(\cdot)}(w) : \mathcal{S} \cap \mathcal{T} \neq \emptyset \right]}.
\end{aligned}$$

It remains to upper bound probability inside the square root, which we denote by p . We only provide a proof sketch as this follows from an almost identical argument

to the remaining game transitions from Game₄ to Game₇ in the main proof. Namely, first observe that p remains the same even if we modify $\mathcal{A}_{\text{O2H}}^{F(\cdot)}(w)$ to set $v_1^* := \Delta$ instead. This is because from $\mathcal{A}$'s point of view, $\Delta := G_r(\mu^*)$ is information theoretically hidden, and Δ is distributed uniformly over R_q . We further modify $\mathcal{A}_{\text{O2H}}^{F(\cdot)}(w)$ to set $(u^*, v_0^*) \xleftarrow{\$} R_q^k \times R_q$ instead. We can show that p only changes negligibly due to the $\text{MLWE}_{q, k+1, n, \chi_r, (\chi_i)_{i \in [k+1]}}$ assumption. That is, if $\mathcal{B}_{\text{O2H}}^{F(\cdot)}(w)$ outputs a set $\mathcal{T}$ that includes μ^* with different probabilities, then it can be used to solve the $\text{MLWE}_{q, k+1, n, \chi_r, (\chi_i)_{i \in [k+1]}}$ problem. Lastly, at this point, since μ^* is information theoretically hidden from $\mathcal{A}$, the probability that $\mathcal{B}_{\text{O2H}}^{F(\cdot)}(w)$ outputs μ^* is at most $\frac{1}{|\mathcal{M}|}$. This

concludes the proof. $\square$

$\square$

E.2. CR-IND Security of RPKE₂

Proof. We consider a sequence of hybrid games, where we denote E_i as the event that $\mathcal{A}$ outputs $b' = 1$ in Game_i . The simulator S_r and some of the critical game transitions are depicted in Fig. 13.

Game₀: This corresponds to the CR-IND game where $b = 0$ is sampled. We thus have

$$\Pr[E_0] = 1 - \Pr[\text{Game}_{\text{RPKE}_2, \mathcal{A}}^{\text{CR-IND}}(1^\lambda) = 1 \mid b = 0].$$

Game₁: This is the same as Game_0 except that the game modifies how the random oracle G is simulated. See Fig. 12 for the exact description. This argument is identical to the one we used to move from Game_0 to Game_1 in the proof of Thm. 9. Thus we have

$$\Pr[E_0] = \Pr[E_1].$$

Game₂: This is the same game as Game_1 except that the game replaces $\mathbf{b}_r$ by a uniform sample in R_q^n . Since the reinforcing secret key $\mathbf{s}_r$ is not used elsewhere than to compute $\mathbf{b}_r$, this reduces to the MLWE problem. There exists an adversary $\mathcal{B}_1$ against the $\text{MLWE}_{q,k,n,\chi,\chi}$ problem such that

$$|\Pr[E_1] - \Pr[E_2]| \leq \text{Adv}_{\mathcal{B}_1}^{\text{MLWE}}(1^\lambda).$$

Game₃: This is the same as Game_2 except that the game modifies how the random oracle G_r is simulated, and how the reinforcing public key $\bar{\mathbf{b}}_r$ is sampled. It first samples $\mathbf{b}_r^* \xleftarrow{\$} R_q^n$ and replaces the oracle G_{Decomp} by an efficient simulator Sim that simulates G_{Decomp} conditioned on $G_{\text{Decomp}}(\bar{\mathbf{b}}_r) = \mathbf{b}_r^*$ where $\bar{\mathbf{b}}_r = \text{Compress}_q(\mathbf{b}_r^*, d_r)$. Due to the complex output distribution of G_{Decomp} that furthermore depends on the input, defining such simulator is more challenging than in previous proofs where we could directly use $2Q$ -wise independent functions. We prove in Lemma 12 that we can define an efficient simulator for this conditioned oracle that succeeds with all but negligible probability $\text{negl}(\lambda)$. We thus have

$$|\Pr[E_2] - \Pr[E_3]| = \text{negl}(\lambda).$$

Game₄: This is the same as Game_3 except that the game modifies how v_0^* in the challenge ciphertext is computed. See Fig. 12 for the exact description. Notice that

$$\begin{aligned} v_0^* - \mu' \cdot \lfloor q/2 \rfloor &= \mathbf{r}^\top \cdot \mathbf{b} + z \\ &= \mathbf{r}^\top \cdot (\mathbf{A} \cdot \mathbf{s} + \mathbf{x}) + z \\ &= \mathbf{u}^* \cdot \mathbf{s} - \mathbf{z}^\top \cdot \mathbf{s} + \mathbf{r}^\top \cdot \mathbf{x} + z, \end{aligned}$$

where the first (resp. last) line is how v_0^* is computed in Game_3 (resp. Game_4). Thus, we have

$$\Pr[E_3] = \Pr[E_4].$$

Game₅: This is the same as Game_4 except that the challenge ciphertext ct^* is sampled by the simulator S_r . See Fig. 12

for the exact description of the game and the simulator. We later show in Lemma 13 that there exists an adversary $\mathcal{B}$ against the $\text{hint-MLWE}_{q,k,n+1,\chi_r,\chi_z,\mathcal{F}}$ problem such that

$$|\Pr[E_4] - \Pr[E_5]| \leq \text{Adv}_{\mathcal{B}}^{\text{hint-MLWE}}(1^\lambda).$$

The remaining games is to undo our modification.

