Can You Sign A Quantum State?

Gorjan Alagic[1,2], Tommaso Gagliardoni, and Christian Majenz[3]

[1] QuICS, University of Maryland, College Park, MD, USA
[2] National Institute of Standards and Technology, Gaithersburg, MD, USA
[3] QuSoft and Institute for Logic, Language and Computation, University of Amsterdam, Amsterdam, Netherlands
galagic@umd.edu; paper.signcrypt2018@gagliardoni.net; c.majenz@uva.nl

Abstract. Cryptography with quantum states exhibits a number of surprising and counterintuitive features. In a 2002 work, Barnum et al. argued informally that these strange features should imply that digital signatures for quantum states are impossible [6].

In this work, we perform the first rigorous study of the problem of signing quantum states. We first show that the intuition of [6] was correct, by proving an impossibility result which rules out even very weak forms of signing quantum states. Essentially, we show that any non-trivial combination of correctness and security requirements results in negligible security. This rules out all quantum signature schemes except those which simply measure the state and then sign the outcome using a classical scheme. In other words, *only classical signature schemes exist.*

We then show a positive result: *it is possible to sign quantum states, provided that they are also encrypted with the public key of the intended recipient.* Following classical nomenclature, we call this notion *quantum signcryption*. Classically, signcryption is only interesting if it provides superior efficiency to simultaneous encryption and signing. Our results imply that, quantumly, it is far more interesting: by the laws of quantum mechanics, it is the only signing method available.

We develop security definitions for quantum signcryption, ranging from a simple one-time two-user setting, to a chosen-ciphertext-secure many-time multi-user setting. We also give secure constructions based on post-quantum public-key primitives. Along the way, we show that a natural hybrid method of combining classical and quantum schemes can be used to "upgrade" a secure classical scheme to the fully-quantum setting, in a wide range of cryptographic settings including signcryption, authenticated encryption, and chosen-ciphertext security.

1 Introduction

The Internet of the future will plausibly include both large-scale quantum computers and high-capacity quantum channels. How will we securely transmit (quantum) data over the resulting "quantum Internet?" Methods based on entanglement (e.g., teleportation) are costly, using many rounds of interaction to build a shared state which must be at least as large as the data itself. Classically, encryption and authentication offer a non-interactive approach with several attractive features: (i.) keys exchanged over public channels, (ii.) a short key suffices for transmitting unlimited amounts of data, and (iii.) security guarantees are maximal for both secrecy and authenticity. Can we encrypt, authenticate, and sign quantum data to the same standard?

Classical digital signatures, for instance, are ubiquitous in everyday classical cryptography, with applications ranging from secure software distribution and email signatures to e-governance and cryptocurrencies. Given their importance in the classical world, it is natural to ask whether it is possible to devise digital signature schemes for quantum data. Unfortunately, this question has been considered in only one previous work [6]. There, the authors only comment that, since any symmetric-key scheme that authenticates quantum states must also encrypt them, quantum digital signatures must be impossible. They suggest that one can use classical public-key cryptography and one-time quantum authentication to build a scheme they call "public-key quantum authentication." Unfortunately, no formal security definitions or proofs are given, and the theory remains undeveloped. In this work, we return to the problem of signing quantum states, with a rigorous and formal approach. We ask: are digital signatures for quantum states really impossible? Does there exist any feasible tradeoff in the security and correctness requirements which makes them possible?

1.1 Summary Of Results In This Work

Impossibility of quantum signatures. We first define quantum signature schemes, leaving as much room as possible for achievable notions. We provide for the possibility of varying levels of correctness and very weak forms of security.

Definition 1 (informal). *A* **quantum signature scheme** *(or QS) is a pair* $(\mathsf{Sign}, \mathsf{Ver})$ *of keyed quantum polynomial time (QPT) algorithms.*

- *A QS is* ***correct for a map*** N *if* $\|N \circ \mathsf{Ver}_{\mathsf{vk}} \circ \mathsf{Sign}_{\mathsf{sk}} - N\| \leq \mathrm{negl}(n)$.
- *A QS is simply* ***correct*** *if it is correct for* $N = \mathbb{1}$.
- *A QS is* ε-***one-time-secure*** *for a map* L *if, for every adversary* $\mathcal{A}$*, the "attack map"* $L \circ \mathsf{Ver}_{\mathsf{pk}} \circ \mathcal{A} \circ \mathsf{Sign}_{\mathsf{sk}}$ *can be simulated (up to* ε *loss) by an algorithm which either applies* L *or discards the input.*

Requiring correctness only for certain maps N, or security only for certain maps L, weakens the signature scheme in that the (secure) further use of the verified state is limited. Note also that we have relaxed the classical requirement that a signature is additional information that accompanies the message; quantumly, this is impossible due to no-cloning. Instead, signing is allowed to transform the message into an arbitrary state. Correctness requires that the message (or some partial information about it) can be recovered at verification time.

We prove two strong impossibility results. First, we show that full correctness implies negligible security, even in a very weak one-time sense which only guarantees the integrity of the outcomes of a single two-outcome measurement. Second, if we fix a pair of measurements and then ask only for correctness and security of their outcomes, then those measurements must commute.

Theorem 1 (informal). *Let* Π *be a quantum signature scheme.*

1. *If* Π *is correct, then for any two-outcome measurement* M*,* Π *is at most* $(1 - \mathrm{negl}(n))$*-one-time secure for* M*.*
2. *If* Π *is correct and* ε*-one-time-secure for a pair* $\{M_0, M_1\}$ *of two-outcome measurements, then* M_0 *and* M_1 *are indistinguishable from* $(1 - \varepsilon)$*-commuting.*

The measurement maps above are viewed as channels which measure and then output the outcome, discarding the post-measurement state. This result shows convincingly that signature schemes can only sign classical information.

Quantum signcryption. Our impossibility results appear to be devastating to prospects for public-key cryptography with quantum data. They seem to imply that authenticated communication requires each pair of parties to share a secret key. For networks with a large number of parties, or where parties frequently come and go, this is an unwieldy and highly inefficient solution.

In the second part of our work, we show that this is in fact not necessary! The key observation is that impossibility can be circumvented *if we also encrypt the message*. We can achieve this by, for each transmission, selecting an intended recipient and encrypting using their public key. Classically, such combined schemes are called signcryption, and are of interest only insofar as they provide efficiency gains over combined encryption and signing [17]. This is in stark contrast to the quantum world: our results show that signcryption provides *the only way* to achieve integrity and authenticity without a pre-shared secret key.

Definition 2 (informal). *A* **quantum signcryption scheme** *(or QSC) is a triple of QPT algorithms:*

1. *(key generation)* $\mathsf{KeyGen}(1^n)$ *: output* $(\mathsf{sdk}, \mathsf{vek}) \leftarrow \{0,1\}^{\mathrm{poly}(n)}$.
2. *(signcrypt)* $\mathsf{SigEnc}_{\mathsf{sdk},\mathsf{vek}} : \mathfrak{D}(\mathcal{H}_M) \to \mathfrak{D}(\mathcal{H}_C)$
3. *(verified decrypt)* $\mathsf{VerDec}_{\mathsf{vek},\mathsf{sdk}} : \mathfrak{D}(\mathcal{H}_C) \to \mathfrak{D}(\mathcal{H}_M \oplus |\bot\rangle\langle\bot|)$

such that $\|\mathsf{VerDec}_{\mathsf{vek}_S,\mathsf{sdk}_R} \circ \mathsf{SigEnc}_{\mathsf{sdk}_S,\mathsf{vek}_R} - \mathbb{1}_M \oplus 0_\bot\| \leq \mathrm{negl}(n)$.

In the envisioned usage, each party on a network first runs key generation, publishing their "verify/encrypt key" vek while keeping their "sign/decrypt key" sdk private. When a sender S wishes to send a state σ_M to a recipient R, they apply $\mathsf{SigEnc}_{\mathsf{sdk}_S,\mathsf{vek}_R}$ using their private and the public key of R. When R receives the signcryption ϱ_C, they apply $\mathsf{VerDec}_{\mathsf{vek}_S,\mathsf{sdk}_R}$ using their private key and the public key of S. Note that "$\oplus 0_\perp$" above indicates that this process of signcryption followed immediately by verified decryption (with the correct keys) never rejects. We remark that the general definition for the multi-user setting will also need to keep track of "user IDs" in order to prevent identity fraud.

Quantum signcryption security. Signcryption security is naturally divided into two settings, according to whether the adversary has access to the private key of one of the parties (the sender or the recipient) or not.

First, *outsider security* ensures that, if S and R are honest parties, then their channel is private and authenticated against adversaries from whom sdk_S and sdk_R are kept secret. Second, *insider security* ensures that, if one party's secret key is compromised, the security guarantees that are due to the other party's secret key still hold. If the sender key is revealed, the receiver should still enjoy full privacy. The other case is quantumly unachievable: impossibility for signatures implies that releasing the receiver's secret key results in a useless scheme. In particular, quantum signcryption cannot provide *non-repudiation.*

Formally defining outsider and insider security for signcryption is a challenge quantumly: they are roughly analogous to, respectively, authenticated encryption and chosen-ciphertext secrecy. Both of these notions are quite troublesome in the quantum setting, since their usual definitions require the ability to make comparisons between previous queries and outputs of the adversaries (e.g., in deciding if the adversary is attempting to decrypt the challenge in $\mathsf{IND\text{-}CCA2}$.) Fortunately, a recent approach [3] shows the way forward on such definitions. The idea is to "split" the standard security game into two experiments: an "unrestricted" experiment and a "cheat-detecting" experiment. More precisely, In the unrestricted experiment, an adversary interacts with a signcryption and a verified-decryption oracle. In the cheat detecting experiment, the modified signcryption oracle signcrypts half of a maximally entangled state and stores the other half together with the plaintext submitted by the adversary. The modified verified-decryption oracle decrypts the ciphertext it is presented with and checks sequetially whether it is in a maximally entangled[4] state with one of the registers stored by the signcryption oracle. If so, it returns the corresponding stored plaintext, otherwise it returns the reject symbol. Following this approach, we formally define outsider signcryption security as follows.

Experiment 1. The **real outsider experiment** $\mathsf{Out\text{-}Real}(\Pi, \mathcal{A}, n, S, R)$:

1: **output** $\mathcal{A}^{\mathsf{SigEnc}_{S,R},\mathsf{VerDec}_{S,R}}(1^n)$.

Experiment 2. The **ideal outsider experiment** $\mathsf{Out\text{-}Ideal}(\Pi, \mathcal{A}, n, S, R)$:

1: define channel $E_{M\to C}$: prepare maximally entangled state $|\phi^+\rangle_{M'M''}$, store (M'', M) in a set $\mathcal{M}$; return $\mathsf{SigEnc}_{S,R}$ applied to M'.
2: define channel $D_{C\to M}$:
   - apply $\mathsf{VerDec}_{S,R}$ to C, place results in M'
   - for each $(M'', M) \in \mathcal{M}$: measure if $M'M''$ are maximally entangled; if yes, return M
   - return $\perp$.
3: **output** $\mathcal{A}^{E,D}(1^n)$.

Security is defined as indistinguishability of the real and ideal worlds.

Definition 3. *A QSC Π is **outsider secure** if for all QPTs $\mathcal{A}$,*

$$|\Pr[\mathsf{Out\text{-}Real}(\Pi, \mathcal{A}, n) \to \mathsf{real}] - \Pr[\mathsf{Out\text{-}Ideal}(\Pi, \mathcal{A}, n) \to \mathsf{real}]| \leq \mathrm{negl}(n).$$

For the relevant case of insider security (i.e., where the sender is the attacker) we define insider security to be $\mathsf{QIND\text{-}CCA2}$ security of the induced public-key quantum encryption scheme. While [3] only defines this notion in the symmetric-key setting, we adapt it here to public-key schemes.

[4] more precisely, whether it is in the standard maximally entangled state used by the modified signcryption oracle

Constructions, security proofs. Finally, we give secure constructions for quantum signcryption. Along the way, we also give several new security proofs of independent interest, for "bootstrapping" classical encryption security properties to the quantum encryption setting. This is done via a generic classical-quantum hybrid construction $\Pi^{\mathsf{Hyb}}[\Pi, \Sigma]$ which combines a classical scheme Π with a quantum scheme Σ, as follows. The key (or keys) for $\Pi^{\mathsf{Hyb}}[\Pi, \Sigma]$ are the same as for Π. To encrypt ϱ, we generate a key k for Σ and encrypt ϱ with k. Furthermore, we "encapsulate" k by encrypting it with Π and attaching the resulting classical encryption to the encryption of ϱ. Decryption first decapsulates the Σ-key, then uses it to decrypt the rest.

Theorem 2 (informal). *Let Σ be one-time quantum authenticating (*cQCA *[3].)*

1. *If Π is post-quantum* IND-CCA2 *private- (resp., public-) key encryption, then $\Pi^{\mathsf{Hyb}}[\Pi, \Sigma]$ is* QIND-CCA2 *private- (resp., public-) key quantum encryption.*
2. *If Π is post-quantum classical authenticated encryption, then $\Pi^{\mathsf{Hyb}}[\Pi, \Sigma]$ is quantum authenticated encryption.*
3. *If Π is post-quantum classical outsider- and insider-secure signcryption, then $\Pi^{\mathsf{Hyb}}[\Pi, \Sigma]$ is outsider- and insider-secure quantum signcryption.*

2 Preliminaries

Notation, conventions. We will largely use the conventions regarding quantum information from [14]. We use $\mathcal{H}_M$ to denote a complex Hilbert space with label M and finite dimension $\dim M$. A quantum register is a physical system whose set of valid states is $\mathfrak{D}(\mathcal{H}_M)$. In our setting, a "quantum register M" is in fact an infinite family of registers $\{M_n\}_{n\in\mathbb{N}}$ consisting of $p(n)$ qubits, where p is some fixed polynomial. The notation τ_M will mean the maximally mixed state (i.e., uniform classical distribution) $\mathbb{1}/\dim M$ on M.

By "QPT" we mean a polynomial-time uniform family of quantum circuits, i.e., an efficient quantum algorithm. Quantum algorithms implement completely positive (CP), trace-preserving (TP) maps, i.e., quantum channels. To indicate that Φ is a channel from register A to B, we will write $\Phi_{A\to B}$. When it helps to clarify notation, we will use $\circ$ to denote composition of operators. We will also often drop tensor products with the identity, e.g., given a map $\Psi_{BC\to D}$, we will write $\Psi \circ \Phi$ to denote the map $\Psi \circ (\Phi \otimes \mathbb{1}_C)$ from AC to D.

From now on, when mentioning a computational security notion for classical schemes, we mean it by default *in the post-quantum sense*, that is, holding against QPT adversaries.

Quantum, classical encryption. We will assume basic familiarity with the usual classical constructions: public-key encryption (PKE), digital signatures (DS), and signcryption (SC.) For the last item, we refer to [5] for the basic theory. For us it will suffice to recall that one can construct a signcryption scheme by combining a DS with a CCA-secure PKE via "encrypt then sign" [5]. The result is a signcryption scheme with outsider-security and insider-security. We adopt the following notation for keys: (i.) dk = secret decryption key, (ii.) ek = public encryption key, (iii.) sk = secret signing key, (iv.) vk = public verification key, (v.) sdk = secret signing and decryption key, and (vi.) vek = public verification and encryption key. For quantum encryption, we follow [2,3].

Definition 4. *A public-key quantum encryption scheme (or PKQE) is a triple of QPT algorithms:*

1. *(key generation)* KeyGen : *on input* 1^n, *outputs* $(\mathsf{ek}, \mathsf{dk}) \xleftarrow{\$} \mathcal{K}_{\mathsf{E}} \times \mathcal{K}_{\mathsf{D}}$
2. *(encryption)* $\mathsf{Enc} : \mathcal{K}_{\mathsf{E}} \times \mathfrak{D}(\mathcal{H}_M) \to \mathfrak{D}(\mathcal{H}_C)$
3. *(decryption)* $\mathsf{Dec} : \mathcal{K}_{\mathsf{D}} \times \mathfrak{D}(\mathcal{H}_C) \to \mathfrak{D}(\mathcal{H}_M \oplus |\bot\rangle\langle\bot|)$

such that $\|\mathsf{Dec}_{\mathsf{dk}} \circ \mathsf{Enc}_{\mathsf{ek}} - \mathbb{1}_M \oplus 0_\bot\|_\diamond \leq \mathrm{negl}(n)$ *for all* $(\mathsf{ek}, \mathsf{dk}) \leftarrow \mathsf{KeyGen}(1^n)$.

All the spaces have $\mathrm{poly}(n)$ bits or qubits; $\mathcal{K}_{\mathsf{E}}$ and $\mathcal{K}_{\mathsf{D}}$ are classical and C and M are quantum. We assume w.l.o.g. that dk also includes ek, and ek also includes n. Setting $\mathsf{ek} = \mathsf{dk}$ in the Definition 4 yields symmetric-key quantum encryption (SKQE), letting Enc and Dec be classical algorithms yields PKE.

Quantum security. We briefly review security for quantum encryption, i.e. encryption for quantum data. The standard one-time quantum authentication notion is DNS, defined by Dupuis et al. [10]. Roughly speaking, it states that for any attack map Λ the "effective map" $\mathbb{E}_k \mathsf{Dec}_k \circ \Lambda \circ \mathsf{Enc}_k$ is equivalent to a combination of the identity and a "discard" map, along with some map on the adversary's private space. It is well-known that DNS only provides for authentication of the plaintext [3]. One can define a stronger notion, called quantum ciphertext authentication (QCA) which prevents any adversarial modification of the ciphertext [3]. This definition places further constraints on the simulator. While DNS and QCA are information-theoretic security notions, it is straightforward to define variants which only require security against QPTs; we denote these by cDNS and cQCA, respectively. All of these notions can be achieved with a simple scheme, as follows [1,10,3]. The key selects a random element C_k of the Clifford group. Encryption maps ϱ to $\mathcal{C}_k \varrho \otimes |0^n\rangle\langle 0^n| C_k^\dagger$. Decryption undoes C_k, and rejects if any of the attached qubits yield a non-zero measurement. For ease of reference, we restate the definitions of DNS and QCA in Supplementary Section C.1.

Basic secrecy (e.g., IND, IND-CPA, IND-CCA1) is straightforward to generalize to the quantum case. We denote these notions (respectively) by QIND, QIND-CPA, and QIND-CCA1 [2,7]. For more advanced notions like IND-CCA2 or authenticated encryption (AE), much more care is needed: the no-cloning theorem [14] prevents classically straightforward tasks such as "remember this query so you can compare it to a later output of the adversary." Fortunately, a recent work showed that a solution is possible [3]. The idea is to compare adversaries in two games: an "unrestricted game," and a "cheat-detecting" game. To simplify exposition, we will focus on the (weaker) *plaintext security* variants of the notions set down in [3]. This means that we only check whether the adversary is attempting to cheat on the level of plaintexts. We attach a "w" to the acronyms to indicate this distinction. In the final portion of the paper, we describe informally how to extend all our results to the full, ciphertext-secure setting.

For chosen-ciphertext security, the unrestricted game is the usual CCA2 game with no restrictions on the adversary's use of the oracles (e.g., decrypting the challenge is allowed.) In the cheat-detecting game, we substitute the challenge with half of a maximally entangled state, and store the other half so we can later test whether the adversary attemps to decrypt the challenge. Security is defined in terms of the advantage of adversaries between these two games. We call this notion QIND-wCCA2. For details, see Appendix B.1.

For authenticated encryption, in the unrestricted (or "real") world, the adversary is given Enc_k and Dec_k oracles and outputs a bit. In the "ideal" world, they are instead given modified oracles E and D. Here E substitutes every input with half of a fresh maximally entangled state, storing the other half. Meanwhile, D uses these stored states to test if the adversary composes E with D; if he does, D ensures the identity results, and otherwise outputs $\perp$. We call this notion wQAE. For details, see Appendix B.2.

3 Impossibility Of Signing Quantum States

3.1 Defining Quantum Signatures

We begin with a general definition of quantum signature schemes. (See Supplementary Section C.2 for a discussion on quantum signature scheme definitions.)

