Tamper-Resilient Quantum Key Encapsulation with Coherent Ciphertexts from qOWFs in the QROM

Shiyu He

College of Computer Science, Chongqing University
Chongqing, China
shiyuhe477@gmail.com

Abstract

Kitagawa et al. (*CRYPTO 2024*) constructed the first tamper-resilient, chosen-ciphertext-secure public-key encryption from quantum-secure one-way functions (qOWFs). However, their construction assumes authenticated classical channels and measures intermediate signed-phase states to produce classical ciphertexts. Therefore we ask whether coherent quantum ciphertexts can support active security and certified deletion over unauthenticated classical and quantum channels. We construct the first tamper-resilient quantum key-encapsulation mechanism (QKEM) from qOWFs with coherent quantum ciphertexts and genuine validity-feedback confidentiality in the quantum random-oracle model. Our scheme, designated *Signed-Fourier QKEM*, retains signed coherence through honest encapsulation to form a state-consuming quantum ciphertext. Aside from an authentic root verification key, all transmitted classical metadata and quantum registers are subject to arbitrary adversarial tampering and entanglement. With one quantum public key per independent signing domain and a challenge serial unused in pre-challenge decapsulation queries, we prove session-key confidentiality under genuine different-header decapsulation, including queries sharing the challenge serial, and genuine same-header validity feedback. We also establish fixed-header key robustness against adaptive decapsulation attacks, bounding wrong-key acceptance by a negligible probability. Furthermore, our scheme provides publicly verifiable computational certified deletion, protecting session-key privacy against subsequent, complete exposure of the receiver's secret state. Finally, composing our QKEM with a classical one-time private data encapsulation mechanism yields header-restricted chosen-ciphertext confidentiality for arbitrary polynomial-length messages.

Index Terms

Quantum key encapsulation, coherent quantum ciphertexts, tamper resilience, certified deletion, quantum random-oracle model.

I. Introduction

In classical cryptography, the black-box separation of Impagliazzo and Rudich [1] established a fundamental barrier: key agreement and public-key encryption cannot be constructed from one-way functions in a black-box manner. Quantum information changes this landscape. Early quantum public-key encryption constructions used quantum-state indistinguishability [2]. Barooti et al. [3] and Coladangelo [4] showed that quantum public keys enable public-key encryption from quantum-secure one-way functions (qOWFs). In this framework, the receiver distributes quantum public keys that an honest sender uses to encrypt classical secrets. Unlike classical public keys, however, unknown quantum states cannot generally be inspected or copied without disturbance. An active adversary can intercept distributed quantum public keys, apply arbitrary completely positive trace-preserving (CPTP) maps, and entangle them with private auxiliary memory. These early constructions assume access to intact quantum public keys.

To counter physical tampering, Kitagawa, Morimae, Nishimaki, and Yamakawa (KMNY) [5] introduced the first tamper-resilient quantum public-key encryption scheme using signed-phase quantum public keys. Despite this breakthrough, existing constructions face two notable limitations. First, KMNY and the robust scheme of Malavolta and Walter [6] assume authenticated classical communication: while quantum states may be manipulated, transmitted classical messages are assumed to be tamper-proof. In realistic hostile networks, transmission channels are entirely unauthenticated, allowing an active adversary to simultaneously tamper with both quantum states and transmitted classical metadata. Second, KMNY's final scheme measures the intermediate signed quantum state to produce a classical ciphertext. This measurement destroys quantum coherence, limiting direct extensions to quantum communication networks and homomorphic encryption settings that require coherent processing of quantum ciphertexts.

Retaining quantum ciphertexts preserves coherence for subsequent quantum processing, supporting their use in quantum communication protocols and, with suitable constructions, homomorphic evaluation on encrypted quantum data. Quantum ciphertexts can also support certified deletion [7]. Classical ciphertexts can be copied and retained for decryption after future key exposure. Certified deletion instead allows a recipient to produce a classical certificate providing evidence that the protected information remains hidden even after subsequent decryption-key exposure. Generic compilers extend certified deletion [8], and further constructions support public verification [9], [10].

However, preserving quantum coherence introduces an additional authentication problem. Consider the signed-phase state

$$|\psi_b\rangle = \frac{|0\rangle_A |\sigma_0\rangle_B + (-1)^b |1\rangle_A |\sigma_1\rangle_B}{\sqrt{2}},$$

where A is the branch register, B stores the signature, and σ_0, σ_1 are valid signatures on the two branch messages. An adversary can apply a Pauli-Z gate to A, obtaining $(Z_A \otimes I_B)\,|\psi_b\rangle = |\psi_{b\oplus 1}\rangle$. Both signatures remain valid, but the phase bit changes. Thus, authenticating the classical branch values does not by itself authenticate their relative phase. Moreover, arbitrary unknown quantum ciphertexts do not admit an exact equality test that preserves their states and entanglement. The SWAP test does not provide such a test, so classical CCA challenge exclusions based on ciphertext equality cannot be applied directly to quantum registers [11].

State consumption does not remove the need for chosen-ciphertext security. An adversary can retain auxiliary memory entangled with a submitted quantum ciphertext and adapt later queries to the receiver's responses. Although decapsulation consumes the submitted ciphertext, acceptance decisions or returned keys may still reveal information. The adversary may also replace the quantum ciphertext while retaining the classical header of the challenge ciphertext and observe whether the receiver accepts.

This leads to the question:

> *Can a quantum key encapsulation mechanism (QKEM) retain coherent quantum ciphertexts while providing confidentiality under adaptive decapsulation feedback, fixed-header key robustness, and publicly verifiable certified deletion from qOWFs in the quantum random-oracle model (QROM)?*

We answer this question by constructing the Signed-Fourier QKEM. We assume only that the receiver's long-term root verification key is authentic; all transmitted classical and quantum information may be modified by the adversary.

A. Our Contributions

We construct Signed-Fourier QKEM, a tamper-resilient quantum key encapsulation mechanism with coherent quantum ciphertexts from quantum-secure one-way functions (qOWFs) in the QROM. Our main contributions are threefold:

a) 1. Stateful Architecture over Unauthenticated Channels: We formulate a stateful issuance model anchored solely by an authentic long-term root verification key vk_{rt}. Each signing domain supplies one quantum public key, with an independent serial ν and signing key pair (sk_ν, vk_ν). The resulting ciphertext $CT = (h, Q)$ binds a classical header h to a coherent quantum ciphertext Q through authenticated Fourier-phase encoding. All transmitted classical and quantum information may be modified and entangled with the adversary's memory.

b) 2. Active Security via Classical Header Anchoring: Because arbitrary unknown quantum ciphertexts cannot generally be compared exactly without disturbance, we use the classical challenge header h^* to determine the response to each decapsulation query:

- **Validity-Feedback Chosen-Ciphertext Attack (VFCCA) Confidentiality:** The challenger returns full decapsulation results for $h \neq h^*$ and the genuine acceptance/rejection bit for $h = h^*$. Every query consumes the submitted quantum ciphertext and induces the corresponding state of the adversary's retained memory.
- **Adaptive Fixed-Header Key Robustness (aFHKR):** Even with adaptive decapsulation access and knowledge of the challenge key K^*, an adversary cannot cause (h^*, Q') to decapsulate to a key $K' \notin \{K^*, \perp\}$ except with negligible probability.

Theorem 1 (Informal: Active Security, see Theorems 30 and 31)**.** *Assuming qOWFs, Signed-Fourier QKEM achieves* VFCCA *confidentiality and* aFHKR *in the QROM. For every QPT adversary $\mathcal{A}$ making polynomially many quantum random-oracle queries Q_H and decapsulation queries Q_D,*

$$\mathbf{Adv}_{\Pi}^{\mathsf{VFCCA}}(\mathcal{A}) \le \mathrm{negl}(\lambda), \qquad \mathbf{Adv}_{\Pi}^{\mathsf{aFHKR}}(\mathcal{A}) \le \min\left\{1,\, C_{\mathrm{col}}\frac{(Q_H + 2Q_D + 5)^3}{2^{2\lambda}}\right\} \le \mathrm{negl}(\lambda),$$

where C_{col} is an absolute constant. The VFCCA guarantee requires one quantum public key per independent signing domain and a challenge serial that has not appeared in any pre-challenge decapsulation query, including rejected queries. Consequently, IND-QKEM-HCCA*,* IND-QKEM-FTCCA*, and* IND-QKEM-CPA *hold.*

c) 3. Publicly Verifiable Certified Deletion: Measuring the signed branch registers in the computational basis produces a publicly verifiable classical deletion certificate. We prove security even when an adversary generates the certificate while retaining arbitrary auxiliary memory.

Theorem 2 (Informal: Certified Deletion, see Theorem 39)**.** *For honestly generated target ciphertexts, Signed-Fourier QKEM achieves publicly verifiable computational certified deletion under receiver-secret exposure: every QPT adversary has negligible advantage in the deletion experiment, which exposes the receiver's entire secret state, including sk_ν, after accepting the deletion certificate.*

d) Extension to Polynomial-Length Messages: Composing our QKEM with a classical data-encapsulation mechanism (DEM) providing one-time privacy against QPT adversaries yields hybrid public-key encryption with CPA and IND-QKEM-HCCA confidentiality for polynomial-length messages in the QROM (Theorems 33 and 35).

B. Related Work

a) Closest Prior Work: Kitagawa et al. [5] and Malavolta and Walter [6] obtain tamper-resilient quantum public-key encryption from qOWFs under authenticated transmission of classical public-key information. Kitagawa et al. also give a CPA-to-CCA compiler. Their construction uses signed quantum branch states and phase encoding before producing classical ciphertexts. Our construction retains the signed coherence in the quantum ciphertext and allows adversarial modification of all transmitted classical and quantum information, assuming only an authentic long-term root verification key. We prove confidentiality with genuine different-header decapsulation and same-header acceptance/rejection feedback, together with adaptive fixed-header key robustness. The retained coherence also enables publicly verifiable certified deletion that protects the session key after subsequent exposure of the receiver's entire secret state. These guarantees are obtained from qOWFs in the QROM.

b) Other Quantum-Ciphertext Constructions: Early quantum public-key encryption appears in Kawachi et al. [2], while subsequent constructions from weak assumptions appear in [3], [4], [12]. Hhan, Morimae, and Yamakawa [13] construct encryption with classical public keys and quantum ciphertexts using non-abelian group actions and distinguishing-to-swapping techniques. Barhoush and Salvail [14] study time-dependent quantum signatures and their applications to authenticated quantum public keys under chosen quantum-key and ciphertext attacks. Practical constructions appear in [15], [16]. Foundations for quantum-data privacy and authentication are developed in [11], [17], [18], while Aaronson, Atia, and Susskind [19] establish the connection between superposition detection and swapping. Certified deletion [7], [8] and publicly verifiable deletion compilers [9], [10] provide the context for our deletion guarantee under receiver-secret exposure.

c) Relation to Impossibility Results: Li et al. [20] prove QROM impossibility results for perfectly correct quantum public-key encryption with classical keys. Their extension to quantum public keys assumes oracle-independent public-key generation from a fixed classical secret key and access to polynomially many copies of the same public key. Our model instead supplies one quantum public key per independent signing domain and permits state-updating issuance. Our claims concern this issuance interface and do not extend to polynomially many copies of a fixed quantum public key.

II. Cryptographic Preliminaries

This section introduces the cryptographic primitives used in the construction and their instantiation from quantum-secure one-way functions. It also collects the quantum random-oracle bounds used to establish confidentiality and key robustness.

A. Notation and Conventions

Let λ be the security parameter, with $\lambda \le \ell = \ell(\lambda) \le 2\lambda$. We write $[n] := \{1, \dots, n\}$, use $\|$ for string concatenation, and define the trace distance as $\Delta(\rho, \sigma) := \frac{1}{2}\|\rho - \sigma\|_1$. A quantum polynomial-time (QPT) algorithm is a uniform polynomial-size quantum circuit family. We denote the single-qubit Hadamard transform by H and the quantum random oracle by H. We use unambiguous injective encodings for tagged tuples and variable-length strings; signatures in quantum registers are canonically padded to $L_{\mathrm{sig}}(\lambda)$ bits. The symbol st denotes any (possibly quantum) state carried across protocol stages. Finally, N, Q_H, Q_D are polynomial bounds satisfying $1 \le N < 2^\lambda$, where Q_H bounds all direct adversarial QRO queries (including local evaluations within public algorithms) and Q_D bounds the adversary's decapsulation queries; queries evaluated internally by reduction oracles are counted separately.

B. Computational Primitives

Our security analysis is performed in the QROM. This subsection introduces the computational primitives used within our construction, together with their instantiations from quantum-secure one-way functions.

a) Signatures against Quantum Adversaries: Root signatures authenticate token metadata, while deterministic phase signatures define the signed branches used in the construction.

A digital signature scheme $\Sigma = (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Vfy})$ consists of PPT algorithms with deterministic verification:

- **Syntax & Correctness:** Key generation outputs $(sk, vk) \leftarrow \mathsf{Gen}(1^\lambda)$, signing outputs $\sigma \leftarrow \mathsf{Sign}_{sk}(m)$, and verification outputs $\mathsf{Vfy}_{vk}(m, \sigma) \in \{0, 1\}$. Correctness requires that for all m,
$$\Pr\big[\mathsf{Vfy}_{vk}(m, \sigma) = 1 : (sk, vk) \leftarrow \mathsf{Gen}(1^\lambda),\ \sigma \leftarrow \mathsf{Sign}_{sk}(m)\big] = 1.$$
- **Canonical Signing:** Σ is *canonical* if Sign is deterministic, yielding a unique signature for each pair (sk, m). Verification is not restricted to accepting only canonical signatures.
- **Security:** We consider the standard notions of existential unforgeability (EUF-CMA) and strong unforgeability (sEUF-CMA) against QPT adversaries with classical signing interfaces.

Deterministic signing provides the reproducible branch signatures used in coherent decoding. Strong unforgeability will be used to bound the accepted weight outside the canonical signature subspace. The reductions prepare signed quantum states from signatures obtained through classical signing queries.

Definition 3 ((s)EUF-CMA Security)**.** Let $\Sigma = (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Vfy})$ be a signature scheme. For a QPT adversary $\mathcal{F}$, define $\mathbf{Exp}_{\Sigma}^{\text{euf-cma}}(\mathcal{F})$:

1. Sample $(sk, vk) \leftarrow \mathsf{Gen}(1^\lambda)$ and initialize $\mathcal{Q} \leftarrow \varnothing$.
2. Run $(m^*, \sigma^*) \leftarrow \mathcal{F}^{\mathcal{O}_{\mathsf{Sign}}}(1^\lambda, vk)$, where $\mathcal{O}_{\mathsf{Sign}}$ is a *classical* signing oracle: on query m, it returns $\sigma \leftarrow \mathsf{Sign}_{sk}(m)$ and records $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(m, \sigma)\}$.
3. The experiment outputs 1 iff $\mathsf{Vfy}_{vk}(m^*, \sigma^*) = 1$ and:
   - **EUF-CMA:** $m^* \notin \{m : \exists \sigma,\ (m, \sigma) \in \mathcal{Q}\}$.
   - **sEUF-CMA:** $(m^*, \sigma^*) \notin \mathcal{Q}$.

We define $\mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma}(\mathcal{F}) := \Pr[\mathbf{Exp}^{\text{seuf-cma}}_{\Sigma}(\mathcal{F}) = 1]$ (analogously for EUF-CMA). Σ is (s)EUF-CMA secure if this advantage is negligible for all QPT $\mathcal{F}$. The strongly unforgeable one-time variant (sOT) restricts $\mathcal{F}$ to at most one classical signing query.

For a one-time signature scheme Γ_{OT}, we denote the corresponding forgery probability by $\mathbf{Adv}^{\mathrm{sOT}}_{\Gamma_{\mathrm{OT}}}(\mathcal{F})$.

b) Quantum-Accessible Pseudorandom Functions: In proof hybrids, replacing the random oracle with a qPRF [21] provides local oracle circuits and their inverses, with the replacement loss bounded by qPRF security.

Definition 4 (qPRF)**.** An efficiently computable function family $F = \{F_k : \mathcal{X}_\lambda \to \mathcal{Y}_\lambda\}_{k \in \mathcal{K}_\lambda}$ is a *quantum-accessible pseudorandom function (qPRF)* if for every QPT distinguisher $\mathcal{C}$ making $q = \mathrm{poly}(\lambda)$ quantum superposition queries to an evaluation oracle $|x\rangle|y\rangle \mapsto |x\rangle|y \oplus \mathcal{O}(x)\rangle$,

$$\mathbf{Adv}^{\mathrm{qPRF}}_{F}(\mathcal{C}) := \left| \Pr_{k \leftarrow \mathcal{K}_\lambda}\left[\mathcal{C}^{|F_k\rangle}(1^\lambda) = 1\right] - \Pr_{R \leftarrow \mathrm{Func}(\mathcal{X}_\lambda, \mathcal{Y}_\lambda)}\left[\mathcal{C}^{|R\rangle}(1^\lambda) = 1\right] \right| \le \mathrm{negl}(\lambda).$$

c) One-Time Private DEMs with Associated Data: The KEM–DEM composition uses the following primitive to encrypt messages with the encapsulated session key.

Definition 5 (OT-Priv DEM)**.** A data encapsulation mechanism $\mathsf{DEM} = (\mathsf{Enc}, \mathsf{Dec})$ with key space $\{0,1\}^\ell$ and associated data space $\mathcal{A}$ satisfies correctness $\mathsf{Dec}_K(a, \mathsf{Enc}_K(a, M)) = M$. DEM provides *one-time privacy (OT-Priv)* if for every QPT adversary $\mathcal{B} = (\mathcal{B}_1, \mathcal{B}_2)$,

$$\mathbf{Adv}^{\mathrm{OTPriv}}_{\mathsf{DEM}}(\mathcal{B}) := \left| \Pr\left[\widehat{\beta} = \beta : \begin{array}{l} (a, M_0, M_1, \mathsf{st}) \leftarrow \mathcal{B}_1(1^\lambda) \text{ s.t. } |M_0| = |M_1|, \\ K \leftarrow \{0,1\}^\ell,\ \beta \leftarrow \{0,1\},\ C^* \leftarrow \mathsf{Enc}_K(a, M_\beta), \\ \widehat{\beta} \leftarrow \mathcal{B}_2(C^*, \mathsf{st}) \end{array} \right] - \frac{1}{2} \right| \le \mathrm{negl}(\lambda).$$

Associated data a is a public input to encryption and decryption; no ciphertext integrity is assumed.

d) Instantiation from Quantum-Secure One-Way Functions:

Definition 6 (qOWF)**.** An efficiently computable function family $f = \{f_\lambda : \{0,1\}^{n(\lambda)} \to \{0,1\}^{m(\lambda)}\}_{\lambda \in \mathbb{N}}$ is a *quantum-secure one-way function (qOWF)* if for every QPT adversary $\mathcal{A}$,

$$\mathbf{Adv}^{\mathrm{qOWF}}_{f}(\mathcal{A}) := \Pr\left[f_\lambda(x') = f_\lambda(x) : x \leftarrow \{0,1\}^{n(\lambda)},\ x' \leftarrow \mathcal{A}(1^\lambda, f_\lambda(x))\right] \le \mathrm{negl}(\lambda).$$

Proposition 7 (Primitives from qOWFs)**.** Assuming the existence of qOWFs, the following primitives exist:

(1) Standard EUF-CMA and canonical sEUF-CMA signature schemes against QPT adversaries with classical signing queries;
(2) Quantum-accessible pseudorandom functions (qPRFs);
(3) PPT DEMs providing OT-Priv security against QPT adversaries for polynomial-length messages.

Proof. By [5, Theorems 3.3 and 3.5], qOWFs imply both strong signatures secure against quantum adversaries with classical signing queries and qPRFs. Applying standard PRF-based derandomization (Lemma 42) yields canonical sEUF-CMA signatures. A separate qPRF instantiates the OT-Priv DEM (Appendix B). All primitive keys are independently sampled, and primitive evaluations make no random-oracle queries. □

C. Quantum Random-Oracle Tools

The following bounds control oracle reprogramming, collision finding, and preimage-fiber sizes.

Quantum access to a random oracle $H : \mathcal{X} \to \mathcal{Y}$ is evaluated via the standard unitary map $U_H : |x\rangle|y\rangle \mapsto |x\rangle|y \oplus H(x)\rangle$ [22]. Throughout this subsection, d denotes an upper bound on oracle queries. Executions issuing fewer than d queries are modeled with d query slots, where inactive slots contribute zero query weight.

a) One-Way-to-Hiding with Intermediate Channels: The one-way-to-hiding framework [23] relates the effect of oracle modifications to the probability of querying the modified inputs. The following form accommodates common intermediate quantum channels, including measurements and register discarding.