Game₆: This is the same as Game_5 except that the game computes the random oracle G_{Decomp} as in the real game. Using an identical argument to move from Game_0 to Game_1 and Game_2 to Game_3 , we have

$$|\Pr[E_6] - \Pr[E_7]| = \text{negl}(\lambda).$$

Game₇: This is the same as Game_6 except that the game samples rpk^* as in the real game. See Fig. 12 for the exact description. Using an identical argument to move from Game_1 to Game_2 , there exists an adversary $\mathcal{B}_3$ against the $\text{MLWE}_{q,k,n,\chi,\chi}$ problem such that

$$|\Pr[E_6] - \Pr[E_7]| \leq \text{Adv}_{\mathcal{B}_3}^{\text{MLWE}}(1^\lambda).$$

Game₈: This is the same as game Game_7 except that the game computes the random oracle G as in the real game. Using an identical argument to move from Game_0 to Game_1 , we have $\Pr[E_7] = \Pr[E_8]$. This corresponds to the C-IND game where $b = 1$ is sampled. We thus have

$$\Pr[E_8] = \Pr[\text{Game}_{\text{RPKE}_2, \mathcal{A}}^{\text{CR-IND}}(1^\lambda) = 1 \mid b = 1].$$

Combining all the bounds together, we arrive at the bound in the theorem statement. It remains to prove the following lemmas to complete the proof.

Lemma 12. *Let $q = \text{poly}(\lambda)$. For any (possibly unbounded) adversary $\mathcal{A}$ making at most Q_{Decomp} (quantum) oracle queries, there is an efficient simulator Sim that samples a function G_{Decomp}^* with polynomial-sized description such that:*

$$\begin{aligned} &\left| \Pr \left[\begin{array}{l} \mathbf{b}_r \xleftarrow{\$} R_q^n, \\ \bar{\mathbf{b}}_r := \text{Compress}_q(\mathbf{b}_r, d_{b_r}), : \mathcal{A}^{G_{\text{Decomp}}}(\bar{\mathbf{b}}_r, \mathbf{b}_r^*) = 1 \end{array} \right] \right. \\ &\quad \left. - \Pr \left[\begin{array}{l} \mathbf{b}_r^* \xleftarrow{\$} R_q^n, \\ \bar{\mathbf{b}}_r := \text{Compress}_q(\mathbf{b}_r^*, d_{b_r}), : \mathcal{A}^{G_{\text{Decomp}}^*}(\bar{\mathbf{b}}_r, \mathbf{b}_r^*) \rightarrow 1 \end{array} \right] \right| = \text{negl}(\lambda), \end{aligned}$$

where the probability is also taken over the randomness of sampling the random oracle G_{Decomp} .

Proof. Let us define the left (resp. right) hand probability in the lemma statement as p_{real} and p_{sim} . Recall that given an input $\bar{\mathbf{b}}_r$, the random oracle G_{Decomp} samples a random element $\mathbf{b}_r^*$ in R_q such that $\text{Compress}_q(\mathbf{b}_r^*) = \bar{\mathbf{b}}_r$, i.e., $\mathbf{b}_r^* \xleftarrow{\$} \mathcal{U}(\{\mathbf{b}' \in R_q^n \mid \text{Compress}_q(\mathbf{b}', d_{b_r}) = \bar{\mathbf{b}}_r\})$. When $\mathbf{b}_r \xleftarrow{\$} R_q^n$, we can see that $\mathbf{b}_r^*$ is also uniformly distributed over R_q^n conditioned on $\bar{\mathbf{b}}_r = \text{Compress}_q(\mathbf{b}_r^*, d_{b_r})$. Put differently, there is an *inefficient* simulator Sim' that first samples $\mathbf{b}_r^*$, computes $\bar{\mathbf{b}}_r = \text{Compress}_q(\mathbf{b}_r^*, d_{b_r})$, and samples a random oracle G'_{Decomp} conditioned on $G'_{\text{Decomp}}(\bar{\mathbf{b}}_r) = \mathbf{b}_r^*$ such that

$$p_{\text{real}} = \Pr \left[\begin{array}{l} \mathbf{b}_r^* \xleftarrow{\$} R_q^n, \\ \bar{\mathbf{b}}_r := \text{Compress}_q(\mathbf{b}_r^*, d_{b_r}), : \mathcal{A}^{G_{\text{Decomp}}}(\bar{\mathbf{b}}_r, \mathbf{b}_r^*) \rightarrow 1 \end{array} \right],$$

$$p_{\text{sim}} = \Pr \left[\begin{array}{l} \mathbf{b}_r^* \xleftarrow{\$} R_q^n, \\ \bar{\mathbf{b}}_r := \text{Compress}_q(\mathbf{b}_r^*, d_{b_r}), : \mathcal{A}^{G'_{\text{Decomp}}}(\bar{\mathbf{b}}_r, \mathbf{b}_r^*) \rightarrow 1 \end{array} \right],$$