Definition 5. *A* **quantum signature scheme** *(or QS) is a triple of QPTs:*

1. *(key generation)* $\mathsf{KeyGen}(1^n)$: *output* $(\mathsf{sk}, \mathsf{vk}) \in \mathcal{K}_\mathsf{S} \times \mathcal{K}_\mathsf{V}$.
2. *(sign)* $\mathsf{Sign} : \mathcal{K}_\mathsf{S} \times \mathfrak{D}(\mathcal{H}_M) \to \mathfrak{D}(\mathcal{H}_C)$
3. *(verify)* $\mathsf{Ver} : \mathcal{K}_\mathsf{V} \times \mathfrak{D}(\mathcal{H}_C) \to \mathfrak{D}(\mathcal{H}_M \oplus |\perp\rangle)$.

Here, we have relaxed the classical requirement that a signature is an additional string that accompanies the message; quantumly, this is impossible due to the no-cloning theorem. Instead, the signing algorithm is allowed to transform the message into an arbitrary state, provided that the message (or some partial information about it) can still be recovered at verification time. The key space $\mathcal{K}_\mathsf{S} \times \mathcal{K}_\mathsf{V}$ is a family of sets of bitstrings of size $\mathrm{poly}(n)$, and that M and C are quantum registers of $\mathrm{poly}(n)$ qubits. In a typical setting, the verification key vk will be public, while the signing key sk stays private.

Since our goal is to show an impossibility result, we define correctness separately and permit schemes with very weak guarantees on message recovery. Recall that a projector P defines a two-outcome measurement channel N via $N(X) = \mathrm{Tr}((\mathbb{1} - P)X)|0\rangle\langle 0| + \mathrm{Tr}(PX)|1\rangle\langle 1|$.

Definition 6. *A QS* $(\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Ver})$ *is* ***correct*** *if it satisfies*

$$\|\mathsf{Ver}_{\mathsf{vk}} \circ \mathsf{Sign}_{\mathsf{sk}} - \mathrm{id}_M \oplus 0_\perp\|_\diamond \leq \mathrm{negl}(n) \tag{1}$$

for all $(\mathsf{sk}, \mathsf{vk}) \leftarrow \mathsf{KeyGen}(1^n)$*. More generally, given a finite set* $\mathcal{N}$ *of two-outcome measurements on* $\mathcal{H}_M \oplus |\perp\rangle$*, we say that a QS is* $\mathcal{N}$***-correct*** *if it satisfies*

$$\|N \circ \mathsf{Ver}_{\mathsf{vk}} \circ \mathsf{Sign}_{\mathsf{sk}} - N \oplus 0_\perp\|_\diamond \leq \mathrm{negl}(n)$$

for all $(\mathsf{sk}, \mathsf{vk}) \leftarrow \mathsf{KeyGen}(1^n)$ *and all* $N \in \mathcal{N}$*.*

Definition 7. *Let* $\mathcal{L}$ *be a finite set of two-outcome measurements on* $\mathcal{H}_M \oplus |\perp\rangle$*. A QS is* ε*-one-time* $\mathcal{L}$***-secure*** *if, for any QPT adversary* $\mathcal{A}$*, for all* $L \in \mathcal{L}$ *there exists a probability* $p \in [0, 1]$ *such that*

$$\|L \circ \mathbb{E}\left[\mathsf{Ver}_{\mathsf{vk}} \circ \mathcal{A}(\mathsf{pk}) \circ \mathsf{Sign}_{\mathsf{sk}}\right] - pL + (1-p)\perp\|_\diamond \leq \varepsilon\,. \tag{2}$$

Here, $\perp$ denotes the reject map $X \mapsto \perp\mathrm{Tr}(X)$. The constraint (2) is essentially a weakened version of the DNS security definition for the authentication of quantum states [9], adapted to the public key case. It is weakened in three ways: (i.) we only ask for computational security, (ii.) the adversary does not hold any side information about the plaintext, and (iii.) security only holds for the selected set of measurements. Note that correctness and security are not required to hold for the same set of measurements. So, we could in principle ask for a QS which is $\mathcal{N}$-correct and $\mathcal{L}$-secure for some $\mathcal{N} \neq \mathcal{L}$[5]; in practice, we are probably primarily interested in $\mathcal{N} \cap \mathcal{L}$. In any case, as we will show next, security cannot be achieved except for trivial choices of $\mathcal{N}$ and $\mathcal{L}$.

3.2 Impossibility Of Quantum Signatures

We begin with a technical lemma, characterizing quantum encryption.

Lemma 1 (generalization of Lemma B.9 in [4]). *Let* $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ *be a PKE with exact (approximate) correctness. Then* Enc *and* Dec *have the following form, for every keypair* $k = (\mathsf{dk}, \mathsf{ek})$*:*

$$\left\|\mathsf{Enc}_{\mathsf{pk}} - V_k\left((\cdot) \otimes (\sigma_k)_T\right)V_k^\dagger\right\|_\diamond \leq \varepsilon$$
$$\left\|\mathsf{Dec}_{\mathsf{sk}}(V_k P_T^{\sigma_k} V_k^\dagger(\cdot)V_k P_T^{\sigma_k} V_k^\dagger) - \mathrm{Tr}_T\left[P_T^{\sigma_k}\left(V_k^\dagger(\cdot)V_k\right)P_T^{\sigma_k}\right]\right\|_\diamond \leq \varepsilon.$$

Here, σ_k *is a state on register* T*,* $P_T^{\sigma_k}$ *and* $\bar{P}_T^{\sigma_k}$ *is an orthogonal projector such that* $\|P_T^{\sigma_k}\sigma_k P_T^{\sigma_k} - \sigma_k\|$ *is negligible,* V_k *is a unitary operator and* ε *is neglible.*

The proof is given in Appendix A. Note that the behaviour of $\mathsf{Dec}_{\mathsf{dk}}$ can be arbitrary outside of the range of $V_k P_T^{\sigma_k} V_k^\dagger$. The standard behavior is to output the reject symbol $\perp$ when presented with an invalid ciphertext. The Lemma gives an information-theoretic characterization of the channels $\mathsf{Enc}_{\mathsf{pk}}, \mathsf{Dec}_{\mathsf{sk}}$; this is why it's ok that V and σ are indexed by the keypair. This does not guarantee that $\mathsf{Enc}_{\mathsf{ek}}$ can be efficiently implemented in this form: knowledge of ek alone is not enough to efficiently implement σ_k and V_k (otherwise it would also be enough to decrypt!). In actual schemes, ek is enough to implement σ_k and V_k *on relevant inputs for encryption* efficiently.

We will now show that a QS cannot be both correct and secure for a pair of measurements unless those measurements commute, in the sense that their sequential application yields the same outcome distribution

[5] Note however that for $\varepsilon = \mathrm{negl}(n)$, $\mathcal{N} \subset \mathcal{L}$, unless $\mathcal{L} = \emptyset$.

no matter the order. Recall that a measurement M also defines an instrument, i.e., a measurement map that keeps the post-measurement state, by

$$\tilde{M}(X) = (\mathbb{1} - P)X(\mathbb{1} - P) \otimes |0\rangle\langle 0| + PXP \otimes |1\rangle\langle 1| \tag{3}$$

In the following theorem, the composition of measurements is understood to mean that the instrument of the first measurement is applied and the the second measurement acts on the post-measurement state of the first one, i.e.

$$(M_1 \circ M_0)_{A\to R_0R_1}(X_A) := (M_1)_{A\to R_1}\left((\tilde{M}_0)_{A\to AR_0}(X_A)\right). \tag{4}$$

The index of the outcome register R is tied to the index of the instrument, not to the order the measurements are performed, i.e.

$$(M_0 \circ M_1)_{A\to R_0R_1}(X_A) = (M_0)_{A\to R_0}\left((\tilde{M}_1)_{A\to AR_1}(X_A)\right). \tag{5}$$

Theorem 3. *Let $\Pi = (\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Ver})$ be a QS and $\{(M_0)_{A\to R_0}, (M_1)_{A\to R_1}\}$ be two-outcome projective measurements with efficiently implementable instruments $\tilde{M}_i$. Suppose Π is $\{M_0, M_1\}$-correct and ε-one-time $\{M_0, M_1\}$-secure. Then*

$$\|M_1 \circ M_0(\varrho) - M_0 \circ M_1(\varrho)\|_1 \leq \varepsilon + \mathrm{negl}(n) \tag{6}$$

for all efficiently preparable quantum states ϱ_{MR}.

We remark that the commutation condition Equation 6 therefore holds whenever $\{M_0, M_1\} \subset \mathcal{N} \cap \mathcal{L}$.

Proof. The plan is as follows. The scheme is correct for M_0, the two-outcome measurement defined by a projector P_0. Using this, one can show that running the verification circuit (without discarding qubits, and delaying measurements), applying the reflection unitary $U = \mathbb{1} - 2P_0$ and running the inverse of the verification circuit, produces a valid signed state. More precisely, if the described attack is applied to a signed state obtained from signing $|\psi\rangle$, the resulting state is a signed state that can be also obtained from signing $U|\psi\rangle$. But if M_0 and M_1 don't commute, then the unitary U changes the outcome of M_1. Therefore the measurements have to commute, up to an error equal to the soundness error ε.

We now explain the details. Assume for contracdiction that the conclusion of the theorem does not hold, i.e. that there exists efficiently prepareable ϱ_{MR} such that

$$\|M_1 \circ M_0(\varrho) - M_0 \circ M_1(\varrho)\|_1 > 2\delta. \tag{7}$$

Let $|\psi\rangle_{MRE}$ be an efficiently prepareable purification of of ϱ_{MR}. Let P_i be the projector for the outcome 0 of M_i, for $i = 0, 1$. The assumption (7) implies WLOG that

$$\left|\|P_0P_1|\psi\rangle\|_2^2 - \|P_1P_0|\psi\rangle\|_2^2\right| > \frac{\delta}{2}. \tag{8}$$

Define the reflection unitary $U_i = \mathbb{1} - 2P_i$. Note that these unitaries are efficiently implementable via a Stinespring dilation of the instruments $\tilde{M}_i$. We would like to show that an application of U_0 to $|\psi\rangle$ changes the outcome of M_1. We calculate

$$\begin{aligned}\left|\|P_1U_0|\psi\rangle\|_2^2 - \|P_1|\psi\rangle\|_2^2\right| &= \left|\|P_1(\mathbb{1} - 2P_0)|\psi\rangle\|_2^2 - \|P_1|\psi\rangle\|_2^2\right| \\ &= 2\left|2\langle\psi|P_0P_1P_0|\psi\rangle - \langle\psi|P_1P_0|\psi\rangle - \langle\psi|P_0P_1|\psi\rangle\right|.\end{aligned} \tag{9}$$

Now we use Equation 8 and rewrite the left hand side as

$$\begin{aligned}
&\left|\|P_0P_1|\psi\rangle\|_2^2 - \|P_1P_0|\psi\rangle\|_2^2\right| = \left|\langle\psi|P_1P_0P_1|\psi\rangle - \langle\psi|P_0P_1P_0|\psi\rangle\right| \\
=&\Bigg|\langle\psi|P_1P_0P_1|\psi\rangle - \frac{1}{2}\langle\psi|P_0P_1|\psi\rangle - \frac{1}{2}\langle\psi|P_1P_0|\psi\rangle \\
&+ \frac{1}{2}\langle\psi|P_0P_1|\psi\rangle + \frac{1}{2}\langle\psi|P_1P_0|\psi\rangle - \langle\psi|P_0P_1P_0|\psi\rangle\Bigg| \\
\leq&\left|\langle\psi|P_1P_0P_1|\psi\rangle - \frac{1}{2}\langle\psi|P_0P_1|\psi\rangle - \frac{1}{2}\langle\psi|P_1P_0|\psi\rangle\right| \\
&+ \left|\frac{1}{2}\langle\psi|P_0P_1|\psi\rangle + \frac{1}{2}\langle\psi|P_1P_0|\psi\rangle - \langle\psi|P_0P_1P_0|\psi\rangle\right|,
\end{aligned}$$

where we have added a zero in the second equality and used the triangle inequality. In particular we obtain the inequality

$$\begin{aligned}
\left|\|P_0P_1|\psi\rangle\|_2^2 - \|P_1P_0|\psi\rangle\|_2^2\right| \leq 2\max\Bigg(&\left|\langle\psi|P_1P_0P_1|\psi\rangle - \frac{1}{2}\langle\psi|P_0P_1|\psi\rangle - \frac{1}{2}\langle\psi|P_1P_0|\psi\rangle\right|, \\
&\left|\langle\psi|P_0P_1P_0|\psi\rangle - \frac{1}{2}\langle\psi|P_0P_1|\psi\rangle - \frac{1}{2}\langle\psi|P_1P_0|\psi\rangle\right|\Bigg)
\end{aligned} \tag{10}$$

Putting Equations (8), (9) as well as its counterpart with 0 and 1 interchanged, and (10) together we arrive at the conclusion that either

$$\left|\|P_1U_0|\psi\rangle\|_2^2 - \|P_1|\psi\rangle\|_2^2\right| > \delta, \tag{11}$$

or this equation holds with 0 and 1 interchanged. Assume WLOG that Equation 11 holds. Consider the adversary $\mathcal{A}$ that applies a Stinespring dilation of the verification algorithm, performs the unitary U_0 and then undoes the verification. This adversary produces a valid signed state by linearity. Indeed, let $(\mathsf{pk}, \mathsf{sk})$ be a given key pair, and V_{sk} and W_{pk} be Stinespring dilation unitaries of the signing and verification algorithms, respectively. Here and in the following a Stinespring dilation unitary of a quantum channel $\Lambda_{R_1\to R_2}$ from a k-qubit register R_1 to an ℓ-qubit register R_2 is a unitary $U^{\Lambda}_{R_1E_1\to R_2E_2}$ such that

$$\Lambda(X_{R_1}) = \mathrm{Tr}_{E_2}\left[U^{\Lambda}_{R_1E_1\to R_2E_2}\left(X_{R_1}\otimes|0\rangle\langle 0|_{E_1}\right)\left(U^{\Lambda}_{R_1E_1\to R_2E_2}\right)^{\dagger}\right].$$

Let furthermore $\tilde{P}_0 = W^{\dagger}_{\mathsf{pk}}P_0W_{\mathsf{pk}}$. By the correctness of the QS for M_0, we have that

$$\left|\|\tilde{P}_0V|\phi\rangle\|_2^2 - \|P_0|\phi\rangle\|_2^2\right| \leq \mathrm{negl}(n),$$

for all $|\phi\rangle$, so in particular

$$\left|\|\tilde{P}_0VP_0|\psi\rangle\|_2 - \|VP_0|\psi\rangle\|_2\right| \leq \mathrm{negl}(n),$$

which implies $\|\tilde{P}_0VP_0|\psi\rangle - VP_0|\psi\rangle\|_2 \leq \mathrm{negl}(n)$ and similarly $\|\tilde{P}_0V(\mathbb{1}-P_0)|\psi\rangle\|_2 \leq \mathrm{negl}(n)$. Together, these inequalities yield

$$\|\tilde{P}_0V|\psi\rangle - VP_0|\psi\rangle\|_2 \leq \mathrm{negl}(n)$$

and therefore

$$\|\tilde{U}_0V|\psi\rangle - VU_0|\psi\rangle\|_2 \leq \mathrm{negl}(n), \tag{12}$$

where $\tilde{U}_0 = W^{\dagger}_{\mathsf{pk}}U_0W_{\mathsf{pk}}$. This shows that the attack U_0 produces a valid ciphertext. Because of correctness of the QS for M_1 we have that

$$\left|\|P_1WVU_0|\psi\rangle\|_2^2 - \|P_1U_0|\psi\rangle\|_2^2\right| \leq \mathrm{negl}(n). \tag{13}$$

Equation 12 implies that verification only rejects with negligible probability when applied to $\tilde{U}_0 V|\psi\rangle$. The ε-one-time security therefore implies that

$$\left| \|P_1 W \tilde{U}_0 V|\psi\rangle\|_2^2 - \|P_1|\psi\rangle\|_2^2 \right| \leq \varepsilon.$$

Combined with Equation 12 we therefore get

$$\left| \|P_1 W V U_0|\psi\rangle\|_2^2 - \|P_1|\psi\rangle\|_2^2 \right| \leq \varepsilon + \mathrm{negl}(n)$$

and hence, using Equation 13,

$$\left| \|P_1 U_0|\psi\rangle\|_2^2 - \|P_1|\psi\rangle\|_2^2 \right| \leq \varepsilon + \mathrm{negl}(n).$$

Via Equation 11, this is a contradiction to assumption (7) as long as $\delta \geq \varepsilon + \mathrm{negl}(n)$, which finishes the proof.

□

Note that for logarithmic-size message space, the conclusion of Theorem 3 is equivalent to an upper bound on the commutator of the projectors forming the measurements M_i, as then all states are efficiently preparable.

Theorem 3 shows that simultaneously signing different properties of a quantum state is only possible if these properties are essentially classical, implying that the best possible security can be achieved by just measuring the M_i sequentially while at most incurring an error equal to the soundness parameter ε plus a negligible function and signing the outcome classically. This classical-message protocol, would, however, destroy any quantum properties of the plaintext state – even if no attack occurs. The following complementary impossibility result indicates that part of this loss of quantum information is unavoidable: If we require *full correctness* for a QS in the sense that the composition of $\mathsf{Sign}_{\mathsf{sk}}$ and $\mathsf{Ver}_{\mathsf{vk}}$ yields the identity channel on the plaintext space, no security can be achieved.

Theorem 4. *Let $\Pi = (\mathsf{KeyGen}, \mathsf{Sign}, \mathsf{Ver})$ be a correct QS, and let M be a non-trivial two-outcome measurement. Then Π is at most $(1 - \mathrm{negl}(n))$-one-time $\{M\}$-secure.*

Proof. We use a similar idea as in the proof of Theorem 3. Just assuming correctness we can construct an attack that begins by applying a Stinespring dilation of the verification algorithm. Subsequently we apply a unitary that changes $|\psi_0\rangle$ to $|\psi_1\rangle$, where the $|\psi_i\rangle$ are efficiently prepareable states such that measuring M on $|\psi_i\rangle$ returns result i with certainty. Finally, the attack undoes the Stinespring dilation of the verification. The result is a valid signed state for $|\psi_1\rangle$ when the attack is applied to a signed state for $|\psi_0\rangle$.