Lemma 8. *Let $G, H : \mathcal{X} \to \mathcal{Y}$ satisfy $G(x) = H(x)$ for all $x \notin S$. Consider two hybrid executions sharing identical initial states and identical intermediate CPTP maps, differing solely in at most $d \ge 1$ oracle evaluations. The intermediate channels*

may be adaptively controlled by private classical registers. If p denotes the probability that a uniformly sampled query slot in the all-H execution is active and queries an element of S, then for any common final test $T \in \{0,1\}$,

$$\big|\Pr[T_G = 1] - \Pr[T_H = 1]\big| \le 2d\sqrt{p}.$$

The statement also holds when the oracle tables (G, H), the set S, and the initial states are jointly sampled.

Proof. Fix S, the oracle tables (G, H), and all initially sampled classical parameters. Intermediate classical measurement records and flags are absorbed into the joint state. We take the expectation over the initial configuration at the end.

Let ρ_i denote the joint density operator immediately preceding the i-th query slot in the all-H execution, with purification $|\rho_i\rangle$. Let

$$\Pi_i := |1\rangle\langle 1|_{\mathrm{act}} \otimes \sum_{x \in S} |x\rangle\langle x| \otimes I$$

project onto the subspace where the i-th slot is active and queries S. The query weight is $w_i := \mathrm{Tr}(\Pi_i \rho_i) = \|\Pi_i |\rho_i\rangle\|^2$. Let $V_{G,i}$ and $V_{H,i}$ denote the corresponding controlled oracle gates, which act as identity on inactive slots. Because $V_{G,i}$ and $V_{H,i}$ coincide on $I - \Pi_i$,

$$\|(V_{G,i} - V_{H,i})|\rho_i\rangle\| = \|(V_{G,i} - V_{H,i})\Pi_i |\rho_i\rangle\| \le 2\sqrt{w_i}. \tag{1}$$

By monotonicity of the trace distance under partial trace and the pure-state bound $\Delta(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|) \le \||\psi\rangle - |\phi\rangle\|$,

$$\Delta(V_{G,i}\rho_i V_{G,i}^\dagger,\, V_{H,i}\rho_i V_{H,i}^\dagger) \le 2\sqrt{w_i}.$$

Consider final hybrid states $\{\sigma^{(k)}\}_{k=0}^d$, where $\sigma^{(k)}$ executes V_H in slots $i \le k$ and V_G in slots $i > k$. The executions defining $\sigma^{(i)}$ and $\sigma^{(i-1)}$ share the same all-H prefix before slot i, entering that slot in state ρ_i, and share the identical continuation. Since intermediate operations are CPTP channels, contractivity under the shared continuation gives $\Delta(\sigma^{(i)}, \sigma^{(i-1)}) \le 2\sqrt{w_i}$. Telescoping across all d steps,

$$\Delta(\rho_G, \rho_H) = \Delta(\sigma^{(0)}, \sigma^{(d)}) \le \sum_{i=1}^d \Delta(\sigma^{(i-1)}, \sigma^{(i)}) \le 2\sum_{i=1}^d \sqrt{w_i}.$$

Averaging over the initial distribution and applying the Cauchy–Schwarz inequality:

$$\mathbb{E}\big[\Delta(\rho_G, \rho_H)\big] \le 2\sum_{i=1}^d \mathbb{E}[\sqrt{w_i}] \le 2\sqrt{d\sum_{i=1}^d \mathbb{E}[w_i]} = 2d\sqrt{\frac{1}{d}\sum_{i=1}^d \mathbb{E}[w_i]} = 2d\sqrt{p},$$

where $p = \frac{1}{d}\sum_{i=1}^d \mathbb{E}[w_i]$ is the probability that a uniform slot hits S. Finally, since test T is common, $|\Pr[T_G = 1] - \Pr[T_H = 1]| \le \mathbb{E}[\Delta(\rho_G, \rho_H)] \le 2d\sqrt{p}$. □

Corollary 9. A final oracle query is excluded from the query count d in Lemma 8 if it never evaluates inputs in S (e.g., when queries in S are intercepted by a prior check). In this case, the final operation acts identically under G and H and contributes no distinguishing advantage.

Proof. Since the final channel $\mathcal{E}$ is identical in both executions, contractivity of trace distance gives $\Delta(\mathcal{E}(\rho_G), \mathcal{E}(\rho_H)) \le \Delta(\rho_G, \rho_H)$. Thus, it incurs no additional hybrid step. □

b) Fiber Regularity: The analysis of header retagging also requires a uniform bound on the number of phase strings consistent with a binding tag. The next lemma provides this bound for the random oracle.

Lemma 10. *Define $H_{\mathrm{bnd}}(\nu, \tau, s) := H(00\|\nu\|\tau\|s)$ with range $\{0,1\}^{2\lambda}$, where $\nu, \tau \in \{0,1\}^\lambda$ and $s \in \{0,1\}^\ell$ with $\lambda \le \ell \le 2\lambda$. For any $(\nu, \tau, \gamma) \in \{0,1\}^\lambda \times \{0,1\}^\lambda \times \{0,1\}^{2\lambda}$, let $L_{\nu,\tau,\gamma} := |\{s \in \{0,1\}^\ell : H_{\mathrm{bnd}}(\nu, \tau, s) = \gamma\}|$. For all $\lambda \ge 3$,*

$$\Pr_H\Big[\max_{\nu,\tau,\gamma} L_{\nu,\tau,\gamma} \ge \lambda\Big] \le \varepsilon_{\mathrm{fib}}(\lambda) := 2^{4\lambda}\left(\frac{e}{\lambda}\right)^\lambda = 2^{-\Omega(\lambda \log \lambda)}.$$

Proof. Fix (ν, τ, γ). For each $s \in \{0,1\}^\ell$, let $X_s := \mathbf{1}[H_{\mathrm{bnd}}(\nu, \tau, s) = \gamma]$. The variables $\{X_s\}$ are independent Bernoulli trials with parameter $2^{-2\lambda}$, so $L_{\nu,\tau,\gamma} \sim \mathrm{Binomial}(2^\ell, 2^{-2\lambda})$. A union bound over all subsets of $\{0,1\}^\ell$ of size λ gives:

$$\Pr[L_{\nu,\tau,\gamma} \ge \lambda] \le \binom{2^\ell}{\lambda}(2^{-2\lambda})^\lambda \le \left(\frac{e \cdot 2^\ell}{\lambda}\right)^\lambda \cdot 2^{-2\lambda^2} = \left(\frac{e \cdot 2^{\ell - 2\lambda}}{\lambda}\right)^\lambda \le \left(\frac{e}{\lambda}\right)^\lambda,$$

where the last inequality follows from $\ell \le 2\lambda$. Taking a union bound over all $2^{4\lambda}$ triples (ν, τ, γ) yields:

$$\Pr_H\Big[\max_{\nu,\tau,\gamma} L_{\nu,\tau,\gamma} \ge \lambda\Big] \le 2^{4\lambda}\left(\frac{e}{\lambda}\right)^\lambda = 2^{4\lambda - \lambda\log_2(\lambda/e)} = 2^{-\Omega(\lambda \log \lambda)}.$$

□

III. State-Consuming QKEM: Foundations and Construction

In this section, we formalize the framework of state-consuming quantum key-encapsulation mechanisms (QKEMs) and present our concrete instantiation. We first introduce the syntax and correctness requirements in Section III-A. Next, Section III-B presents the Signed-Fourier QKEM construction, followed by its correctness and efficiency analysis in Section III-C. Finally, Section III-D formalizes the active security models, including validity-feedback CCA confidentiality and adaptive fixed-header key robustness.

A. QKEM Syntax

Definition 11. A quantum key-encapsulation mechanism $\Pi = (\mathsf{SKGen}, \mathsf{PKGen}, \mathsf{Encap}, \mathsf{Decap})$ with classical header space $\mathcal{H}$, quantum payload space $\mathcal{D}(\mathcal{H}_{\mathrm{pay}})$, and key space $\{0,1\}^\ell$ consists of four efficient algorithms:

- $\mathsf{SKGen}(1^\lambda)$: Outputs root verification key vk_{rt} and initial receiver state $\mathsf{st}_{\mathrm{R}}^{(0)}$.
- $\mathsf{PKGen}(\mathsf{st}_{\mathrm{R}})$: Takes receiver state, outputs token $pk_\nu = (\nu, vk_\nu, \pi_\nu, \rho_\nu)$ with a fresh serial $\nu \in \{0,1\}^\lambda$, payload $\rho_\nu \in \mathcal{D}(\mathcal{H}_{\mathrm{pay}})$, and updated state $\mathsf{st}'_{\mathrm{R}}$.
- $\mathsf{Encap}^H(vk_{\mathrm{rt}}, pk')$: Takes root key and candidate token register, outputs ciphertext $CT = (h, Q) \in \mathcal{H} \times \mathcal{D}(\mathcal{H}_{\mathrm{pay}})$ and key $K \in \{0,1\}^\ell$, or fail.
- $\mathsf{Decap}^H(\mathsf{st}_{\mathrm{R}}, CT)$: Consumes quantum payload Q and outputs session key $K' \in \{0,1\}^\ell \cup \{\perp\}$.

Classical parsers $\mathsf{sid}_{\mathrm{pk}}(pk)$, $\mathsf{sid}_{\mathrm{ct}}(h)$, and $\mathsf{hdr}(CT)$ extract the token serial, ciphertext serial, and header, returning $\perp$ on parsing failure; honest encapsulation satisfies $\mathsf{sid}_{\mathrm{ct}}(\mathsf{hdr}(CT)) = \mathsf{sid}_{\mathrm{pk}}(pk_\nu)$.

a) Correctness: Π is correct if for all polynomial $N = N(\lambda) < 2^\lambda$ and every $j \in [N]$,

$$\Pr\left[K' = K : \begin{array}{l} (vk_{\mathrm{rt}}, \mathsf{st}_{\mathrm{R}}^{(0)}) \leftarrow \mathsf{SKGen}(1^\lambda), \\ (pk_i, \mathsf{st}_{\mathrm{R}}^{(i)}) \leftarrow \mathsf{PKGen}(\mathsf{st}_{\mathrm{R}}^{(i-1)}) \quad (i \in [N]), \\ (CT, K) \leftarrow \mathsf{Encap}^H(vk_{\mathrm{rt}}, pk_j), \\ K' \leftarrow \mathsf{Decap}^H(\mathsf{st}_{\mathrm{R}}^{(N)}, CT) \end{array}\right] \geq 1 - \mathrm{negl}(\lambda).$$

Encapsulation failure makes the event false. The probability is over a uniformly sampled random oracle H and the randomness of the algorithms. If the probability equals 1, Π achieves perfect correctness.

B. The Signed-Fourier QKEM Construction

Let $\Sigma_{\mathrm{rt}} = (\mathsf{Gen}_{\mathrm{rt}}, \mathsf{Sign}_{\mathrm{rt}}, \mathsf{Vfy}_{\mathrm{rt}})$ be an EUF-CMA signature scheme, and let $\Sigma_{\mathrm{ph}} = (\mathsf{Gen}_{\mathrm{ph}}, \mathsf{Sign}_{\mathrm{ph}}, \mathsf{Vfy}_{\mathrm{ph}})$ be a canonical sEUF-CMA signature scheme with deterministic signing $\mathsf{Sign}_{\mathrm{ph}}$. Fix $\lambda \leq \ell \leq 2\lambda$ and a random oracle

$$H : \{0,1\}^{2+2\lambda+\ell} \longrightarrow \{0,1\}^{2\lambda}.$$

Define

$$H_{\mathrm{bnd}}(\nu, \tau, s) := H(00\|\nu\|\tau\|s), \qquad H_{\mathrm{kdf}}(\nu, \tau, s) := \mathrm{first}_\ell\big(H(01\|\nu\|\tau\|s)\big).$$

Each serial ν identifies one issued token and an independently generated phase-signing key pair (sk_ν, vk_ν). The quantum registers of tokens and ciphertexts are $A = A_1 \cdots A_\ell$, consisting of branch qubits, and $B = B_1 \cdots B_\ell$, consisting of signature registers of $L_{\mathrm{sig}}(\lambda)$ qubits each. For $i \in [\ell]$, $b \in \{0,1\}$, and $z \in \{0,1\}^\ell$, define

$$\sigma_{\nu,i,b} := \mathsf{Sign}_{\mathrm{ph},sk_\nu}(\mathsf{phase}\|\nu\|i\|b), \qquad \boldsymbol{\sigma}_{\nu,z} := (\sigma_{\nu,1,z_1}, \ldots, \sigma_{\nu,\ell,z_\ell}).$$

a) Projective Coherent Validation: For verification key vk_ν, public validation is performed via the measurement $\{\Pi_\nu^{\mathrm{val}}, I - \Pi_\nu^{\mathrm{val}}\}$, where

$$\begin{aligned} \Pi_\nu^{\mathrm{val}} &:= \bigotimes_{i=1}^{\ell} \Pi_{\nu,i}^{\mathrm{val}}, \\ \Pi_{\nu,i}^{\mathrm{val}} &:= \sum_{b \in \{0,1\}} \sum_{\sigma:\, \mathsf{Vfy}_{\mathrm{ph},vk_\nu}(\mathsf{phase}\|\nu\|i\|b,\sigma)=1} |b\rangle\langle b|_{A_i} \otimes |\sigma\rangle\langle\sigma|_{B_i}. \end{aligned} \tag{2}$$

This measurement preserves superpositions within the accepting subspace. Parsing failures return fail in Encap and $\perp$ in Decap, and all rejected quantum inputs are discarded. The complete scheme Π is given in Figure 1.

$\underline{\mathsf{SKGen}(1^\lambda)}$:

1. Sample $(sk_{\mathrm{rt}}, vk_{\mathrm{rt}}) \leftarrow \mathsf{Gen}_{\mathrm{rt}}(1^\lambda)$. Initialize $\mathsf{ctr} := 0$ and $\mathsf{DB} := \varnothing$.
2. Return $(vk_{\mathrm{rt}}, \mathsf{st}_{\mathrm{R}} := (sk_{\mathrm{rt}}, \mathsf{ctr}, \mathsf{DB}))$.

$\underline{\mathsf{PKGen}(\mathsf{st}_{\mathrm{R}})}$:

1. Increment ctr and encode it as $\nu \in \{0,1\}^\lambda$, issuing at most $N < 2^\lambda$ tokens.
2. Sample $(sk_\nu, vk_\nu) \leftarrow \mathsf{Gen}_{\mathrm{ph}}(1^\lambda)$. Compute $\sigma_{\nu,i,b} := \mathsf{Sign}_{\mathrm{ph},sk_\nu}(\mathtt{phase}\|\nu\|i\|b)$ for $i \in [\ell]$ and $b \in \{0,1\}$.
3. Prepare the token state
$$|\Psi_\nu\rangle_{AB} := \bigotimes_{i=1}^{\ell} \frac{|0\rangle_{A_i} |\sigma_{\nu,i,0}\rangle_{B_i} + |1\rangle_{A_i} |\sigma_{\nu,i,1}\rangle_{B_i}}{\sqrt{2}}.$$
4. Compute the root certificate $\pi_\nu \leftarrow \mathsf{Sign}_{\mathrm{rt},sk_{\mathrm{rt}}}(\mathtt{token}\|\nu\|vk_\nu)$, and store $\mathsf{DB}[\nu] := (sk_\nu, vk_\nu)$.
5. Return $(pk_\nu, \mathsf{st}'_{\mathrm{R}})$, where $pk_\nu := (\nu, vk_\nu, \pi_\nu, |\Psi_\nu\rangle_{AB})$ and $\mathsf{st}'_{\mathrm{R}}$ contains the updated counter and database.

$\underline{\mathsf{Encap}^H(vk_{\mathrm{rt}}, pk')}$:

1. Parse $pk' = (\nu, vk_\nu, \pi_\nu, Q)$. Return fail if $\mathsf{Vfy}_{\mathrm{rt},vk_{\mathrm{rt}}}(\mathtt{token}\|\nu\|vk_\nu, \pi_\nu) \neq 1$.
2. Measure $\{\Pi_\nu^{\mathrm{val}}, I - \Pi_\nu^{\mathrm{val}}\}$ on Q; return fail upon rejection.
3. Sample $s \leftarrow \{0,1\}^\ell$ and $\tau \leftarrow \{0,1\}^\lambda$. Apply $Z_A^s := \bigotimes_{i=1}^{\ell} Z_{A_i}^{s_i}$ to Q, yielding Q_s.
4. Compute $\gamma := H_{\mathrm{bnd}}(\nu, \tau, s)$ and $K := H_{\mathrm{kdf}}(\nu, \tau, s)$.
5. Return $(CT := ((\nu, \tau, \gamma), Q_s), K)$.

$\underline{\mathsf{Decap}^H(\mathsf{st}_{\mathrm{R}}, CT)}$:

1. Parse $CT = (h, Q)$ with $h = (\nu, \tau, \gamma)$. Retrieve $(sk_\nu, vk_\nu) := \mathsf{DB}[\nu]$; return $\perp$ if absent.
2. Measure $\{\Pi_\nu^{\mathrm{val}}, I - \Pi_\nu^{\mathrm{val}}\}$ on Q; return $\perp$ upon rejection.
3. Using sk_ν, apply coherent erasure
$$U_{\mathrm{ers}}^{(\nu)} : |z\rangle_A |\boldsymbol{\sigma}\rangle_B \longmapsto |z\rangle_A |\boldsymbol{\sigma} \oplus \boldsymbol{\sigma}_{\nu,z}\rangle_B .$$
4. Measure B in the computational basis; return $\perp$ if the outcome is non-zero.
5. Apply $\mathsf{H}^{\otimes\ell}$ to A and measure in the computational basis, obtaining candidate phase string s'.
6. If $H_{\mathrm{bnd}}(\nu, \tau, s') \neq \gamma$, return $\perp$.
7. Return $K := H_{\mathrm{kdf}}(\nu, \tau, s')$.

Fig. 1. The Signed-Fourier QKEM Π.

C. Correctness and Concrete Efficiency

Theorem 12. If Σ_{rt} and Σ_{ph} have perfect correctness, then for every honestly generated token pk_ν and state st_{R} retaining $\mathsf{DB}[\nu]$, Π satisfies perfect correctness. When per-signature verification errors are bounded by $\varepsilon_{\mathrm{rt}}^{\mathrm{cor}}$ and $\varepsilon_{\mathrm{ph}}^{\mathrm{cor}}$, correctness holds except with probability at most $N\varepsilon_{\mathrm{rt}}^{\mathrm{cor}} + 2N\ell\varepsilon_{\mathrm{ph}}^{\mathrm{cor}}$.

Proof. For an honest token pk_ν, $\mathsf{Vfy}_{\mathrm{rt},vk_{\mathrm{rt}}}$ accepts and $\Pi_\nu^{\mathrm{val}} |\Psi_\nu\rangle_{AB} = |\Psi_\nu\rangle_{AB}$. Applying Z_A^s during encapsulation yields

$$|\Psi_{\nu,s}\rangle_{AB} = 2^{-\ell/2} \sum_{z \in \{0,1\}^\ell} (-1)^{z \cdot s} |z\rangle_A |\boldsymbol{\sigma}_{\nu,z}\rangle_B .$$

In decapsulation, $|\Psi_{\nu,s}\rangle$ satisfies $\Pi_\nu^{\mathrm{val}} |\Psi_{\nu,s}\rangle = |\Psi_{\nu,s}\rangle$. Applying coherent signature erasure $U_{\mathrm{ers}}^{(\nu)}$ with sk_ν yields

$$U_{\mathrm{ers}}^{(\nu)} |\Psi_{\nu,s}\rangle_{AB} = \left(2^{-\ell/2} \sum_{z \in \{0,1\}^\ell} (-1)^{z \cdot s} |z\rangle_A \right) \otimes |0\rangle_B = (\mathsf{H}^{\otimes\ell} |s\rangle_A) \otimes |0\rangle_B .$$

Measuring B returns $0^{\ell L_{\mathrm{sig}}(\lambda)}$ deterministically, and measuring $\mathsf{H}^{\otimes\ell} A$ yields $s' = s$, which satisfies $H_{\mathrm{bnd}}(\nu, \tau, s') = \gamma$ and decapsulates to $K = H_{\mathrm{kdf}}(\nu, \tau, s)$. A union bound over the N root certificates and $2N\ell$ branch signatures establishes the bound for imperfect signatures. □

a) Complexity: Each quantum token and ciphertext payload contains $\ell(1 + L_{\mathrm{sig}}(\lambda))$ qubits; classical headers h contain 4λ bits. Issuance evaluates 2ℓ phase signatures and one root certificate. Encapsulation performs one root verification, ℓ coherent verifications for Π_ν^{val}, ℓ phase gates, and at most two queries to H. Decapsulation requires validation, coherent evaluation of ℓ deterministic signing circuits for $U_{\mathrm{ers}}^{(\nu)}$, ℓ Hadamard gates, and at most two queries to H. Receiver storage is $O(N)$ classical keys.