Game₁, Game₂ 1 : $(\text{seed}_A, \mathbf{A}) \xleftarrow{\$} \{0, 1\}^{256} \times R_q^{n \times k}$ 2 : Sample $2Q_G$ -wise independent function G s.t. $G(\text{seed}_A) = \mathbf{A}$ 3 : $(\text{pk}^*, \text{sk}^*) \xleftarrow{\$} \text{KeyGen}_2(1^\lambda)$ 4 : $\text{parse}(\text{seed}_A, \mathbf{b}) \leftarrow \text{pk}^*$ 5 : $\mathbf{A} := G(\text{seed}_A)$ 6 : $(\mathbf{s}_r, \mathbf{x}_r) \xleftarrow{\$} \chi^k \times \chi^n$ / If Game ₁ 7 : $\mathbf{b}_r := \mathbf{A} \cdot \mathbf{s}_r + \mathbf{x}_r$ / If Game ₁ 8 : $\mathbf{b}_r \xleftarrow{\$} R_q^n$ / If Game ₂ 9 : $\overline{\mathbf{b}}_r := \text{Compress}_q(\mathbf{b}_r, d_{b_r})$ 10 : $(\text{rpk}^*, \text{rsk}^*) := (\overline{\mathbf{b}}_r, \mathbf{s}_r)$ / If Game ₁ 11 : $\text{rpk}^* := \overline{\mathbf{b}}_r$ / If Game ₂ 12 : $(\mu', \text{state}) \xleftarrow{\$} \mathcal{A}^{G(\cdot), G_{\text{Decomp}}, G_r(\cdot)}(\text{pk}^*, \text{sk}^*, \text{rpk}^*)$ 13 : $\mu_r^* \xleftarrow{\$} \mathcal{U}(\mathcal{M})$ 14 : $\mathbf{b}_r \xleftarrow{\$} G_{\text{Decomp}}(\overline{\mathbf{b}}_r, d_{b_r})$ 15 : $(\mathbf{r}, \mathbf{z}, z, z_r) \xleftarrow{\$} \chi_r^n \times \chi_r^k \times \chi_z \times \chi_r$ 16 : $\Delta := G_r(\mu')$ 17 : $\mathbf{u}^* := \mathbf{r}^\top \cdot \mathbf{A} + \mathbf{z}^\top$ 18 : $v_0^* := \mathbf{r}^\top \cdot \mathbf{b} + z + \mu' \cdot \lfloor q/2 \rfloor$ 19 : $v_1^* := \mathbf{r}^\top \cdot \mathbf{b}_r^* + z_r + \mu_r^* \cdot \lfloor q/2 \rfloor + \Delta$ 20 : $\overline{\mathbf{u}}^* := \text{Compress}_q(\mathbf{u}, d_u)$ 21 : $\overline{v}_0^* := \text{Compress}_q(v_0, d_{v_0})$ 22 : $\overline{v}_1^* := \text{Compress}_q(v_1, d_{v_1})$ 23 : $\text{ct}^* := (\overline{\mathbf{u}}^*, \overline{v}_0^*, \overline{v}_1^*)$ 24 : $b' \xleftarrow{\$} \mathcal{A}^{G(\cdot), G_{\text{Decomp}}, G_r(\cdot)}(\text{ct}^*, \text{state})$ 25 : return b' Simulator $\mathcal{S}_r(1^\lambda, \text{pk}^*, \text{sk}^*, \text{rpk}^*, \mu')$ 1 : $\text{parse}(\text{seed}_A, \mathbf{b}) \leftarrow \text{pk}^*$ 2 : $\text{parse}(\text{seed}_A, \mathbf{s}) \leftarrow \text{sk}^*$ 3 : $\mathbf{A} := G(\text{seed}_A)$ 4 : $\mathbf{x} := \mathbf{b} - \mathbf{A} \cdot \mathbf{s}$ 5 : $(\mathbf{r}, \mathbf{z}, z) \xleftarrow{\$} \chi_r^n \times \chi_r^k \times \chi_z$ 6 : $\mathbf{u}^* \xleftarrow{\$} R_q^k$ 7 : $v_0^* := \mathbf{u}^* \cdot \mathbf{s} - \mathbf{z}^\top \cdot \mathbf{s} + \mathbf{r}^\top \cdot \mathbf{x} + z + \mu' \cdot \lfloor q/2 \rfloor$ 8 : $v_1^* \xleftarrow{\$} R_q$ 9 : $\overline{\mathbf{u}}^* := \text{Compress}_q(\mathbf{u}, d_u)$ 10 : $\overline{v}_0^* := \text{Compress}_q(v_0, d_{v_0})$ 11 : $\overline{v}_1^* := \text{Compress}_q(v_1, d_{v_1})$ 12 : $\text{ct}^* := (\overline{\mathbf{u}}^*, \overline{v}_0^*, \overline{v}_1^*)$ 13 : return ct^*	Game₃, Game₄ / Identical up to Game ₂ , Line 4 5 : $\mathbf{b}_r^* \xleftarrow{\$} R_q^n$ 6 : $\overline{\mathbf{b}}_r := \text{Compress}_q(\mathbf{b}_r^*, d_{b_r})$ 7 : $G_{\text{Decomp}} \xleftarrow{\$} \text{Sim}(\overline{\mathbf{b}}_r, \mathbf{b}_r^*)$ 8 : $\mathbf{A} := G(\text{seed}_A)$ 9 : $\text{rpk}^* := \overline{\mathbf{b}}_r$ 10 : $(\mu', \text{state}) \xleftarrow{\$} \mathcal{A}^{G(\cdot), G_{\text{Decomp}}, G_r(\cdot)}(\text{pk}^*, \text{sk}^*, \text{rpk}^*)$ 11 : $\mu_r^* \xleftarrow{\$} \mathcal{U}(\mathcal{M})$ 12 : $(\mathbf{r}, \mathbf{z}, z, z_r) \xleftarrow{\$} \chi_r^n \times \chi_r^k \times \chi_z \times \chi_r$ 13 : $\Delta := G_r(\mu')$ 14 : $\mathbf{u}^* := \mathbf{r}^\top \cdot \mathbf{A} + \mathbf{z}^\top$ 15 : $v_0^* := \mathbf{r}^\top \cdot \mathbf{b} + z + \mu' \cdot \lfloor q/2 \rfloor$ / If Game ₃ 16 : $v_0^* := \mathbf{u}^* \cdot \mathbf{s} - \mathbf{z}^\top \cdot \mathbf{s} + \mathbf{r}^\top \cdot \mathbf{x} + z + \mu' \cdot \lfloor q/2 \rfloor$ / If Game ₄ 17 : $v_1^* := \mathbf{r}^\top \cdot \mathbf{b}_r^* + z_r + \mu_r^* \cdot \lfloor q/2 \rfloor + \Delta$ / Identical to Game ₂ , Line 20 onward Game₅ / Identical up to Game ₄ , Line 7 8 : $\mathbf{b}_r^* \xleftarrow{\$} R_q^n$ 9 : $\overline{\mathbf{b}}_r := \text{Compress}_q(\mathbf{b}_r^*, d_{b_r})$ 10 : $G_{\text{Decomp}} \xleftarrow{\$} \text{Sim}(\overline{\mathbf{b}}_r, \mathbf{b}_r^*)$ 11 : $\mathbf{A} := G(\text{seed}_A)$ 12 : $\text{rpk}^* := \overline{\mathbf{b}}_r$ 13 : $(\mu', \text{state}) \xleftarrow{\$} \mathcal{A}^{G(\cdot), G_{\text{Decomp}}, G_r(\cdot)}(\text{pk}^*, \text{sk}^*, \text{rpk}^*)$ 14 : $\text{ct}^* \xleftarrow{\$} \mathcal{S}_r(1^\lambda, \text{pk}^*, \text{sk}^*, \text{rpk}^*, \mu')$ 15 : $b' \xleftarrow{\$} \mathcal{A}^{G(\cdot), G_{\text{Decomp}}, G_r(\cdot)}(\text{ct}^*, \text{state})$ 16 : return b'
--	--