The details are as follows. First note that Lemma 1 does not use the fact that the public key is public, or that the secret key is secret, it only uses the correctness of the scheme. Therefore we can apply it to the present QS as well to conclude that the quantum channel implemented by the verification algorithm fulfills the equation

$$\left\| \mathsf{Sign}_{\mathsf{sk}} - V_k \left((\cdot) \otimes (\sigma_k)_T \right) V_k^\dagger \right\|_\diamond \leq \mathrm{negl}(n) \tag{14}$$

$$\left\| \mathsf{Ver}_{\mathsf{vk}}(V_k P_T^{\sigma_k} V_k^\dagger (\cdot) V_k P_T^{\sigma_k} V_k^\dagger) - \mathrm{Tr}_T \left[P_T^{\sigma_k} \left(V_k^\dagger (\cdot) V_k \right) P_T^{\sigma_k} \right] \right\|_\diamond \leq \mathrm{negl}(n), \tag{15}$$

where $k = \mathsf{vk}, \mathsf{sk}$.[6] In particular, Equation 15 implies that there exists a Stinespring dilation unitary $W^{\mathsf{Ver}_{\mathsf{vk}}}$ of $\mathsf{Ver}_{\mathsf{vk}}$ such that

$$\left\| W^{\mathsf{Ver}_{\mathsf{vk}}} V_k P_T^{\sigma_k} V_k^\dagger (\cdot) V_k P_T^{\sigma_k} V_k^\dagger \left(W^{\mathsf{Ver}_{\mathsf{vk}}} \right) - P_T^{\sigma_k} V_k^\dagger (\cdot) V_k P_T^{\sigma_k} \right\|_\diamond \leq \mathrm{negl}(n). \tag{16}$$

Let $\hat{W}^{\mathsf{Ver}_{\mathsf{vk}}}_{CE_1 \to M E_2}$ be an efficiently implementable Stinespring dilation unitary of $\mathsf{Ver}_{\mathsf{vk}}$. We consider the following attack. Let $|\psi_i\rangle_M$, $i = 0, 1$ be efficiently prepareable pure states such that M returns i with certainty

[6] note that the unitary V_k is not necessarily efficiently implementable, i.e. it does not make sense to say it "only depends on the verification key". The same holds for the state σ_k.

when applied to $|\psi_i\rangle\langle\psi_i|_M$, and let U_i be preparation unitaries for $|\psi_i\rangle_M$, i.e. $|\psi_i\rangle_M = U_i|0\rangle_M$. The attack map is now given by

$$\mathcal{A}(X) = \mathrm{Tr}_{E_1}\left[U_{\mathcal{A}} X U_{\mathcal{A}}^\dagger\right]$$
$$U_{\mathcal{A}} = \left(\hat{W}^{\mathsf{Ver}_{\mathsf{vk}}}\right)^\dagger U_1 U_0^\dagger \hat{W}^{\mathsf{Ver}_{\mathsf{vk}}}.$$

Let us define $P := V_k P_T^{\sigma_k} V_k^\dagger$. According to Equation 16, we have

$$\left\| P\mathcal{A}\left(P(\cdot)P\right)P - P\tilde{U}_{\mathcal{A}}\left(P(\cdot)P\right)\tilde{U}_{\mathcal{A}}^\dagger P\right\|_\diamond \leq \mathrm{negl}(n) \tag{17}$$
$$\tilde{U}_{\mathcal{A}} = V_k U_1 U_0^\dagger V_k^\dagger$$

When applying this attack to a ciphertext for $|\psi_0\rangle\langle\psi_0|_M$ we calculate, using the symbol $\approx$ for equality up to negligible difference in trace norm,

$$\begin{aligned}
&P\mathcal{A}\left(P\mathsf{Sign}_{\mathsf{sk}}\left(|\psi_0\rangle\langle\psi_0|_M\right)P\right)P\\
&\approx P\tilde{U}_{\mathcal{A}}\left(P\mathsf{Sign}_{\mathsf{sk}}\left(|\psi_0\rangle\langle\psi_0|_M\right)P\right)\tilde{U}_{\mathcal{A}}^\dagger P\\
&\approx P\tilde{U}_{\mathcal{A}} V_k\left(|\psi_0\rangle\langle\psi_0|_M \otimes (\sigma_k)_T\right)V_k^\dagger \tilde{U}_{\mathcal{A}}^\dagger P\\
&= PV_k\left(\left(U_1 U_0^\dagger|\psi_0\rangle\langle\psi_0|_M U_0 U_1^\dagger\right)\otimes(\sigma_k)_T\right)V_k^\dagger P\\
&= V_k\left(|\psi_1\rangle\langle\psi_1|_M \otimes (\sigma_k)_T\right)V_k^\dagger\\
&\approx \mathsf{Sign}_{\mathsf{sk}}\left(|\psi_1\rangle\langle\psi_1|_M\right).
\end{aligned}$$

Here we have used Equation 17 in the second line, and Equation 14 in the third and last lines. In particular, the calculation above implies that

$$\begin{aligned}
&P\mathcal{A}\left(P\mathsf{Sign}_{\mathsf{sk}}\left(|\psi_0\rangle\langle\psi_0|_M\right)P\right)P\\
&\approx \mathcal{A}\left(\mathsf{Sign}_{\mathsf{sk}}\left(|\psi_0\rangle\langle\psi_0|_M\right)\right).
\end{aligned}$$

We have thus, up to negligible difference, transformed a signature of $|\psi_0\rangle$ into a signature of $|\psi_1\rangle$ and therefore have changed the outcome of M from 0 to 1 up to negligible probability. □

Some variants. Here we briefly discuss why some a-priori plausible ways to circumvent the above impossibility results cannot work.

First, one might think of circumventing our impossibility result by defining a weak form of signatures where the receiver can choose to *either* recover the original message *or* check the authenticity (but not both.) Such "weak quantum signatures" may be enough for interesting applications. However, any useful formulation of weak signatures is likely to imply strong signatures, via quantum error correction: first encode the message, then weak-sign the qubits; during verification, check a random subset of the qubits (at most half the code distance) and then use decoding to recover the message.

Second, one might try quantum keys. In [11], Gottesman and Chuang investigate the possibility of using quantum keys to achieve information-theoretically secure digital signatures for classical data. In their model, the public key is a quantum state, a copy of which can be requested by anybody, including the adversary. A signature scheme like that for quantum data must fail, by the same arguments as in the proofs above: by correctness, an attack using the reflection corresponding to one of the signed measurements will implement the desired attack and leave the copy of the public key undisturbed. This implies impossibility even in a weaker model where the adversary has to hand back the public key.

4 Quantum Signcryption: Basic Ideas

As we saw in the previous section, signing states for public verification is impossible. We thus must resort to using secret verification. This might seem to imply that we have to use secret-key authentication schemes, which come with a dramatic cost: all pairs of parties wishing to communicate must share a key secret to that pair. As it turns out, there is a far better option: *signcryption.* It is natural to define a notion of *quantum signcryption* (analogous to classical signcryption [17,5]), which is a scheme for *simultaneously* signing and encrypting a quantum state (for a particular recipient.)

4.1 Definition, Some Basic Facts, And A Construction.

The basic notion. Consider m parties who require pairwise authenticated and encrypted communication. This can be solved using symmetric-key authenticated encryption, but at a cost of a quadratic number of key-exchange executions. Moreover, each party would need to store m keys, and adding any new parties would require another round of m key-exchange executions and an additional stored key for all parties.

With signcryption, the aim is to achieve the same goal but with drastically reduced resource requirements. Each party needs only to run key generation privately, publish their public key, and keep their single private key. Sending a message is now a matter of "signcrypting" with the sender's private key ("sign") and the receiver's public key ("encrypt.") Receiving a message requires the opposite "verified-decrypting" with the sender's public key ("verify") and the receiver's private key ("decrypt.") Note that adding new parties becomes trivial.

We define signcryption for quantum states as follows.

Definition 8. *A* **quantum signcryption scheme** *(or QSC) is a triple of QPT algorithms:*

1. *(key generation)* $\mathsf{KeyGen}(1^n)$: *output* $(\mathsf{sdk}, \mathsf{vek}) \leftarrow \mathcal{K}_{\mathsf{SD}} \times \mathcal{K}_{\mathsf{VE}}$
2. *(signcrypt)* $\mathsf{SigEnc} : \mathcal{K}_{\mathsf{SD}} \times \mathcal{K}_{\mathsf{VE}} \times \mathfrak{D}(\mathcal{H}_M) \to \mathfrak{D}(\mathcal{H}_C)$
3. *(verified decrypt)* $\mathsf{VerDec} : \mathcal{K}_{\mathsf{VE}} \times \mathcal{K}_{\mathsf{SD}} \times \mathfrak{D}(\mathcal{H}_C) \to \mathfrak{D}(\mathcal{H}_M \oplus |\bot\rangle\langle\bot|)$

such that

$$\|\mathsf{VerDec}_{\mathsf{vek}_S,\mathsf{sdk}_R} \circ \mathsf{SigEnc}_{\mathsf{sdk}_S,\mathsf{vek}_R} - \mathbb{1}_M \oplus 0_\bot\|_\diamond \leq \mathrm{negl}(n) \tag{18}$$

for all $(\mathsf{sdk}_S, \mathsf{vek}_S), (\mathsf{sdk}_R, \mathsf{vek}_R) \leftarrow \mathsf{KeyGen}(1^n)$.

The key spaces $\mathcal{K}_{\mathsf{SD}}$ and $\mathcal{K}_{\mathsf{VE}}$ are classical and of size $\mathrm{poly}(n)$, while the registers C and M are quantum registers consisting of at most $\mathrm{poly}(n)$ qubits. Note that we have adopted the convention that, in the subscripts of SigEnc and VerDec, the sender key always goes first. We will sometimes write $\mathsf{SigEnc}_{S,R}$ and $\mathsf{VerDec}_{S,R}$ to simplify this notation. As usual, we assume w.l.o.g. that sdk also includes vek, and vek also includes n.

A QSC is used to transmit messages as follows. First, a sender S selects a message (placing it in register M) and a receiver R. Then S applies $\mathsf{SigEnc}_{\mathsf{sdk}_S,\mathsf{vek}_R}$ to M, using their secret key sdk_S and the receiver's public key vek_R. The resulting register C is sent to R, who applies $\mathsf{VerDec}_{\mathsf{vek}_S,\mathsf{sdk}_R}$ to C, using their secret key sdk_R and the sender's public key vek_S. Correctness (i.e., (18)) requires that the overall channel implemented by this honest process should be $\mathbb{1}_M$ along with an "accept" output (indicated by "$\oplus 0_\bot$".)

Signatures and encryption from signcryption. Any quantum signcryption scheme trivially yields a quantum signature scheme (Definition 5), as follows.

Proposition 1. *Let* $\Pi = (\mathsf{KeyGen}, \mathsf{SigEnc}, \mathsf{VerDec})$ *be a signcryption scheme (QSC.) Then the following is a signature scheme (QS.)*

- $\mathsf{KeyGen}'(1^n)$*:* $(\mathsf{sdk}_S, \mathsf{vek}_S) \leftarrow \mathsf{KeyGen}(1^n)$ *and* $(\mathsf{sdk}_R, \mathsf{vek}_R) \leftarrow \mathsf{KeyGen}(1^n)$*; output signing key* $\mathsf{sk} := (\mathsf{sdk}_S, \mathsf{vek}_R)$ *and verification key* $\mathsf{vk} := (\mathsf{vek}_S, \mathsf{sdk}_R)$.
- $\mathsf{Sign}'_{\mathsf{sk}} := \mathsf{SigEnc}_{\mathsf{sk}}$ *and* $\mathsf{Ver}'_{\mathsf{vk}} := \mathsf{VerDec}_{\mathsf{vk}}$.

As we will later show, this proposition and the impossibility results of Section 3 together imply that certain security conditions which are achievable by classical signcryption cannot be fulfilled when signcrypting quantum states.

Any QSC also trivially yields a public-key quantum encryption scheme (Definition 4) – simply by swapping the role of the public and secret keys above.

Proposition 2. *Let $\Pi = (\mathsf{KeyGen}, \mathsf{SigEnc}, \mathsf{VerDec})$ be a signcryption scheme (QSC.) Then the following is a public-key encryption scheme (PKQE.)*

- $\mathsf{KeyGen}'(1^n)$*:* $(\mathsf{sdk}_S, \mathsf{vek}_S) \leftarrow \mathsf{KeyGen}(1^n)$ *and* $(\mathsf{sdk}_R, \mathsf{vek}_R) \leftarrow \mathsf{KeyGen}(1^n)$*; output secret key* $\mathsf{dk} := (\mathsf{vek}_S, \mathsf{sdk}_R)$ *and public key* $\mathsf{ek} := (\mathsf{sdk}_S, \mathsf{vek}_R)$.
- $\mathsf{Enc}'_{\mathsf{ek}} := \mathsf{SigEnc}_{\mathsf{ek}}$ *and* $\mathsf{Dec}'_{\mathsf{dk}} := \mathsf{VerDec}_{\mathsf{dk}}$.

A basic construction. We now define a generic "hybrid" construction of quantum signcryption. We will use a classical signcryption scheme to signcrypt a random key, and then use that key to encrypt the quantum state with a (usually one-time secure) symmetric-key quantum encryption scheme.

Construction 1. Let Π^{SC} be a classical signcryption scheme, and Π^{SKQE} a symmetric-key quantum encryption scheme. Define a quantum signcryption scheme $\Pi^{\mathsf{Hyb}}_{\mathsf{QSC}}[\Pi^{\mathsf{SC}}, \Pi^{\mathsf{SKQE}}] = (\mathsf{KeyGen}, \mathsf{SigEnc}, \mathsf{VerDec})$ as follows:

1. $\mathsf{KeyGen}(1^n)$: output $(\mathsf{sdk}, \mathsf{vek}) \leftarrow \mathsf{KeyGen}^{\mathsf{SC}}(1^n)$;
2. $\mathsf{SigEnc}_{S,R}$: on input ϱ_M, generate $k \leftarrow \mathsf{KeyGen}^{\mathsf{SKQE}}(1^n)$ and output $(\mathsf{SigEnc}^{\mathsf{SC}}_{S,R}(k), \mathsf{Enc}^{\mathsf{SKQE}}_k(\varrho))$;
3. $\mathsf{VerDec}_{S,R}$: on input (c, σ_C), set $k = \mathsf{VerDec}^{\mathsf{SC}}_{S,R}(c)$; if $k = \perp$, output $\perp$ and otherwise output $\mathsf{Dec}^{\mathsf{SKQE}}_k(\sigma)$.

We denote the analogous construction of public- (resp., symmetric-) key quantum encryption by $\Pi^{\mathsf{Hyb}}_{\mathsf{PKQE}}$ (resp., $\Pi^{\mathsf{Hyb}}_{\mathsf{SKQE}}$). A special case of this hybrid construction was proposed by Barnum et al. [6] (Section 5.1), but without any security definitions or proofs. We will later formally examine the security of this construction in various settings, and for various choices of Π^{SC} and Π^{SKQE}.

4.2 One-Time Signcryption Security

We now consider the simplest nontrivial setting of quantum signcryption: a single sender S needs to send a single signcrypted state to a single receiver R. Our goal here is to describe the most basic setting of interest, with definitions and security notions that require only minimal familiarity with previous work on quantum encryption and authentication. We will build up to the full setting in Section 5.

Outsider security. In the first, so-called "outsider security" setting, an untrusted third-party adversary $\mathcal{A}$ attacks the channel between S and R. We assume that $\mathcal{A}$ can convince S to signcrypt any message, leaving $\mathcal{A}$ free to manipulate the resulting signcryption before convincing R to unsigncrypt.

Experiment 3. Let $\Pi = (\mathsf{KeyGen}, \mathsf{SigEnc}, \mathsf{VerDec})$ be a quantum signcryption scheme. An ***outsider attack*** (in the one-time, two-user setting) by a QPT $\mathcal{A}$ proceeds as follows.

1. *(Setup.)* Generate $(\mathsf{sdk}_S, \mathsf{vek}_S) \leftarrow \mathsf{KeyGen}(1^n)$ and $(\mathsf{sdk}_S, \mathsf{vek}_S) \leftarrow \mathsf{KeyGen}(1^n)$. Give vek_S and vek_R to $\mathcal{A}$.
2. *(Signcrypt.)* $\mathcal{A}$ prepares a state ϱ_{MB} and $\mathsf{SigEnc}_{S,R}$ is applied to register M, yielding registers C and B.
3. *(Attack, unsigncrypt.)* $\mathcal{A}$ applies a channel $\Lambda_{CB\to CB}$ (possibly depending on $\mathsf{vek}_S, \mathsf{vek}_R$), and $\mathsf{VerDec}_{S,R}$ is applied to register C.

We refer to $\mathcal{A}$ as an ***outside attacker***, and define the effective attack map

$$\tilde{\Lambda}^{(\mathcal{A},S,R)}_{MB\to MB} := \mathsf{VerDec}_{S,R} \circ \Lambda \circ \mathsf{SigEnc}_{S,R}\,.$$

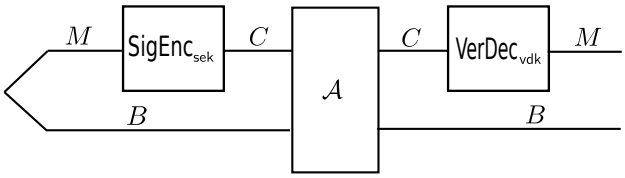

Fig. 1. A one-time outsider attack.

Our security definition will be based on the standard DNS security notion for quantum authentication [10,8], adapted to the particulars of our setting. Specifically, we will require that any (efficient) outside attack amounts to a choice between implementing the "identity map" or the "discard map" on the plaintext space. The former results in a reject.

Definition 9. *A QSC Π is (one-time, two-user)* ***outsider-secure*** *if, for every QPT outside attacker $\mathcal{A}$, there exists a QPT simulator $\mathcal{S}$ which implements a map $\Phi_{MB\to MB}$ of the form*

$$\Phi : \sigma_{MB} \longmapsto \Phi^{\mathsf{acc}}_{B\to B}(\sigma_{MB}) + |\bot\rangle\langle\bot| \otimes \Phi^{\mathsf{rej}}_{B\to B}(\sigma_B)$$

which is indistinguishable from $\tilde{\Lambda}^{(\mathcal{A},S,R)}$, and where Φ^{acc} and Φ^{rej} act on B only.

A strengthening. Note that both SigEnc and VerDec receive, in addition to their input, a private parameter (i.e., sdk) and a public parameter (i.e., vek.) It is thus reasonable to define a stronger notion of security, where the adversary can set these parameters for the oracle calls occurring in the "Signcrypt" and "Attack, unsigncrypt" phases of Experiment 3. Indeed, this models a very plausible real-world attack. Since our goal here is to develop the one-time setting as a minimal first stepping stone, we will examine this stronger security later, when we address the multi-user setting in Section 5.

Insider security. Note that an outside attacker knows all but two pieces of information: the private key of S, and the private key of R. In "insider security," we consider the case where one of these private keys is known to the attacker (if both are known, security is impossible.)

Insider security of sender. First, consider the case of an adversary in possession of sdk_R. The private information is now only sdk_S, and the public information (i.e., the information accessible by the adversary) is $\mathsf{sdk}_R, \mathsf{vek}_R, \mathsf{vek}_S$. This means the adversary can unsigncrypt, so we can't expect S to maintain any secrecy. Can S expect unforgeability? If they could, then this would yield (as in Proposition 1) an unforgeable quantum signature scheme with secret signing key $(\mathsf{sdk}_S, \mathsf{vek}_R)$ and public verification key $(\mathsf{vek}_S, \mathsf{sdk}_R)$.) However, the impossibility results of Section 3 state that this induced signature scheme cannot satisfy even the most minimal security requirements. It follows that quantum signcryption cannot fulfill any insider security guarantees in this case.

Recall that *non-repudiation* is a property of classical digital signatures. It means simply that, due to the unforgeability property, the recipient R of a document m signed by a sender S can present the document and signature to a third party (e.g., the judge) as evidence that m was indeed signed by S. If quantum signcryption is used instead, the impossibility results of Section 3 tell us that R will be able to forge signcryptions; so, this property is impossible to achieve quantumly.

Insider security of receiver. It remains to consider insider security where the adversary has sdk_S. The private information is now only sdk_R. By Proposition 2, Π is now a public-key quantum encryption scheme, with secret decryption key $(\mathsf{sdk}_R, \mathsf{vek}_S)$ and public encryption key $(\mathsf{sdk}_S, \mathsf{vek}_R)$.

Definition 10. *A QSC Π is (one-time, two-user)* ***insider-secure*** *if the public-key quantum encryption scheme induced by making the sender key public (i.e., via Proposition 2) satisfies* QIND-CPA.

As with outsider security, we will not (yet) concern ourselves with adversaries who can select the key parameters of the SigEnc and VerDec calls.

We remark that the above definitions of one-time security translate directly to the classical setting, by using standard notions of classical unforgeability and secrecy (see, e.g., [12]) and plugging them into Definition 9 and Definition 10.

Achieving one-time outsider and insider security. We now show that the hybrid scheme $\Pi^{\mathsf{Hyb}}_{\mathsf{QSC}}$ achieves security in this setting, when equipped with appropriate component schemes.

Theorem 5. *Select the following:*

- *A classical signcryption scheme Π^{SC} with one-time outsider and insider security (defined as above);*
- *A* cDNS*-secure SKQE scheme Π^{SKQE} (Definition 18).*

Then $\Pi^{\mathsf{Hyb}}_{\mathsf{QSC}}[\Pi^{\mathsf{SC}}, \Pi^{\mathsf{SKQE}}]$ (Construction 1) is a quantum signcryption scheme which is one-time outsider-secure and insider-secure.