D. Security Models

We define confidentiality under public-key tampering, its extensions with decapsulation access, and fixed-header key robustness. Write $CT = (h, Q)$ for a ciphertext and $CT^* = (h^*, Q^*)$ for the challenge ciphertext. Throughout, $\mathcal{A}$ is a uniform QPT adversary and $N = N(\lambda)$ is polynomially bounded. Classical fields use canonical encodings. Quantum inputs ρ_{QE} may be arbitrary states on the prescribed registers, where Q is submitted and E is retained by $\mathcal{A}$. Encapsulation and decapsulation failures are denoted by fail and $\perp$, respectively.

1) Confidentiality under Public-Key Tampering: The CPA experiment allows the adversary to modify the issued public-key tokens and select the candidate used for challenge encapsulation. No decapsulation oracle is provided.

Definition 13 (CPA confidentiality)**.** For a QPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, define $\mathbf{Exp}_{\Pi}^{\mathrm{CPA}}(\mathcal{A})$ as follows:

1. Sample a random oracle H and $b \leftarrow \{0,1\}$. Generate
$$(vk_{\mathrm{rt}}, \mathsf{st}_{\mathrm{R}}^{(0)}) \leftarrow \mathsf{SKGen}(1^{\lambda}),$$
and issue N tokens:
$$(pk_i, \mathsf{st}_{\mathrm{R}}^{(i)}) \leftarrow \mathsf{PKGen}(\mathsf{st}_{\mathrm{R}}^{(i-1)}) \qquad (i \in [N]).$$
2. Run
$$(pk^*, E) \leftarrow \mathcal{A}_1^{H}(vk_{\mathrm{rt}}, pk_1, \ldots, pk_N).$$
3. Execute $(CT^*, K_0) \leftarrow \mathsf{Encap}^{H}(vk_{\mathrm{rt}}, pk^*)$. If encapsulation returns fail, terminate and output an independent uniform bit. Otherwise, sample $K_1 \leftarrow \{0,1\}^{\ell}$.
4. Run
$$\widehat{b} \leftarrow \mathcal{A}_2^{H}(CT^*, K_b, E).$$
Output 1 if $\widehat{b} = b$, and 0 otherwise.

Define
$$\mathbf{Adv}_{\Pi}^{\mathrm{CPA}}(\mathcal{A}) := \left| \Pr[\mathbf{Exp}_{\Pi}^{\mathrm{CPA}}(\mathcal{A}) = 1] - \frac{1}{2} \right|.$$
The scheme is CPA-secure if this advantage is negligible for every QPT adversary $\mathcal{A}$.

2) Confidentiality with Validity Feedback: Active security additionally allows the adversary to submit ciphertexts and observe decapsulation results. In classical CCA experiments, the challenger compares ciphertext strings to exclude direct decapsulation of the challenge. For $CT = (h, Q)$, the classical header can be compared directly, whereas arbitrary quantum ciphertexts cannot generally be compared exactly without disturbing them.

In Signed-Fourier, the header $h = (\nu, \tau, \gamma)$ is linked to the phase string s encoded in Q through $\gamma = H_{\mathrm{bnd}}(\nu, \tau, s)$. The serial ν identifies the signing domain used for validation and decoding. Decapsulation recovers a candidate phase string s' from Q and accepts only if $H_{\mathrm{bnd}}(\nu, \tau, s') = \gamma$. The header thus specifies the conditions for accepting the quantum ciphertext and deriving its key.

We use the header to determine the feedback on each query. Queries with $h \neq h^*$ receive genuine decapsulation results, including those that retain ν^* but change other header fields. For $h = h^*$, returning the key would allow direct recovery of the challenge key by submitting the original challenge ciphertext in an honest execution. The oracle therefore performs genuine decapsulation but returns only its acceptance bit. This defines validity-feedback CCA (VFCCA), our main confidentiality notion.

The active experiments require that the challenge serial has not appeared in any canonically parsed pre-challenge decapsulation query, including rejected queries. Every query consumes the submitted quantum ciphertext. The adversary retains its auxiliary memory in the state induced by the actual decapsulation outcome.

Definition 14 (Validity-feedback CCA)**.** For a QPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, define $\mathbf{Exp}_{\Pi}^{\mathrm{VFCCA}}(\mathcal{A})$ as follows:

1. Sample a random oracle H and $b \leftarrow \{0,1\}$. Generate
$$(vk_{\mathrm{rt}}, \mathsf{st}_{\mathrm{R}}^{(0)}) \leftarrow \mathsf{SKGen}(1^{\lambda}),$$
and issue N tokens:
$$(pk_i, \mathsf{st}_{\mathrm{R}}^{(i)}) \leftarrow \mathsf{PKGen}(\mathsf{st}_{\mathrm{R}}^{(i-1)}) \qquad (i \in [N]).$$
Set $\mathsf{st}_{\mathrm{R}} := \mathsf{st}_{\mathrm{R}}^{(N)}$ and $S := \varnothing$.
2. Run
$$(pk^*, E) \leftarrow \mathcal{A}_1^{H, \mathcal{O}_{\mathrm{pre}}}(vk_{\mathrm{rt}}, pk_1, \ldots, pk_N).$$
On each query (h, Q), the pre-challenge oracle records $\mathsf{sid}_{\mathrm{ct}}(h)$ in S whenever h parses canonically, and returns
$$K' \leftarrow \mathsf{Decap}^{H}(\mathsf{st}_{\mathrm{R}}, (h, Q)).$$
The serial is recorded even if decapsulation rejects.
3. If $\mathsf{sid}_{\mathrm{pk}}(pk^*) = \perp$ or $\mathsf{sid}_{\mathrm{pk}}(pk^*) \in S$, terminate and output an independent uniform bit. Otherwise, execute
$$(CT^*, K_0) \leftarrow \mathsf{Encap}^{H}(vk_{\mathrm{rt}}, pk^*).$$
If encapsulation returns fail, terminate and output an independent uniform bit. Otherwise, write $CT^* = (h^*, Q^*)$, set $\nu^* := \mathsf{sid}_{\mathrm{ct}}(h^*)$, and sample $K_1 \leftarrow \{0,1\}^{\ell}$.

4. Run
$$\widehat{b} \leftarrow \mathcal{A}_2^{H,\mathcal{O}_{\mathrm{post}}}(CT^*, K_b, E).$$
On each query (h, Q), the post-challenge oracle executes
$$K' \leftarrow \mathsf{Decap}^H(\mathsf{st}_{\mathrm{R}}, (h, Q))$$
once, consuming Q, and returns
$$\begin{cases} \mathbf{1}[K' \neq \perp] \in \{0, 1\}, & \text{if } h = h^*, \\ K', & \text{if } h \neq h^*. \end{cases}$$
5. Output 1 if $\widehat{b} = b$, and 0 otherwise.

Define
$$\mathbf{Adv}_{\Pi}^{\mathrm{VFCCA}}(\mathcal{A}) := \left| \Pr[\mathbf{Exp}_{\Pi}^{\mathrm{VFCCA}}(\mathcal{A}) = 1] - \frac{1}{2} \right|.$$
The scheme is VFCCA-secure if this advantage is negligible for every QPT adversary $\mathcal{A}$.

VFCCA is our primary active confidentiality notion, providing full decapsulation for $h \neq h^*$ and validity feedback for $h = h^*$. To distinguish the roles of same-header feedback and same-serial decapsulation, we define two restricted variants. HCCA excludes queries with $h = h^*$ but permits decapsulation under other headers with serial ν^*. FTCCA further excludes all queries with serial ν^*.

In both models, an excluded query is handled by discarding Q and returning $\perp$, without executing decapsulation:
$$\mathcal{F}_Q(\rho_{QE}) := |\perp\rangle\langle\perp|_C \otimes \mathrm{Tr}_Q(\rho_{QE}),$$
where C stores the response and E is the adversary's retained memory.

Definition 15 (HCCA and FTCCA)**.** The HCCA and FTCCA experiments are identical to Definition 14, except for the post-challenge oracle. On a query (h, Q):

- **HCCA:** If $h = h^*$, apply $\mathcal{F}_Q$. Otherwise, return $\mathsf{Decap}^H(\mathsf{st}_{\mathrm{R}}, (h, Q))$.
- **FTCCA:** If $\mathsf{sid}_{\mathrm{ct}}(h) = \nu^*$, apply $\mathcal{F}_Q$. Otherwise, return $\mathsf{Decap}^H(\mathsf{st}_{\mathrm{R}}, (h, Q))$.

For $X \in \{\mathrm{HCCA}, \mathrm{FTCCA}\}$, define
$$\mathbf{Adv}_{\Pi}^{X}(\mathcal{A}) := \left| \Pr[\mathbf{Exp}_{\Pi}^{X}(\mathcal{A}) = 1] - \frac{1}{2} \right|.$$
The scheme is X-secure if this advantage is negligible for every QPT adversary $\mathcal{A}$.

VFCCA security implies HCCA security, which in turn implies FTCCA security. Appendix C separates FTCCA from HCCA and key robustness from confidentiality.

The preceding notions protect session-key confidentiality. We next define adaptive fixed-header key robustness (aFHKR): even given the true challenge key K^* and adaptive decapsulation access, an adversary should have negligible probability of causing acceptance of a different key under h^*.

Definition 16 (Adaptive fixed-header key robustness)**.** For a QPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, define $\mathbf{Exp}_{\Pi}^{\mathrm{aFHKR}}(\mathcal{A})$ as follows:

1. Sample a random oracle H. Generate
$$(vk_{\mathrm{rt}}, \mathsf{st}_{\mathrm{R}}^{(0)}) \leftarrow \mathsf{SKGen}(1^\lambda),$$
and issue N tokens:
$$(pk_i, \mathsf{st}_{\mathrm{R}}^{(i)}) \leftarrow \mathsf{PKGen}(\mathsf{st}_{\mathrm{R}}^{(i-1)}) \qquad (i \in [N]).$$
Set $\mathsf{st}_{\mathrm{R}} := \mathsf{st}_{\mathrm{R}}^{(N)}$.
2. Run
$$(pk^*, E) \leftarrow \mathcal{A}_1^H(vk_{\mathrm{rt}}, pk_1, \ldots, pk_N),$$
and execute $(CT^*, K^*) \leftarrow \mathsf{Encap}^H(vk_{\mathrm{rt}}, pk^*)$. If encapsulation returns fail, terminate and output 0. Otherwise, write $CT^* = (h^*, Q^*)$.
3. Run
$$\widetilde{Q} \leftarrow \mathcal{A}_2^{H,\mathcal{O}_{\mathrm{dec}}}(CT^*, K^*, E).$$
On each query (h, Q), the oracle executes
$$K' \leftarrow \mathsf{Decap}^H(\mathsf{st}_{\mathrm{R}}, (h, Q))$$
once, consuming Q. If $h = h^*$ and $K' \notin \{K^*, \perp\}$, terminate the experiment and output 1. Otherwise, return K'.

4. Execute
$$K' \leftarrow \mathsf{Decap}^H(\mathsf{st}_{\mathrm{R}}, (h^*, \widetilde{Q})).$$
Output 1 if $K' \notin \{K^*, \perp\}$, and 0 otherwise.

Define
$$\mathbf{Adv}^{\mathrm{aFHKR}}_{\Pi}(\mathcal{A}) := \Pr[\mathbf{Exp}^{\mathrm{aFHKR}}_{\Pi}(\mathcal{A}) = 1].$$
Let Q_D bound the queries to $\mathcal{O}_{\mathrm{dec}}$, excluding the final decapsulation test. Restricting $Q_D = 0$ defines FHKR and its advantage $\mathbf{Adv}^{\mathrm{FHKR}}_{\Pi}(\mathcal{A})$. The scheme satisfies aFHKR (resp. FHKR) if the corresponding advantage is negligible for every QPT adversary.

IV. Security of Authenticated Signed Coherence

In this section, we first use signature unforgeability to bound the noncanonical components of validated states. We then prove single-coordinate phase hiding and extend it to exact phase-recovery hardness in the QROM. Finally, we analyze how decapsulation acceptance or rejection affects the adversary's retained state. These results are used in the subsequent VFCCA security proof.

A. Public-Key Authentication and Canonicalization

An adversary may modify quantum public keys and entangle them with its retained memory. We bound the accepted noncanonical mass, allowing validated states to be replaced by their canonical projections with a quantified error.

a) QRO Replacement via qPRFs: For a QPT experiment $\mathcal{E}$ with an efficiently checkable winning event, let q bound all oracle queries, including those needed to check that event. Replacing a random oracle H by an independently keyed qPRF F_k gives [21]
$$\left|\Pr[\mathcal{E}^H = 1] - \Pr[\mathcal{E}^{F_k} = 1]\right| \le \mathbf{Adv}^{\mathrm{qPRF}}_{F}(\mathcal{C}; q). \tag{3}$$
The distinguisher $\mathcal{C}$ runs $\mathcal{E}$, forwards its oracle queries, and outputs its winning bit. After this replacement, the signature-security reductions sample k locally and implement $U_{F_k}|x, y\rangle = |x, y \oplus F_k(x)\rangle$ and its inverse via polynomial-size quantum circuits. Since $U^{\dagger}_{F_k} = U_{F_k}$, both operations are available without external oracle queries.

Lemma 17. *Let $q_{\mathrm{rt}} \le Q_H + 2Q_D$ bound all QRO queries before challenge candidate selection. The probability $\varepsilon_{\mathrm{rt}}$ of accepting a root certificate for an unissued pair (ν^*, vk^*) satisfies*
$$\varepsilon_{\mathrm{rt}} \le \mathbf{Adv}^{\mathrm{qPRF}}_{F}(\mathcal{C}_{\mathrm{rt}}; q_{\mathrm{rt}}) + \mathbf{Adv}^{\mathrm{EUF}}_{\Sigma_{\mathrm{rt}}}(\mathcal{B}_{\mathrm{rt}}),$$
for uniform QPT reductions $\mathcal{C}_{\mathrm{rt}}$ and $\mathcal{B}_{\mathrm{rt}}$. The latter makes exactly N classical signing queries.

Proof. Apply (3). In the keyed experiment, $\mathcal{B}_{\mathrm{rt}}$ queries its signing oracle on $\mathtt{token}\|\nu_j\|vk_j$ for every $j \in [N]$ and simulates pre-challenge decapsulation using locally generated phase keys. Any accepted unissued pair (ν^*, vk^*) yields the EUF-CMA forgery $(\mathtt{token}\|\nu^*\|vk^*, \pi^*)$ against Σ_{rt}. □

b) Canonical Subspace Projection: Fix a phase-signing key pair for the following argument. For coordinate $\xi = (\nu, i)$, let Π^{val}_{ξ} be the public validation projector and let $\sigma_{\xi,b} := \mathsf{Sign}_{\mathrm{ph},sk}(\mathtt{phase}\|\nu\|i\|b)$ be the canonical signature for branch $b \in \{0, 1\}$. For a coordinate set $\mathcal{I}$, define
$$P_{\mathrm{val}} := \left(\bigotimes_{\xi\in\mathcal{I}} \Pi^{\mathrm{val}}_{\xi}\right) \otimes I_E, \qquad P_{\mathrm{can}} := \left(\bigotimes_{\xi\in\mathcal{I}} \left[\sum_{b\in\{0,1\}} |b\rangle\langle b|_{A_\xi} \otimes |\sigma_{\xi,b}\rangle\langle\sigma_{\xi,b}|_{B_\xi}\right]\right) \otimes I_E.$$
Under perfect correctness, $P_{\mathrm{can}} \le P_{\mathrm{val}}$. Thus $P_{\mathrm{nc}} := P_{\mathrm{val}} - P_{\mathrm{can}}$ projects onto the valid but noncanonical subspace.

Lemma 18. *Let ρ be the output of a QPT algorithm given one honest signed state for each coordinate in $\mathcal{I}$. Assume all auxiliary operations are simulatable without sk or additional signing queries under that key. Let $\eta := \mathrm{Tr}(P_{\mathrm{nc}}\rho)$. Then:*

(1) *There exists a uniform QPT sEUF-CMA forger $\mathcal{B}$ making at most $2|\mathcal{I}|$ classical signing queries such that $\eta \le \mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma_{\mathrm{ph}}}(\mathcal{B})$.*

(2) *Canonical projection satisfies $\Delta(P_{\mathrm{val}}\rho P_{\mathrm{val}}, P_{\mathrm{can}}\rho P_{\mathrm{can}}) \le \sqrt{\eta} + \eta/2$. For every effect $0 \le M \le I$,*
$$|\mathrm{Tr}(MP_{\mathrm{val}}\rho P_{\mathrm{val}}) - \mathrm{Tr}(MP_{\mathrm{can}}\rho P_{\mathrm{can}})| \le \sqrt{\eta} + \eta \le 2\sqrt{\eta}.$$

Proof. The forger queries both branch signatures at every coordinate to prepare the honest input. It runs the adversary, performs public validation, and measures the branch and signature registers. With probability η, the result contains a valid signature $\sigma^* \ne \sigma_{\xi,b}$ on a queried message $\mathtt{phase}\|\nu\|i\|b$, giving a strong forgery.

For the metric bounds, write $\omega := P_{\mathrm{val}}\rho P_{\mathrm{val}}$. Since $P_{\mathrm{val}} = P_{\mathrm{can}} + P_{\mathrm{nc}}$,
$$\omega - P_{\mathrm{can}}\rho P_{\mathrm{can}} = P_{\mathrm{nc}}\rho P_{\mathrm{nc}} + P_{\mathrm{can}}\rho P_{\mathrm{nc}} + P_{\mathrm{nc}}\rho P_{\mathrm{can}}.$$

The first term has trace norm η. By Schatten–Hölder,

$$\|P_{\mathrm{can}}\rho P_{\mathrm{nc}}\|_1 \le \|P_{\mathrm{can}}\rho^{1/2}\|_2\|\rho^{1/2}P_{\mathrm{nc}}\|_2 \le \sqrt{\mathrm{Tr}(P_{\mathrm{can}}\rho)\eta} \le \sqrt{\eta}.$$

The other cross-term is its adjoint. Hence $\|\omega - P_{\mathrm{can}}\rho P_{\mathrm{can}}\|_1 \le \eta + 2\sqrt{\eta}$, and the trace distance satisfies

$$\Delta(\omega, P_{\mathrm{can}}\rho P_{\mathrm{can}}) = \tfrac{1}{2}\|\omega - P_{\mathrm{can}}\rho P_{\mathrm{can}}\|_1 \le \sqrt{\eta} + \eta/2.$$

For the effect bound, set $X := \omega - P_{\mathrm{can}}\rho P_{\mathrm{can}}$. Since $\mathrm{Tr}\, X = \eta$, the inequality $|\,\mathrm{Tr}(MX)| \le (\|X\|_1 + |\,\mathrm{Tr}\, X|)/2$ gives $|\,\mathrm{Tr}(MX)| \le \sqrt{\eta} + \eta \le 2\sqrt{\eta}$. □

B. Single-Coordinate Signed-Phase Hiding

We convert phase distinguishing into extraction of both branch signatures, and then into a signature forgery. The canonicalization bound quantifies the gap between the actual validated state and its canonical component.

For $\xi = (\nu, i)$ and $\beta \in \{0,1\}$, define $|\phi_\xi^\beta\rangle := 2^{-1/2}(|0\rangle_A |\sigma_{\xi,0}\rangle_B + (-1)^\beta |1\rangle_A |\sigma_{\xi,1}\rangle_B)$.