Figure 13. Game₁ to Game₅ for the proof of CR-IND security, and the description of simulator $\mathcal{S}_r$.

where denote the right hand probability as p' .

It remains to prove that we can replace Sim' with an efficient algorithm sampling an efficiently computable G'_{Decomp} . The core idea relies on the observation that we can sample G'_{Decomp} only with access to a uniform sampler over $\mathbb{Z}_q$, using rejection sampling. Indeed, we can equivalently (independently) sample the i -th coordinate of $\mathbf{b}_r^*$, denoted as $(\overline{\mathbf{b}}_r)_i$, by sampling an element c in $\mathbb{Z}_q$ and restarting as long as $\text{Compress}_q(c) \neq (\overline{\mathbf{b}}_r)_i$.

As a first step, we consider a simulator Sim'' that samples a random oracle $G_{\text{Decomp}}^{(\kappa)}$ which samples its outputs with the above method, however limiting the number of attempts per coordinate to κ – setting $\perp$ in case of failure. Furthermore, we consider that $G_{\text{Decomp}}^{(\kappa)}$ calls a random oracle on its input for each pair of coordinate and attempt number instead of a sampler (with an efficient description) over $\mathbb{Z}_q$. We implement the condition $G'_{\text{Decomp}}(\overline{\mathbf{b}}_r) = \mathbf{b}_r^*$ by conditioning the first random oracle for each coordinate i to take as value $(\mathbf{b}_r^*)_i$ on $(\overline{\mathbf{b}}_r)_i$. Then, since the depth of the circuit necessary to implement the rejection sampling with κ attempts is $\text{poly}(\lambda, \kappa)$, we can aggregate the results of the individual quantum oracles into a final result with an efficient quantum circuit.