Proof. We provide a brief sketch; for the full proof see Supplemental Section C.3. Insider security follows from the QIND-CPA security of the hybrid PKQE construction (i.e. Theorem 6) and the fact that cDNS implies QIND [6,4]. For outsider security, observe that if the attack map Λ modifies the classical part c of the ciphertext, the decryption function will reject with overwhelming probability due to the unforgeability of Π^{SC}. Assuming that Λ doesn't change c, we can replace c by an encryption c' of 0 by the IND-CPA security of Π^{SC}, and swapping c and c' back right before decryption. This modified experiment can be simulated using the cDNS oracles: we (i.) apply Λ_{CB}, then (ii.) discard register C^{SC}. Π^{SKQE} is cDNS secure, so for every attack there exists a simulator that either does not attack or rejects. Applying this fact to the attack described above, we obtain a simulator for $\Pi^{\mathsf{Hyb}}_{\mathsf{QSC}}$ as mandated by Definition 9 when the input state

$$\varrho'_{MBC^{\mathsf{Cl}}K} = \varrho_{MB} \otimes \mathsf{SigEnc}_{\mathsf{sdk}_S, \mathsf{vek}_R}\left(\tilde{k}\right) \otimes (\mathsf{sdk}_S, \mathsf{vek}_S, \mathsf{sdk}_R, \mathsf{vek}_R)$$

with $(\mathsf{sdk}_i, \mathsf{vek}_i) \leftarrow \mathsf{KeyGen}^{\mathsf{SC}}$ for $i = S, R$ and $\tilde{k} \leftarrow \mathsf{KeyGen}^{\mathsf{SKQE}}$ is supplied. □

To get an explicit instantiation, we can let Π^{SC} be a classical signcryption scheme constructed from Lamport signatures and standard post-quantum IND-CPA public-key LWE encryption via "encrypt-then-sign," and let Π^{SKQE} to be the scheme $\mathsf{Enc}_k : \varrho \mapsto C_k\left(\varrho \otimes |0^n\rangle\langle 0^n|\right) C_k^\dagger$ where $\{C_k\}_k$ is the Clifford group.

5 Quantum Signcryption: Full Security

We now describe in detail how to "upgrade" from the (one-time, two-user) setting of the previous section, to the full setting. This will involve three steps. First, in Section 5.1 we will describe how to upgrade to many-time security. The adversary will now have oracle access to $\mathsf{SigEnc}_{S,R}$ and $\mathsf{VerDec}_{S,R}$, but secrecy and authenticity must still be preserved. The usual classical security games for these notions do not make sense with quantum data. A recent approach of [3] shows how to get around these issues for symmetric-key encryption; we will extend this approach to our setting. Second, in Section 5.2 we will upgrade to multiple users. This means keeping track of "IDs" for each user, and adding some constraints, e.g., to prevent identity fraud. Here we will give a generic transformation (following [5]) which turns any two-user secure QSC into a multi-user secure QSC simply by attaching IDs to the plaintext before signcrypting. Finally, in Section 5.3 we will describe how to upgrade all of the above to ensure ciphertext authenticity.

5.1 Upgrading To The Many-Time Setting

We begin by upgrading from the one-time setting to the many-time setting, while keeping the number of users at two, i.e., one sender S and one receiver R.

Outsider security. As in the one-time setting, the attack here is launched by a third-party adversary $\mathcal{A}$ against a sender-receiver pair (S, R). In addition to knowing the public information $\mathsf{vek}_S, \mathsf{vek}_R$, $\mathcal{A}$ will now also be able to observe some of the transmissions from S to R. It is also reasonable to expect that $\mathcal{A}$ will have some knowledge about (or even control over) what is being transmitted. As usual, we model this by being as generous as possible to the adversary; specifically, we give $\mathcal{A}$ oracle access to both $\mathsf{SigEnc}_{S,R}$ and $\mathsf{VerDec}_{S,R}$. As it turns out, we will be able to use the approach of [3] to define security, thus avoiding the usual issues with quantum no-cloning and measurement. As signcryption provides secrecy, authenticity and integrity, it can be seen as a kind of public-key analogue of authenticated encryption. Due to the similar security goals, we use a plaintext-based public-key version of the real vs. ideal approach developed in [3] for the definition of quantum authenticated encryption.

Let $\Pi = (\mathsf{KeyGen}, \mathsf{SigEnc}, \mathsf{VerDec})$ be a quantum signcryption scheme with message register M and signcryption register C. In addition, select a security parameter n and an oracle QPT adversary $\mathcal{A}$. Let $|\phi^+\rangle$ denote some choice of maximally entangled state, and Π^+ the projector $|\phi^+\rangle\langle\phi^+|$ onto that state.

Experiment 4. The **real outsider experiment** $\mathsf{Out\text{-}Real}(\Pi, \mathcal{A}, n)$:

1: $(\mathsf{sdk}_S, \mathsf{vek}_S) \leftarrow \mathsf{KeyGen}(1^n)$ and $(\mathsf{sdk}_R, \mathsf{vek}_R) \leftarrow \mathsf{KeyGen}(1^n)$;
2: **output** $\mathcal{A}^{\mathsf{SigEnc}_{S,R}, \mathsf{VerDec}_{S,R}}(1^n)$.

Experiment 5. The **ideal outsider experiment** $\mathsf{Out\text{-}Ideal}(\Pi, \mathcal{A}, n)$:

1: $(\mathsf{sdk}_S, \mathsf{vek}_S) \leftarrow \mathsf{KeyGen}(1^n)$ and $(\mathsf{sdk}_R, \mathsf{vek}_R) \leftarrow \mathsf{KeyGen}(1^n)$;
2: define channel $E_{M \to C}$:
   (1) prepare $|\phi^+\rangle_{M'M''}$, store (M'', M) in a set $\mathcal{M}$;
   (2) apply $\mathsf{SigEnc}_{S,R}$ to M'; **return** result.
3: define channel $D_{C \to M}$:
   (1) apply $\mathsf{VerDec}_{S,R}$ to C, place results in M';
   (2) **for each** $(M'', M) \in \mathcal{M}$ **do**:
   (3) apply $\{\Pi^+, \mathbb{1} - \Pi^+\}$ to $M'M''$;
   (4) **if** outcome is 0: **return** M;
   (5) **end for**
   (6) **return** $|\bot\rangle\langle\bot|$;
4: **output** $\mathcal{A}^{E,D}(1^n)$.

Definition 11. *A quantum signcryption scheme Π is (many-time, two-user)* ***outsider secure*** *if for all QPT adversaries $\mathcal{A}$,*

$$|\Pr[\mathsf{Out\text{-}Real}(\Pi, \mathcal{A}, n) \to \mathsf{real}] - \Pr[\mathsf{Out\text{-}Ideal}(\Pi, \mathcal{A}, n) \to \mathsf{real}]| \leq \mathrm{negl}(n).$$

Insider Security. As before, the remaining conditions to consider amount to allowing the adversary access to the private key of either S or R (but not both) which we can simply view as one of the two parties (S or R) attacking the other. Also as before, the impossibility results for quantum signatures (Section 3) imply that one cannot expect any security in the case where R is the adversary. It remains to consider the case where S is the adversary, and ask if the secrecy of R can be preserved. Here we ask for the strongest notion of secrecy, which is QIND-CCA2. This notion was recently defined in [3] in the symmetric-key case; we show how to adapt it to the public-key case in Section B.1 (where we also define the corresponding cDNS-based "weak" variant QIND-wCCA2).

Definition 12. *A quantum signcryption scheme Π is* ***insider-secure*** *(many-time, two-user) if the public-key quantum encryption scheme induced by making the sender key public (i.e., via Proposition 2) satisfies* QIND-CCA2.

Security of the classical-quantum hybrid. We will show security for the hybrid quantum signcryption scheme by leveraging results about the security of the hybrid approach in two other settings: public-key chosen-ciphertext-secure encryption, and symmetric-key authenticated encryption. The public-key hybrid scheme was first introduced in [7]. In Appendix B, we prove that this hybrid construction preserves security, in the following sense.

Theorem 6. *Let Π be a PKE, and Σ a SKQE. Then $\Pi^{Hyb}_{PKQE}[\Pi, \Sigma]$ is a public-key quantum encryption scheme (PKQE). Moreover, we have:*

1. *If Π is* IND-CPA *and Σ is* QIND*, then $\Pi^{Hyb}_{PKQE}[\Pi, \Sigma]$ is* QIND-CPA*.*
2. *If Π is* IND-CCA1 *and Σ is* QIND*, then $\Pi^{Hyb}_{PKQE}[\Pi, \Sigma]$ is* QIND-CCA1*.*
3. *If Π is* IND-CCA2 *and Σ is* cDNS*, then $\Pi^{Hyb}_{PKQE}[\Pi, \Sigma]$ is* QIND-wCCA2*.*
4. *If Π is* IND-CCA2 *and Σ is* cQCA*, then $\Pi^{Hyb}_{PKQE}[\Pi, \Sigma]$ is* QIND-CCA2*.*

In the symmetric-key case, we need authenticated encryption. Classically this is chosen-ciphertext secrecy plus unforgeability of ciphertexts. Quantumly, we need the aforementioned notion of QAE (and in our case, for ease of exposition in this work, also a cDNS-based weakened version, wQAE) defined in the real/ideal approach as described in Appendix B.2. There we also show the following.

Theorem 7. *Let Π be a classical symmetric-key authenticated encryption scheme, and let Σ be a* cDNS*-(resp.,* cQCA*-) secure symmetric-key quantum encryption scheme. Then $\Pi^{Hyb}_{PKQE}[\Pi, \Sigma]$ is a* wQAE*-(resp.,* QAE*-) secure symmetric-key quantum encryption scheme.*

Achieving Many-Time, Two-User Security. With the above generic results on the classical-quantum hybrid in hand, we can now show security of the quantum signcryption hybrid. It turns out that the hybrid scheme can achieve both outsider and insider (sender) security even in the many-time setting, provided that the classical component is strengthened sufficiently. This is proven as Theorem 11 and Corollary 2 in Appendix B.

Theorem 8. *Select the following:*

- *A many-time, two-user outsider- and insider-secure signcryption Π^{SC};*
- *and a* cDNS*-secure symmetric-key quantum encryption Π^{SKQE}.*

Then $\Pi^{Hyb}_{QSC}[\Pi^{SC}, \Pi^{SKQE}]$ (Construction 1) is a quantum signcryption scheme which is (many-time, two-user) outsider-secure and insider-secure.

5.2 Upgrading To The Multi-User Setting

In the multi-user setting there are different users, each with unique "IDs." We denote the ID of a party P by $\mathsf{ID}_P \in \mathcal{I}$. We assume that there is an efficient public lookup for the map $\mathsf{ID}_P \mapsto \mathsf{vek}_P$ (this is usually achieved by a root of trust, PKI, or similar). In this setting we need to change the syntax of signcryption in order to account for the adversary's ability to *spoof sender's and/or receiver's identitites*. In order to protect the scheme against these attacks we need to bind the users' identities to the corresponding keys, and impose extra correctness conditions based on the match between identities and keys. There are different ways to achieve this; in this work, we opt for the simplified notation below. In particular, we assume that every user P has a single keypair (it is possible to remove this restriction by adding a *key index* to the output of KeyGen and managing the index together with the identities. This is straightforward but beyond the scope of this work.)

Definition 13. *A* (multi-user) quantum signcryption scheme (or muQSC) *is a triple of QPT algorithms:*

- *(key generation)* $\mathsf{KeyGen} : (1^n, \mathsf{ID}_P \in \mathcal{I}) \mapsto (\mathsf{sdk}_P, \mathsf{vek}_P) \in \mathcal{K}_{\mathsf{SD}} \times \mathcal{K}_{\mathsf{VE}}$

- *(signcrypt)* $\mathsf{SigEnc} : \mathcal{K}_{\mathsf{SD}} \times \mathcal{K}_{\mathsf{VE}} \times \mathfrak{D}(\mathcal{H}_M) \to \mathfrak{D}(\mathcal{H}_C)$
- *(unsigncrypt)* $\mathsf{VerDec} : \mathcal{K}_{\mathsf{SD}} \times \mathcal{K}_{\mathsf{VE}} \times \mathfrak{D}(\mathcal{H}_C) \to \mathcal{I}^2 \otimes \mathfrak{D}(\mathcal{H}_M \oplus |\bot\rangle\langle\bot|)$

such that:

1. $\|\mathsf{VerDec}_{\mathsf{sdk}_R,\mathsf{vek}_S} \circ \mathsf{SigEnc}_{\mathsf{sdk}_S,\mathsf{vek}_R} - \{S,R\} \otimes (\mathbb{1}_M \oplus 0_\bot)\|_\diamond \leq \mathrm{negl}(n)$
2. *For* $\mathsf{VerDec}_{\mathsf{sdk}_R,\mathsf{vek}_S}(X) = (S', R', Y)$*, if* $Y \neq \bot$*, then* $(R,S) = (R',S')$

for all $(\mathsf{sdk}_P, \mathsf{vek}_P) \in \mathbf{supp}\,\mathsf{KeyGen}(1^n, \mathsf{ID}_P)$*, and where we assume w.l.o.g. that* sdk_P *also includes* vek_P*, and* vek_P *also includes* 1^n *and* ID_P*, for every* $(\mathsf{sdk}_P, \mathsf{vek}_P) \leftarrow \mathsf{KeyGen}(1^n, \mathsf{ID}_P), \forall\ P \in \mathcal{I}$*.*

Note that the output of $\mathsf{VerDec}_{\mathsf{sdk}_R,\mathsf{vek}_S}$ includes user identities for sender and receiver. Correctness demands that such identities are the ones corresponding to the keys used during a "correct" usage of SigEnc. In theory nothing forbids a malicious adversary to modify a ciphertext in such a way that the resulting identities output by VerDec are different. In this sense, an "identity fraud attack" for signcryption in the multi-user setting results in a forgery attack.

Outsider security. The scenario here is similar to outsider security for the two-user setting: an external adversary $\mathcal{A}$ mounts an attack against a sender/receiver pair (S,R). However, there are two fundamental differences:

- there are many users in $\mathcal{I}$, and hence many possible (S,R) pairs; and
- beyond the "usual" attack scenarios covered by the two-user case, $\mathcal{A}$ might be able to attack the signcryption scheme by performing *identity fraud* (i.e., spoofing the identity of the sender or the receiver of a signcrypted message).

Dealing with the latter is easy: as previously discussed in Definition 13, in order to mount an identity fraud attack $\mathcal{A}$ has to modify the identities output by the VerDec algorithm. This means that such attacks are actually a special case of *message forgeries*, which are already covered by the real vs. ideal approach in the outsider security scenario of Section 5.1.

The other issue is more subtle. There might be insecure schemes which produce weaker keys for certain user IDs, and the adversary might exploit this weakness. Or there might be schemes where the compromise of a certain number of users (and hence knoweldge of their secret keys) also compromises the security of other users. In order to deal with all these scenarios, we will adopt a very conservative approach:

1. the adversary is given the possibility of choosing S and R among $\mathcal{I}$; and
2. the adversary receives the secret keys of *all the other users*.

We define the modified experiments M-Out-Real and M-Out-Ideal accordingly (see Supplemental Section D, and security is given as usual in terms of indistinguishability of the two.

Definition 14. *A multi-user quantum signcryption scheme* Π *is* ***outsider secure*** *if for all QPT* $\mathcal{A}$*,*

$$|\Pr[\mathsf{M\text{-}Out\text{-}Real}(\Pi, \mathcal{A}, n) \to \mathsf{real}] - \Pr[\mathsf{M\text{-}Out\text{-}Ideal}(\Pi, \mathcal{A}, n) \to \mathsf{real}]| \leq \mathrm{negl}(n).$$

Insider security. In addition to picking a sender and receiver of his choice, the adversary is now allowed to access the private key of either of the two (but not both). Also as before, given the impossibility results for quantum signatures (Section 3), it makes sense to consider the case where S is the adversary, and ask if the secrecy of R can be preserved. Again, we ask for the strongest notion of secrecy. However, given the multi-user setting, we must now consider a multi-user version of QIND-CCA2, where the adversary gets to pick the target R to attack, and receives all other user's secret keys. We omit the details here, as the resulting definition follows a similar approach as in the outsider security case.

Definition 15. *A multi-user quantum signcryption scheme* Π *is* ***insider-secure*** *if the public-key quantum encryption scheme induced by making the sender key public (i.e., via Proposition 2) satisfies (the multi-user version of)* QIND-CCA2*.*

Achieving Many-Time, Multi-User Security. The multi-user definitions can then be fulfilled in a way analogous to the classical case. Here, An, Dodis and Rabin [5] describe a simple generic transformation from a two-user secure scheme to a multi-user secure scheme, which proceeds as follows. Whenever signcrypting any plaintext, we first attach to the plaintext the IDs of both the sender and the intended receiver. When applying verified decryption, we check whether the IDs attached to the plaintext are correct. If so, we output the plaintext, otherwise reject. Arguing multi-user security of this transformation then reduces in a straightforward way to two-user security, along the same line of argument as is done in [5].

5.3 Upgrading To Ciphertext Authentication

As discussed above, the security notions we have developed for signcryption provide for *plaintext security*. For instance, in the outsider security case, this means that the adversary cannot perform any attack which modifies the underlying plaintext. The advantage of this approach is that the security games and definitions are rather simple to state and describe, and intuitive. However, this leaves open the possibility of adversaries that modify ciphertexts without being noticed. As it turns out, we can address this case as well, roughly following the route followed by [3] in the private-key encryption setting. Here we only briefly describe the modifications, as they are essentially identical to those in [3].

First, we observe that the characterization lemma (Lemma 1) shows that all quantum encryption schemes follow a simple "attach auxiliary state, then apply keyed isometry V_k" form. This form is also efficient in all schemes we are aware of (we state this formally as Condition 1), but may *in principle* be inefficient.

Next, with this characterization in hand, we can replace the "trap setting" step of each relevant security game (e.g., Step 1 in the Out-Ideal experiment) as follows. We still encrypt half of a maximally-entangled state ϕ^+ and store the input and half of ϕ^+. In addition, we now also store the classical randomness used to sample the aforementioned auxiliary state (prior to applying V_k.)

We then correspondingly adjust the second "cheat detection" step of each relevant security game (e.g., Step 2 in the Out-Ideal experiment), as follows. Given a ciphertext, we first undo the isometry $V_k^\dagger$. We then check if the auxiliary state agrees with the stored randomness. If it does, we additionally perform the entanglement check as before. If both checks say "yes," we supply the stored state. Otherwise we output $\perp$. (In the insider-security case, it's slightly different: "yes" now means that the adversary is attempting to decrypt the challenge, so we simply terminate and output "cheat.")

This transformation yields enhanced cheat-detection games for both outsider-security (analogous to QAE) and insider-security (analogous to QIND-CCA2.) Security is then still defined in terms of the advantage of adversaries at distinguishing the test game from the corresponding cheat-detection game for outsider security, and the winning advantage in the test game over the cheat-detection game for insider security, respectively. Provided that the relevant construction now uses QCA [3] instead of cDNS quantum authentication as a building block, the security proofs carry over essentially unchanged.

6 Acknowledgements

The authors would like to thank Yfke Dulek, Christopher Portmann, Christian Schaffner, and Yi-Kai Liu for helpful feedback about this work. Part of this work was done while T.G. was supported by IBM Research Zurich. T.G. acknowledges financial support from the H2020 FutureTPM project (agreement 779391). GA acknowledges support from NSF grant CCF-1763736.