Definition 19 (Validated Signed-Phase Hiding)**.** Fix a target coordinate $\xi^* = (\nu^*, i^*)$ in a polynomial-size set $\mathcal{L}$. For a signature scheme Σ and QPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, define $\mathbf{Exp}^{\mathsf{VSPH}}_{\Sigma,\xi^*,\mathcal{L}}(\mathcal{A})$:

1. Sample $(sk, vk) \leftarrow \mathsf{Gen}(1^\lambda)$ and $\beta \leftarrow \{0,1\}$. Prepare $\{|\phi_\xi^0\rangle\}_{\xi\in\mathcal{L}}$ using sk, and give these states and vk to $\mathcal{A}_1$.
2. $\mathcal{A}_1$ outputs candidate registers AB and retains auxiliary memory E, possibly entangled with AB.
3. Measure AB with $\{\Pi^{\mathrm{val}}_{\xi^*}, I - \Pi^{\mathrm{val}}_{\xi^*}\}$. On rejection, terminate and output an independent uniform bit.
4. On acceptance, apply Z^β to the A_{ξ^*} subregister of AB and return AB together with the retained memory E to $\mathcal{A}_2$.
5. Run $\widehat{\beta} \leftarrow \mathcal{A}_2(AB, E)$. Output 1 if $\widehat{\beta} = \beta$, and 0 otherwise.

Define $\mathbf{Adv}^{\mathsf{VSPH}}_{\Sigma,\xi^*}(\mathcal{A}) := |\Pr[\mathbf{Exp}^{\mathsf{VSPH}}_{\Sigma,\xi^*,\mathcal{L}}(\mathcal{A}) = 1] - 1/2|$. Validated signed-phase hiding requires this advantage to be negligible for every target $\xi^* \in \mathcal{L}$ and every uniform QPT adversary with $\mathcal{L}$ an experiment parameter fixed for the advantage definition.

For an extractor $\mathcal{T}$ given the same initial input, define

$$\mathbf{Succ}^{\mathsf{TwoSig}}_{\Sigma,\xi^*}(\mathcal{T}) := \Pr\big[\mu_0 = \sigma_{\xi^*,0} \wedge \mu_1 = \sigma_{\xi^*,1} : (\mu_0,\mu_1) \leftarrow \mathcal{T}(vk, \{|\phi_\xi^0\rangle\}_{\xi\in\mathcal{L}})\big].$$

Lemma 20 (Phase Distinguishing to Strong Forgery)**.** *Consider Definition 19 under the conditions of Lemma 18. All auxiliary interfaces must be phase-independent and simulatable without the target signing key or additional target signing queries. Those used by $\mathcal{A}_2$ must admit efficient unitary dilations and inverses available to the reduction.*

Let ε be the phase advantage on the subnormalized canonical accepted component, weighted by its mass without conditioning on acceptance. Let η be the accepted noncanonical mass. There exists a uniform QPT sEUF-CMA forger $\mathcal{F}$ making at most $2|\mathcal{L}| - 1$ classical signing queries such that

$$\varepsilon^2 \le 2\mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma}(\mathcal{F}) + \sqrt{\eta} + \eta. \tag{4}$$

Proof. **Step 1: Relating distinguishing to a transition amplitude.** Dilate $\mathcal{A}_2$ to a unitary V acting on AB, auxiliary memory and workspace E, and a decision qubit D_{out} initialized to $|0\rangle$. For the analysis, include any purifying environment in E; V acts trivially on that environment. Set $O := V^\dagger(Z_{D_{\mathrm{out}}} \otimes I)V$.

Fix the setup randomness for this calculation. The purified, subnormalized canonical state is $|\psi_\beta\rangle = c_0 |x_0\rangle + (-1)^\beta c_1 |x_1\rangle$, where $|x_b\rangle := |b\rangle_A |\sigma_{\xi^*,b}\rangle_B |\Omega_b\rangle_E |0\rangle_{D_{\mathrm{out}}}$ are orthonormal and $|c_0|^2 + |c_1|^2 \le 1$. Let r_β be the probability mass of output 1 on this component. Assigning an independent uniform output to the missing mass gives $\varepsilon = \frac{1}{2}|r_1 - r_0|$ for this setup. Since $Z_{D_{\mathrm{out}}} = I - 2\,|1\rangle\langle 1|_{D_{\mathrm{out}}}$,

$$\varepsilon = \frac{1}{4}\,|\langle\psi_0|\,O\,|\psi_0\rangle - \langle\psi_1|\,O\,|\psi_1\rangle| = \big|\mathrm{Re}\big(c_0^* c_1\,\langle x_0|\,O\,|x_1\rangle\big)\big| \le |c_0 c_1|\,|\,\langle x_0|\,O\,|x_1\rangle\,|. \tag{5}$$

Step 2: Extracting both signatures. The extractor $\mathcal{T}$ runs $\mathcal{A}_1$, performs public validation, and measures A. It halts unless validation accepts and the outcome is 0. It measures B to record the first signature, initializes the fresh ancillas of V, including D_{out}, applies O, and measures B again. It preserves the adversary's retained memory.

On the canonical component, the first branch has mass $|c_0|^2$ and yields $\sigma_{\xi^*,0}$. Let Π_{σ_1} project onto states with signature register $\sigma_{\xi^*,1}$. Since $|x_1\rangle \in \mathrm{Im}(\Pi_{\sigma_1})$, $O = O^\dagger$, and $|c_1|^2 \le 1$,

$$r_{\mathrm{can}} \ge |c_0|^2\|\Pi_{\sigma_1} O\,|x_0\rangle\,\|^2 \ge |c_0|^2|\,\langle x_1|\,O\,|x_0\rangle\,|^2 \ge \varepsilon^2.$$

The extractor implements O by running V, applying $Z_{D_{\mathrm{out}}}$, and running $V^\dagger$. In the keyed oracle setting, qPRF calls are implemented locally, with $U^\dagger_{F_k} = U_{F_k}$. The triangle inequality and Jensen's inequality preserve $r_{\mathrm{can}} \ge \varepsilon^2$ after averaging, now with both quantities taken over the full experiment.

Step 3: Comparing canonical and actual extraction. Let ρ be the subnormalized validated state. The remaining extraction procedure defines a success effect $0 \le M \le I$, so $p := \mathrm{Tr}(M\rho) = \mathbf{Succ}^{\mathsf{TwoSig}}_{\Sigma,\xi^*}(\mathcal{T})$ and $r_{\mathrm{can}} = \mathrm{Tr}(MP_{\mathrm{can}}\rho P_{\mathrm{can}})$. Lemma 18

gives $|p - r_{\mathrm{can}}| \le \sqrt{\eta} + \eta$, hence $\varepsilon^2 \le p + \sqrt{\eta} + \eta$. The bound holds after averaging by concavity of the square root. The canonical projection is used only for this comparison.

Step 4: Turning extraction into forgery. Represent the complete extraction success event on the issued input by an effect $0 \le M_{\mathrm{in}} \le I$. Decompose the input into its orthogonal target-branch components $|\varphi_0\rangle + |\varphi_1\rangle$. The coherent success probability is $p = \langle\varphi_0 + \varphi_1| M_{\mathrm{in}} |\varphi_0 + \varphi_1\rangle$. Measuring the target branch first gives $q := \langle\varphi_0| M_{\mathrm{in}} |\varphi_0\rangle + \langle\varphi_1| M_{\mathrm{in}} |\varphi_1\rangle$. Positivity implies

$$2\,|\mathrm{Re}\,\langle\varphi_0| M_{\mathrm{in}} |\varphi_1\rangle| \le 2\sqrt{\langle\varphi_0| M_{\mathrm{in}} |\varphi_0\rangle \langle\varphi_1| M_{\mathrm{in}} |\varphi_1\rangle} \le q,$$

so $p \le 2q$, also after averaging.

The forger samples $a \leftarrow \{0,1\}$, requests the signature on $m_{\xi^*,a}$, and requests both signatures at every other coordinate. These $2|\mathcal{L}| - 1$ queries prepare the branch-measured honest input, on which it runs $\mathcal{T}$. Whenever both canonical signatures are extracted, μ_{1-a} is a valid signature on the unqueried message $m_{\xi^*,1-a}$. Thus $\mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma}(\mathcal{F}) \ge q \ge p/2$. Combining this with Step 3 proves (4). □

Theorem 21 (Single-Coordinate Phase Hiding)**.** *If Σ_{ph} is canonical and sEUF-CMA secure, then for every uniform QPT adversary $\mathcal{A}$ in Definition 19, there exist uniform QPT forgers $\mathcal{B}_0, \mathcal{B}_1$ making at most $2|\mathcal{L}|$ and $2|\mathcal{L}| - 1$ classical signing queries, respectively, such that*

$$\mathbf{Adv}^{\mathsf{VSPH}}_{\Sigma_{\mathrm{ph}},\xi^*}(\mathcal{A}) \le 2\sqrt{\mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma_{\mathrm{ph}}}(\mathcal{B}_0)} + \sqrt{2\mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma_{\mathrm{ph}}}(\mathcal{B}_1) + 2\sqrt{\mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma_{\mathrm{ph}}}(\mathcal{B}_0)}} \le \mathrm{negl}(\lambda). \tag{6}$$

Auxiliary interfaces must satisfy the simulation and inversion conditions of Lemma 20. For QROM applications, the qPRF replacement loss is accounted for separately.

Proof. Lemma 18 gives $\eta \le \mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma_{\mathrm{ph}}}(\mathcal{B}_0)$. Set $e := \sqrt{\eta} + \eta \le 2\sqrt{\eta}$, and let ε be the phase advantage on the subnormalized canonical component. Canonical projection gives $\mathbf{Adv}^{\mathsf{VSPH}}_{\Sigma_{\mathrm{ph}},\xi^*}(\mathcal{A}) \le \varepsilon + e$. Lemma 20 gives $\varepsilon^2 \le 2\mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma_{\mathrm{ph}}}(\mathcal{B}_1) + e$. Consequently, $\mathbf{Adv}^{\mathsf{VSPH}}_{\Sigma_{\mathrm{ph}},\xi^*}(\mathcal{A}) \le e + \sqrt{2\mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma_{\mathrm{ph}}}(\mathcal{B}_1) + e}$. Substituting $e \le 2\sqrt{\mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma_{\mathrm{ph}}}(\mathcal{B}_0)}$ proves (6). The query counts follow from the two reductions. □

C. Lifting to the Full Fourier Phase

We extend single-coordinate phase hiding to the full ℓ-bit phase string by changing one comparison bit at a time.

Definition 22 (Fourier Hybrids)**.** Let $\mathcal{O}$ be a phase-independent auxiliary interface simulatable without the target signing key or additional target signing queries, with an efficient unitary dilation and inverse available to the reduction. Each domain $j \in [N]$ uses an independent signing key pair, and one quantum public key is issued per domain.

For $t \in \{0, \dots, \ell\}$, define $\mathsf{Hyb}^{\mathcal{O}}_t(\mathcal{A})$:

1. Generate N honest quantum public keys $pk_1, \dots, pk_N$ and give them to $\mathcal{A}^{\mathcal{O}}_1$.
2. $\mathcal{A}^{\mathcal{O}}_1$ outputs metadata (ν^*, vk^*) and a quantum register $Q^* = A_1 B_1 \cdots A_\ell B_\ell$, retaining memory E. If the metadata matches no issued domain, output 0. Otherwise, let J be the unique matching domain.
3. Measure Q^* with $\{\Pi^{\mathrm{val}}_J, I - \Pi^{\mathrm{val}}_J\}$. On rejection, output 0.
4. Sample independent uniform $s, u \leftarrow \{0,1\}^\ell$ and apply $Z^s = \bigotimes_{i=1}^{\ell} Z^{s_i}_{A_i}$ to Q^*, obtaining Q^*_s.
5. Set $Y^{(t)} := (u_1, \dots, u_t, s_{t+1}, \dots, s_\ell)$ and output $b \leftarrow \mathcal{A}^{\mathcal{O}}_2(Q^*_s, Y^{(t)}, E)$.

Write $p_t := \Pr[\mathsf{Hyb}^{\mathcal{O}}_t(\mathcal{A}) = 1]$ and define

$$\Delta^{(j)}_{\mathrm{F}}(\mathcal{A}) := \left|\Pr[\mathsf{Hyb}^{\mathcal{O}}_0(\mathcal{A}) = 1 \wedge J = j] - \Pr[\mathsf{Hyb}^{\mathcal{O}}_\ell(\mathcal{A}) = 1 \wedge J = j]\right|.$$

Lemma 23. *For every $j \in [N]$ and $i \in [\ell]$, there exists a uniform QPT adversary $\mathcal{B}_{j,i}$ such that*

$$\left|\Pr[\mathsf{Hyb}^{\mathcal{O}}_{i-1}(\mathcal{A}) = 1 \wedge J = j] - \Pr[\mathsf{Hyb}^{\mathcal{O}}_i(\mathcal{A}) = 1 \wedge J = j]\right| \le \mathbf{Adv}^{\mathsf{VSPH}}_{\Sigma_{\mathrm{ph}},(\nu_j,i)}(\mathcal{B}_{j,i}).$$

Proof. The reduction receives vk_j and honest signed states for $\mathcal{L}_j = \{(\nu_j, k) : k \in [\ell]\}$. It generates the other domains locally and simulates $\mathcal{O}$ using the assumed dilation.

If $J \ne j$, or validation under $\Pi^{\mathrm{val}}_{j,-i} := \bigotimes_{k \ne i} \Pi^{\mathrm{val}}_{j,k}$ fails, it makes a dummy challenge submission and outputs an independent uniform guess. Otherwise, it submits (A_i, B_i) to its single-coordinate challenger. The commuting validation projectors ensure that joint validation agrees with the hybrid experiment.

After receiving the modulated register, it independently samples s_k for $k \ne i$, u_k for $k < i$, and $y \leftarrow \{0,1\}$. It applies Z^{s_k} at each coordinate $k \ne i$ and runs $\mathcal{A}^{\mathcal{O}}_2(Q^*_s, \widetilde{Y}, E)$ with $\widetilde{Y} := (u_1, \dots, u_{i-1}, y, s_{i+1}, \dots, s_\ell)$. It outputs y if $\mathcal{A}_2$ outputs 1, and $1 - y$ otherwise.

Let $A_{j,i}$ and $B_{j,i}$ be the probabilities of $J = j$, full validation accepting, and $\mathcal{A}_2 = 1$, conditioned on $y = \beta$ and $y \ne \beta$, respectively. The first condition gives Hyb_{i-1}. Independent uniform y gives Hyb_i, whose output probability on domain j is

$(A_{j,i}+B_{j,i})/2$. The adjacent gap is therefore $|A_{j,i}-B_{j,i}|/2$. The reduction guesses correctly with probability $1/2+(A_{j,i}-B_{j,i})/2$, so its VSPH advantage equals this gap. □

Decomposing over the domain variable J and telescoping across coordinates yields

$$|p_0 - p_\ell| \le \sum_{j=1}^{N} \Delta_{\mathrm{F}}^{(j)}(\mathcal{A}) \le \sum_{j=1}^{N}\sum_{i=1}^{\ell} \mathbf{Adv}^{\mathsf{VSPH}}_{\Sigma_{\mathrm{ph}},(\nu_j,i)}(\mathcal{B}_{j,i}) \le \mathrm{negl}(\lambda), \tag{7}$$

by Theorem 21 and the polynomial bounds on N and ℓ.

D. Phase-Recovery Hardness in the QROM

We now use full-phase hiding to bound the exact recovery probability of the phase string s used to derive the session key $K = H_{\mathrm{kdf}}(\nu, \tau, s)$.

Fix a target domain $j \in [N]$. Both stages of $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ have quantum access to H and access to an interface $\mathcal{O}$ satisfying Definition 22. The recovery experiment proceeds as follows:

1. Give $\mathcal{A}_1^{H,\mathcal{O}}$ honest quantum public keys for all N domains. It outputs metadata (ν^*, vk^*) and Q^*, retaining memory E.
2. Match the metadata to an issued domain J. If no match exists or $J \neq j$, terminate with failure. Measure Q^* under Π_j^{val}; on rejection, terminate with failure. Otherwise, sample uniform $s \leftarrow \{0,1\}^\ell$, apply $Z^s = \bigotimes_{i=1}^{\ell} Z_{A_i}^{s_i}$ to Q^*, obtaining Q_s^*.
3. Run $\widehat{s} \leftarrow \mathcal{A}_2^{H,\mathcal{O}}(Q_s^*, E)$. The experiment succeeds if $\widehat{s} = s$.

In security reductions, all non-target domains are simulated locally. Let Acc denote validation acceptance and define $p_{\mathrm{rec}}^{(j)}(\mathcal{A}) := \Pr[\mathsf{Acc} \wedge J = j \wedge \widehat{s} = s]$. Failure branches contribute zero; this probability is not conditioned on acceptance.

Corollary 24 (QROM Phase-Recovery Bound)**.** *Let q bound all queries to H, including simulation of $\mathcal{O}$ and the terminal success check. There exist uniform QPT reductions $\mathcal{C}, \mathcal{B}_0, \mathcal{B}_1$ such that*

$$p_{\mathrm{rec}}^{(j)}(\mathcal{A}) \le 2^{-\ell} + \mathbf{Adv}_F^{\mathrm{qPRF}}(\mathcal{C}; q) + 4\ell \left(\mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma_{\mathrm{ph}}}(\mathcal{B}_0)^{1/4} + \sqrt{\mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma_{\mathrm{ph}}}(\mathcal{B}_1)} \right), \tag{8}$$

where $\mathcal{B}_0$ and $\mathcal{B}_1$ make at most 2ℓ and $2\ell - 1$ classical signing queries, respectively.

Proof. By (3), replace H by F_k at loss at most $\mathbf{Adv}_F^{\mathrm{qPRF}}(\mathcal{C}; q)$. In the keyed experiment, sample independent uniform $u \leftarrow \{0,1\}^\ell$ and define $p_i := \Pr[\mathsf{Acc} \wedge J = j \wedge \widehat{s} = (u_1, \ldots, u_i, s_{i+1}, \ldots, s_\ell)]$ for $i \in \{0, \ldots, \ell\}$. The comparison string is used only to test $\widehat{s}$ and is never given to the recovery algorithm. Thus p_0 is its recovery probability, while $p_\ell = 2^{-\ell} \Pr[\mathsf{Acc} \wedge J = j] \le 2^{-\ell}$.

For each adjacent pair, an outer distinguisher runs the recovery algorithm without revealing the comparison string and tests whether its output matches that string. Lemma 23 and the extraction analysis in the proof of Lemma 20 give $|p_{i-1} - p_i| \le 2\sqrt{\eta_i} + \sqrt{r_i + 2\sqrt{\eta_i}}$, where η_i is the accepted noncanonical mass and r_i is the actual two-signature extraction probability.

Uniformly sampling $i \leftarrow [\ell]$ defines $\mathcal{B}_0, \mathcal{B}_1$. Write $a_0 := \mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma_{\mathrm{ph}}}(\mathcal{B}_0)$ and $a_1 := \mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma_{\mathrm{ph}}}(\mathcal{B}_1)$. Then $\ell^{-1}\sum_i \eta_i \le a_0$ and $\ell^{-1}\sum_i r_i \le 2a_1$. By Jensen's inequality,

$$\begin{aligned} \sum_{i=1}^{\ell} |p_{i-1} - p_i| &\le 2\ell\sqrt{a_0} + \ell\sqrt{2a_1 + 2\sqrt{a_0}} \\ &\le 2\ell\sqrt{a_0} + \ell\sqrt{2a_1} + \sqrt{2}\,\ell a_0^{1/4} \\ &\le 4\ell(a_0^{1/4} + \sqrt{a_1}). \end{aligned}$$

Telescoping gives $p_0 \le p_\ell + \sum_i |p_{i-1} - p_i|$. Adding the qPRF replacement loss proves (8). □

This bounds phase recovery under the stated auxiliary-interface conditions. We next analyze how decapsulation acceptance or rejection in VFCCA affects the adversary's retained state and its ability to recover the phase string.

E. Authenticated Dephasing and Branch Preservation

We establish a dephased reference execution in which changing a canonical signed branch is bounded by signature unforgeability.

For a pre-phase state ρ_{ABE} independent of s,

$$2^{-\ell} \sum_{s \in \{0,1\}^\ell} Z_A^s \rho_{ABE} Z_A^s = \sum_z P_z \rho_{ABE} P_z, \qquad P_z := |z\rangle\langle z|_A \otimes I_{BE}. \tag{9}$$

The (z, z') block acquires $(-1)^{s\cdot(z \oplus z')}$, whose average vanishes unless $z = z'$. This describes the phase-averaged state; the actual challenge need not be measured.