Let's now evaluate the probability that the sampling of $G_{\text{Decomp}}^{(\kappa)}$ succeeds on all possible inputs, as it then has the same distribution as G_{Decomp} . One coordinate sampling attempt succeeds with probability at least $1/q$. After κ attempts, the probability of still failing for one coordinate is thus $(1 - 1/q)^\kappa \leq e^{-\kappa/q}$. By union-bound the probability of failing on any of d coordinates is bounded by $d \cdot e^{-\kappa/q}$.

Going further, since there are at most $2^{d \cdot d_{v_1}}$ possible inputs for this function, the probability of successful sampling of $G_{\text{Decomp}}^{(\kappa)}$ is at least $1 - 2^{d \cdot d_{v_1}} \cdot d \cdot e^{-\kappa/q}$ by union bound. We set $\kappa = \omega(q \cdot d \log d \cdot d_{v_1})$ so that this probability is overwhelming. In that case, we deduce

$$\left| p' - \Pr \left[\begin{array}{l} \mathbf{b}_r^* \xleftarrow{\$} R_q^n, \\ \mathbf{b}_r := \text{Compress}_q(\mathbf{b}_r^*, d_{b_r}), : \mathcal{A}^{G'_{\text{Decomp}}(\overline{\mathbf{b}}_r, \mathbf{b}_r^*)} \rightarrow 1 \end{array} \middle| \begin{array}{l} G_{\text{Decomp}}^{(\kappa)} \xleftarrow{\$} \text{Sim}''(\overline{\mathbf{b}}_r, \mathbf{b}_r^*) \end{array} \right] \right| = \text{negl}(\lambda),$$

where denote the right hand probability as p'' .

Lastly, we define $\text{Sim}(\overline{\mathbf{b}}_r, \mathbf{b}_r^*)$ which works as Sim'' but replaces each random oracle H over $\mathbb{Z}_q$ (with a conditioning $H(c) = c'$ on one input) using a $2Q_{\text{Decomp}}$ -wise independent function. This oracle is G_{Decomp}^* , which clearly has a polynomial-sized description. As in previous proofs, one way to sample such function would be to first sample a random $2Q_{\text{Decomp}}$ -wise independent function H' (which is possible since q is a prime, cf. [81, Theorem 6.1]). Then simply define H as computing H' and adding $c' - H(c)$.

Using [81, Theorem 6.1], a $2Q_{\text{Decomp}}$ -wise independent function is perfectly indistinguishable from a random oracle to an adversary making at most $2Q_{\text{Decomp}}$ oracle queries. Hence, we have

$$p'' = p_{\text{sim}}.$$

□

Lemma 13. *There exists an adversary $\mathcal{B}$ against the hint-MLWE $_{q,k+1,n,\chi_r,\chi_z,\mathcal{F}}$ problem such that*

$$|\Pr[E_4] - \Pr[E_5]| \leq \text{Adv}_{\mathcal{B}}^{\text{hint-MLWE}}(1^\lambda).$$

Proof. The proof follows immediately from the hint-MLWE problem. Assume $\mathcal{B}$ is given as challenge $(\mathbf{B}, \mathbf{d}, [\mathbf{x}^\top \mid -\mathbf{s}^\top \mid 0], h)$, where $\mathbf{B} \in R_q^{(k+1) \times n}$, $\mathbf{d} \in R_q^{k+1}$, and $(\mathbf{s}, \mathbf{x}) \xleftarrow{\$} \chi^k \times \chi^n$. $\mathcal{B}$ then sets the last row of $\mathbf{B}$ as $\mathbf{b}_r^{*\top}$ and the remaining matrix as $\mathbf{A}^\top$. Similarly, it sets the last entry of $\mathbf{d}$ as v_1 and the remaining column vector as $\mathbf{u}^*$. It then simulates the Game_4 challenger with the provided $\mathbf{A}$. When the adversary $\mathcal{A}$ outputs (μ', state) , $\mathcal{B}$ uses the provided $\mathbf{u}^*$, and computes $v_0^* := \mathbf{u}^* \cdot \mathbf{s} + h + \mu' \cdot \lfloor q/2 \rfloor$ and $v_1^* := v_1 + \mu_r^* \cdot \lfloor q/2 \rfloor + \Delta$. It then sets the challenge ciphertext ct^* as in Game_4 and completes the simulation. $\mathcal{B}$ then outputs whatever output by $\mathcal{A}$ as its guess.

It can be checked that $\mathcal{B}$ perfectly simulates Game_4 (resp. Game_5) if $\mathbf{d}$ is a valid (resp. invalid). Thus, we arrive at the bound in the lemma statement. Concretely, if $\mathbf{d} = \begin{bmatrix} \mathbf{A}^\top \\ \mathbf{b}_r^{*\top} \end{bmatrix} \cdot \mathbf{r} + \begin{bmatrix} \mathbf{z} \\ z_r \end{bmatrix}$, then it can be checked that $h = -\mathbf{z}^\top \cdot \mathbf{s} + \mathbf{r}^\top \cdot \mathbf{x} + z$ as desired, where $z \xleftarrow{\$} \chi_z$. Note that the random oracle G_r can be perfectly simulated by a $2Q_{G_r}$ -wise independent function as we did for G . □

□

E.3. Statistical Disjointness of $\overline{\text{RPKE}}_2$

This is a simple consequence of the counting argument.