References

1. D. Aharonov, M. Ben-Or, and E. Eban. Interactive proofs for quantum computations. In *Innovations in Computer Science - ICS 2010, Tsinghua University, Beijing, China, January 5-7, 2010. Proceedings*, pages 453–469, 2010.
2. G. Alagic, A. Broadbent, B. Fefferman, T. Gagliardoni, C. Schaffner, and M. S. Jules. Computational security of quantum encryption. In *Information Theoretic Security - 9th International Conference, ICITS 2016, Tacoma, WA, USA, August 9-12, 2016, Revised Selected Papers*, pages 47–71, 2016.
3. G. Alagic, T. Gagliardoni, and C. Majenz. Unforgeable quantum encryption. In J. B. Nielsen and V. Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018*, pages 489–519, Cham, 2018. Springer International Publishing.
4. G. Alagic and C. Majenz. Quantum non-malleability and authentication. In *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part II*, pages 310–341, 2017.
5. J. H. An, Y. Dodis, and T. Rabin. On the security of joint signature and encryption. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 83–107. Springer, 2002.
6. H. Barnum, C. Crépeau, D. Gottesman, A. D. Smith, and A. Tapp. Authentication of quantum messages. In *43rd Symposium on Foundations of Computer Science (FOCS 2002), 16-19 November 2002, Vancouver, BC, Canada, Proceedings*, pages 449–458, 2002.
7. A. Broadbent and S. Jeffery. Quantum homomorphic encryption for circuits of low t-gate complexity. In *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part II*, pages 609–629, 2015.
8. A. Broadbent and E. Wainewright. Efficient simulation for quantum message authentication. In *Information Theoretic Security - 9th International Conference, ICITS 2016, Tacoma, WA, USA, August 9-12, 2016, Revised Selected Papers*, pages 72–91, 2016.
9. F. Dupuis, J. B. Nielsen, and L. Salvail. Secure two-party quantum evaluation of unitaries against specious adversaries. In *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings*, pages 685–706, 2010.
10. F. Dupuis, J. B. Nielsen, and L. Salvail. Actively secure two-party evaluation of any quantum operation. In *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, pages 794–811, 2012.
11. D. Gottesman and I. Chuang. Quantum digital signatures. *arXiv preprint quant-ph/0105032*, 2001.
12. J. Katz and Y. Lindell. *Introduction to Modern Cryptography, Second Edition*. CRC Press, 2014.
13. D. Kretschmann, D. Schlingemann, and R. F. Werner. The information-disturbance tradeoff and the continuity of stinespring's representation. *IEEE transactions on information theory*, 54(4):1708–1717, 2008.
14. M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, New York, NY, USA, 10th edition, 2011.
15. T. Shrimpton. A characterization of authenticated-encryption as a form of chosen-ciphertext security. *IACR Cryptology ePrint Archive*, 2004:272, 2004.
16. A. J. Winter. Coding theorem and strong converse for quantum channels. *IEEE Trans. Information Theory*, 45(7):2481–2485, 1999.
17. Y. Zheng. Digital signcryption or how to achieve cost(signature & encryption) ¡¡ cost(signature) + cost(encryption). In B. S. Kaliski, editor, *Advances in Cryptology — CRYPTO '97*, pages 165–179, Berlin, Heidelberg, 1997. Springer Berlin Heidelberg.

A Characterization Of Quantum Encryption

In this section, we will prove Lemma 1 which requires the following

Theorem 9 (Theorem I in [13]). *Let $T^{(i)}_{A\to A}$, $i = 1, 2$ be quantum channels with the same input and output systems, and let $U^{(i)}_{A\to AE}$ be Stinespring dilation isometries of $T^{(i)}$. then there exists a unitary $V_{E\to E}$ such that*

$$\|VU^{(1)} - U^{(2)}\|^2_\infty \leq \|T^{(1)} - T^{(2)}\|_\diamond \leq 2\|VU^{(1)} - U^{(2)}\|_\infty.$$

Proof. (of Lemma 1.) To simplify notation, we fix a key pair k and omit all key subscripts. We indicate input and output systems for channels, and support systems for operators, using subscripts as necessary. Let $(U)_{M\to CE}$, and $(W)_{C\to MF}$, be Stinespring dilation isometries of Enc, and Dec, respectively, i.e.,

$$\mathsf{Enc}(X) = \mathrm{Tr}_E UXU^\dagger \qquad \text{and} \qquad \mathsf{Dec}(Y) = \mathrm{Tr}_F WYW^\dagger.$$

Now, ε-approximate correctness together with Theorem 9 implies that there exists a pure state $|\phi\rangle_{EF}$ such that

$$\|(W)_{C\to MF}(U)_{M\to CE} - \mathbb{1}_M \otimes |\phi\rangle_{EF}\|_\infty \leq \sqrt{\varepsilon}.$$

Let $\hat{W}_{CR\to MF}$ be a unitary such that $\hat{W}_{CR\to MF}|0\rangle_R = W_{C\to MF}$. Such a unitary exists without loss of generality (If C does not divide MF, we can jus pick a bigger F). By the unitary invariance of the operator norm, we get

$$\left\|(U)_{M\to CE}\otimes|0\rangle_R - \left(\hat{W}^\dagger\right)_{MF\to CR}|\phi\rangle_{EF}\right\|_\infty \leq \sqrt{\varepsilon}.$$

Therefore we get

$$\left\|(U)_{M\to CE} - \langle 0|_R\left(\hat{W}^\dagger\right)_{MF\to CR}|\phi\rangle_{EF}\right\|_\infty \leq \sqrt{\varepsilon}. \tag{19}$$

by the submultiplicativity of the operator norm. Let

$$\langle 0|_R\left(\hat{W}^\dagger\right)_{MF\to CR}|\phi\rangle_{EF} = U^{(1)}_{M\to CE}D_M U^{(2)}_M$$

be the singular value decomposition of the second matrix. Equation 19 implies $\|\mathbb{1}_M - D_M\|_\infty \leq \sqrt{\varepsilon}$, and hence

$$\begin{aligned}&\left\|\langle 0|_R\left(\hat{W}^\dagger\right)_{MF\to CR}|\phi\rangle_{EF} - \langle 0|_R\left(\hat{W}^\dagger\left(U^{(2)}\right)^\dagger D^{-1}U^{(2)}\right)_{MF\to CR}|\phi\rangle_{EF}\right\|_\infty\\ \leq &\left\|U^{(1)}_{M\to CE}D_M U^{(2)}_M - U^{(1)}_{M\to CE}U^{(2)}_M\right\|_\infty \leq \sqrt{\varepsilon}\end{aligned} \tag{20}$$

and

$$\left(\tilde{V}\right)_{M\to CE} = \langle 0|_R\left(\hat{W}^\dagger\left(U^{(2)}\right)^\dagger D^{-1}U^{(2)}\right)_{MF\to CR}|\phi\rangle_{EF}$$

is an isometry. Combining Equations (19) and (20), we arrive at

$$\left\|(U)_{M\to CE} - \left(\tilde{V}\right)_{M\to CE}\right\|_\infty \leq 2\sqrt{\varepsilon}.$$

Defining $\tilde{\mathsf{Enc}}(X) = \mathrm{Tr}_E \tilde{V}X\tilde{V}^\dagger$, we conclude that

$$\left\|\mathsf{Enc} - \tilde{\mathsf{Enc}}\right\|_\diamond \leq \max_{\varrho_{MS}}\left\|U\varrho U^\dagger - \tilde{V}\varrho\tilde{V}^\dagger\right\|_\diamond \leq 4\sqrt{\varepsilon},$$

where the first inequality is due to the definition of the diamond norm and the fact that the trace norm is non-increasing under partial trace, and the second inequality is a double application of Hölder's inequality. $\tilde{\mathsf{Enc}}$ has the form we want, although this fact is still quite hidden. To show it, we define

$$(A)_{MF\to C} = \langle 0|_R\left(\hat{W}^\dagger\left(U^{(2)}\right)^\dagger D^{-1}U^{(2)}\right)_{MF\to CR},$$

so that we can write

$$\tilde{\mathsf{Enc}}(X) = (A)_{MF\to C}\left(X_M \otimes (\phi)_F\right)\left(A^\dagger\right)_{C\to MF}.$$

A is not, in general, an isometry. $\tilde{\mathsf{Enc}}$ is, however, trace preserving, implying $A^\dagger A = \mathbb{1}_M \otimes \kappa_F$ such that $\mathrm{Tr}(\kappa_F\phi_F) = 1$. Setting $T = F$, $\sigma_T = \sqrt{\kappa}_F\phi_F\sqrt{\kappa}_F$ and letting V be a completion of $A\kappa^{-1/2}$ to a unitary shows now that $\tilde{\mathsf{Enc}}$ has the desired form. For notational convenience, define $\delta = \varepsilon + 4\sqrt{\varepsilon}$.

To show the form of the decryption map, observe that, again by Theorem 9, there exists a possibly different Stinespring dilation isometry $W'_{C\to CT}$ for Dec and a quantum state $|\phi'_k\rangle_{ET}$ such that

$$\|W'_{C\to CT}V_{MT\to C}|\phi\rangle_{TE} - \mathbb{1}_M \otimes |\phi'\rangle_{ET}\|_\infty \leq \sqrt{\delta}$$

In particular, there is a unitary U_T such that $\|U_T|\phi\rangle_{TE} - |\phi'\rangle_{TE}\|_2 \leq \sqrt{\delta}$, so

$$\|W'_{C\to CT}V'_{MT\to C}|\phi'\rangle_{TE} - \mathbb{1}_M \otimes |\phi'\rangle_{TE}\|_\infty \leq 2\sqrt{\delta}, \tag{21}$$

with $V' = VU^\dagger$. Let $|\phi'\rangle_{TE} = \sum_i \sqrt{q}_i|\gamma_i\rangle_T \otimes |\eta_i\rangle_E$ be the Schmidt decomposition of $|\phi'\rangle$. Equation 21 implies that

$$q_i\,\|W'_{C\to CT}V'_{MT\to C}|\gamma_i\rangle_T - \mathbb{1}_M \otimes |\gamma_i\rangle_T\|_\infty \leq 2\sqrt{\delta}. \tag{22}$$

Let $P_T = \sum_{i:q_i\geq(4\delta)^{\frac{1}{6}}} |\gamma_i\rangle\langle\gamma_i|$. Then we get

$$\begin{aligned}
&\left\|W'_{C\to CT}V'_{MT\to C}P_T(V'^\dagger)_{C\to MT} - P_T(V'^\dagger)_{C\to MT}\right\|_\infty \\
\leq &\sum_{i:q_i\geq(4\delta)^{\frac{1}{6}}} \left\|W'_{C\to CT}V'_{MT\to C}|\gamma_i\rangle\langle\gamma_i|_T(V'^\dagger)_{C\to MT} - \mathbb{1}_M \otimes |\gamma_i\rangle\langle\gamma_i|_T(V'^\dagger)_{C\to MT}\right\|_\infty \\
\leq &(4\delta)^{\frac{1}{6}},
\end{aligned}$$

where we have used Equation 22 and the fact that $\mathrm{rk}P \leq \frac{1}{(4\delta)^{\frac{1}{6}}}$. Using that W' is a Stinespring dilation isometry for Dec shows the second claimed inequality. □

By viewing σ_{sk} in its eigenbasis, we see that the characterization can be enhanced so it only uses classical randomness.

Corollary 1. *Let $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ be a PKQE. Then for every keypair $k := (\mathsf{ek}, \mathsf{dk})$, there exists a probability distribution $p_k : \{0,1\}^t \to [0,1]$ and a family of quantum states $|\psi^{(k,r)}\rangle_T$ such that Enc_k is equivalent to the following algorithm: (i.) sample $r \in \{0,1\}^t$ according to p_k; (ii.) apply the map $X_M \mapsto V_k\left(X_M \otimes |\psi^{(k,r)}\rangle\langle\psi^{(k,r)}|_T\right)V_k^\dagger$. Here V_k and T are as in Lemma 1, and t is the number of qubits in T.*

Importantly, even if Enc_k is a polynomial-time algorithm, the functionally-equivalent algorithm provided by Corollary 1 may not be. We thus define the following.

Condition 1. *Let Π be a PKQE, and let p_k, $|\psi^{(k,r)}\rangle$ and V_k be as in Corollary 1. We say Π* satisfies Condition 1 *if there exist QPTs for (i.) sampling from p_k, (ii.) preparing $|\psi^{(k,r)}\rangle$, and (iii.) implementing V_k on inputs of the form $\varrho \otimes |\psi^{(k,r)}\rangle\langle\psi^{(k,r)}|$, and this holds for all but a negligible fraction of k and r.*

We stress that we are not aware of any known or trivially defined construction of PKQE where Condition 1 does not hold.

B Security Proofs

We now give several new security proofs for quantum encryption, and a proof of many-time security for hybrid quantum signcryption, i.e., Construction 1.

B.1 Strong Secrecy For Hybrid Public-Key Encryption

For PKQE schemes, the notions of QIND-CPA and QIND-CCA1 are defined in [2], where the adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ consists of two parts (pre-challenge and post-challenge), and is playing against a challenger $\mathcal{C}$, which is a fixed algorithm determined only by the security game and the scheme.

Theorem 10. *Let Π^{PKE} be an* IND-CPA *(resp.,* IND-CCA1*) PKE. Let further Π^{SKQE} be a* QIND *SKQE. Then $\Pi^{Hyb}_{PKQE}[\Pi^{PKE}, \Pi^{SKQE}]$ is a* QIND-CPA *(resp.,* QIND-CCA1*) PKQE.*

Proof. We will prove the CPA case; the CCA1 case is essentially identical. For a contradiction, let $\mathcal{A}$ be a QPT machine which wins the QIND-CPA game against $\Pi^{\mathsf{Hyb}}_{\mathsf{PKQE}}[\Pi^{\mathsf{PKE}}, \Pi^{\mathsf{SKQE}}]$ with probability $1/2 + \delta$ where δ is non-negligible in n. Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be the pre-challenge and post-challenge algorithms of $\mathcal{A}$. Define an adversary $\mathcal{A}'$ against Π^{PKE} as follows. The pre-challenge algorithm $\mathcal{A}'_1$ accepts ek, runs $\varrho \leftarrow \mathcal{A}_1(\mathsf{ek})$, and outputs $m \leftarrow \mathsf{KeyGen}^{\mathsf{SKQE}}(1^n)$. The post-challenge algorithm $\mathcal{A}'_2$, on input c, proceeds as follows: (i.) flip a coin $b' \xleftarrow{\$} \{0,1\}$, (ii.) if $b' = 0$, set $t = \mathcal{A}_2(c, \mathsf{Enc}_m^{\Pi^{\mathsf{SKQE}}}(\varrho))$; if $b' = 1$ set $t = \mathcal{A}_2(c, \mathsf{Enc}_m^{\Pi^{\mathsf{SKQE}}}(\mathbb{1}/d_M))$; (iii.) if $t = b'$ set $b_{\mathsf{out}} = 0$; otherwise set $b_{\mathsf{out}} \xleftarrow{\$} \{0,1\}$, and (iv.)output b_{out}.

Now consider $\mathcal{A}'$ in the IND-CPA game, and let b be the challenge bit. If $b = 0$, then the execution of $\mathcal{A}$ is exactly simulating the QIND-CPA game, and so $\Pr[b_{\mathsf{out}} = b] \geq 1/2 + \delta$. If $b = 1$, then the encryption of ϱ is done with an independent key; by QIND, $\Pr[t = b'] = 1/2$ which implies that $\Pr[b_{\mathsf{out}} = b] = 1/2$. We thus have that, overall, $\Pr[b_{\mathsf{out}} = b] \geq 1/2 \cdot (1/2 + \delta) + 1/2 \cdot 1/2 \geq 1/2 + \delta/2$, meaning that $\mathcal{A}'$ would break the IND-CPA security of Π^{PKE}. □

In this section, we adopt techniques from [3] in order to define quantum CCA2 security for PKQE. Recall that in the present work, for ease of exposition, we use integrity check techniques based on the simpler and more familiar DNS notion of plaintext authentication instead of the more complicated notion of QCA from [3]. This, in particular, has the advantage that the analysis in this work applies to all quantum encryption schemes (see [3] for details). On the other hand, this approach requires the repeated application of the Gentle Measurement Lemma [16], and can therefore only be fulfilled by schemes with plaintext spaces of large dimension, with a loss in the tightness of the resulting reductions. We discuss in Section 5.3 how to lift our results to the QCA-based setting.

We begin by defining a cDNS-style version of QIND-CCA2, which we call *weak* CCA2 *(*QIND-wCCA2*)*. Here, intuitively, the decryption oracle refuses to decrypt any ciphertext that decrypts to the challenge plaintext. While not providing the best possible quantum counterpart to IND-CCA2, this notion fits the simplified framework of this article, and the upgrade to QIND-CCA2 is straightforward.

As in the case of QIND-CCA2, we define QIND-wCCA2 in terms of the advantage gap of adversaries between two games. The first game is the same as QCCA2-Test from [3] for the symmetric-key case, except that the Enc oracle is replaced by the public key; in this game there are no restrictions on the use of $\mathsf{Dec}_{\mathsf{dk}}$ by $\mathcal{A}_2$. In the cheat-detection game QwCCA2-Fake, instead, the adversary is declared to cheat whenever he replays the challenge state (and declared to cheat anyway with probability 1/2 otherwise). The only difference in respect to QCCA2-Fake from [3] (beyond the replacement of the Enc oracle with the public key), is that the challenge-replay test is done cDNS-style (detecting plaintext replay) instead of cQCA-style (detecting ciphertext replay).

Experiment 6. The QCCA2-Test$(\Pi, \mathcal{A}, n)$ experiment:

1: $\mathcal{C}$ runs $(\mathsf{dk}, \mathsf{ek}) \leftarrow \mathsf{KeyGen}(1^n)$ and flips a coin $b \xleftarrow{\$} \{0,1\}$;
2: $\mathcal{A}_1$ receives ek and oracle access to $\mathsf{Dec}_{\mathsf{dk}}$;
3: $\mathcal{A}_1$ prepares a side register S, and sends to $\mathcal{C}$ a challenge register M;

4: $\mathcal{C}$ puts into C either $\mathsf{Enc}_{\mathsf{ek}}(M)$ (if $b = 0$) or $\mathsf{Enc}_{\mathsf{ek}}(\tau_M)$ (if $b = 1$);
5: $\mathcal{A}_2$ receives registers C and S and oracle access to $\mathsf{Dec}_{\mathsf{dk}}$;
6: $\mathcal{A}_2$ outputs a bit b'. **If** $b' = b$, **output** win; otherwise **output** rej.

Experiment 7. The QwCCA2-Fake$(\Pi, \mathcal{A}, n)$ experiment:
1: $\mathcal{C}$ runs $(\mathsf{dk}, \mathsf{ek}) \leftarrow \mathsf{KeyGen}(1^n)$;
2: $\mathcal{A}_1$ receives ek and oracle access to $\mathsf{Dec}_{\mathsf{dk}}$;
3: $\mathcal{A}_1$ prepares a side register S, and sends to $\mathcal{C}$ a challenge register M;
4: $\mathcal{C}$ discards M, prepares $|\phi^+\rangle_{M'M''}$, and stores M''; then $\mathcal{C}$ encrypts M' (using ek) and puts the resulting ciphertext into C';
5: $\mathcal{A}_2$ receives registers C' and S and oracle access to D_{dk}, where D_{dk} is defined as follows. On input a register C:
(1) $\mathcal{C}$ applies $\mathsf{Dec}_{\mathsf{dk}}$ to C, places result in M;
(2) $\mathcal{C}$ applies $\{\Pi^+, \mathbb{1} - \Pi^+\}$ to MM'';
(3) **if** the outcome is 1 **then**: abort and **output** cheat; **else return** M;
6: $\mathcal{C}$ draws a bit b at random. **If** $b = 1$, **output** cheat; if $b = 0$ **output** rej.

As usual, we define QIND-wCCA2 in terms of the advantage gap of adversaries between the two games.