Lemma 25 (Authenticated Dephasing and Branch Preservation)**.** *Assume quantum public-key preparation satisfies Lemma 18. All operations before the final submission must be simulatable without the target signing key, additional target signing queries, or additional quantum public keys under that key.*

Let ω be the subnormalized validated state. In the reference execution, project onto P_{can}, dephase the branch register, and privately record z. Let ω_t be the subnormalized joint state of the final submission, retained memory E, and recorded z. Let P_{diff} project onto valid canonical submissions with branch $z' \neq z$.

For at most d oracle queries, including simulation and terminal event checks, there exist uniform QPT reductions $\mathcal{C}_0, \mathcal{C}_1, \mathcal{B}_0, \mathcal{B}_1$ such that:

(1) ***Noncanonical mass.*** *The mass $\eta := \mathrm{Tr}((I - P_{\mathrm{can}})\omega)$ satisfies*

$$\eta \leq \mathbf{Adv}_F^{\mathrm{qPRF}}(\mathcal{C}_0; d) + \mathbf{Adv}_{\Sigma_{\mathrm{ph}}}^{\mathrm{sEUF}}(\mathcal{B}_0). \tag{10}$$

(2) ***Branch mismatch.*** *The mass $\mu := \mathrm{Tr}(P_{\mathrm{diff}}\omega_t)$ satisfies*

$$\mu \leq \mathbf{Adv}_F^{\mathrm{qPRF}}(\mathcal{C}_1; d) + 2\ell\, \mathbf{Adv}_{\Sigma_{\mathrm{ph}}}^{\mathrm{sEUF}}(\mathcal{B}_1), \tag{11}$$

where $\mathcal{B}_1$ makes at most $2\ell - 1$ classical signing queries.

(3) ***Expectation change.*** *After fixing the baseline oracle and public parameters, suppose the pre-phase state, continuation channels, and final effect $0 \leq M \leq I$ are independent of s, with M supported on the canonical subspace. For uniform s, canonical projection and removal of branch mismatches in the dephased reference execution change the phase-averaged expectation by at most $2\sqrt{\eta} + 2\sqrt{\mu} + \mu$.*

In locally simulated settings, (1) and (2) hold without the qPRF terms. The expectation bound is negligible under qPRF security and sEUF-CMA security of Σ_{ph}.

Proof. **Step 1: Bounding noncanonical mass.** Replace H by an independently keyed qPRF F_k. Equation (3) bounds the change in noncanonical mass by $\mathbf{Adv}_F^{\mathrm{qPRF}}(\mathcal{C}_0; d)$. The signature reduction samples k locally and applies Lemma 18, giving $\eta^{F_k} \leq \mathbf{Adv}_{\Sigma_{\mathrm{ph}}}^{\mathrm{sEUF}}(\mathcal{B}_0)$.

Step 2: Bounding branch mismatch. In the keyed execution, let μ_i be the probability of a valid canonical submission with $z_i' \neq z_i$. Then $\mu^{F_k} \leq \sum_i \mu_i$. The measurement-loss argument in the proof of Lemma 20 gives $\mu_i \leq 2q_i$, where q_i is the mismatch probability when coordinate i of the honest quantum public key is measured at issuance.

The forger samples $i \leftarrow [\ell]$ and $c \leftarrow \{0,1\}$. It requests the signature on $m_{(\nu,i),c}$ and both signatures at every other coordinate, using $2\ell - 1$ queries to prepare the branch-measured honest input.

At reference dephasing, it measures the candidate branch z. If $z_i \neq c$ and public validation accepts, measuring B_i yields a valid signature on the unqueried message $m_{(\nu,i),1-c}$. When $z_i = c$, all canonical signatures for z are known. It projects B onto $|\boldsymbol{\sigma}_{\nu,z}\rangle$ and, on success, continues. A later valid submission with $z_i' = 1 - c$ again yields a signature on the unqueried message.

Thus $\mathbf{Adv}_{\Sigma_{\mathrm{ph}}}^{\mathrm{sEUF}}(\mathcal{B}_1) \geq \ell^{-1} \sum_i q_i \geq \mu^{F_k}/(2\ell)$. The qPRF comparison experiment generates the signing keys locally, so its mismatch test is efficiently checkable. Applying (3) adds the loss $\mathbf{Adv}_F^{\mathrm{qPRF}}(\mathcal{C}_1; d)$ and proves (2).

Step 3: Bounding the expectation change. Under the independence conditions, (9) evaluates the phase-averaged expectation in the dephased reference execution. Canonical projection changes the expectation by at most $\sqrt{\eta} + \eta \leq 2\sqrt{\eta}$. Removing the mismatch component of weight μ changes the subnormalized state in trace norm by at most $2\sqrt{\mu} + \mu$, by the gentle measurement bound. Adding these errors proves (3). □

F. Consuming Validity Instruments and State Decoupling

In VFCCA, same-header queries return an acceptance bit and may change the adversary's retained state. Let r denote the phase string decoded during decapsulation, and let s be the phase string encoded in the challenge. We bound the contribution from the rejection branch and characterize the adversary's residual state conditioned on acceptance with decoded phase $r = s$.

For an issued serial ν, let $\boldsymbol{\sigma}_{\nu,z}$ denote the tuple of canonical signatures specified by the branch string $z \in \{0,1\}^\ell$. Set $D := 2^\ell$ and $|\chi_z\rangle := |z\rangle_A |\boldsymbol{\sigma}_{\nu,z}\rangle_B$. The canonical Fourier state for decoded phase r is

$$|\Phi_r\rangle := D^{-1/2} \sum_{z \in \{0,1\}^\ell} (-1)^{r \cdot z} |\chi_z\rangle . \tag{12}$$

For a header $h = (\nu, \tau, \gamma)$, the acceptance effect is

$$E_h^H := \sum_{\substack{r \in \{0,1\}^\ell: \\ H_{\mathrm{bnd}}(\nu,\tau,r) = \gamma}} |\Phi_r\rangle\langle\Phi_r| . \tag{13}$$

Set $E_h^H := 0$ if $\nu \notin \mathsf{DB}$ or h is ill-formed. The states in this expression use the signing domain identified by ν.

Pulling decoded outcome r back through coherent erasure and the Hadamard transform gives $|\Phi_r\rangle\langle\Phi_r|$. Validation and the zero-test restrict acceptance to the canonical subspace. Let Q be the submitted quantum ciphertext, E the adversary's retained memory, and C the classical response register. Summing over outcomes passing the binding check gives

$$\mathcal{V}_h^H(\rho_{QE}) = |1\rangle\langle 1|_C \otimes \mathrm{Tr}_Q[(E_h^H \otimes I_E)\rho_{QE}] + |0\rangle\langle 0|_C \otimes \mathrm{Tr}_Q[((I - E_h^H) \otimes I_E)\rho_{QE}]. \tag{14}$$

This channel is completely positive and trace preserving, and independent of H_{kdf}. Write its subnormalized rejection branch as $\mathcal{R}_h(\rho) := |0\rangle\langle 0|_C \otimes \mathrm{Tr}_Q[((I - E_h^H) \otimes I_E)\rho]$. For comparison, define $\mathcal{F}_Q^{(0)}(\rho) := |0\rangle\langle 0|_C \otimes \mathrm{Tr}_Q(\rho)$, which discards Q and always returns 0.

Lemma 26 (Validity Feedback and Phase Decoupling)**.** (1) ***Rejection-branch domination.*** *The map $\mathcal{F}_Q^{(0)} - \mathcal{R}_h$ is completely positive. Replacing the subnormalized rejection branch by $\mathcal{F}_Q^{(0)}$ upper-bounds subsequent event probabilities under a common completely positive continuation. The same holds with 0 relabelled as $\perp$.*

(2) ***Accepted-phase decoupling.*** *Fix a baseline oracle H_0, receiver keys, and public parameters, including the challenge header $h^* = (\nu^*, \tau^*, \gamma^*)$. Sample uniform $s \leftarrow \{0,1\}^\ell$ independently of these parameters and the pre-phase state. Consider a simulated prefix using H_0 whose only dependence on s is the challenge modulation Z^s. The terminal test uses h^* and $H_s := H_0[00\|\nu^*\|\tau^*\|s \mapsto \gamma^*]$.*
Let η and μ be the noncanonical and branch-mismatch masses for this prefix, as defined in the preceding subsection. Let S be a classical register recording s, with uniform state $U_S := 2^{-\ell} I_S$. Let $\rho_{SE}^{r=s}$ be the subnormalized joint output associated with acceptance under h^* and decoded phase $r = s$. There exists an s-independent subnormalized state ω_E, with $\mathrm{Tr}(\omega_E) \le 1$, such that

$$\|\rho_{SE}^{r=s} - U_S \otimes \omega_E\|_1 \le 4\sqrt{\eta} + 2\sqrt{\mu}. \tag{15}$$

(3) ***Recovery from the reference state.*** *With the same fixed parameters and oracle H_s, suppose a continuation starts from $U_S \otimes \omega_E$. Apart from access to H_s and its consuming validity channel, it receives no information about s. Let d and Q_D bound its direct QRO queries and consuming validity queries, respectively. Its phase-recovery probability is at most*

$$b_{\mathrm{ind}} := 2^{-\ell} + 2d \cdot 2^{-\ell/2} + Q_D 2^{-\ell}. \tag{16}$$

Proof. **(1) Rejection-branch domination.** The difference map evaluates as

$$(\mathcal{F}_Q^{(0)} - \mathcal{R}_h)(\rho) = |0\rangle\langle 0|_C \otimes \mathrm{Tr}_Q[(E_h^H \otimes I_E)\rho].$$

This map is completely positive because $\mathrm{Tr}_Q[(E_h^H \otimes I_E)\rho] = \mathrm{Tr}_Q[(\sqrt{E_h^H} \otimes I_E)\rho(\sqrt{E_h^H} \otimes I_E)]$. Composition with a common completely positive continuation preserves this ordering, giving the event-probability bound. This ordering comparison applies to subnormalized branches, not states renormalized by conditioning on rejection.

(2) Accepted-phase decoupling. We work with the canonical Fourier states for serial ν^*. Purify the canonical pre-phase state as $|\Psi_0\rangle = \sum_z |\chi_z\rangle\,|\Omega_z\rangle$, where $|\Omega_z\rangle$ is independent of s and $\sum_z \|\Omega_z\|^2 \le 1$. After modulation, $|\Psi_s\rangle = \sum_z (-1)^{s\cdot z} |\chi_z\rangle\,|\Omega_z\rangle$.

Let V be an s-independent contraction dilating the subsequent prefix and projecting the final submitted register onto the canonical subspace. Write $V|\Psi_s\rangle = \sum_{y,z} (-1)^{s\cdot z} |\chi_y\rangle\, e_{y,z}$, where the vectors $e_{y,z}$ are independent of s. The dephased reference execution gives $\sum_{y\neq z} \|e_{y,z}\|^2 = \mu$, and contractivity gives $\sum_{y,z} \|e_{y,z}\|^2 \le 1$.

Projection onto $|\Phi_s\rangle$ leaves the memory vector $v_s := D^{-1/2} \sum_{y,z} (-1)^{s\cdot(z\oplus y)} e_{y,z}$. Since $v_s = (\langle\Phi_s|_Q \otimes I)V|\Psi_s\rangle$, we have $\|v_s\| \le 1$. The diagonal terms form $w := D^{-1/2} \sum_z e_{z,z}$, which is independent of s and satisfies $\|w\|^2 \le \sum_z \|e_{z,z}\|^2 \le 1$.

Grouping the remaining terms by $a := y \oplus z \neq 0$ and using character orthogonality gives

$$\mathbb{E}_s \|v_s - w\|^2 = D^{-1} \sum_{a\neq 0} \left\| \sum_z e_{z\oplus a, z} \right\|^2 \le \sum_{a\neq 0, z} \|e_{z\oplus a, z}\|^2 = \mu. \tag{17}$$

Therefore $\mathbb{E}_s \| |v_s\rangle\langle v_s| - |w\rangle\langle w| \|_1 \le 2\mathbb{E}_s \|v_s - w\| \le 2\sqrt{\mu}$. Tracing out the purifying environment in $|w\rangle\langle w|$ gives an s-independent state ω_E with trace at most 1. Thus the canonical output is within $2\sqrt{\mu}$ in trace norm of $U_S \otimes \omega_E$.

By Lemma 18, canonical projection changes the input in trace norm by at most $2\sqrt{\eta} + \eta \le 4\sqrt{\eta}$. Contractivity under the common prefix and terminal test, followed by the triangle inequality, proves (15).

(3) Recovery from the reference state. Let $m := d + Q_D$. For $k \in \{0, \dots, m\}$, define Hyb_k to use H_0 in the first k query slots and H_s thereafter, including the corresponding validity channels. In Hyb_m, the view is independent of s, so recovery succeeds with probability at most $2^{-\ell}\,\mathrm{Tr}(\omega_E) \le 2^{-\ell}$.

Adjacent hybrids share an all-H_0 prefix before the changed query and a common continuation afterwards. The entering state is therefore independent of s. For a direct QRO query, the oracles differ only at $00\|\nu^*\|\tau^*\|s$. The average query weight there is at most $2^{-\ell}$, giving trace distance at most $2 \cdot 2^{-\ell/2}$.

For a consuming validity query, the channels can differ only when $(\nu, \tau) = (\nu^*, \tau^*)$ and the decoded phase is s. Let ρ be the subnormalized input on this matching branch. The average mass of that outcome is $2^{-\ell} \sum_r \mathrm{Tr}[(|\Phi_r\rangle\langle\Phi_r| \otimes I_E)\rho] \le 2^{-\ell}$.

The channels leave the same subnormalized memory state for that outcome and can differ only in the classical response. Its mass bounds their output trace distance.

Contractivity under the common continuation and summation over query slots give (16). Averaging over the fixed parameters completes the proof. □

Applying this unconditional recovery bound to the actual accepted $r = s$ output additionally accumulates the trace-distance error from (15).

V. From Signed Coherence to Active QKEM Security

Building on the phase-hiding and decoupling bounds developed above, we prove CPA confidentiality, extend it to FTCCA and VFCCA, and then establish adaptive fixed-header key robustness.

A. CPA Confidentiality

We first prove key confidentiality without decapsulation queries. By O2H, key distinguishing reduces to phase recovery, whose probability is bounded in Section IV-D.

Definition 27 (Selective challenge experiments)**.** Sample $J \leftarrow [N]$, $s \leftarrow \{0,1\}^{\ell}$, $\tau \leftarrow \{0,1\}^{\lambda}$, and independent uniform $Y_{\mathrm{b}}, Y_{\mathrm{k}} \leftarrow \{0,1\}^{2\lambda}$. Set $x_{\mathrm{b}} := 00\|\nu_J\|\tau\|s$ and $x_{\mathrm{k}} := 01\|\nu_J\|\tau\|s$. For a uniform random oracle H_0, define

$$H^{(0)} := H_0[x_{\mathrm{b}} \mapsto Y_{\mathrm{b}}, x_{\mathrm{k}} \mapsto Y_{\mathrm{k}}], \qquad H^{(1)} := H_0[x_{\mathrm{b}} \mapsto Y_{\mathrm{b}}].$$

On an accepted candidate with metadata (ν_J, vk_J), apply Z^s to its quantum register, obtaining Q_s, and return

$$CT^* = ((\nu_J, \tau, Y_{\mathrm{b}}), Q_s), \qquad K_{\mathrm{test}} := \mathrm{first}_{\ell}(Y_{\mathrm{k}}). \tag{18}$$

A mismatched candidate or encapsulation failure causes the experiment to terminate and output an independent uniform bit.

Let $\mathsf{Hy}^{(b)}$ use $H^{(b)}$ throughout. The cases $b = 0$ and $b = 1$ correspond to a real test key and an independent uniform test key. The common ideal execution $\mathsf{Hy}_{\mathrm{ideal}}$ uses H_0 throughout with the same challenge tuple (18).

Theorem 28 (CPA confidentiality)**.** *Assume EUF-CMA security of Σ_{rt}, canonical sEUF-CMA security of Σ_{ph}, and qPRF security of F. Then the Signed-Fourier QKEM satisfies*

$$\mathbf{Adv}_{\Pi}^{\text{IND-QKEM-CPA}}(\mathcal{A}) \le \varepsilon_{\mathrm{rt}} + 2Nd\sqrt{p_0} \le \mathrm{negl}(\lambda),$$

where $d = \max\{1, Q_H\}$ and p_0 is the probability that a uniformly selected query slot in $\mathsf{Hy}_{\mathrm{ideal}}$ hits $\{x_{\mathrm{b}}, x_{\mathrm{k}}\}$. Inactive or unreached slots contribute zero.

Proof. First terminate neutrally whenever a certificate for an unissued pair verifies. Lemma 17 bounds the resulting change in advantage by $\varepsilon_{\mathrm{rt}}$. Guessing $J \leftarrow [N]$ retains each remaining contribution to the distinguishing gap with probability $1/N$. Hence $\mathbf{Adv}_{\Pi}^{\text{IND-QKEM-CPA}}(\mathcal{A}) \le \varepsilon_{\mathrm{rt}} + N\mathbf{Adv}_{\mathrm{sel}}^{\text{IND-QKEM-CPA}}(\mathcal{A})$.

Before the challenge, the ideal query state is independent of (τ, s), so its probability of hitting either programmed point is at most $2^{-(\lambda+\ell)}$.

For post-challenge queries, construct a recovery algorithm $\mathcal{R}_0$ that samples $J, \tau, Y_{\mathrm{b}}, Y_{\mathrm{k}}$ and a query slot $i \leftarrow [d]$. It generates non-target domains locally and embeds its phase-recovery challenge in domain J. On an accepted target candidate, it submits the candidate quantum register to its challenger, receives Q_s^*, and forms (18) without knowing s.

At an active post-challenge slot i, it measures the query input. If the result is $00\|\nu_J\|\tau\|u$ or $01\|\nu_J\|\tau\|u$ for $u \in \{0,1\}^{\ell}$, it outputs u; otherwise it fails. Pre-challenge, inactive, and unreached slots count as failure. Every post-challenge hit yields $u = s$, so

$$p_0 \le 2^{-(\lambda+\ell)} + p_{\mathrm{rec}}(\mathcal{R}_0). \tag{19}$$

Here $p_{\mathrm{rec}}(\mathcal{R}_0)$ is the joint recovery probability averaged over J.

Apply Lemma 8 between each $\mathsf{Hy}^{(b)}$ and $\mathsf{Hy}_{\mathrm{ideal}}$. Each gap is at most $2d\sqrt{p_0}$. The distinguishing gap between $\mathsf{Hy}^{(0)}$ and $\mathsf{Hy}^{(1)}$ is bounded by the sum of the two individual gaps. The triangle inequality gives a real-versus-random gap of at most $4d\sqrt{p_0}$, hence $\mathbf{Adv}_{\mathrm{sel}}^{\text{IND-QKEM-CPA}}(\mathcal{A}) \le 2d\sqrt{p_0}$.

The ideal simulation requires no target signing key. Corollary 24, averaged over J, gives $p_{\mathrm{rec}}(\mathcal{R}_0) \le 2^{-\ell} + \mathrm{negl}(\lambda)$. Combining these inequalities proves the theorem. □

B. FTCCA Confidentiality

We next allow decapsulation queries under other serials. These queries can be simulated without the target signing key, so the CPA argument continues to apply.

Theorem 29 (FTCCA confidentiality)**.** *Under the assumptions of Theorem 28, the Signed-Fourier QKEM satisfies*

$$\mathbf{Adv}_{\Pi}^{\text{IND-QKEM-FTCCA}}(\mathcal{A}) \le \varepsilon_{\text{rt}} + 2Nd\sqrt{p_0} \le \text{negl}(\lambda),$$

where $d = \max\{1, Q_H + 2Q_D\}$. Here p_0 is the query-hit probability in the FTCCA common ideal execution, including oracle calls made by the decapsulation simulation. It satisfies $p_0 \le 2^{-(\lambda+\ell)} + p_{\text{rec}}(\mathcal{R}_0)$ for a uniform QPT recovery algorithm $\mathcal{R}_0$.

Proof. As in the CPA proof, exclude unissued certificates at cost ε_{rt} and guess $J \leftarrow [N]$. In the selective active experiment, terminate neutrally upon any canonically parsed pre-challenge query with serial ν_J, including rejected queries, or upon selection of different candidate metadata.