Lemma 14. *Let $B > 0$ such that $\Pr[x \xleftarrow{\$} \chi_r : |x| > B] \leq \text{negl}(\lambda)$, $c := \max\{n, k\}$, and $T := \max\{(2 \cdot B + 1)^2, (2 \cdot \lfloor \frac{q}{2^{d_{v_0}+1}} \rfloor + 1)^2\}$. Then, our $\overline{\text{RPKE}}_2$ is statistical disjoint assuming $\left(\frac{T^c}{q^k}\right)^d = \text{negl}(\lambda)$.*

Proof. Recall from [Thm. 9](#) that the simulator $\mathcal{S}$ takes as input $(1^\lambda, \text{rpk}, \mu_r)$, samples random $(\mathbf{u}, v_0, v_1) \xleftarrow{\$} R_q^k \times R_q \times R_q$, and outputs a simulated ciphertext ct as $(\overline{\mathbf{u}}, \overline{v_0}, \overline{v_1}) := (\text{Compress}_q(\mathbf{u}, d_u), \text{Compress}_q(v_0, d_{v_0}), \text{Compress}_q(v_1, d_{v_1}))$. In contrast, $\mathbf{u}$ in a valid ciphertext generated using $\overline{\text{RPKE}}_2$ is set as $\mathbf{u} = \mathbf{r}^\top \cdot \mathbf{A} + \mathbf{z}^\top$ for $(\mathbf{r}, \mathbf{z}) \xleftarrow{\$} \chi_r^n \times \chi_r^k$. If simulated and valid ciphertexts agree, then we must have

$$\mathbf{u} + \delta_{\text{sim}} = \mathbf{A}^\top \cdot \mathbf{r} + \mathbf{z} + \delta_{\text{val}}, \quad (1)$$

where δ_{sim} and δ_{val} are offsets due to the decompression, whose infinity norm is smaller than $\lfloor \frac{q}{2^{d_{v_0}+1}} \rfloor$ due to [Lemma 1](#). Using a simple counting argument, $\mathbf{A}^\top \cdot \mathbf{r} + \mathbf{z} + \delta_{\text{val}} - \delta_{\text{sim}}$ can take at most $(2 \cdot B + 1)^{nd} \cdot (2 \cdot B + 1)^{kd}$. $(2 \cdot \lfloor \frac{q}{2^{d_{v_0}+1}} \rfloor + 1)^{2kd} \leq T^{cd}$. Assuming $(T/q)^{cd} = \text{negl}(\lambda)$, the probability that [Eq. \(1\)](#) holds for $\mathbf{u} \xleftarrow{\$} R_q^k$ is negligible as desired. □

E.4. (Half-)Correctness of RPKE_2

Correctness can be shown as for ML-KEM [19], although we additionally need to take into account rounding terms for the reinforcing public key.

Lemma 15. *Our RPKE_2 is half-correct assuming*

$$\Pr [\|\mathbf{r}^\top \cdot \mathbf{x} - (\mathbf{z} + \boldsymbol{\delta})^\top \cdot \mathbf{s} + z + \delta_0\|_\infty \leq q/4] = 1 - \text{negl}(\lambda),$$

where the probability is taken over the randomness to sample $(\mathbf{s}, \mathbf{x}) \stackrel{\$}{\leftarrow} \chi^k \times \chi^n$, $(\mathbf{r}, \mathbf{z}) \stackrel{\$}{\leftarrow} \chi_r^n \times \chi_r^k$, $z \stackrel{\$}{\leftarrow} \chi_z$, and $(\delta, \delta_0) \stackrel{\$}{\leftarrow} \chi_{\text{round}}(d_u) \times \chi_{\text{round}}(d_{v_0})$. Here, $\chi_{\text{round}}(c)$ is the distribution $\{\delta = \text{Decompress}_q(\text{Compress}_q(\mathbf{b}, c), c) - \mathbf{b} \mid \mathbf{b} \leftarrow \mathcal{U}(R_q)\}$.

Moreover, it is correct further assuming the above holds where we replace d_{v_0} by $d_{v_1} \stackrel{s}{\leftarrow} \chi_{\text{round}}(d_{v_1})$, sample $\mathbf{z}$ from $\chi_{\mathbf{r}}$, and add a term $\mathbf{r}^\top \cdot \delta_{\text{rpk}}$ where δ_{rpk} is sampled from $\{\mathbf{b}' - \mathbf{b} \mid \mathbf{b}, \mathbf{b}' \leftarrow \mathcal{U}(R_q) \wedge \text{Compress}_q(\mathbf{b}, d_{v_1}) = \text{Compress}_q(\mathbf{b}', d_{v_1})\}$.