Definition 16. *A PKQE Π is* QIND-wCCA2 *if, for all QPT adversaries $\mathcal{A}$,*

$$\Pr[\mathsf{QCCA2\text{-}Test}(\Pi, \mathcal{A}, n) \to \mathsf{win}] - \Pr[\mathsf{QwCCA2\text{-}Fake}(\Pi, \mathcal{A}, n) \to \mathsf{cheat}] \leq \mathrm{negl}(n)\,.$$

Clearly, because an adversary for QIND-wCCA2 is also an adversary for QIND-CCA1 with additional power and additional constraints on the scheme itself, we see that QIND-wCCA2 implies QIND-CCA1 and hence also QIND-CPA. Moreover, one can show that the hybrid construction behaves as expected:

Theorem 11. *Let Π^{PKE} be an* IND-CCA2 *PKE scheme, and let Π^{SKQE} be a* cDNS*-secure SKQE scheme. Then $\Pi^{Hyb}_{PKQE}[\Pi^{PKE}, \Pi^{SKQE}]$ is a* QIND-wCCA2*-secure PKQE scheme.*

Proof (Sketch). (For details, see Supplemental Section C.4.) We follow the same strategy as the proof of QIND-CPA security, i.e., Theorem 10. Specifically, given an adversary $\mathcal{A}$ which can distinguish the "test" and "fake" games, we will build an adversary $\mathcal{A}'$ against the CCA2-secure classical public-key scheme. This adversary proceeds as follows:
(1.) pass the input (1^n and the public key ek) to start $\mathcal{A}$;
(2.) simulate each decryption query in the obvious way: query the PKE Dec oracle on the classical part of the input, then decrypt the quantum part using the resultant plaintext as the SKQE key.
(3.) When $\mathcal{A}$ outputs the challenge ϱ, send a fresh one-time key k as our IND-CCA2 challenge, receive the challenge PKE ciphertext c back.
(4.) Now flip a coin b to decide whether, for the remainder of the game, we will simulate the "test" game ($b = 0$) or the "fake" game ($b = 1$) with $\mathcal{A}$. In the case "test", we will use the challenge ciphertext $(c, \mathsf{Enc}^{\mathsf{QCA}}_k(\varrho))$. It is clear that, in this case, we can faithfully simulate the rest of the game.
(5.) if $\mathcal{A}$ correctly guesses b, we output 0. If not, we output a fair coin.

We argue that $\mathcal{A}'$ wins the PKE IND-CCA2 game with non-negligible advantage over random guessing. First, if the PKE challenge bit is 0 (i.e., undisturbed challenge), then we are faithfully simulating $\mathcal{A}$ in either the "test" or the "fake" games. We will thus gain precisely the advantage of $\mathcal{A}$ over random guessing, in this case. It remains to check that, if the PKE challenge bit is 1 (i.e., discarded challenge), then $\mathcal{A}'$ does *no worse* than random guessing. This is done by reduction to cDNS security. Note that, in this case, the SKQE encryption provided to $\mathcal{A}$ is performed with a key k which is independent of all other random variables in the game (since c is an encryption of a random string, and not k.) If the adversary can nonetheless distinguish the "test" game from the "fake" game, then (by the definition of these games) this implies that he cannot be simulated by an "ignore or discard" channel demanded by the cDNS security definition. □

B.2 Plaintext QAE For Hybrid Private-Key Encryption

In this section we show that the hybrid construction for symmetric-key schemes, $\Pi^{\mathsf{Hyb}}_{\mathsf{SKQE}}$, can be used to lift the security of a SKQE scheme from cQCA to QAE, by combining it with a classical AE scheme. We will show this for a weak, plaintext-based version of QAE starting from a cDNS SKQE scheme. The general result follows from the discussion in Section 5.3.

We start from the QAE-Real and QAE-Ideal games defined in [3] for QAE, but in the latter game we replace the cQCA-style quantum ciphertext replay check by a cDNS-style plaintext replay check. We call the new gamewQAE-Ideal; a SKQE is now defined to be *weak quantum authenticated encryption* (wQAE) secure if no QPT adversary can reliably distinguish the two.

Experiment 8. The experiment $\mathsf{QAE\text{-}Real}(\Pi, \mathcal{A}, n)$:

1: $k \leftarrow \mathsf{KeyGen}(1^n)$;
2: **output** $\mathcal{A}^{\mathsf{Enc}_k, \mathsf{Dec}_k}(n)$.

Experiment 9. The experiment $\mathsf{wQAE\text{-}Ideal}(\Pi, \mathcal{A}, n)$:

1: $k \leftarrow \mathsf{KeyGen}(1^n)$;
2: define channel $E_{M\to C}$ as follows:
   (1) prepare $|\phi^+\rangle_{M'M''}$, store (M'', M) in a set $\mathcal{M}$;
   (2) apply Enc_k to M'; **return** result.
3: define channel $D_{C\to M}$ as follows:
   (1) apply Dec_k to C, place results in M';
   (2) **for each** $(M'', M) \in \mathcal{M}$ **do**:
   (3) apply $\{\Pi^+, \mathbb{1} - \Pi^+\}$ to $M'M''$;
   (4) **if** outcome is 0: **return** M;
   (5) **end for**
   (6) **return** $|\bot\rangle\langle\bot|$;
4: **output** $\mathcal{A}^{E,D}(1^n)$.

A SKQE is now defined to be *weak quantum authenticated encryption (*wQAE*)* secure if for all QPT adversaries, the two experiments are indistinguishable up to a negligible advantage.

Definition 17. *A SKQE Π is* wQAE *if for all QPT adversaries $\mathcal{A}$ it holds:*

$$|\Pr[\mathsf{QAE\text{-}Real}(\Pi, \mathcal{A}, n) \to \mathsf{real}] - \Pr[\mathsf{wQAE\text{-}Ideal}(\Pi, \mathcal{A}, n) \to \mathsf{real}]| \leq \mathrm{negl}(n).$$

Analogously we say that a classical scheme is wAE secure if it fulfills the classical restriction of the above definition (see [3]).

The following theorem shows that the hybrid construction for the symmetric-key case provides wQAE security when applied to a AE secure SKES and a cDNS secure SKQE. We provide a proof sketch below, the full proof can be found in the Supplemental Section C.6.

Theorem 12. *Let Π^{Cl} be an* AE *secure SKES and Π^{Qu} a* cDNS *secure SKQE. Then $\Pi^{Hyb}_{SKQE}[\Pi^{\mathsf{Cl}}, \Pi^{\mathsf{Qu}}]$ is a* wQAE *secure SKQE.*

Proof. (sketch.) We begin by defining a hybrid game Hybrid 0, modifying the encryption and decryption routine. For encryption, we replace the encrypted one-time key k' used for $\mathsf{Enc}^{\mathsf{Qu}}$ by an encryption of a freshly sampled key $k'' \leftarrow \mathsf{KeyGen}^{\mathsf{Qu}}$ and store (k', k'') in a database $\mathcal{S}$. For decryption, we decrypt the classical part of the ciphertext and check whether the result is equal to one of the k'' in $\mathcal{S}$. If not, return $\bot$, if so, use the corresponding k' to decrypt the quantum ciphertext and return the result. Suppose now that there exists an adversary $\mathcal{A}$ that can distinguish QAE-Real and Hybrid 0. Then we can build an adversary $\mathcal{A}$' that distinguishes the real and ideal worlds, AE-Real and AE-Ideal, in the real vs ideal characterization of AE by Shrimpton [15]. This is done by simulating the QAE-Real game played by $\mathcal{A}$, assuming we are in the

AE-Real world. It turns out that this results in $\mathcal{A}$' simulating Hybrid 0 when in the AE-Ideal world. Hence QAE-Real and Hybrid 0 are indistinguishable by the AE security of Π^{Cl}.

We continue to show that the experiments Hybrid 0 and wQAE-Ideal are indistinguishable due to the cDNS security of Π^{Qu}. We proceed using a standard hybrid method over the encryption queries of $\mathcal{A}$. We define Hybrid i to replace the quantum plaintext by half of a maximally entangled state for the first i queries, storing the other half together with the pair (k', k'') and the plaintext. For the remaining queries it behaves like Hybrid 0. The decryption routine behaves like in Hybrid 0, unless the presented one-time key matches one from the first i queries, then it performs the entanglement check measurement as in wQAE.

Suppose now that $\mathcal{A}$ can distinguish Hybrid i and Hybrid $(i+1)$. Then we can build a cDNS adversary $\mathcal{A}$', using the cDNS oracles for the i-th encryption call and corresponding decryption call. Note that this adversary exactly simulates Hybrid i, but its Broadbent-Waynewright simulator [8] simulates Hybrid $(i+1)$. (It is easy to see that if there exists a simulator as required by cDNS security, then the Broadbent-Waynewright simulator works as well.) Therefore we have built a successful adversary against the cDNS security of Π^{Qu}, a contradiction.

We conclude that Hybrid 0 and Hybrid q are indistinguishable. It remains to show that Hybrid q and wQAE-Ideal are indistinguishable. This follows by Gentle Measurement [16]: every time the entanglement test is applied to the current plaintext and an unrelated purification, the state remains unperturbed. □

B.3 Many-Time, Two-User Security For Hybrid Signcryption

We now prove outsider security of $\Pi^{\mathsf{Hyb}}_{\mathsf{QSC}}\left[\Pi^{\mathsf{SC}}, \Pi^{\mathsf{SKQE}}\right]$ (Construction 1) if Π^{SC} is an outsider-secure classical signcryption scheme, and Π^{SKQE} is a cDNS-secure SKQE. We will do this by showing that the (cDNS-style version of) outsider security of a QSC (resp., SC) scheme is equivalent to the wQAE (resp., wAE) security of a derived SKQE (resp., SKES) that is equivalent to the hybrid scheme $\Pi^{\mathsf{Hyb}}_{\mathsf{SKQE}}$. Then, we will use Theorem 12 to conclude the proof. We first define a generic way to obtain a SKQE from a QSC, or a SKES from a SC, in a way that preserves secrecy and unforgeability.

Construction 2. Let $\Pi = (\mathsf{KeyGen}, \mathsf{SigEnc}, \mathsf{VerDec})$ be a QSC (SC). We define the SKQE (SKES) $\Pi^\sharp = (\widehat{\mathsf{KeyGen}}, \mathsf{Enc}, \mathsf{Dec})$ in the following way:

1. $\widehat{\mathsf{KeyGen}}$ runs KeyGen twice to obtain $k = (\mathsf{sdk}, \mathsf{vek}, \mathsf{sdk}', \mathsf{vek}')$;
2. Enc_k runs $\mathsf{SigEnc}_{\mathsf{sdk},\mathsf{vek}'}$ and appends vek and vek′, i.e., $\mathsf{Enc}_k(X) = \left(\mathsf{SigEnc}_{\mathsf{sdk},\mathsf{vek}'}, \mathsf{vek}, \mathsf{vek}'\right)$;
3. Dec_k checks whether the second and third part of a ciphertext are equal to vek and vek′. If one of them is not, it outputs $\perp$, otherwise it runs $\mathsf{VerDec}_{\mathsf{sdk}',\mathsf{vek}}$ on the first part of the ciphertext.

Lemma 2. *Let $\Pi = (\mathsf{KeyGen}, \mathsf{SigEnc}, \mathsf{VerDec})$ be a QSC (resp., SC). Then $\Pi^\sharp$ is* wQAE *(resp.,* AE*) if and only if Π is many-time, outsider secure (resp., many-time outsider secure in the strong unforgeability sense).*

Proof. Let us first look at the classical case. If Π is an outsider secure SC (in the strong unforgeability sense), then it is by definition strongly unforgeable and IND-CCA2 secure. (Recall that, in the symmetric-key case, strong unforgeability is called integrity of ciphertexts.) Now note that these properties carry over to $\Pi^\sharp$ using the following argument. Given a forging or IND-CCA2 adversary $\mathcal{A}$ against Π, we can build an adversary $\mathcal{A}$' against $\Pi^\sharp$ by making one arbitrary extra query at the beginning to obtain the public keys, and then running $\mathcal{A}$, ignoring the second and third parts of all ciphertexts. Whenever $\mathcal{A}$ submits a ciphertext to the verified decryption oracle or as a forgery, $\mathcal{A}$' appends the public keys before submission. $\mathcal{A}$' has the same output distribution as $\mathcal{A}$, i.e. the winning probability of the two games is equal. On the other hand, let $\mathcal{A}$ now be an adversary against $\Pi^\sharp$. We construct an adversary $\mathcal{A}$' against Π in the following way. $\mathcal{A}$' runs $\mathcal{A}$, relaying any encryption queries to the signcryption oracle and appending the public keys to the result. The same is done in the challenge phase of the IND-CCA2 game, here the challenge plaintext submitted by $\mathcal{A}$ is forwarded as the challenge in the IND-CCA2 game for Π, and when the challenge ciphertext is received, the public keys are appended. For decryption queries, or forgeries, $\mathcal{A}$' checks whether the second and third

parts of the ciphertext are equal to the public keys before relaying it to the decryption oracle, or submitting it as a forgery, respectively. The observation that AE is equivalent to integrity of ciphertext and IND-CCA2 finishes the proof in the classical case.

Let us now turn to the quantum case. Here the proof is even simpler. The two adversary tranformations above simulate the Out-Real and Out-Ideal games for Π using oracles provided by the QAE-Real and wQAE-Ideal games for $\Pi^\sharp$, respectively, and vice versa. □

As a corollary, we obtain outsider security of the hybrid construction with an outsider-secure SC and a cDNS secure SKQE.

Corollary 2. *Let Π^{Cl} be a many-time, outsider secure SC and let Π^{Qu} be a* cDNS*-secure SKQE. Then $\Pi^{Hyb}_{QSC}[\Pi^{\mathsf{Cl}}, \Pi^{\mathsf{Qu}}]$ is a many-time outsider-secure QSC.*

Proof. This follows immediately from Theorem 12 and Lemma 2, together with the observation that:

$$\left(\Pi^{\mathsf{Hyb}}_{\mathsf{QSC}}[\Pi^{\mathsf{Cl}}, \Pi^{\mathsf{Qu}}]\right)^\sharp = \Pi^{\mathsf{Hyb}}_{\mathsf{QSC}}\left[\left(\Pi^{\mathsf{Cl}}\right)^\sharp, \Pi^{\mathsf{Qu}}\right]$$

□

C SUPPLEMENTARY MATERIAL

C.1 Definition Of DNS And QCA Security

Here we recall the formal definition of DNS one-time quantum authentication [10]. Given an attack map $\Lambda_{CB\to C\tilde{B}}$ on a scheme $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ (where the adversary holds B and $\tilde{B}$), we define the "effective attack map" by $\Lambda^{\Pi}_{MB\to M\tilde{B}} := \mathbb{E}_{k\leftarrow\mathsf{KeyGen}(1^n)}\left[\mathsf{Dec}_k \circ \Lambda \circ \mathsf{Enc}_k\right]$.

Definition 18 ([10]). *A SKQE $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ is* DNS*-authenticating if, for all CP-maps $\Lambda_{CB\to C\tilde{B}}$, there exist CP-maps $\Lambda^{\mathsf{acc}}_{B\to\tilde{B}}$ and $\Lambda^{\mathsf{rej}}_{B\to\tilde{B}}$ that sum to a TP map, such that:*

$$\left\|\Lambda^{\Pi}_{MB\to M\tilde{B}} - \left(\mathrm{id}_M \otimes \Lambda^{\mathsf{acc}}_{B\to\tilde{B}} + |\bot\rangle\langle\bot|_M \otimes \Lambda^{\mathsf{rej}}_{B\to\tilde{B}}\right)\right\|_\diamond \leq \mathrm{negl}(n)\,.$$

We also recall the stronger definition of QCA security [3].

Definition 19. *A SKQE $\Pi = (\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ is* ciphertext authenticating*, or* QCA*, if for all CP-maps $\Lambda_{CB\to C\tilde{B}}$, there exists a CP-map $\Lambda^{\mathsf{rej}}_{B\to\tilde{B}}$ such that:*

$$\left\|\Lambda^{\Pi}_{MB\to M\tilde{B}} - \left(\mathrm{id}_M \otimes \Lambda^{\mathsf{acc}}_{B\to\tilde{B}} + |\bot\rangle\langle\bot|_M \otimes \Lambda^{\mathsf{rej}}_{B\to\tilde{B}}\right)\right\|_\diamond \leq \mathrm{negl}(n),$$

and $\Lambda^{\mathsf{acc}}_{B\to\tilde{B}} + \Lambda^{\mathsf{rej}}_{B\to\tilde{B}}$ is TP. Here $\Lambda^{\mathsf{acc}}_{B\to\tilde{B}}$ is given by:

$$\Lambda^{\mathsf{acc}}_{B\to\tilde{B}}(Z_B) = \mathbb{E}_{k,r}\left[\langle\Phi_{k,r}|V_k^\dagger \Lambda\left(\mathsf{Enc}_{k;r}\left(\phi^+_{MM'}\otimes Z_B\right)\right)V_k|\Phi_{k,r}\rangle\right]$$

where $|\Phi_{k,r}\rangle = |\phi^+\rangle_{MM'}\otimes|\psi^{(k,r)}\rangle_T$.

C.2 Defining Quantum Signature Schemes

In this section we discuss how to properly define in a formal way what a quantum signature scheme should be. A first attempt would be to translate "quantumly" in the most natural way the usual definition of classical schemes; that is, trying to define the following.

Definition 20 (Quantum Signature - Wrong Definition). *A* quantum signature scheme (or QS) *with signing-verifying key space $\{\mathcal{S}_n\}_n \times \{\mathcal{V}_n\}_n$ (where $\{\mathcal{S}_n\}_n$ and $\{\mathcal{V}_n\}_n$ are family of spaces of bitstrings of size polynomial in n), message space $\mathcal{P}$ and signature space $\mathcal{C}$ (both being complex Hilbert spaces of fixed finite dimension) is a triple of QPT algorithms:*

1. *(key generation)* KeyGen *: on input 1^n, outputs signing-verifying key pair $(\mathsf{sk},\mathsf{vk})\in\mathcal{S}_n\times\mathcal{V}_n$, where we assume WLOG that* sk *includes* vk*, and* vk *includes n.*
2. *(Sign)* Sign*: on input a signing key $\mathsf{sk}\in\mathcal{S}_n$ and a quantum state $\mu\in\mathfrak{D}(\mathcal{P})$, outputs a quantum state ("quantum signature") $\sigma\in\mathfrak{D}(\mathcal{C})$; we write this as $\sigma\leftarrow\mathsf{Sign}_{\mathsf{sk}}(\mu)$.*
3. *(Verify)* Ver*: on input a verification key $\mathsf{vk}\in\mathcal{V}_n$, a quantum state $\mu\in\mathfrak{D}(\mathcal{P})$, and a quantum signature $\sigma\in\mathfrak{D}(\mathcal{C})$, outputs a verification bit $b\in\{\mathsf{acc},\mathsf{rej}\}$; we write this as $b\leftarrow\mathsf{Ver}_{\mathsf{vk}}(\mu,\sigma)$.*

Moreover, the following correctness property must hold:

$$\forall(\mathsf{sk},\mathsf{vk})\leftarrow\mathsf{KeyGen}(1^n),\forall\mu\in\mathfrak{D}(\mathcal{P})\implies\Pr[\mathsf{Ver}_{\mathsf{vk}}(\mu,\mathsf{Sign}_{\mathsf{sk}}(\mu))\to\mathsf{rej}]\leq\mathrm{negl}(n).$$

Looking at the above definition, it is clear that there is an issue: the Sign and Ver procedures consume the message. This is a problem in the quantum setting, given the no-cloning theorem. A natural "fix" would be to make sure that both the Sign and Ver algorithms output an unmodified copy of the message state. This would lead to the following.