By the freshness condition of the selective experiment, a domain queried before the challenge cannot be the eventual challenge domain. Every eligible contribution is therefore retained with probability $1/N$. For $X \in \{\text{FTCCA}, \text{HCCA}, \text{VFCCA}\}$, this gives

$$\mathbf{Adv}_{\Pi}^{X}(\mathcal{A}) \le \varepsilon_{\text{rt}} + N\mathbf{Adv}_{\text{sel}}^{X}(\mathcal{A}). \tag{20}$$

Non-target decapsulations use locally generated keys and require at most two oracle queries each. Pre-challenge target queries cause neutral termination, and post-challenge target queries are answered by $\mathcal{F}_Q$ without oracle evaluation. The simulation therefore requires no target signing key and uses at most d oracle queries.

Counting all internal oracle calls, with inactive or unreached slots contributing zero, the query-measurement and O2H arguments in the proof of Theorem 28 give $p_0 \le 2^{-(\lambda+\ell)} + p_{\text{rec}}(\mathcal{R}_0)$ and $\mathbf{Adv}_{\text{sel}}^{\text{IND-QKEM-FTCCA}}(\mathcal{A}) \le 2d\sqrt{p_0}$. The simulated interface satisfies the conditions of Corollary 24, so $p_{\text{rec}}(\mathcal{R}_0) \le 2^{-\ell} + \text{negl}(\lambda)$. Equation (20) completes the proof. □

C. VFCCA Confidentiality

We now allow acceptance or rejection feedback under the challenge header and full decapsulation under other headers. We first bound phase recovery during this interaction, then use that bound to prove key confidentiality.

Theorem 30 (VFCCA confidentiality)**.** *Under the assumptions of Theorem 28, the Signed-Fourier QKEM satisfies*

$$\mathbf{Adv}_{\Pi}^{\text{VFCCA}}(\mathcal{A}) \le \varepsilon_{\text{rt}} + NQ_H\sqrt{2^{-(\lambda+\ell)} + p_{\text{rec,sel}}^{\text{feedback}}(\mathcal{R})} \le \text{negl}(\lambda). \tag{21}$$

Here $p_{\text{rec,sel}}^{\text{feedback}}(\mathcal{R})$ is the joint phase-recovery probability in the selective experiment below, including target matching, freshness, and challenge validation. Aborts count as failure.

Proof. Work in the selective experiment with target serial $\nu^* = \nu_J$. Independently of H_0, the receiver keys, and the pre-phase state, sample uniformly $s \in \{0,1\}^\ell$, $\tau^* \in \{0,1\}^\lambda$, and $\gamma^* \in \{0,1\}^{2\lambda}$. Let $h^* := (\nu^*, \tau^*, \gamma^*)$ and $H_s := H_0[00\|\nu^*\|\tau^*\|s \mapsto \gamma^*]$. The recovery algorithm receives the challenge quantum ciphertext and header and has the VFCCA interface under H_s, but receives no additional information about s. Any test key supplied in this recovery experiment is sampled independently.

When $Q_D = 0$, only the all-rejection case below remains. We first analyze $Q_D \ge 1$ and then apply the same oracle-comparison argument to $Q_D = 0$.

A target query is a post-challenge submission with serial ν^*. Let r denote the phase string decoded during decapsulation. Let $\mathcal{F}_{h,Q}$ discard Q and return 0 when $h = h^*$, or $\perp$ otherwise. Partition recovery success by the first accepted target query, with $t = 0$ denoting no target acceptance. All branch probabilities include query reachability.

All target queries reject. Let P_0 replace every target query by $\mathcal{F}_{h,Q}$ and run to completion. Lemma 26 upper-bounds the all-rejection contribution by the corresponding event in P_0. Replacing H_s by H_0 incurs error a_0. If r_0 is the recovery probability in this all-H_0 execution, the contribution is at most $r_0 + a_0$.

A target query accepts. For $t \in [Q_D]$, let P_t replace earlier target queries by $\mathcal{F}_{h,Q}$ and stop before processing query t. Rejection-branch domination upper-bounds the subnormalized branch reaching this first acceptance. Replace H_s by H_0 in the prefix at error a_t, then apply the same H_s terminal test and continuation to both prefix states. These terminal operations are used only in the analysis.

Let η_t, μ_t be the noncanonical and branch-mismatch masses of the all-H_0 prefix. For a submitted header $h' = (\nu^*, \tau', \gamma')$, every accepting outcome outside $h' = h^*$ and $r = s$ also passes the binding check under H_0. If $\tau' \ne \tau^*$, the binding input is unchanged. If $\tau' = \tau^*$ and $\gamma' \ne \gamma^*$, the outcome $r = s$ rejects, while all other inputs are unchanged. If $h' = h^*$ and $r \ne s$, the input is again unchanged.

By Lemma 25, canonical projection and removal of branch mismatches cost at most $2\sqrt{\eta_t} + 2\sqrt{\mu_t} + \mu_t$. The baseline prefix depends on s only through the challenge modulation. Averaging over s gives the dephased reference execution, in which a preserved canonical branch has a uniform Fourier-decoding outcome conditional on H_0 and the submitted header. Lemma 10

bounds each relevant binding fiber by λ, except with probability $\varepsilon_{\mathrm{fib}}(\lambda)$. Thus acceptance outside $h' = h^*$ and $r = s$ has probability at most

$$\beta_t := 2\sqrt{\eta_t} + 2\sqrt{\mu_t} + \mu_t + \lambda 2^{-\ell} + \varepsilon_{\mathrm{fib}}(\lambda). \tag{22}$$

On acceptance with $h' = h^*$ and $r = s$, Lemma 26 replaces the remaining subnormalized state by $U_S \otimes \omega_E$, with ω_E independent of s, at trace-norm error $\xi_t := 4\sqrt{\eta_t} + 2\sqrt{\mu_t}$.

The continuation argument in the proof of Lemma 26 extends to different-header queries returning keys. Compare H_s and H_0 one query at a time, using an all-H_0 prefix before the changed query. The entering state is independent of s. Both worlds perform the same decoding measurement. Their responses can differ only when $(\nu, \tau) = (\nu^*, \tau^*)$ and the decoded phase is s, because only that binding entry changes and the KDF table remains fixed. The average mass of this outcome is at most $2^{-\ell}$, whether the response is an acceptance bit or a key. Each decapsulation query therefore changes the final success probability by at most $2^{-\ell}$.

Each direct QRO query contributes at most $2 \cdot 2^{-\ell/2}$. In the all-H_0 continuation, the view is independent of s, so recovery succeeds with probability at most $2^{-\ell}$. Consequently, recovery from the reference state is bounded by $b_{\mathrm{ind}} := 2^{-\ell} + 2d \cdot 2^{-\ell/2} + Q_D 2^{-\ell}$. The contribution of first acceptance at t is at most $a_t + \beta_t + \xi_t + b_{\mathrm{ind}}$.

Combining the recovery bounds. For $t \in \{0, \ldots, Q_D\}$, let p_t be the average probability of querying the programmed point in the all-H_0 execution P_t, assigning zero to inactive or unreached slots. Lemma 8 gives $a_t \le 2d\sqrt{p_t}$.

Construct $\mathcal{R}_1$ by sampling $t^* \leftarrow \{0, \ldots, Q_D\}$ and $i^* \leftarrow [d]$. It runs P_0 under H_0 if $t^* = 0$, and the stopped prefix P_{t^*} otherwise. At an active post-challenge slot i^*, it measures the query input and outputs u if the result is $00\|\nu^*\|\tau^*\|u$ or $01\|\nu^*\|\tau^*\|u$ for $u \in \{0,1\}^\ell$. All other cases, including pre-challenge slots, count as failure.

Before the challenge, the hit probability is at most $2^{-(\lambda+\ell)}$. After the challenge, every hit recovers s. Writing $r_1 := p_{\mathrm{rec}}(\mathcal{R}_1)$, Jensen's inequality gives

$$\sum_{t=0}^{Q_D} a_t \le 2d \sum_{t=0}^{Q_D} \sqrt{p_t} \le 2(Q_D+1)d \sqrt{\frac{1}{Q_D+1} \sum_{t=0}^{Q_D} p_t} \le 2(Q_D+1)d\sqrt{2^{-(\lambda+\ell)} + r_1}.$$

Both baseline recovery algorithms filter target queries and simulate non-target queries locally. Corollary 24, averaged over the target domain, gives $r_0, r_1 \le 2^{-\ell} + \mathrm{negl}(\lambda)$.

Set $\bar{\eta} := Q_D^{-1} \sum_{t=1}^{Q_D} \eta_t$ and $\bar{\mu} := Q_D^{-1} \sum_{t=1}^{Q_D} \mu_t$. The stopped baseline prefixes require no target signing key or target-secret oracle. Lemma 25, applied with a uniformly sampled stopping index, bounds both averages by negligible quantities.

Summing all contributions and applying Jensen's inequality to the square-root terms gives, for $Q_D \ge 1$,

$$\begin{aligned} p^{\mathrm{feedback}}_{\mathrm{rec,sel}}(\mathcal{R}) &\le r_0 + 2(Q_D+1)d\sqrt{2^{-(\lambda+\ell)} + r_1} \\ &\quad + Q_D \left(6\sqrt{\bar{\eta}} + 4\sqrt{\bar{\mu}} + \bar{\mu} + \varepsilon_{\mathrm{fib}}(\lambda) + \frac{\lambda + Q_D + 1}{2^\ell} + \frac{2d}{2^{\ell/2}} \right) \le \mathrm{negl}(\lambda). \end{aligned} \tag{23}$$

For $Q_D = 0$, the all-rejection argument gives $p^{\mathrm{feedback}}_{\mathrm{rec,sel}}(\mathcal{R}) \le r_0 + a_0$. The same query-measurement reduction, with stopping index fixed to $t = 0$, gives $a_0 \le 2d\sqrt{2^{-(\lambda+\ell)} + r_1}$. Corollary 24 bounds $r_0, r_1 \le 2^{-\ell} + \mathrm{negl}(\lambda)$ in the all-H_0 executions, so recovery is negligible in this case as well.

From recovery to confidentiality. Keep the binding table fixed and reprogram only $x_{\mathrm{k}} := 01\|\nu^*\|\tau^*\|s$ to an independent uniform value Y. We verify that every VFCCA response channel is unchanged.

For a query $h' = (\nu', \tau', \gamma')$, if $\nu' \ne \nu^*$ or $\tau' \ne \tau^*$, all candidate KDF inputs differ from x_{k}. If $\nu' = \nu^*$ and $\tau' = \tau^*$ but $h' \ne h^*$, then $\gamma' \ne \gamma^*$. The outcome $r = s$ fails the binding check, and all other outcomes use different KDF inputs. If $h' = h^*$, the oracle returns only the acceptance bit, whose channel (14) is independent of the KDF table. Thus the responses and induced auxiliary states are identical in both worlds.

The reference execution programs only the binding point and supplies an independent uniform test key. Only direct QRO queries can expose the changed KDF value. Let p be the average probability of querying x_{k} in this execution. Lemma 8, after converting the distinguishing gap to guessing advantage, gives $\mathbf{Adv}^{\mathrm{VFCCA}}_{\mathrm{sel}}(\mathcal{A}) \le Q_H\sqrt{p}$. If $Q_H = 0$, the advantage is zero directly.

Pre-challenge hits contribute at most $2^{-(\lambda+\ell)}$. Measuring a selected post-challenge query recovers s whenever it hits, giving $p \le 2^{-(\lambda+\ell)} + p^{\mathrm{feedback}}_{\mathrm{rec,sel}}(\mathcal{R})$. The reduction supplies the independent test key locally and preserves the VFCCA interface. Applying (20) and the recovery bound above proves (21). □

D. Adaptive Fixed-Header Key Robustness

Finally, we prove that even an adversary given the challenge key cannot make the same header accept a different key. Such an acceptance would yield a collision in the binding hash.

Theorem 31 (Adaptive fixed-header key robustness)**.** *For every QPT adversary $\mathcal{A}$ making at most Q_H direct QRO queries and Q_D decapsulation queries,*

$$\mathbf{Adv}^{\mathrm{aFHKR}}_{\Pi}(\mathcal{A}) \le \min\left\{1,\, C_{\mathrm{col}}\frac{(Q_H+2Q_D+5)^3}{2^{2\lambda}}\right\} \le \mathrm{negl}(\lambda),$$

where $C_{\mathrm{col}} > 0$ is a universal constant. Setting $Q_D = 0$ gives the corresponding FHKR bound.

Proof. Construct a collision-finding algorithm $\mathcal{C}$ that generates the receiver keys and quantum public keys honestly and simulates the aFHKR experiment. During encapsulation, it records the sampled phase s^*, so $h^* = (\nu^*, \tau^*, \gamma^*)$ satisfies $\gamma^* = H_{\mathrm{bnd}}(\nu^*, \tau^*, s^*)$ and $K^* = H_{\mathrm{kdf}}(\nu^*, \tau^*, s^*)$.

Suppose an intermediate decapsulation query or the final test query under h^* accepts $K' \notin \{K^*, \perp\}$. Let r be the decoded phase. Since $r = s^*$ would give $K' = K^*$, we have $r \neq s^*$. Acceptance implies

$$H(00\|\nu^*\|\tau^*\|r) = \gamma^* = H(00\|\nu^*\|\tau^*\|s^*).$$

The two distinct inputs form a collision. Every aFHKR win therefore yields a collision in H.

The simulation uses two oracle queries for encapsulation, at most two per intermediate decapsulation, and at most two for the final test. Including direct queries, its total is at most $q := Q_H + 2Q_D + 4$. The quantum collision bound [24] for a uniform random function with 2λ-bit output gives

$$\Pr[\mathcal{C}^H \text{ finds a collision}] \le \min\left\{1,\, C_{\mathrm{col}}\frac{(q+1)^3}{2^{2\lambda}}\right\}.$$

Substitution proves the claim. □

VI. QKEM–DEM Hybrid Encryption

This section combines the QKEM with a classical DEM to encrypt arbitrary-length messages and proves CPA and HCCA confidentiality of the resulting hybrid scheme.

A. Construction

Let $\Pi = (\mathsf{SKGen}, \mathsf{PKGen}, \mathsf{Encap}, \mathsf{Decap})$ be a QKEM and DEM a classical one-time private DEM with key space $\{0,1\}^\ell$. The hybrid scheme inherits $\mathsf{SKGen}, \mathsf{PKGen}$ and uses:

- $\mathsf{Enc}^H_{\mathrm{hyb}}(vk_{\mathrm{rt}}, pk', a, M)$: Run $(CT_{\mathrm{kem}}, K) \leftarrow \mathsf{Encap}^H(vk_{\mathrm{rt}}, pk')$. On failure, return fail. Otherwise, write $CT_{\mathrm{kem}} = (h, Q)$, compute $C \leftarrow \mathsf{DEM.Enc}_K(a\|h, M)$, and return $(CT_{\mathrm{kem}}, a, C)$.
- $\mathsf{Dec}^H_{\mathrm{hyb}}(\mathsf{st}_{\mathrm{R}}, (CT_{\mathrm{kem}}, a, C))$: Write $CT_{\mathrm{kem}} = (h, Q)$ and run $K' \leftarrow \mathsf{Decap}^H(\mathsf{st}_{\mathrm{R}}, CT_{\mathrm{kem}})$. Return $\perp$ if $K' = \perp$; otherwise, return $\mathsf{DEM.Dec}_{K'}(a\|h, C)$.

Here a is associated data, $a\|h$ is injectively encoded, and DEM makes no queries to H.

a) Correctness: For honestly generated $(vk_{\mathrm{rt}}, \mathsf{st}_{\mathrm{R}}, pk)$,

$$\Pr\left[\mathsf{Dec}^H_{\mathrm{hyb}}\left(\mathsf{st}_{\mathrm{R}}, \mathsf{Enc}^H_{\mathrm{hyb}}(vk_{\mathrm{rt}}, pk, a, M)\right) = M\right] \ge 1 - \mathrm{negl}(\lambda).$$

B. CPA Confidentiality

We prove that the hybrid ciphertext hides which of two equal-length messages is encrypted, in the setting without decryption queries.

Definition 32 (Hybrid CPA experiment)**.** For a uniform QPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, define $\mathbf{Exp}^{\mathrm{CPA}}_{\Pi_{\mathrm{hyb}}}(\mathcal{A})$:

1. Generate H, the receiver setup, and $pk_1, \ldots, pk_N$ as in Definition 14.
2. Run $(pk^*, a^*, M_0, M_1, E) \leftarrow \mathcal{A}_1^H(vk_{\mathrm{rt}}, pk_1, \ldots, pk_N)$, with $|M_0| = |M_1|$.
3. Sample $\beta \leftarrow \{0,1\}$ and compute $CT^* \leftarrow \mathsf{Enc}^H_{\mathrm{hyb}}(vk_{\mathrm{rt}}, pk^*, a^*, M_\beta)$. If pk^* is invalid or encryption returns fail, output an independent uniform bit.
4. Run $\widehat{\beta} \leftarrow \mathcal{A}_2^H(CT^*, E)$ and output $\mathbf{1}[\widehat{\beta} = \beta]$.

The CPA distinguishing advantage of $\mathcal{A}$ is defined as $\mathbf{Adv}^{\mathrm{CPA}}_{\Pi_{\mathrm{hyb}}}(\mathcal{A}) := |\Pr[\mathbf{Exp}^{\mathrm{CPA}}_{\Pi_{\mathrm{hyb}}}(\mathcal{A}) = 1] - 1/2|$.

Theorem 33 (Hybrid CPA confidentiality)**.** *Assume CPA confidentiality of Π, qPRF security of F, and one-time privacy of* DEM *against QPT adversaries with adversarially chosen associated data. For every uniform QPT adversary $\mathcal{A}$, there exist uniform QPT reductions $\mathcal{B}_{\mathrm{kem}}, \mathcal{C}_{\mathrm{qprf}}, \mathcal{B}_{\mathrm{dem}}$ such that*

$$\mathbf{Adv}^{\mathrm{CPA}}_{\Pi_{\mathrm{hyb}}}(\mathcal{A}) \le 2\mathbf{Adv}^{\mathrm{CPA}}_{\Pi}(\mathcal{B}_{\mathrm{kem}}) + \mathbf{Adv}^{\mathrm{qPRF}}_{F}(\mathcal{C}_{\mathrm{qprf}}; q_{\mathrm{hyb}}) + \mathbf{Adv}^{\mathrm{OTPriv}}_{\mathsf{DEM}}(\mathcal{B}_{\mathrm{dem}}).$$

For the Signed-Fourier QKEM, $q_{\mathrm{hyb}} \le Q_H + 2$, where Q_H bounds the adversary's direct QRO queries.

Proof. Let G_0 be the real experiment. Let G_1 replace the key used for DEM encryption by an independent uniform $K_{\mathrm{rnd}} \leftarrow \{0,1\}^{\ell}$. Let G_2 additionally replace H by an independently keyed qPRF F_k. Write $g_i := \Pr[G_i = 1]$. All games and reductions preserve aborts by outputting an independent uniform bit.

To compare G_0 and G_1, $\mathcal{B}_{\mathrm{kem}}$ gives $\mathcal{A}_1$ the setup and quantum public keys from its QKEM challenger. It forwards the resulting candidate pk^* and receives $(CT^*_{\mathrm{kem}}, K_b)$, where K_b is real or random. Writing $CT^*_{\mathrm{kem}} = (h^*, Q^*)$, it samples $\beta \leftarrow \{0,1\}$, computes $C^* \leftarrow \mathsf{DEM.Enc}_{K_b}(a^* \| h^*, M_\beta)$, and gives the hybrid ciphertext to $\mathcal{A}_2$. It guesses that K_b is real exactly when $\widehat{\beta} = \beta$. Thus $|g_0 - g_1| \le 2\mathbf{Adv}^{\mathrm{CPA}}_{\Pi}(\mathcal{B}_{\mathrm{kem}})$.

For G_1 and G_2, the qPRF distinguisher generates the receiver keys and quantum public keys locally, samples K_{rnd}, and forwards every oracle query to its challenger. Equation (3) gives $|g_1 - g_2| \le \mathbf{Adv}^{\mathrm{qPRF}}_{F}(\mathcal{C}_{\mathrm{qprf}}; q_{\mathrm{hyb}})$. For Signed-Fourier, these queries comprise at most Q_H direct queries and two encapsulation queries.