Proof. We first provide the proof for half-correctness. Observe that we have

$$\begin{aligned} & v_0 - \mathbf{u}^\top \cdot \mathbf{s} \\ &= \mathbf{r}^\top \cdot (\mathbf{A} \cdot \mathbf{s} + \mathbf{x}) + z + \mu \cdot \lfloor q/2 \rfloor - (\mathbf{r} \cdot \mathbf{A}^\top + \mathbf{z}^\top)^\top \cdot \mathbf{s} \\ &= \mathbf{r}^\top \cdot \mathbf{x} - \mathbf{z}^\top \cdot \mathbf{s} + z + \mu \cdot \lfloor q/2 \rfloor. \end{aligned}$$

Due to rounding and [Lemma 1](#), we further have $\hat{\mathbf{u}} := \text{Decompress}_q(\bar{\mathbf{u}}, d_u) = \mathbf{u} + \delta$ and $\hat{v}_0 := \text{Decompress}_q(\bar{v}_0, d_{v_0}) = v_0 + \delta_0$, where δ and δ_0 are from the lemma statement. Plugging this into the above equation, it can be checked that RHalfDec_2 correctly decodes μ under the assumption made in the lemma.

The case for correctness follows straightforwardly considering also the reinforcing ciphertext, and the compression plus decompression of the reinforcing public key which add a term δ_{rpk} to $\mathbf{A} \cdot \mathbf{s} + \mathbf{x}$ in v_1 uniformly sampled from $\{\mathbf{b}' - \mathbf{b} \mid \text{Compress}_q(\mathbf{b}, d_{v_1}) = \text{Compress}_q(\mathbf{b}', d_{v_1})\}$. $\square$

Appendix F.

(Reinforced) KEM-based Noise Protocols

The Noise Protocol Framework [4] gives key exchange protocols based on Diffie–Hellman. WireGuard is based on the Noise “IKpsk2” pattern. There exists a preliminary specification [24] and prior academic work [13] which show how to convert DH-based Noise to KEM-based protocols. However, they did not show how to handle PSK in PQ Noise exchanges. Use of pre-shared keys triggers special handling of DH keys in classic Noise, which leaves the question on how to handle this for KEM-based approaches. Additionally, since [13] was published, Cremers et al. [2] have introduced the concept of binding properties for KEMs. By including public keys in the C_i key derivation chain, the derived K_i will bind them (achieving *MAL-BIND-K-PK*, [2, Sec. 4.8]). Otherwise, an adversary is able to set up two sessions that have the same shared key but disagree on identities (c.f. Sec. B). We also show how to extend Noise protocols with reinforced KEM.

KEM Noise patterns. In classic Noise, the `e` and `s` messages indicate the transmission of ephemeral and static public keys. To complete ephemeral key agreement, both parties include `e` tokens in their messages, after which they trigger the key agreement computation by the `ee` token. For KEM-based key exchange, we will instead need to compute and transmit a ciphertext message. As done in [13], [24], we replace the responder's `e` and `ee` tokens by a single `ekem` token, which will indicate the encapsulation, transmission, and decapsulation of the ciphertext and the inclusion of the shared secret into the key schedule. For static key exchange, there is a similar `skem` token. Fig. 14 shows an example.

-> e, s	-> e, s
<- e, ee, se	<- ekem, skem

Figure 14. Translation to PQ tokens in the IN pattern.

Key schedule computations. Noise attempts to ensure that different sessions compute unique keys. In the description of the PSK extension to classic Noise, [4, Sec. 9.2] states that in PSK handshakes, all ephemeral public keys are included in the key schedule to ensure that PSK-derived session keys are randomized. This reasoning can be extended to the use of KEMs. In addition, without relying on binding properties of the KEM [2], we can not assume that any computed shared secrets are unique to a particular public key, which could lead to similar results as with shared PSK keys. So to ensure randomized keys and guard against misbinding, we require that e and s tokens are followed by `MixKey` operations which update the C_i chaining key with the transmitted public key. Because sharing ciphertexts will immediately result in a (sufficiently random) shared secret being included in the chaining key, we do not include ciphertexts into the chaining key; ciphertexts are still included in computations of H_i and thus authenticated. The key schedule computations are shown in Fig. 15. Note that both [13], [24] only update H_i for the public keys and ciphertexts. There are other possible solutions to bind the K_i , such as deriving K_i from both C_i and H_i , but we leave exploring this to future work.

F.1. Noise Patterns for KEM Key Exchange

Below, we give the updated descriptions of the operations to process the `e`, `ekem`, `s`, and `skem` tokens in Noise handshake patterns. The precise details of the referenced




Figure 15. PQ Noise key computations with KEM (left) and RKEM (right)

functions are given in [4, §5]; these new rules update the definition of `WriteMessage` specifically. We encrypt ciphertexts encapsulated to static keys following [13]; we differ mainly in the calls to `MixKey()` for all public keys.

Symbol $\mathbf{e}$: Generate $(\text{epk}, \text{esk}) \xleftarrow{\$} \text{KEM.KeyGen}$, transmit epk , and call `MixHash(epk)` and `MixKey(epk)` to update H_i and C_i , respectively.