Definition 21 (Quantum Signature - Another Wrong Definition). *A* quantum signature scheme (or QS) *with signing-verifying key space* $\{\mathcal{S}_n\}_n \times \{\mathcal{V}_n\}_n$ *(where* $\{\mathcal{S}_n\}_n$ *and* $\{\mathcal{V}_n\}_n$ *are family of spaces of bitstrings of size polynomial in* n*), message space* $\mathcal{P}$ *and signature space* $\mathcal{C}$ *(both being complex Hilbert spaces of fixed finite dimension) is a triple of QPT algorithms:*

1. *(key generation)* KeyGen *: on input* 1^n*, outputs signing-verifying key pair* $(\mathsf{sk}, \mathsf{vk}) \in \mathcal{S}_n \times \mathcal{V}_n$*, where we assume WLOG that* sk *includes* vk*, and* vk *includes* n*.*
2. *(Sign)* Sign*: on input a signing key* $\mathsf{sk} \in \mathcal{S}_n$ *and a quantum state* $\mu \in \mathfrak{D}(\mathcal{P})$*, outputs two quantum states: a quantum message* $\mu' \in \mathfrak{D}(\mathcal{P})$ *and a quantum signature* $\sigma \in \mathfrak{D}(\mathcal{C})$*; we write this as* $(\sigma, \mu') \leftarrow \mathsf{Sign}_{\mathsf{sk}}(\mu)$*.*
3. *(Verify)* Ver*: on input a verification key* $\mathsf{vk} \in \mathcal{V}_n$*, a quantum state* $\mu \in \mathfrak{D}(\mathcal{P})$*, and a quantum signature* $\sigma \in \mathfrak{D}(\mathcal{C})$*, outputs a quantum message* $\mu' \in \mathfrak{D}(\mathcal{P})$ *and a verification bit* $b \in \{\mathsf{acc}, \mathsf{rej}\}$*; we write this as* $(\mu', b) \leftarrow \mathsf{Ver}_{\mathsf{vk}}(\mu, \sigma)$*.*

Moreover, the following correctness properties must hold:

$$\forall (\mathsf{sk}, \mathsf{vk}) \leftarrow \mathsf{KeyGen}(1^n), \forall \mu \in \mathfrak{D}(\mathcal{P}) \implies$$

1. $\Pr[\mathsf{Ver}_{\mathsf{vk}}(\mathsf{Sign}_{\mathsf{sk}}(\mu)) \to (., \mathsf{rej})] \leq \mathrm{negl}(n)$
2. $\mathsf{Sign}_{\mathsf{sk}}(\mu) \to (\mu', .) \implies \|\mu - \mu'\|_\diamond \leq \mathrm{negl}(n)$
3. $\mathsf{Ver}_{\mathsf{vk}}(\mu, \sigma) \to (\mu', .) \implies \|\mu - \mu'\|_\diamond \leq \mathrm{negl}(n)$

It is easy to notice that this definition also has a lot of problems. First of all, condition 2. above is too strong, as it implies that the signature cannot be (in any noticeable way) correlated with the message state itself. This would make the goal of achieving any reasonable security notion hopeless. Condition 3. is also unnecessarily strong, as it requires that the original state is recovered even when verification fails. This finally leads to the following definition, which we adopted in Section 3.1.

Definition 22 (Restated Definition 5).
A **quantum signature scheme** *(or QS) is a triple of QPT algorithms:*

1. *(key generation)* $\mathsf{KeyGen}(1^n)$ *: output* $(\mathsf{sk}, \mathsf{vk}) \in \mathcal{K}_\mathsf{S} \times \mathcal{K}_\mathsf{V}$*.*
2. *(sign)* $\mathsf{Sign} : \mathcal{K}_\mathsf{S} \times \mathfrak{D}(\mathcal{H}_M) \to \mathfrak{D}(\mathcal{H}_C)$
3. *(verify)* $\mathsf{Ver} : \mathcal{K}_\mathsf{V} \times \mathfrak{D}(\mathcal{H}_C) \to \mathfrak{D}(\mathcal{H}_M \oplus |\perp\rangle)$ *.*

Moreover, the following correctness properties must hold:

$$\|\mathsf{Ver}_{\mathsf{vk}} \circ \mathsf{Sign}_{\mathsf{sk}} - \mathrm{id}_M \oplus 0_\perp\|_\diamond \leq \mathrm{negl}(n)$$

for all $(\mathsf{sk}, \mathsf{vk}) \leftarrow \mathsf{KeyGen}(1^n)$*.*

C.3 Proof Of Theorem 5

First, insider security follows from the $\mathsf{QIND\text{-}CPA}$ security of the hybrid PKQE construction (i.e. Theorem 6) and the fact that cDNS implies QIND.

Let $\mathcal{A}$ be an outside attacker that prepares a quantum state $\varrho_{MB}^{(\mathsf{vek}_R, \mathsf{vek}_S)}$ upon input $(\mathsf{vek}_R, \mathsf{vek}_S)$, and attacks the ciphertext it receives with an attack channel $\Lambda_{CB \to CB}$. We can assume without loss of generality that Λ is independent of vek_R and vek_S. This is because any dependency can be removed by redefining $\varrho_{MB}^{(\mathsf{vek}_R, \mathsf{vek}_S)}$ to include a copy of the public keys in the B register and having Λ read it from there. Here $C = C^{\mathsf{Qu}} C^{\mathsf{Cl}}$ is the ciphertext space of the scheme $\Pi_{\mathsf{QSC}}^{\mathsf{Hyb}}[\Pi^{\mathsf{SC}}, \Pi^{\mathsf{SKQE}}]$, where C^{Qu} and C^{Cl} are a quantum and a classical register, respectively. Let $\sigma_{CB}^{(\mathsf{vek}_R, \mathsf{vek}_S)} = \mathsf{SigEnc}_{\mathsf{vek}_R, \mathsf{sdk}_S}\left(\varrho_{MB}^{(\mathsf{vek}_R, \mathsf{vek}_S)}\right)$ be the ciphertext $\mathcal{A}$ receives, and let c^{Cl} be it's classical part. Let further $\hat{\sigma}_{CB}^{(\mathsf{vek}_R, \mathsf{vek}_S)} = \Lambda_{CB \to CB}\left(\sigma_{CB}^{(\mathsf{vek}_R, \mathsf{vek}_S)}\right)$, and let $\hat{c}^{\mathsf{Cl}}$ be the classical part[7] of $\hat{\sigma}_{CB}^{(\mathsf{vek}_R, \mathsf{vek}_S)}$. We first observe that if $\hat{c}^{\mathsf{Cl}} \neq c^{\mathsf{Cl}}$, $\mathsf{VerDec}_{\mathsf{sdk}_R, \mathsf{vek}_S}$ rejects with overwhelming

[7] For simplicity of exposition we assume here that C^{Cl} is a classical register that can only hold classical information

probability due to the unforgeability of the classical signcryption scheme. Let now $\tilde{c}^{\mathsf{Cl}} = \mathsf{SigEnc}^{\mathsf{SC}}_{\mathsf{vek}_R,\mathsf{sdk}_S}(\tilde{k})$ with $\tilde{k} \leftarrow \mathsf{KeyGen}^{\mathsf{SKQE}}(1^n)$. We note that $\sigma_{CB}^{(\mathsf{vek}_R,\mathsf{vek}_S)}$ and $\tilde{\sigma}_{CB}^{(\mathsf{vek}_R,\mathsf{vek}_S)}$ are indistinguishable for $\mathcal{A}$, where $\tilde{\sigma}_{CB}^{(\mathsf{vek}_R,\mathsf{vek}_S)}$ is equal to $\sigma_{CB}^{(\mathsf{vek}_R,\mathsf{vek}_S)}$ except that c^{Cl} is replaced with $\tilde{c}^{\mathsf{Cl}}$. In particular, $\tilde{c}^{\mathsf{Cl}}$ is unchanged with the same probability when $\Lambda_{CB\to CB}$ is applied to $\hat{\sigma}_{CB}^{(\mathsf{vek}_R,\mathsf{vek}_S)}$ as c^{Cl} is when $\Lambda_{CB\to CB}$ is applied to $\sigma_{CB}^{(\mathsf{vek}_R,\mathsf{vek}_S)}$, up to negligible difference. Suppose now that $\mathcal{A}$ breaks the outsider security of $\Pi^{\mathsf{Hyb}}_{\mathsf{QSC}}[\Pi^{\mathsf{SC}},\Pi^{\mathsf{SKQE}}]$, i.e. suppose that there is no simulator as required by Definition 9. Then it is easy to check that the attack map $\Lambda'_{C^{\mathsf{Qu}}B'}$ with side information register $B' = BC^{\mathsf{Cl}}K$ that

1. applies Λ_{CB}, and
2. discards register C^{Cl}

breaks cDNS security of Π^{SKQE}: When the input state

$$\varrho'_{MBC^{\mathsf{Cl}}K} = \varrho_{MB} \otimes \mathsf{SigEnc}_{\mathsf{sdk}_S,\mathsf{vek}_R}\left(\tilde{k}\right) \otimes (\mathsf{sdk}_S, \mathsf{vek}_S, \mathsf{sdk}_R, \mathsf{vek}_R)$$

with $(\mathsf{sdk}_i, \mathsf{vek}_i) \leftarrow \mathsf{KeyGen}^{\mathsf{SC}}$ for $i = S, R$ and $\tilde{k} \leftarrow \mathsf{KeyGen}^{\mathsf{SKQE}}$ is supplied, this attack simulates the attack $\mathcal{A}$ against $\Pi^{\mathsf{Hyb}}_{\mathsf{QSC}}[\Pi^{\mathsf{SC}},\Pi^{\mathsf{SKQE}}]$. □

C.4 Proof Of Theorem 11

Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be an adversary playing QCCA2-Test. We will go through a game hopping, where the winning probability of the adversary always increases (or decreases at most negligibly), up to an experiment which is equivalent to QwCCA2-Fake. First of all, we restate the two experiments explicitly.

Experiment 10. The QCCA2-Test$(\Pi, \mathcal{A}, n)$ experiment:

1: $\mathcal{C}$ runs $(\mathsf{dk}, \mathsf{ek}) \leftarrow \mathsf{KeyGen}(1^n)$ and flips a coin $b \xleftarrow{\$} \{0,1\}$;
2: $\mathcal{A}_1$ receives ek and oracle access to $\mathsf{Dec}_{\mathsf{dk}}$;
3: $\mathcal{A}_1$ prepares a side register S, and sends to $\mathcal{C}$ a challenge register M;
4: $\mathcal{C}$ puts into C either $\mathsf{Enc}_{\mathsf{ek}}(M)$ (if $b = 0$) or $\mathsf{Enc}_{\mathsf{ek}}(\tau_M)$ (if $b = 1$);
5: $\mathcal{A}_2$ receives registers C and S and oracle access to $\mathsf{Dec}_{\mathsf{dk}}$;
6: $\mathcal{A}_2$ outputs a bit b'. **If** $b' = b$, **output** win; otherwise **output** rej.

Experiment 11. The QwCCA2-Fake$(\Pi, \mathcal{A}, n)$ experiment:

1: $\mathcal{C}$ runs $(\mathsf{dk}, \mathsf{ek}) \leftarrow \mathsf{KeyGen}(1^n)$;
2: $\mathcal{A}_1$ receives ek and oracle access to $\mathsf{Dec}_{\mathsf{dk}}$;
3: $\mathcal{A}_1$ prepares a side register S, and sends to $\mathcal{C}$ a challenge register M;
4: $\mathcal{C}$ discards M, prepares $|\phi^+\rangle_{M'M''}$, and stores M''; then $\mathcal{C}$ encrypts M' (using ek) and puts the resulting ciphertext into C';
5: $\mathcal{A}_2$ receives registers C' and S and oracle access to D_{dk}, where D_{dk} is defined as follows. On input a register C:
(1) $\mathcal{C}$ applies $\mathsf{Dec}_{\mathsf{dk}}$ to C, places result in M;
(2) $\mathcal{C}$ applies $\{\Pi^+, \mathbb{1} - \Pi^+\}$ to MM'';
(3) **if** the outcome is 1 **then**: abort and **output** cheat; **else return** M;
6: $\mathcal{C}$ draws a bit b at random. **If** $b = 1$, **output** cheat; if $b = 0$ **output** rej.

We start now with the game-hopping.

Game 0: this is just QCCA2-Test.

Game 1: as Game 0, except that $\mathcal{C}$ will prepare two entangled registers $|\phi^+\rangle_{M'M''}$; then if $b = 1$ instead of encrypting the maximally mixed state τ_M, he will discard M and return to $\mathcal{A}$ the encryption of M' instead. Clearly the winning probability of $\mathcal{A}$ is unaffected by this, otherwise $\mathcal{A}$ could distinguish τ_M from $\phi^+_{M'}$ without access to $\phi^+_{M''}$.

Game 2: as Game 1, but instead of oracle access to $\mathsf{Dec}_{\mathsf{dk}}$, $\mathcal{A}_2$ gets access to a modified oracle D'_{dk} which, on input a ciphertext register C, does the following:

1. applies $\mathsf{Dec}_{\mathsf{dk}}$ to C, places result in M;
2. applies $\{\Pi^+, \mathbb{1} - \Pi^+\}$ to MM'';
3. **if** the outcome is 1 **then**: abort game and **output** win; **else return** M;

Notice that D'_{dk} is the same as D_{dk} in QwCCA2-Fake, except it aborts with win instead of cheat whenever it detects a challenge plaintext replay. Clearly, the probability of $\mathcal{A}$ of winning this game does not decrease in respect to the previous game, hence so far we have:

$$\Pr\left[\mathcal{A} \text{ wins Game 2}\right] \geq \Pr\left[\mathcal{A} \text{ wins QCCA2-Test}\right].$$

Game 3: as Game 2, but $\mathcal{C}$ "blinds" the classical part of the challenge ciphertext, replacing it with a different one, and D'_{dk} is modified in order to undo the blinding whenever queried on the fake ciphertext, making the substitution transparent to the adversary. The experiment looks as follows.

Experiment 12. The Game 3$(\Pi, \mathcal{A}, n)$ experiment:

1: $\mathcal{C}$ runs $(\mathsf{dk}, \mathsf{ek}) \leftarrow \mathsf{KeyGen}(1^n)$, prepares entangled registers $|\phi^+\rangle_{M'M''}$, and flips a bit $b \xleftarrow{\$} \{0,1\}$;
2: $\mathcal{A}_1$ receives ek and oracle access to $\mathsf{Dec}_{\mathsf{dk}}$;
3: $\mathcal{A}_1$ prepares a side register S, and sends to $\mathcal{C}$ a challenge register M;
4: **if** $b = 1$, **then**: $\mathcal{C}$ discards M and replaces it with M';
5: $\mathcal{C}$ encrypts M (using ek), puts the resulting ciphertext into C, and records the classical part (PKE ciphertext) t of the resulting ciphertext;
6: $\mathcal{C}$ generates a fresh one-time key $k' \leftarrow \mathsf{KeyGen}^{\mathsf{SKQE}}(1^n)$;
7: $\mathcal{C}$ encrypts $t' \leftarrow \mathsf{Enc}^{\mathsf{PKE}}_{\mathsf{ek}}(k')$, records t', and replaces the classical t with t' in the ciphertext register C;
8: $\mathcal{A}_2$ receives registers C and S and oracle access to D^t_{dk}, where D^t_{dk} is defined as follows. On input a register C:
   (1) measure the classical-subsystem part of C: **if** the outcome is t' **then**: replace the classical-subsystem part of C with $|t\rangle\langle t|$;
   (2) applies $\mathsf{Dec}_{\mathsf{dk}}$ to C, places result in M;
   (3) applies $\{\Pi^+, \mathbb{1} - \Pi^+\}$ to MM'';
   (4) **if** the outcome is 1 **then**: abort game and **output** win; **else return** M;
9: $\mathcal{A}_2$ outputs a bit b'. **If** $b' = b$, **output** win; otherwise **output** fail.

Notice that, since in the algorithm Dec of Π^{Hyb} the first step is to measure the classical part of the register C, the new measurement introduced by D^t_{dk} does not disturb $\mathcal{A}$'s behavior. Moreover, the substitution of t with t' is undetectable for $\mathcal{A}$, otherwise we could build a reduction $\mathcal{B}$ against the IND-CCA2 security of Π^{PKE}, in the following way:

1. $\mathcal{B}$ plays the IND-CCA2 game against $\mathcal{C}^{\mathsf{CCA2}}$ (for a secrept bit b), simulating $\mathcal{C}$ for $\mathcal{A}$ in the obvious way (forwarding ek, simulating Dec by decrypting classical keys with his own $\mathsf{Dec}^{\mathsf{PKE}}_{\mathsf{dk}}$ oracle and then doing the SKQE decryption himself).
2. During the challenge from $\mathcal{A}$, $\mathcal{B}$ generates a fresh k and uses it as a challenge to $\mathcal{C}^{\mathsf{CCA2}}$, receiving back a ciphertext t;
3. Now flip a coin $b' \xleftarrow{\$} \{0,1\}$ to decide whether for the rest of the game $\mathcal{B}$ will try to simulate Game 2 ($b' = 0$) or Game 3 ($b' = 0$) for $\mathcal{A}$. In the first case just return to $\mathcal{A}$ the correctly formed ciphertext register C, and simulate decryption queries on t-parts by defining their decryption as k. In the latter case instead, generate a fake encryption t', modify C accordingly before returning it to $\mathcal{A}$, and then simulate a "blinded" oracle D^t_{dk} by decrypting t'-part queries as k.
4. Finally, look at $\mathcal{A}$'s output. If adver guesses b' correctly, then $\mathcal{B}$ guesses $b = 0$, otherwise guesses b at random.

The reduction works because if $b = 0$, then we are simulating for $\mathcal{A}$ correctly either Game 2 or Game 3 (depending on b'). By assumption, in this case $\mathcal{A}$ should be able to guess correctly b' with non-negligible advantage, otherwise his guess will be unrelated to b. Therefore, we have:

$$\Pr[\mathcal{A} \text{ wins Game 3}] \geq \Pr[\mathcal{A} \text{ wins Game 2}] - \text{negl}.$$

Also notice that, at this point, $\mathcal{A}$ has no information whatsoever about the key k used to encrypt the quantum part of the challenge, because the encryption of k was blinded with another, unrelated one by the modified oracle D^t_{dk}. This observation will be important in the next game hop, because the definition of DNS security (and of the computational variant cDNS) argues about adversaries averaged over the secret key used.

Game 4: as Game 3, but this time $\mathcal{C}$ *always* replaces the challenge plaintext register M with the entangled half $\phi^+_{M'}$, regardless of b. We show that this replacement cannot be efficiently detected. Clearly, if $b = 1$ nothing changes for $\mathcal{A}$. If $b = 0$ instead, the quantum part of the challenge ciphertext received back by $\mathcal{A}$ is now an encryption (through Π^{SKQE}) of the register M' for a secret key k unknown to $\mathcal{A}$ by the reasoning made in the last game.

Then, consider any query to D^t_{dk} performed by $\mathcal{A}_2$. Such query can be seen as a convex combination of: a query state which does not depend on the challenge (and hence does not change the win probability), and the output of an attack map acting on the challenge ciphertext and an internal adversarial state S. By the cDNS security of Π^{SKQE}, such an attack map is (up to a negligible factor) a convex combination of two maps: the one which replaces the underlying plaintext with $|\bot\rangle\langle\bot|$ (and hence independent from the challenge plaintext, so again does not change the win probability), and the one which acts as the identity, and leaves the underlying plaintext untouched. However, this latter map would produce a query which is detected as $\phi^+_{M'}$ by D^t_{dk}, and would thus cause the adversary to win. Hence we have:

$$\Pr[\mathcal{A} \text{ wins Game 6}] \geq \Pr[\mathcal{A} \text{ wins Game 5}] - \text{negl}.$$

Game 5: as Game 4, but we remove the "t blinding" feature from the simulated oracle D^t_{dk}. That is, we replace D^t_{dk} back with D'_{dk}. This replacement is undetectable for the same reasoning as in Game 3.

Game 6: this is like Game 5, except that:

1. D'_{dk} is replaced by D_{dk} (i.e., aborts with cheat instead of win);
2. $\mathcal{C}$ does not flip the random bit b; and
3. on $\mathcal{A}$'s output, regardless of the outcome, we declare $\mathcal{A}$ to cheat or to rej with 50/50 probability.