In G_2, the DEM key is uniform and independent of the QKEM ciphertext and the adversary's state before DEM encryption. The reduction $\mathcal{B}_{\mathrm{dem}}$ generates the receiver keys, quantum public keys, and qPRF key locally, then runs $\mathcal{A}_1$ and encapsulates to pk^*. It submits $(a^* \| h^*, M_0, M_1)$ to its one-time DEM challenger, uses the returned ciphertext as C^*, and runs $\mathcal{A}_2$. It outputs $\widehat{\beta}$, giving $|g_2 - 1/2| \le \mathbf{Adv}^{\mathrm{OTPriv}}_{\mathsf{DEM}}(\mathcal{B}_{\mathrm{dem}})$.

Combining the three bounds with $|g_0 - 1/2| \le |g_0 - g_1| + |g_1 - g_2| + |g_2 - 1/2|$ proves the claim. □

C. HCCA Confidentiality

We now allow adaptive decryption queries. After the challenge, decryption queries carrying the challenge QKEM header $h = h^*$ are rejected; queries with $h \neq h^*$ are decrypted normally.

Definition 34 (Hybrid HCCA experiment)**.** The experiment $\mathbf{Exp}^{\mathrm{HCCA}}_{\Pi_{\mathrm{hyb}}}(\mathcal{A})$ extends Definition 32 as follows:

(1) Initialize $S := \varnothing$. Give $\mathcal{A}_1$ access to $\mathsf{Dec}^{H}_{\mathrm{hyb}}(\mathsf{st}_{\mathrm{R}}, \cdot)$. For each query $((h, Q), a, C)$, record $\mathsf{sid}_{\mathrm{ct}}(h)$ in S whenever h parses canonically, even if decryption returns $\perp$.

(2) Before challenge encryption, require $\mathsf{sid}_{\mathrm{pk}}(pk^*) \neq \perp$ and $\mathsf{sid}_{\mathrm{pk}}(pk^*) \notin S$. Otherwise, terminate and output an independent uniform bit.

(3) Write the challenge QKEM ciphertext as $CT^*_{\mathrm{kem}} = (h^*, Q^*)$. Give $\mathcal{A}_2$ a decryption oracle that, on input $((h, Q), a, C)$, discards Q and returns $\perp$ if $h = h^*$; otherwise, it returns $\mathsf{Dec}^{H}_{\mathrm{hyb}}(\mathsf{st}, ((h, Q), a, C))$.

Every decryption query consumes the submitted quantum ciphertext. The HCCA distinguishing advantage of $\mathcal{A}$ is defined as $\mathbf{Adv}^{\mathrm{HCCA}}_{\Pi_{\mathrm{hyb}}}(\mathcal{A}) := |\Pr[\mathbf{Exp}^{\mathrm{HCCA}}_{\Pi_{\mathrm{hyb}}}(\mathcal{A}) = 1] - 1/2|$.

Theorem 35 (Hybrid HCCA confidentiality)**.** *Assume HCCA confidentiality of* Π*, qPRF security of* F*, and one-time privacy of* DEM *against QPT adversaries with adversarially chosen associated data. For every uniform QPT adversary* $\mathcal{A}$*, there exist uniform QPT reductions* $\mathcal{B}_{\mathrm{kem}}, \mathcal{C}_{\mathrm{qprf}}, \mathcal{B}_{\mathrm{dem}}$ *such that*

$$\mathbf{Adv}^{\mathrm{HCCA}}_{\Pi_{\mathrm{hyb}}}(\mathcal{A}) \le 2\mathbf{Adv}^{\mathrm{HCCA}}_{\Pi}(\mathcal{B}_{\mathrm{kem}}) + \mathbf{Adv}^{\mathrm{qPRF}}_{F}(\mathcal{C}_{\mathrm{qprf}}; q_{\mathrm{hyb}}) + \mathbf{Adv}^{\mathrm{OTPriv}}_{\mathsf{DEM}}(\mathcal{B}_{\mathrm{dem}}).$$

For the Signed-Fourier QKEM, $q_{\mathrm{hyb}} \le Q_H + 2Q_D + 2$*, where* Q_H *and* Q_D *bound direct QRO and decryption queries, respectively.*

Proof. Let G_0 be the real experiment. Let G_1 replace only the key used for challenge DEM encryption by an independent uniform $K_{\mathrm{rnd}} \leftarrow \{0,1\}^{\ell}$. Decryption queries still use the keys returned by genuine QKEM decapsulation. Let G_2 additionally replace H by an independently keyed qPRF F_k. Write $g_i := \Pr[G_i = 1]$. All games and reductions preserve aborts by outputting an independent uniform bit.

To compare G_0 and G_1, $\mathcal{B}_{\mathrm{kem}}$ gives $\mathcal{A}_1$ the setup and quantum public keys from its QKEM challenger. For each pre-challenge decryption query, it forwards the QKEM ciphertext to its decapsulation oracle. If the returned key is $\perp$, it returns $\perp$; otherwise, it evaluates DEM decryption locally. Forwarding these queries preserves the serial set S, including serials from rejected queries.

For the challenge, it submits pk^* and receives $(CT^*_{\mathrm{kem}}, K_b)$, where K_b is real or random. Writing $CT^*_{\mathrm{kem}} = (h^*, Q^*)$, it samples $\beta \leftarrow \{0,1\}$, computes $C^* \leftarrow \mathsf{DEM.Enc}_{K_b}(a^* \| h^*, M_\beta)$, and gives the hybrid ciphertext to $\mathcal{A}_2$.

After the challenge, it discards Q and returns $\perp$ on queries with $h = h^*$. For $h \neq h^*$, it obtains the key from its decapsulation oracle and evaluates DEM decryption locally. This preserves both the classical responses and the induced state of the adversary's retained memory. It guesses that K_b is real exactly when $\widehat{\beta} = \beta$, giving $|g_0 - g_1| \le 2\mathbf{Adv}^{\mathrm{HCCA}}_{\Pi}(\mathcal{B}_{\mathrm{kem}})$.

For G_1 and G_2, the qPRF distinguisher generates the receiver keys and quantum public keys locally, samples K_{rnd}, and forwards all oracle queries to its challenger. It simulates hybrid decryption using its receiver keys, while discarding and rejecting same-header queries. Equation (3) gives $|g_1 - g_2| \le \mathbf{Adv}^{\mathrm{qPRF}}_{F}(\mathcal{C}_{\mathrm{qprf}}; q_{\mathrm{hyb}})$. For Signed-Fourier, encapsulation uses at most two oracle calls and each decryption uses at most two, so $q_{\mathrm{hyb}} \le Q_H + 2Q_D + 2$.

In G_2, the challenge DEM key is uniform and independent of the QKEM ciphertext and the adversary's state before DEM encryption. The reduction $\mathcal{B}_{\mathrm{dem}}$ generates the receiver keys, quantum public keys, and qPRF key locally. It runs $\mathcal{A}_1$, forms $CT^*_{\mathrm{kem}} = (h^*, Q^*)$, and submits $(a^* \| h^*, M_0, M_1)$ to its one-time DEM challenger. It uses the returned ciphertext as C^*.

The reduction answers permitted decryption queries using its local QKEM keys and then DEM decryption. Same-header queries are discarded and rejected. Thus every query is answered without knowing the DEM challenge key. The reduction outputs $\widehat{\beta}$, yielding $|g_2 - 1/2| \le \mathbf{Adv}^{\mathrm{OTPriv}}_{\mathsf{DEM}}(\mathcal{B}_{\mathrm{dem}})$.

Combining $|g_0 - 1/2| \le |g_0 - g_1| + |g_1 - g_2| + |g_2 - 1/2|$ with the three bounds above proves the claim. □

VII. Publicly Verifiable Certified Deletion

We extend the QKEM with certified deletion so that a party holding the quantum ciphertext can provide publicly verifiable evidence that it has relinquished the ability to recover the session key. After certificate acceptance, the key remains computationally hidden even if the receiver's secret state is later exposed.

A. Syntax and Correctness

Definition 36 (Certified Deletion Syntax)**.** For an honest ciphertext $CT = ((\nu, \tau, \gamma), |\Phi_s\rangle)$ and public verification record $v_{\mathrm{del}} = (\nu, vk_\nu, \pi_\nu)$, define:

- $\mathsf{Del}(CT)$: Measure A, B in the computational basis and output $c = (z, \widetilde{\sigma}_1, \ldots, \widetilde{\sigma}_\ell)$, where $z \in \{0,1\}^\ell$ and $\widetilde{\sigma}_i \in \{0,1\}^{L_{\mathrm{sig}}(\lambda)}$.
- $\mathsf{VDel}(v_{\mathrm{del}}, c)$: Return 1 if $\mathsf{Vfy}_{\mathrm{rt}, vk_{\mathrm{rt}}}(\mathtt{token} \| \nu \| vk_\nu, \pi_\nu) = 1$ and

$$\mathsf{Vfy}_{\mathrm{ph}, vk_\nu}(\mathtt{phase} \| \nu \| i \| z_i, \widetilde{\sigma}_i) = 1 \quad \text{for all } i \in [\ell]. \tag{24}$$

Return 0 otherwise. Here vk_{rt} is a global public root verification key.

a) Correctness: For honestly generated ciphertexts,

$$\Pr[\mathsf{VDel}(v_{\mathrm{del}}, \mathsf{Del}(CT)) = 1] \ge 1 - \mathrm{negl}(\lambda). \tag{25}$$

B. Security under Receiver Key Exposure

We prove that an accepted deletion certificate ensures session-key confidentiality even after the receiver's secret state is exposed.

Definition 37 (Certified Deletion Experiment under Receiver Key Exposure)**.** For a uniform QPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, define $\mathbf{Exp}^{\mathrm{del}}_{\Pi}(\mathcal{A})$:

1. Sample H and $b \leftarrow \{0,1\}$. Generate $(vk_{\mathrm{rt}}, \mathsf{st}^{(0)}_{\mathrm{R}}) \leftarrow \mathsf{SKGen}(1^\lambda)$ and $(pk_\nu, \mathsf{st}_{\mathrm{R}}) \leftarrow \mathsf{PKGen}(\mathsf{st}^{(0)}_{\mathrm{R}})$. Compute $(CT^*, K_0) \leftarrow \mathsf{Encap}^H(vk_{\mathrm{rt}}, pk_\nu)$, record v_{del}, and sample independent $K_1 \leftarrow \{0,1\}^\ell$. If encapsulation returns fail, terminate and output an independent uniform bit.
2. Run $(c, E) \leftarrow \mathcal{A}^H_1(CT^*, v_{\mathrm{del}})$. If $\mathsf{VDel}(v_{\mathrm{del}}, c) = 0$, terminate and output an independent uniform bit.
3. Run $\widehat{b} \leftarrow \mathcal{A}^H_2(K_b, \mathsf{st}_{\mathrm{R}}, E)$, exposing the full receiver state, including sk_ν. Output $\mathbf{1}[\widehat{b} = b]$.

The certified deletion advantage is $\mathbf{Adv}^{\mathrm{del}}_{\Pi}(\mathcal{A}) := |\Pr[\mathbf{Exp}^{\mathrm{del}}_{\Pi}(\mathcal{A}) = 1] - 1/2|$. The public key vk_{rt} is available to both stages, and E may include the certificate and public transcript.

For the following lemma, use a random oracle H_0 throughout and sample the challenge tag γ uniformly and independently of s. We first bound the accepted state's dependence on s, then relate this ideal experiment to the real experiment.

Lemma 38. *Consider the ideal experiment above, with uniform s independent of H_0, the receiver keys, the pre-phase state, and the public parameters, including γ. Use D, U_S, and $|\chi_x\rangle$ from Section IV-F, with S recording s.*

Let $z(c)$ be the branch string in certificate c. Represent certificate generation by operators A_c, retaining a purifying environment, and define

$$\delta := \mathbb{E}\left[D^{-1} \sum_x \sum_{\substack{c:\, \mathsf{VDel}(v_{\mathrm{del}}, c) = 1 \\ z(c) \ne x}} \| A_c \, |\chi_x\rangle \, \|^2 \right]. \tag{26}$$

The expectation includes setup, H_0, and the remaining classical parameters. There exist uniform QPT reductions $\mathcal{C}_{\mathrm{del}}, \mathcal{F}_{\mathrm{del}}$ such that $\delta \le \mathbf{Adv}^{\mathrm{qPRF}}_F(\mathcal{C}_{\mathrm{del}}) + \mathbf{Adv}^{\mathrm{EUF}}_{\Sigma_{\mathrm{ph}}}(\mathcal{F}_{\mathrm{del}})$, where $\mathcal{F}_{\mathrm{del}}$ makes ℓ classical signing queries.

Let $\rho^{\mathrm{acc}}_{SE, \mathsf{st}_{\mathrm{R}}}$ be the subnormalized state after certificate acceptance, including the subsequently exposed receiver state. There exists an s-independent subnormalized state $\omega_{E, \mathsf{st}_{\mathrm{R}}}$ with trace at most 1 such that

$$\left\| \rho^{\mathrm{acc}}_{SE, \mathsf{st}_{\mathrm{R}}} - U_S \otimes \omega_{E, \mathsf{st}_{\mathrm{R}}} \right\|_1 \le 2\sqrt{\delta}. \tag{27}$$

Proof. **Step 1: Bounding branch mismatch.** Replace H_0 by an independently keyed qPRF. The forger samples $x \leftarrow \{0,1\}^\ell$, requests the signature on $m_{(\nu,i),x_i}$ for every $i \in [\ell]$, and prepares $|\chi_x\rangle$. It generates the root authentication record locally and runs $\mathcal{A}_1$.

If the certificate verifies and $z(c) \neq x$, some coordinate j satisfies $z(c)_j = 1 - x_j$. Its certificate signature is a valid forgery on the unqueried message $m_{(\nu,j),1-x_j}$. The forger stops before receiver-key exposure and does not need sk_ν. The qPRF comparison experiment generates the signing keys locally and can check this event. Equation (3) gives the stated bound on δ.

Step 2: Bounding phase dependence. Fix setup, H_0, and the remaining classical parameters; write δ_{fix} for their mismatch mass. The operators A_c are independent of s and satisfy $\sum_c A_c^\dagger A_c \le I$. For each certificate, write the actual output vector as $v_{s,c} := D^{-1/2} \sum_x (-1)^{s\cdot x} A_c \,|\chi_x\rangle$ and retain only its matching-branch term $w_{s,c} := D^{-1/2}(-1)^{s\cdot z(c)} A_c \,|\chi_{z(c)}\rangle$.

The density operator $|w_{s,c}\rangle\langle w_{s,c}|$ is independent of s, since its phase is global. Fourier orthogonality gives

$$\mathbb{E}_s \sum_{c:\,\mathsf{VDel}(v_{\mathrm{del}},c)=1} \|v_{s,c} - w_{s,c}\|^2 = D^{-1} \sum_{c:\,\mathsf{VDel}(v_{\mathrm{del}},c)=1} \sum_{x\neq z(c)} \|A_c \,|\chi_x\rangle\,\|^2 = \delta_{\mathrm{fix}}.$$

Measurement completeness gives $\sum_{c:\,\mathsf{VDel}(v_{\mathrm{del}},c)=1} \|v_{s,c}\|^2 \le 1$ and $\sum_{c:\,\mathsf{VDel}(v_{\mathrm{del}},c)=1} \|w_{s,c}\|^2 \le D^{-1} \sum_x \sum_c \|A_c \,|\chi_x\rangle\,\|^2 \le 1$. Using $\| \,|v\rangle\langle v| - |w\rangle\langle w|\,\|_1 \le (\|v\| + \|w\|)\|v - w\|$ and Cauchy–Schwarz yields

$$\mathbb{E}_s \sum_{c:\,\mathsf{VDel}(v_{\mathrm{del}},c)=1} \| \,|v_{s,c}\rangle\langle v_{s,c}| - |w_{s,c}\rangle\langle w_{s,c}|\,\|_1 \le 2\sqrt{\delta_{\mathrm{fix}}}.$$

Retain the certificate and receiver secrets, trace out the purifying environment, and average over the fixed parameters. The reference state remains independent of s, and Jensen's inequality gives $\mathbb{E}[2\sqrt{\delta_{\mathrm{fix}}}] \le 2\sqrt{\delta}$, proving (27). The same argument bounds the average distance conditional on H_0 and the classical parameters, so the bound also applies to subsequent queries to H_0. □

Theorem 39 (Certified Deletion under Receiver Key Exposure)**.** *Under the assumptions of Theorem 28, the Signed-Fourier QKEM achieves publicly verifiable computational certified deletion for honestly generated target ciphertexts. For every uniform QPT adversary making at most q direct QRO queries across both stages,*

$$\mathbf{Adv}^{\mathrm{del}}_{\Pi}(\mathcal{A}) \le 2q\sqrt{2D^{-1} + \varepsilon_{\mathrm{pre}} + 2\sqrt{\delta}} \le \mathrm{negl}(\lambda), \tag{28}$$

where $D = 2^\ell$, δ is the mismatch mass in (26), and $\varepsilon_{\mathrm{pre}}$ is the negligible excess over D^{-1} in the recovery bound of Corollary 24.

Proof. Fix the honest target domain ν and set $x_{\mathrm{b}} := 00\|\nu\|\tau\|s$ and $x_{\mathrm{k}} := 01\|\nu\|\tau\|s$. Use the programmed experiments $H^{(0)}, H^{(1)}$ from Definition 27, with this single target domain. They represent the real-key and random-key experiments and expose the same receiver state after certificate acceptance. Their common ideal experiment uses H_0 throughout, with an independently sampled binding tag and test key.

Assume $q \ge 1$. Let p be the probability that a uniformly selected query slot in the ideal experiment hits $\{x_{\mathrm{b}}, x_{\mathrm{k}}\}$, assigning zero to inactive or unreached slots. Lemma 8 bounds each programmed experiment's gap from the ideal experiment by $2q\sqrt{p}$. The real-versus-random gap is therefore at most $4q\sqrt{p}$, giving $\mathbf{Adv}^{\mathrm{del}}_{\Pi}(\mathcal{A}) \le 2q\sqrt{p}$.

Before certificate submission. A reduction measures the selected query input and, if it parses as $00\|\nu\|\tau\|u$ or $01\|\nu\|\tau\|u$, outputs u. Every hit recovers s. The reduction stops before key exposure and requires no target signing key. Corollary 24 bounds this contribution to p by $D^{-1} + \varepsilon_{\mathrm{pre}}$.

After certificate acceptance. Lemma 38 replaces the subnormalized accepted state, including the exposed receiver secrets, by an s-independent reference state at average trace-norm error at most $2\sqrt{\delta}$. The subsequent execution uses H_0, so the query input in the reference execution remains independent of uniform s. Its probability of hitting either hidden point is at most D^{-1}. Contractivity bounds the change in this event probability by $2\sqrt{\delta}$, giving a contribution of at most $D^{-1} + 2\sqrt{\delta}$.

Adding the two contributions yields $p \le 2D^{-1} + \varepsilon_{\mathrm{pre}} + 2\sqrt{\delta}$. Substituting this into $\mathbf{Adv}^{\mathrm{del}}_{\Pi}(\mathcal{A}) \le 2q\sqrt{p}$ proves (28). For $q = 0$, the compared experiments have no differing oracle calls, so the advantage is zero. □

VIII. Conclusion

We construct the Signed-Fourier QKEM from quantum-secure one-way functions in the QROM, achieving VFCCA confidentiality and adaptive fixed-header key robustness under quantum public-key tampering. The construction uses signed Fourier states both to protect the session key during adaptive decapsulation and to support publicly verifiable certified deletion. For honestly generated ciphertexts, an accepted deletion certificate ensures computational key confidentiality even after receiver-secret exposure. A classical one-time private DEM extends the scheme to polynomial-length messages with CPA and HCCA confidentiality. Removing the random oracle and supporting reusable quantum public keys remain directions for future work.