Symbol $\mathbf{ekem}$: Generate $(\text{ct}, \text{ss}) \xleftarrow{\$} \text{KEM.Encaps}(\text{epk})$ and transmit ct , or receive and decapsulate ct to get the shared secret $\text{ss} \leftarrow \text{KEM.Decaps}(\text{esk}, \text{ct})$. Next, call `MixHash(ct)` to update H_i . Update the key schedule by calling `MixKey(ss)`.

Symbol $\mathbf{s}$: Encrypt and transmit static public key spk by calling `EncryptAndHash(spk)`, which also updates H_i , and call `MixKey(spk)` to update C_i .

Symbol $\mathbf{skem}$: Generate $(\text{ct}, \text{ss}) \xleftarrow{\$} \text{KEM.Encaps}(\text{spk})$ and encrypt and transmit ct by calling `EncryptAndHash(ct)`, or decrypt and decapsulate ct to get the shared secret $\text{ss} \leftarrow \text{KEM.Decaps}(\text{ssk}, \text{ct})$. Note that this implicitly updates the hash transcript. Update the key schedule by calling `MixKey(ss)`.

F.2. Extending Post-Quantum Noise With Reinforced KEM

We define the following symbols for reinforced KEM:

Symbol $\mathbf{r}$: Generate $(\text{rpk}, \text{rsk}) \xleftarrow{\$} \text{RKeyGen}(\text{spk})$, transmit rpk , and call `MixHash(rpk)` and `MixKey(rpk)` to update H_i and C_i , respectively.

Symbol $\mathbf{rkem}$: Call $(\text{ct}, \text{ss}) \xleftarrow{\$} \text{KEM.Encaps}(\text{spk}, \text{rpk})$ and encrypt and transmit ct by calling `EncryptAndHash(ct)`, or decrypt and decapsulate ct to get the shared secret $\text{ss} \leftarrow \text{RDec}(\text{ssk}, \text{rsk}, \text{ct})$. Note that this implicitly updates the hash H_i . Update the key schedule by calling `MixKey(ss)`.

We can define the `rkem` extension that transforms Noise patterns by the following rules. An example is shown in Fig. 16, and the key schedule computations in Fig. 15.

- 1) Any pair of `ekem`, `skem` or `skem`, `ekem` tokens can be replaced by a single `rkem` token.
- 2) If the previous replacement was made, the `e` token which should be present in the prior message can be replaced by a `r` token.

```

-> e, s          -> r, s
<- ekem, skem   <- rkem

```

Figure 16. Translation to RKEM tokens in the `pqIN` pattern.

Contents

1	Introduction	1
1.1	Our Contribution	2
1.2	Related Work	3
1.3	Concurrent and Independent Work . .	3

2	Preliminaries	4
2.1	Notation	4
2.2	Lattices	4
2.3	Rounding	4
2.4	Quantum Tools	5
2.5	General Cryptographic Primitives . .	5
3	Reviewing PQ WireGuard by Hülsing et al.	6
3.1	Construction Overview	6
3.2	Security Properties	7
4	Computational Model for PQ WireGuard	8
4.1	Syntax and Execution Environment .	8
4.2	Security Guarantees	8
5	PQ WireGuard from Reinforced KEM	10
5.1	Introducing Reinforced KEM	10
5.2	Building PQ WireGuard from RKEM	11
5.3	Security of Our RKEM-based PQ WireGuard	12
6	Generic Construction of Reinforced KEM	14
6.1	Reinforced Double-Message PKE . .	14
6.2	OW-Secure RPKE ₂ to CCA Secure RKEM	15
7	Our Lattice-based Reinforced KEM: Rebar	17
7.1	OW-(R)CPA Security	17
7.2	Ciphertext (R-)Indistinguishability . .	18
7.3	Remaining Properties	18
8	Implementation and Experiments	18
8.1	Parameters and Instantiation of Our RKEM	18
8.2	Integration in WireGuard and Experiments	19
Appendix A:	Omitted Details of Security Model	23
A.1	Subtle Issue in Prior Computational Model	23
A.2	Key Indistinguishability	23
Appendix B:	UKS Attack Without Proper Key Schedule	24
Appendix C:	Key-Indistinguishability of Our PQ WireGuard	24
Appendix D:	Omitted Details of Generic Construction of RKEM	30
D.1	Proof of T Transform	30
D.2	Proof of U_m^X Transform	30
Appendix E:	Omitted Details of Lattice-based $\overline{\text{RPKE}}_2$	34
E.1	C-IND Security of RPKE ₂	34
E.2	CR-IND Security of RPKE ₂	36
E.3	Statistical Disjointness of $\overline{\text{RPKE}}_2$. .	38
E.4	(Half-)Correctness of $\overline{\text{RPKE}}_2$	39

Appendix F: (Reinforced) KEM-based Noise Protocols	39
F.1 Noise Patterns for KEM Key Exchange	39
F.2 Extending Post-Quantum Noise With Reinforced KEM	40