Notice that this final game is equivalent to QwCCA2-Fake. In fact, by now $\mathcal{A}$'s output is completely unrelated to b. Combining all inequalities, we have shown:

$$\Pr[\mathcal{A} \text{ cheats in QwCCA2-Fake}] \geq \Pr[\mathcal{A} \text{ wins QCCA2-Test}] - \text{negl},$$

which proves finally that Π^{Hyb} is QIND-wCCA2 secure. □

C.5 "Full" Ciphertext (cQCA-Based) Version Of CCA2 Security For Public-Key Quantum Encryption

We show here the "correct" (strongest) formulation of quantum CCA2 security for the public-key setting by using the framework from Section B.1 and extending in a natural way the work done in [3]. We recall that, in particular, this formulation has the advantage that it closely matches "in spirit" the classical version of IND-CCA2 by disallowing replay of the challenge ciphertext in a strict quantum sense, where detection of a replay is done by exploiting QCA-based quantum ciphertext integrity. The drawback is that the use of such formulation in actual security proofs relies on a characterization of quantum encryption schemes given by a decomposition of the encryption/decryption procedure in terms of operators V_k and $\Pi_{k,r}$, where k (which in our case can be thought w.l.o.g. as the pair $(\mathsf{dk}, \mathsf{ek})$) is the underlying QCA authentication key (see [3] for details). One can prove that such decomposition always exists, but the necessary characterization here requires that this decomposition is efficient (Condition 1). There is currently no known counterexample of a quantum encryption scheme (either symmetric- or public-key) where this is not the case. Regardless, as shown in [3], this approach has many advantages, even if strictly speaking it can only be used on quantum encryption schemes which fulfill the characterization.

We start by recalling again the unrestricted quantum CCA2 experiment.

Experiment 13. The QCCA2-Test$(\Pi, \mathcal{A}, n)$ experiment:

1: $\mathcal{C}$ runs $(\mathsf{dk}, \mathsf{ek}) \leftarrow \mathsf{KeyGen}(1^n)$ and flips a coin $b \xleftarrow{\$} \{0,1\}$;
2: $\mathcal{A}_1$ receives ek and access to oracle $\mathsf{Dec}_{\mathsf{dk}}$;
3: $\mathcal{A}_1$ prepares a side (state) register S, and sends $\mathcal{C}$ a challenge register M;
4: $\mathcal{C}$ puts into C either $\mathsf{Enc}_{\mathsf{ek}}(M)$ (if $b = 0$) or $\mathsf{Enc}_{\mathsf{ek}}(\tau_M)$ (if $b = 1$);
5: $\mathcal{A}_2$ receives registers C and S and oracle access to $\mathsf{Dec}_{\mathsf{dk}}$;
6: $\mathcal{A}_2$ outputs a bit b'. **If** $b' = b$, **output** win; otherwise **output** fail.

Notice that in this game there are no restrictions on the use of $\mathsf{Dec}_{\mathsf{dk}}$ by $\mathcal{A}_2$. In particular, $\mathcal{A}_2$ is free to decrypt the challenge. In the second game, the challenge plaintext is replaced by half of a maximally entangled state, and $\mathcal{A}$ only gains an advantage over guessing if he cheats, i.e., if he tries to decrypt the challenge.

Experiment 14. The QCCA2-Fake$(\Pi, \mathcal{A}, n)$ experiment:

1: $\mathcal{C}$ runs $k \leftarrow \mathsf{KeyGen}(1^n)$;
2: $\mathcal{A}_1$ receives n and access to oracles Enc_k and Dec_k;
3: $\mathcal{A}_1$ prepares a side register S, and sends $\mathcal{C}$ a challenge register M;
4: $\mathcal{C}$ discards M, prepares $|\phi^+\rangle_{M'M''}$ and fresh randomness r, and stores (M'', r); then $\mathcal{C}$ encrypts the M' register and sends the resulting ciphertext C' to $\mathcal{A}_2$;
5: $\mathcal{A}_2$ receives registers C' and S and oracles Enc_k and D_k, where D_k is defined as follows. On input a register C:
(1) $\mathcal{C}$ applies $V_k^\dagger$ to C, places results in MT;
(2) $\mathcal{C}$ applies $\{P_T^{\sigma_k}, \mathbb{1} - P_T^{\sigma_k}\}$ to T;
(3) **if** outcome is 0 **then**:
(4) $\quad$ $\mathcal{C}$ applies $\{\Pi_{k,r}, \mathbb{1} - \Pi_{k,r}\}$ to T;
(5) $\quad$ **if** outcome is 0 **then**:
(6) $\quad\quad$ $\mathcal{C}$ applies $\{\Pi^+, \mathbb{1} - \Pi^+\}$ to MM'';
(7) $\quad\quad$ **if** outcome is 0: **output** cheat;
(8) $\quad$ **end if**
(9) **else**
(10) $\quad$ apply the default map for invalid ciphertexts, i.e., $\hat{D}_k$ to M.
(11) **end if**
(12) **return** M;
6: $\mathcal{C}$ draws a bit b at random. **If** $b = 1$, **output** cheat; if $b = 0$ **output** reject.

We now define QIND-CCA2 in terms of the advantage gap of adversaries between the above two games.[8]

Definition 23. *A PKQE Π is* QIND-CCA2 *if, for all QPT adversaries $\mathcal{A}$,*

$$\Pr[\mathsf{QCCA2\text{-}Test}(\Pi, \mathcal{A}, n) \to \mathsf{win}] - \Pr[\mathsf{QCCA2\text{-}Fake}(\Pi, \mathcal{A}, n) \to \mathsf{cheat}] \leq \mathrm{negl}(n)\,.$$

The omission of absolute values in the above is intentional. Indeed, an adversary can artificially inflate his cheating probability by querying the decryption oracle on the challenge and then ignoring the result. What he should not be able to do (against a secure scheme) is make his win probability larger than his cheating probability.

Separation and implication results follow the same strategy of [3], which we refer the reader to. We only recall here the following.

Proposition 3. *Let Π be a* QIND-CCA2 *secure PKQES. Then it is in particular also* QIND-CCA1 *and* QIND-CPA.

[8] The interface that the two games provide to the adversary differ slightly in that the adversary is not asked to output a bit in the end of the QCCA2-Fake game. This is not a problem as the games have the same interface until the second one terminates.

Moreover, following a similar strategy as in Theorem 11, we have the following.

Theorem 13. *Select the following:*

- *An* IND-CCA2 *PKE scheme* Π^{PKE}*;*
- *A* cQCA *secure* SKQE Π^{SKQE}*.*

Then $\Pi^{Hyb}_{PKQE}[\Pi^{PKE}, \Pi^{SKQE}]$ *is a* QIND-CCA2 *PKQE.*

C.6 Proof Of Theorem 12

We have to show that the games QAE-Ideal and wQAE-Real using the scheme $\Pi^{\mathsf{Hyb}}[\Pi^{\mathsf{Cl}}, \Pi^{\mathsf{Qu}}]$ are indistinguishable for any QPT adversary $\mathcal{A}$. We begin by defining a hybrid game in the following way:

Experiment 15. The experiment Hybrid 0:

1: $k \leftarrow \mathsf{KeyGen}(1^n)$;
2: define channel $E_{M\to C}$ as follows:
(1) run $\mathsf{Enc}^{\mathsf{Hyb}}_k$ on M, put result in registers $C^{\mathsf{Qu}}C'^{\mathsf{Cl}}$;
(2) sample $k'' \leftarrow \mathsf{KeyGen}^{\mathsf{Qu}}$ place $\mathsf{Enc}^{\mathsf{Cl}}_k(k'')$ in C^{Cl};
(3) store (C'^{Cl}, k'') in a set $\mathcal{S}$;
(4) **return** $(C^{\mathsf{Qu}}, C^{\mathsf{Cl}})$.
3: define channel $D_{C\to M}$ as follows:
(1) **for each** $(C'^{\mathsf{Cl}}, C''^{\mathsf{Cl}}) \in \mathcal{S}$ **do**:
(2) compute $\tilde{k}$ by running $\mathsf{Dec}^{\mathsf{Cl}}_k$ on C^{Cl};
(3) **if** $\tilde{k} = k''_j$: **then**
(4) apply $\mathsf{Dec}^{\mathsf{Hyb}}$ to $(C^{\mathsf{Qu}}, C'^{\mathsf{Cl}})$, place result in M **output** M;
(5) **end if**
(6) **end for**
(7) **return** $|\perp\rangle\langle\perp|$.
4: **output** $\mathcal{A}^{E,D}(1^n)$.

Suppose now first there exists an adversary $\mathcal{A}$ that can distinguish QAE-Real and Hybrid 0. Then we can build an adversary $\mathcal{A}$' that distinguishes the real and ideal worlds, AE-Real and AE-Ideal, in the real vs ideal characterization of AE by Shrimpton [15]. $\mathcal{A}$' runs $\mathcal{A}$, answering its queries in the following way. On an encryption query with register M, sample $k' \leftarrow \mathsf{KeyGen}^{\mathsf{Qu}}$ and encrypt M with $\mathsf{Enc}^{\mathsf{Qu}}_{k'}$ to obtain C^{Qu}. Now send k' to the encryption oracle and put the result into C^{Cl}. Return $(C^{\mathsf{Qu}}, C^{\mathsf{Cl}})$. For a decryption query with register $C = (C^{\mathsf{Qu}}, C^{\mathsf{Cl}})$, decrypt the contents of C^{Cl} using the decrytion oracle and use the result k' to decrypt the contents of C^{Qu} into M. Return M. It is easy to see that $\mathcal{A}$' uses the experiment AE-Real for Π^{Cl} to make $\mathcal{A}$ play QAE-Real, and the experiment AE-Ideal for Π^{Cl} to make $\mathcal{A}$ play Hybrid 0. Hence the indistinguishability between QAE-Real and Hybrid 0 follows from the AE security of Π^{Cl}.

We continue to show that the experiments Hybrid 0 and wQAE-Ideal are indistinguishable due to the cDNS security of Π^{Qu}. First observe that Π^{Cl} is in particular IND-CPA secure, and therefore randomized. More precisely it holds that for $c \leftarrow \mathsf{Enc}^{\mathsf{Cl}}_k$ and $c' \leftarrow \mathsf{Enc}^{\mathsf{Cl}}_k$, the probability of $c = c'$ is negligible. Let us therefore assume that the ciphertexts in $\mathcal{S}$ in Hybrid 0 are all distinct. It follows that for each decryption oracle call, there is at most one entry in $(C'^{\mathsf{Cl}}, C''^{\mathsf{Cl}}) \in \mathcal{S}$ such that $C^{\mathsf{Cl}} = C''^{\mathsf{Cl}}$. Let $\mathcal{A}$ be an adversary against the wQAE security of Π^{Hyb}, and assume that it makes q encryption queries. We build q hybrid experiments as follows:

Experiment 16. The experiment Hybrid i:

1: $k \leftarrow \mathsf{KeyGen}(1^n)$;
2: define a stateful channel $E_{M\to C}$ as follows:
(1) let $j-1$ be the number of times $E_{M\to C}$ has been called before;
(2) **If** $j \leq i$ move content of M to M'_j, prepare $|\phi^+\rangle_{M_j M''_j}$, store (M'_j, M''_j) in $\mathcal{M}$;
(3) run $\mathsf{Enc}^{\mathsf{Hyb}}_k$ on M, put result in registers $C^{\mathsf{Qu}}C'^{\mathsf{Cl}}_j$;

(4) sample $k''_j \leftarrow \mathsf{KeyGen}^{\mathsf{Qu}}$ place $\mathsf{Enc}^{\mathsf{Cl}}_k(k''_j)$ in C^{Cl};
(5) store $(C'^{\mathsf{Cl}}_j, k''_j)$ in a set $\mathcal{S}$;
(6) **return** $(C^{\mathsf{Qu}}, C^{\mathsf{Cl}})$.

3: define channel $D_{C\to M}$ as follows:
(1) **for each** $(C'^{\mathsf{Cl}}_i, C''^{\mathsf{Cl}}_i) \in \mathcal{S}$ **do**:
(2) compute $\tilde{k}$ by running $\mathsf{Dec}^{\mathsf{Cl}}_k$ on C^{Cl};
(3) **if** $\tilde{k} = k''_j$: **then**
(4) apply $\mathsf{Dec}^{\mathsf{Hyb}}$ to $(C^{\mathsf{Qu}}, C'^{\mathsf{Cl}})$, place result in M ;
(5) **if** $j \leq i$: **then**
(6) measure $|\phi^+\rangle\langle\phi^+|$ vs. $\mathbb{1} - |\phi^+\rangle\langle\phi^+|$ on MM''_j; **if** outcome is 0 swap M and M'_j, **else** prepare $|\perp\rangle\langle\perp|_M$;
(7) **end if**
(8) **output** M.
(9) **end if**
(10) **end for**
(11) **return** $|\perp\rangle\langle\perp|$.

4: **output** $\mathcal{A}^{E,D}(1^n)$.

Suppose now that $\mathcal{A}$ can distinguish $\mathsf{Hybrid}\ i$ and $\mathsf{Hybrid}\ (i+1)$. Then we can build a cDNS adversary $\mathcal{A}$' in the obvious way: $\mathcal{A}$' simulates $\mathsf{Hybrid}\ i$ with $\mathcal{A}$ until before the $(i+1)$-th query to prepare an input state. The attack map consists of continuing to simulate $\mathsf{Hybrid}\ i$ with $\mathcal{A}$ until there is a decryption query with a decrypted one-time key $\tilde{k}$ matching k''_{i+1}. If no such query occurs, prepare a random ciphertext and ignore the output of the decryption. Otherwise, continue to simulate $\mathsf{Hybrid}\ i$ with $\mathcal{A}$ to the end. Now note that this adversary exactly simulates $\mathsf{Hybrid}\ i$, but its Broadbent-Waynewright simulator [8] simulates[9] $\mathsf{Hybrid}\ (i+1)$.

Therefore we have built a successful adversary against the cDNS security of Π^{Qu}, a contradiction. We conclude that $\mathsf{Hybrid}\ 0$ and $\mathsf{Hybrid}\ q$ are indistinguishable. It remains to show that $\mathsf{Hybrid}\ q$ and $\mathsf{wQAE\text{-}Ideal}$ are indistinguishable. This follows by an application of the Gentle Measurement Lemma [16]: while $\mathsf{Hybrid}\ q$ first checks which register M'' should be used together with the decrypted quantum plaintext to measure $|\phi^+\rangle\langle\phi^+|$ vs. $\mathbb{1} - |\phi^+\rangle\langle\phi^+|$ by means of comparing $\tilde{k}$ and the k''s, $\mathsf{wQAE\text{-}Ideal}$ just tries them one by one. For the ones that do not fit, though, this measurement yields "not maximally entangled" with overwhelming probability, and the gentle measurement lemma implies that the decrypted plaintext is not disturbed (see also the discussion in Section 5.1). □

C.7 Explanatory example for Condition 1

Condition 1 is weaker than the analogous condition from [3] where it is required that V_k is efficiently implementable *on all inputs*. The results of that paper also hold for all schemes satisfying this weaker Condition 1. As discussed in [3], all known SKQE schemes satisfy Condition 1.

For example, let's look at why $\Pi^{\mathsf{Hyb}}_{\mathsf{PKQE}}$ satisfies Condition 1, provided that the same holds for the underlying SKQE. Let Π_Q be the SKQE and Π^{Cl} the classical PKE. Suppose Π^{Qu} satisfies

$$\mathsf{Enc}_{k;r}(X_M) = V_k \left(X_M \otimes |\psi^{(k,r)}\rangle\langle\psi^{(k,r)}|_T \right) V_k^\dagger.$$

for some efficiently-prepareable $|\psi^{(k,r)}\rangle$ and an V_k. Consider the following way of implementing encryption of the combined scheme $\Pi^{\mathsf{Hyb}}_{\mathsf{SKQE}}[\Pi^{\mathsf{Qu}}, \Pi^{\mathsf{Cl}}]$. Here r_C, k, r_Q denote uniformly random classical strings which are (respectively) the randomness for Π^{Cl} encryption, a key for Π^{Qu}, and randomness for Π^{Qu}.

$$\begin{aligned}\varrho &\longmapsto \varrho \otimes |k\rangle\langle k|_K \otimes |\mathsf{Enc}^C_{\mathsf{ek};r_C}(k)\rangle\langle \mathsf{Enc}^C_{\mathsf{ek};r_C}(k)| \otimes |\psi_{k,r_Q}\rangle\langle\psi_{k,r_Q}| \\ &\longmapsto \mathrm{Tr}_K \left(|k\rangle\langle k|_K \otimes |\mathsf{Enc}^C_{\mathsf{ek};r_C}(k)\rangle\langle \mathsf{Enc}^C_{\mathsf{ek};r_C}(k)| \otimes V_k(\varrho \otimes |\psi_{k,r_Q}\rangle\langle\psi_{k,r_Q}|)V_k^\dagger \right)\end{aligned}$$

[9] It is easy to see that if there exists a simulator as required by cDNS security, then the Broadbent-Waynewright simulator works as well.

Here, the first step attaches the auxiliary state to ϱ, and the second step implements a unitary V_k^{Π} (on relevant inputs.) We emphasize that both steps are efficiently implementable. Note that, despite the fact that (for these particular inputs) we implemented V_k^{Π} by tracing out a register, it is still the case that V_k^{Π} is a unitary operator. Indeed, V_k^{Π} is invertible because k can be extracted from $\mathsf{Enc}^C_{\mathsf{ek};r_C}(k)$; of course, to do this efficiently one needs dk.

D The multiuser outsider security experiments

In this section, we proved the definition of the multi-user outsider security experiments. In the following, we will describe the adversary as a two-stage adversary $(\mathcal{A}_1, \mathcal{A}_2$.

Experiment 17. The **multi-user real outsider experiment** $\mathsf{M\text{-}Out\text{-}Real}(\Pi, \mathcal{A}, n)$:

1: **for** every $\mathsf{ID}_P \in \mathcal{I}$ **do:** $(\mathsf{sdk}_P, \mathsf{vek}_P) \leftarrow \mathsf{KeyGen}(1^n, \mathsf{ID}_P)$.
2: $\mathcal{A}$ gets as input 1^n, $\mathcal{I}$, and the list of all public keys $\{\mathsf{vek}_P\}$ for all $\mathsf{ID}_P \in \mathcal{I}$;
3: $\mathcal{A}$ outputs two identities $S, R \in \mathcal{I}$;
4: $\mathcal{A}$ receives the list of all secret keys $\{\mathsf{sdk}_P\}$ for all $\mathsf{ID}_P \in \mathcal{I} \setminus \{S, R\}$;
5: **return** the output of $\mathcal{A}^{\mathsf{SigEnc}_{S,R},\mathsf{VerDec}_{S,R}}$.

Experiment 18. The **multi-user ideal outsider experiment** $\mathsf{M\text{-}Out\text{-}Ideal}(\Pi, \mathcal{A}, n)$:

1: **for** every $\mathsf{ID}_P \in \mathcal{I}$ **do:** $(\mathsf{sdk}_P, \mathsf{vek}_P) \leftarrow \mathsf{KeyGen}(1^n, \mathsf{ID}_P)$.
2: $\mathcal{A}$ gets as input 1^n, $\mathcal{I}$, and the list of all public keys $\{\mathsf{vek}_P\}$ for all $\mathsf{ID}_P \in \mathcal{I}$;
3: $\mathcal{A}$ outputs two identities $S, R \in \mathcal{I}$;
4: $\mathcal{A}$ receives the list of all secret keys $\{\mathsf{sdk}_P\}$ for all $\mathsf{ID}_P \in \mathcal{I} \setminus \{S, R\}$;
5: define channel $E_{M\to C}$:
   (1) prepare $|\phi^+\rangle_{M'M''}$, store (M'', M) in a set $\mathcal{M}$;
   (2) apply $\mathsf{SigEnc}_{S,R}$ to M'; **return** result.
6: define channel $D_{C\to M}$:
   (1) apply $\mathsf{VerDec}_{S,R}$ to C, place results in M';
   (2) **for each** $(M'', M) \in \mathcal{M}$ **do**:
   (3) apply $\{\Pi^+, \mathbb{1} - \Pi^+\}$ to $M'M''$;
   (4) **if** outcome is 0: **return** M;
   (5) **end for**
   (6) **return** $|\bot\rangle\langle\bot|$;
7: **return** the output of $\mathcal{A}^{E,D}$.