References

[1] R. Impagliazzo and S. Rudich, "Limits on the provable consequences of one-way permutations," in *Proceedings of the 21st Annual ACM Symposium on Theory of Computing*, 1989. [Online]. Available: https://doi.org/10.1145/73007.73012

[2] A. Kawachi, T. Koshiba, H. Nishimura, and T. Yamakami, "Computational indistinguishability between quantum states and its cryptographic application," in *Advances in Cryptology – EUROCRYPT 2005*, 2005. [Online]. Available: https://doi.org/10.1007/11426639_16

[3] K. Barooti, A. B. Grilo, L. Huguenin-Dumittan, G. Malavolta, O. Sattath, Q.-H. Vu, and M. Walter, "Public-key encryption with quantum keys," in *Theory of Cryptography – TCC 2023*, 2023. [Online]. Available: https://doi.org/10.1007/978-3-031-48624-1_8

[4] A. Coladangelo, "Quantum trapdoor functions from classical one-way functions," Cryptology ePrint Archive, Paper 2023/282, 2023. [Online]. Available: https://eprint.iacr.org/2023/282

[5] F. Kitagawa, T. Morimae, R. Nishimaki, and T. Yamakawa, "Quantum public-key encryption with tamper-resilient public keys from one-way functions," in *Advances in Cryptology – CRYPTO 2024*, 2024. [Online]. Available: https://doi.org/10.1007/978-3-031-68394-7_4

[6] G. Malavolta and M. Walter, "Robust quantum public-key encryption with applications to quantum key distribution," in *Advances in Cryptology – CRYPTO 2024*, 2024. [Online]. Available: https://doi.org/10.1007/978-3-031-68394-7_5

[7] A. Broadbent and R. Islam, "Quantum encryption with certified deletion," Cryptology ePrint Archive, Paper 2020/1423, 2020. [Online]. Available: https://eprint.iacr.org/2020/1423

[8] J. Bartusek and D. Khurana, "Cryptography with certified deletion," in *Advances in Cryptology – CRYPTO 2023*, 2023. [Online]. Available: https://doi.org/10.1007/978-3-031-38554-4_7

[9] F. Kitagawa, R. Nishimaki, and T. Yamakawa, "Publicly verifiable deletion from minimal assumptions," in *Theory of Cryptography – TCC 2023*, 2023. [Online]. Available: https://doi.org/10.1007/978-3-031-48624-1_9

[10] J. Bartusek, D. Khurana, G. Malavolta, A. Poremba, and M. Walter, "Weakening assumptions for publicly-verifiable deletion," in *Theory of Cryptography – TCC 2023*, 2023. [Online]. Available: https://doi.org/10.1007/978-3-031-48624-1_7

[11] G. Alagic, T. Gagliardoni, and C. Majenz, "Unforgeable quantum encryption," in *Advances in Cryptology – EUROCRYPT 2018*, 2018. [Online]. Available: https://doi.org/10.1007/978-3-319-78372-7_16

[12] K. Barooti, G. Malavolta, and M. Walter, "A simple construction of quantum public-key encryption from quantum-secure one-way functions," Cryptology ePrint Archive, Paper 2023/306, 2023. [Online]. Available: https://eprint.iacr.org/2023/306

[13] M. Hhan, T. Morimae, and T. Yamakawa, "From the hardness of detecting superpositions to cryptography: Quantum public key encryption and commitments," in *Advances in Cryptology – EUROCRYPT 2023*, 2023. [Online]. Available: https://doi.org/10.1007/978-3-031-30545-0_22

[14] M. Barhoush and L. Salvail, "How to sign quantum messages," *Quantum*, 2026. [Online]. Available: https://doi.org/10.22331/q-2026-01-22-1980

[15] N. Rodrigues, W. O. Krawec, B. Lackey, D. Mukhopadhyay, and B. Wang, "Quantum public key encryption for NISQ devices," arXiv:2509.17758, 2025. [Online]. Available: https://arxiv.org/abs/2509.17758

[16] K. Sudo, K. Fujii, K. Hara, K. Ito, M. Tezuka, and Y. Yoshida, "Practical quantum public key encryption from one-way functions," in *Proceedings of the 12th ACM ASIA Public-Key Cryptography Workshop*, 2025. [Online]. Available: https://doi.org/10.1145/3709015.3728674

[17] G. Alagic, A. Broadbent, B. Fefferman, T. Gagliardoni, C. Schaffner, and M. St. Jules, "Computational security of quantum encryption," in *Information Theoretic Security – ICITS 2016*, 2016. [Online]. Available: https://doi.org/10.1007/978-3-319-49175-2_3

[18] H. Barnum, C. Crépeau, D. Gottesman, A. Smith, and A. Tapp, "Authentication of quantum messages," in *43rd Annual IEEE Symposium on Foundations of Computer Science*, 2002. [Online]. Available: https://doi.org/10.1109/SFCS.2002.1181969

[19] S. Aaronson, Y. Atia, and L. Susskind, "On the hardness of detecting macroscopic superpositions," arXiv:2009.07450, 2020. [Online]. Available: https://arxiv.org/abs/2009.07450

[20] L. Li, Q. Li, X. Li, and Q. Liu, "Cryptomania v.s. minicrypt in a quantum world," in *Advances in Cryptology – CRYPTO 2026*, 2026. [Online]. Available: https://doi.org/10.1007/978-3-032-35409-9_1

[21] M. Zhandry, "How to construct quantum random functions," in *53rd Annual IEEE Symposium on Foundations of Computer Science*, 2012. [Online]. Available: https://doi.org/10.1109/FOCS.2012.37

[22] D. Boneh, Ö. Dagdelen, M. Fischlin, A. Lehmann, C. Schaffner, and M. Zhandry, "Random oracles in a quantum world," in *Advances in Cryptology – ASIACRYPT 2011*, 2011. [Online]. Available: https://doi.org/10.1007/978-3-642-25385-0_3

[23] A. Ambainis, M. Hamburg, and D. Unruh, "Quantum security proofs using semi-classical oracles," in *Advances in Cryptology – CRYPTO 2019*, 2019. [Online]. Available: https://doi.org/10.1007/978-3-030-26951-7_10

[24] M. Zhandry, "A note on the quantum collision and set equality problems," *Quantum Information and Computation*, 2015. [Online]. Available: https://doi.org/10.26421/QIC15.7-8-2

APPENDIX A
REDUCTION INTERFACES AND QUERY ACCOUNTING

We formalize the oracle access, algorithmic complexity, and query accounting for each reduction invoked in our security theorems. Whenever an auxiliary random oracle is replaced by an independently keyed qPRF instance F_k, its unitary U_{F_k} and inverse $U_{F_k}^\dagger = U_{F_k}$ are evaluated locally as explicit polynomial-size circuits, incurring zero external oracle queries. External phase challenges are issued at most once per execution.

- **Single-Coordinate Extraction (proof of Lemma 20):**
  - *Target Secret Access:* Receives the signed states indexed by $\mathcal{L}_j$; possesses no knowledge of sk_j and issues zero signing queries.
  - *Locally Executed Circuits:* Public coordinate validation $\Pi_{\xi^*}^{\mathrm{val}}$, unitary dilation V, reflection $Z_{D_{\mathrm{out}}}$, and inverse $V^\dagger$ (using local F_k gates and retaining dilation environments).
  - *Query Accounting:* Let q_V bound the number of local applications of U_{F_k} in one execution of V, including those used to simulate auxiliary interfaces. Implementing $O = V^\dagger(Z_{D_{\mathrm{out}}} \otimes I)V$ requires at most $2q_V$ such applications, since $U_{F_k}^\dagger = U_{F_k}$. This count covers only the implementation of O. The extraction procedure, given the issued input states, makes no external signing or random-oracle queries.
- **Concrete Phase-Recovery Bounds (Corollary 24):**
  - *Target Secret Access:* Phase challenger applies Z^s once. Reduction $\mathcal{C}$ generates target keys locally. Forgers $\mathcal{B}_0, \mathcal{B}_1$ receive the target verification key and a classical signing interface, and generate the required signed-state inputs using the listed signing queries; neither knows sk_J.
  - *Locally Executed Circuits:* Algorithm $\mathcal{C}$ executes the final recovery test; $\mathcal{B}_i$ execute public tests and local F_k evaluations with explicit inverses.
  - *Query Accounting:* $\mathcal{C}$ makes at most q external oracle queries, where q is the total query bound in Corollary 24, including auxiliary-interface simulation and the terminal success check; $\mathcal{B}_1$ makes at most $2\ell - 1$ classical signing queries.
- **O2H Query-Reduction to Phase Recovery (proof of Theorem 28):**

- *Target Secret Access:* Receives a single target token; challenger applies Z^s once. No access to sk_J or target-domain decapsulation.
- *Locally Executed Circuits:* Runs stages $\mathcal{A}_1, \mathcal{A}_2$; simulates non-target decapsulation via locally known keys; filters target-serial queries.
- *Query Accounting:* Evaluates at most $Q_H + 2Q_D$ active oracle gates (formally padded to $d = \max\{1, Q_H + 2Q_D\}$); issues at most 1 external phase challenge.

- **First-Acceptance Recovery (bound (23)):** The stopped-prefix reductions internally sample the target-query position and oracle slot, discard preceding target payloads, and simulate non-target queries using local keys. They use at most d oracle slots and one external phase challenge. The terminal acceptance test and continuation under H_s are used only in the analysis. The canonicality and branch-preservation forgers use at most 2ℓ and $2\ell - 1$ classical signing queries, respectively.
- **Computational Deletion (Theorem 39):** The mismatch forger requests one branch signature per coordinate, totaling ℓ classical signing queries. It executes only the pre-certificate stage with a local qPRF. The pre-certificate recovery reduction stops before secret exposure. O2H compares the complete executions with receiver-secret disclosure as a common channel.
- **Adaptive Collision Finder (Theorem 31):**
  - *Target Secret Access:* Generates all receiver secrets locally; issues zero signing or phase-challenger queries.
  - *Locally Executed Circuits:* Executes full $\mathcal{A}_1, \mathcal{A}_2$ environments with genuine encapsulation and decapsulation algorithms.
  - *Query Accounting:* At most 2 QRO queries for challenge encapsulation (one each to H_{bnd} and H_{kdf} if validation succeeds), at most 2 queries per intermediate decapsulation query, and at most 2 queries for final replacement testing. Including at most Q_H direct queries by the adversary, the total is at most $Q_H + 2Q_D + 4$ QRO queries.

a) Prefix Budget Consistency: In all reductions, non-target decapsulation queries are simulated using locally sampled receiver secrets, consuming at most two QRO queries per invocation. In stopped prefixes, target queries consume the payload and return the prescribed public failure response without invoking the QRO. Consequently, $d = \max\{1, Q_H + 2Q_D\}$ bounds the simulated CPA, FTCCA, and all-failure stopped prefixes. The final acceptance effects are not executed by the no-secret recovery reductions. Challenge hash evaluations are supplied by the coupling in these prefixes; the separate collision and hybrid-encryption reductions count their actual challenge evaluations explicitly.

APPENDIX B
FOUNDATIONS AND PRIMITIVE CLOSURES FROM QOWFS

This appendix supports Proposition 7. We derive strong signatures from qOWFs, derandomize signing, and construct a one-time private DEM. All primitive keys are independently sampled, and primitive evaluations make no random-oracle queries.

Lemma 40 (Straight-Line Simulation)**.** *Consider a QPT reduction that runs a QPT adversary without rewinding or cloning and exchanges only classical messages with it. If the simulation reproduces the target experiment's initial joint state and implements the same stateful classical interfaces, then the final joint classical–quantum state of the transcript and the adversary's retained workspace is identical to that in the target experiment.*

Proof. The initial joint states coincide by assumption. Track the transcript, the adversary's workspace, and the private interface state. At each round, apply the adversary's prescribed operation and the same classical message processing and interface update as in the target experiment. These operations implement identical channels in both executions, so induction gives identical joint states after every round. Tracing out the private interface state proves the claim. □

Lemma 41 (Strongly Unforgeable Signatures from qOWFs)**.** *Assuming quantum-secure one-way functions, there exist sEUF-CMA signature schemes secure against QPT adversaries making classical signing queries.*

Proof. This is the strong-signature existence result stated in [5, Theorem 3.3], with security against QPT adversaries using classical signing queries. It supplies the scheme Σ_{S} used below. Strong unforgeability also implies ordinary EUF-CMA security, so an independent instance supplies Σ_{rt}. Restricting the signing oracle to one query gives the one-time variant when needed. □

Lemma 42 (Canonical Strong Signatures via qPRFs)**.** *Let Σ_{S} be sEUF-CMA secure against QPT adversaries with adaptive classical access to the signing oracle, with signing randomness in $\{0,1\}^{r(\lambda)}$. Let F^{can} be an independent qPRF with $r(\lambda)$-bit output. Add an independently sampled key k_{can} to the signing key, leave verification unchanged, and define*

$$\mathsf{Sign}^{\mathrm{can}}_{(sk,k_{\mathrm{can}})}(m) := \mathsf{Sign}^{\Sigma_{\mathrm{S}}}_{sk}\big(m; F^{\mathrm{can}}_{k_{\mathrm{can}}}(\mathsf{can}\|m)\big).$$

The resulting scheme Σ_{can} is deterministic. For every QPT adversary $\mathcal{A}$ making at most q_{sig} classical signing queries, including repetitions, there exist QPT reductions $\mathcal{B}, \mathcal{C}$ such that

$$\mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma_{\mathrm{can}}}(\mathcal{A}) \le \mathbf{Adv}^{\mathrm{sEUF}}_{\Sigma_{\mathrm{S}}}(\mathcal{B}) + \mathbf{Adv}^{\mathrm{qPRF}}_{F^{\mathrm{can}}}(\mathcal{C}; q_{\mathrm{sig}}).$$

The scheme preserves perfect correctness when Σ_{S} is perfectly correct.

Proof. Replace $F^{\mathrm{can}}_{k_{\mathrm{can}}}$ by a random function. The distinguisher $\mathcal{C}$ generates the base signing key pair locally, uses its challenge-oracle response at $\mathtt{can}\|m$ as the signing randomness for each new message m, and caches signatures for repeated queries. It outputs $\mathcal{A}$'s strong-forgery success bit. Thus, it makes at most q_{sig} oracle queries, and the replacement loss is at most $\mathbf{Adv}^{\mathrm{qPRF}}_{F^{\mathrm{can}}}(\mathcal{C}; q_{\mathrm{sig}})$.

In the random-function experiment, distinct messages receive independent uniform signing tapes. The reduction $\mathcal{B}$ exactly simulates this interface by querying its Σ_{S} signing oracle once per distinct message and caching the response, as justified by Lemma 40. The sets of message–signature pairs obtained by $\mathcal{B}$ and returned to $\mathcal{A}$ coincide. Hence any strong forgery by $\mathcal{A}$ is also a strong forgery against Σ_{S}. Adding the replacement loss proves the bound.

Perfect correctness is preserved because every qPRF output is a valid signing tape for Σ_{S}. □

a) One-Time Private DEM from a qPRF: Quantum-secure one-way functions imply qPRFs [5, Theorem 3.5]. Let $G = \{G_k\}_{k\in\{0,1\}^\lambda}$ be an independent qPRF with λ-bit output. For a uniform session key $K \in \{0,1\}^\ell$ with $\ell \geq \lambda$, use $k_{\mathrm{dem}} := \mathrm{first}_\lambda(K)$. Fix polynomial bounds on message and associated-data lengths, and use injective encodings into the input domain of G.

For a message M, let $t := \lceil |M|/\lambda \rceil$ and zero-pad M to obtain λ-bit blocks $M_1, \ldots, M_t$. For associated data a, compute

$$C_j := M_j \oplus G_{k_{\mathrm{dem}}}(\mathtt{dem}\|a\|\|M\|\|j), \qquad j \in [t],$$

and output $(|M|, C_1, \ldots, C_t)$. Decryption recomputes the pads, XORs them with the ciphertext blocks, and truncates the result to $|M|$ bits, giving perfect correctness.

For any OT-Priv adversary $\mathcal{A}$, a QPT distinguisher $\mathcal{B}$ simulates the experiment using its challenge oracle in place of $G_{k_{\mathrm{dem}}}$ and outputs $\mathcal{A}$'s success bit. It makes one oracle query per block, hence at most $t_{\max}$ queries, where $t_{\max}$ is a fixed polynomial upper bound on the number of challenge-message blocks.

In the random-function experiment, conditioned on the adversary's classical choices, distinct pad inputs yield independent uniform pads, independent of its retained quantum memory. For equal-length challenge messages with the same associated data, the joint ciphertext–memory state is identical for both challenge bits. Thus the success probability is exactly $1/2$, giving

$$\mathbf{Adv}^{\mathrm{OTPriv}}_{\mathsf{DEM}}(\mathcal{A}) \leq \mathbf{Adv}^{\mathrm{qPRF}}_{G}(\mathcal{B}; t_{\max}).$$

Together, Lemmas 41 and 42, qPRF existence, and this DEM construction establish Proposition 7.

Appendix C
Proofs for Security-Notion Separations

The following counterexamples separate the notions defined in Section III-D.

Proposition 43 (Separation Counterexamples)**.** *The following implications do not hold:*

(1) HCCA $\not\Longrightarrow$ FHKR.

(2) $\{\mathrm{FTCCA}, \mathrm{aFHKR}\} \not\Longrightarrow$ HCCA.

(3) FHKR $\not\Longrightarrow$ CPA.

Proof. **(1) HCCA does not imply FHKR.** Start with a correct HCCA-secure QKEM Π. Append a flag qubit initialized to $|0\rangle$ to each quantum ciphertext. Decapsulation measures the flag. On outcome 0, it runs $\Pi.\mathsf{Decap}$ on the remaining quantum ciphertext. On outcome 1, it discards that register and returns 0^ℓ. Honest correctness is unchanged.

An HCCA reduction attaches $|0\rangle$ to the challenge. For queries, it measures the flag locally. Flag-zero queries are forwarded to the base decapsulation oracle; flag-one queries have the public response 0^ℓ. All same-header post-challenge queries are discarded and rejected, as required by HCCA.

For a pre-challenge flag-one query, the reduction discards the submitted quantum register and sends an independently prepared dummy register under the same header to its base oracle. It ignores the reply and returns 0^ℓ. The dummy register has the prescribed base register space and is efficiently preparable. This records the queried serial without adding more than one base query per simulated query. Discarding the submitted register also preserves the correct reduced state of the adversary's memory. The simulation is therefore exact.

For FHKR, select an honest challenge quantum public key and flip the challenge ciphertext's flag to $|1\rangle$. The replacement retains the challenge header and decapsulates to 0^ℓ. CPA security of the base scheme, implied by HCCA, gives $\Pr[K^* = 0^\ell] \leq 2^{-\ell} + \mathrm{negl}(\lambda)$: otherwise testing whether the supplied key equals 0^ℓ would distinguish real and random keys. The replacement therefore wins FHKR with probability at least $1 - 2^{-\ell} - \mathrm{negl}(\lambda)$.

(2) FTCCA and aFHKR do not imply HCCA. Start with a correct QKEM Π satisfying FTCCA and aFHKR. Replace every honest header h_{raw} by $(\mathtt{norm}, h_{\mathrm{raw}})$. Also accept headers $(\mathtt{relay}, h_{\mathrm{raw}})$. Both modes retain the base serial and invoke the same base decapsulation on (h_{raw}, Q).

An FTCCA reduction removes the mode field and forwards queries to the base oracle. Both modes under the challenge serial are filtered, and pre-challenge serial recording is unchanged. The simulation preserves FTCCA security.

For aFHKR, the honest challenge header has normal mode. The reduction forwards all intermediate queries to base decapsulation after removing the mode field. Any accepted wrong key under the fixed normal challenge header is also an accepted wrong key under the fixed base challenge header. The final test is handled in the same way. Thus aFHKR is preserved.

In HCCA, select an honest challenge quantum public key. Given the challenge $((\texttt{norm}, h^*_{\mathrm{raw}}), Q^*)$, submit $((\texttt{relay}, h^*_{\mathrm{raw}}), Q^*)$. The header differs from the challenge header, so the query is permitted and returns the true challenge key whenever honest decapsulation succeeds. Comparing it with the supplied test key gives advantage at least $(1 - 2^{-\ell})/2 - \mathrm{negl}(\lambda)$. Hence the modified scheme is not HCCA-secure.

(3) FHKR does not imply CPA. Construct a QKEM that samples a uniform session key K and includes it in the classical ciphertext header. Its quantum ciphertext can be a fixed state. Decapsulation discards the quantum register and returns the key in the header.

The scheme is correct and has perfect FHKR: a fixed challenge header fixes the decapsulation key, regardless of the submitted quantum state. In the CPA experiment, comparing the exposed header key with the supplied test key distinguishes real and random keys with advantage $(1 - 2^{-\ell})/2$. □

ACKNOWLEDGMENTS

I thank Tianqiao Hu for her help and companionship throughout my long and often solitary research journey. We are the family we have chosen for each other. Her warmth has carried me this far, and I look forward to all that lies ahead for us, together. I use AI to help organize the writing framework of research content and look up references. Of course, all parts that use AI are manually verified.