



# Unforgeable Quantum Encryption

Gorjan Alagic<sup>1,2(✉)</sup>, Tommaso Gagliardoni<sup>3</sup>, and Christian Majenz<sup>4,5</sup>

<sup>1</sup> Joint Center for Quantum Information and Computer Science,  
University of Maryland, College Park, MD, USA  
galagic@umd.edu

<sup>2</sup> National Institute of Standards and Technology,  
Gaithersburg, MD, USA

<sup>3</sup> IBM Research, Zurich, Switzerland  
tog@zurich.ibm.com

<sup>4</sup> Institute for Logic, Language and Computation, University of Amsterdam,  
Amsterdam, Netherlands  
c.majenz@uva.nl

<sup>5</sup> Centrum for Wiskunde en Informatica, Amsterdam, Netherlands

**Abstract.** We study the problem of encrypting and authenticating quantum data in the presence of adversaries making adaptive chosen plaintext and chosen ciphertext queries. Classically, security games use string copying and comparison to detect adversarial cheating in such scenarios. Quantumly, this approach would violate no-cloning. We develop new techniques to overcome this problem: we use entanglement to detect cheating, and rely on recent results for characterizing quantum encryption schemes. We give definitions for (i) ciphertext unforgeability, (ii) indistinguishability under adaptive chosen-ciphertext attack, and (iii) authenticated encryption. The restriction of each definition to the classical setting is at least as strong as the corresponding classical notion: (i) implies INT-CTXT, (ii) implies IND-CCA2, and (iii) implies AE. All of our new notions also imply QIND-CPA privacy. Combining one-time authentication and classical pseudorandomness, we construct symmetric-key quantum encryption schemes for each of these new security notions, and provide several separation examples. Along the way, we also give a new definition of one-time quantum authentication which, unlike all previous approaches, authenticates ciphertexts rather than plaintexts.

## 1 Introduction

Given the rapid development of quantum information processing, it is reasonable to conjecture that future communication networks will include at least some large-scale quantum computers and high-capacity quantum channels. What will secure communication look like on the resulting “quantum Internet”? For instance, how will we transmit quantum messages securely over a completely insecure channel? One approach is via interactive and information-theoretically secure methods, e.g., combining entanglement distillation with teleportation.

In this work, we will instead consider the non-interactive, highly efficient approach which dominates the current classical Internet. A natural goal here is to achieve, in the quantum setting, all the basic features that are enjoyed by classical encryption: (i) a single small key suffices for transmitting an essentially unlimited amount of data, (ii) these keys can be exchanged over public channels, and (iii) the security guarantees are as strong as possible. Previous work has shown how to achieve both (i) and (ii), but only for secrecy against chosen-plaintext and non-adaptive chosen-ciphertext attacks [3, 14]. Authentication or adaptive chosen-ciphertext security for such schemes has, as yet, not been considered. In fact, at the time of writing, there is not even a definition for *two-time quantum authentication*, much less for quantum analogues of EUF-CMA or IND-CCA2. The aim of this work is to address this problem.

The security definitions we seek do not yet exist due to a number of technical obstacles, all of which can be traced to quantum no-cloning and the destructiveness of quantum measurements. These obstacles make it difficult even just to formulate the basic security notion, much less to prove reductions or to construct secure schemes. In unforgeability, for example, no-cloning makes it impossible to record the adversary’s queries and check whether the final output is a fresh forgery. In adaptive chosen-ciphertext security, no-cloning makes it impossible to record the challenge ciphertext and ensure that the adversary doesn’t “cheat” by simply decrypting it (and thus win against any scheme). Moreover, due to the destructiveness of quantum measurement, it is unclear if one can *both* perform cheat-detection *and* answer non-cheating queries correctly.

In this work, we overcome these obstacles, and present the first definitions of multiple-query unforgeability and adaptive chosen-ciphertext indistinguishability for quantum encryption schemes, thereby solving a longstanding open problem [3, 12, 20]. While our definitions are inherently quantum in nature, we are able to show that they are in fact natural analogues of well-known classical security definitions, such as INT-CTXT and IND-CCA2. The strongest security notion we define is called *quantum authenticated encryption* (or QAE) and corresponds to the strongest form of security normally studied in the classical setting. A secret-key scheme satisfying QAE is unforgeable and indistinguishable even against adversaries that can make adaptive encryption and decryption queries.

In an effort to explore this new landscape, we prove several theorems which relate our new notions to each other and to established quantum and classical security definitions. We also show how to satisfy each of our new security notions with explicit, efficient constructions. In particular, we show that combining a post-quantum secure pseudorandom function with a unitary 2-design yields the strongest form of secret-key quantum encryption defined thus far, i.e., QAE.

**Related Work.** Computationally-secure quantum encryption has garnered significant interest in the past few years, beginning with basic security notions like QIND-CPA and QIND-CCA1 [3, 14], and then with more advanced concepts such as quantum fully-homomorphic encryption (QFHE) [14, 17]. For authentication, uncloneability, and non-malleability, the one-time setting has received

considerable attention (see, e.g., [5, 6, 15, 19, 21, 23, 24, 27].) We will make use of the authentication definition of [19], a characterization lemma of [5], and a simulation adversary of [15]. For classical notions of unforgeability and chosen-ciphertext security, see e.g. [25].

## 1.1 Our Approach

**The Problem.** We begin by outlining the technical difficulties in some further detail. Let us consider many-time authentication for symmetric-key encryption schemes first. In the classical setting, secure many-time authentication is defined in terms of *unforgeability*. A scheme is unforgeable if no adversary, even if granted the black-box power to authenticate with our secret key, can generate a fresh and properly authenticated message (i.e., a forgery). Translating this idea to the quantum setting presents immediate technical difficulties. First, no-cloning prevents us from recording the adversary’s previous queries. Second, even if the first problem is surmounted, the nature of measurement might make it difficult to reliably identify whether the adversary’s output is indeed fresh. For example, we might need many copies of the adversary’s query, as well as many copies of their final output.

A similar problem occurs for secrecy. The current state-of-the-art is the so-called QIND-CCA1 model. In this model, the transmitted state (the “challenge”) remains secret even to adversaries with the black-box power to both encrypt and non-adaptively decrypt with our secret key. Our experience in the classical world tells us that this model is too weak, because real-world adversaries can sometimes gain *adaptive* access to decryption (e.g., in WEP and early versions of SSL [8]). Classically, this is addressed using the so-called IND-CCA2 model, where the adversary is allowed adaptive decryption queries *but cannot use them on the challenge* (without this caveat, security becomes impossible). Here again, the quantum setting presents numerous technical difficulties: no-cloning prevents us from recording the challenge, and the nature of measurement makes it difficult to tell if the adversary is attempting to decrypt the challenge.

Recall that the strongest form of classical security, so-called “authenticated encryption” (or AE) is defined to be IND-CCA2 together with unforgeability of ciphertexts [25]. Achieving a comparable quantum notion thus seems to require solving all of the above problems.

Using classical intuition, one might attempt a solution as follows: consider only pure-state plaintexts, and demand that the final forgery is orthogonal to the previous queries (or, in CCA2, that decryption queries are orthogonal to the challenge). This may seem promising at first, but a closer look reveals numerous issues; for example: (i) quantum states are in general not pure, and may include side registers kept by the adversary, (ii) this idea charges the adversary with adhering to very strict demands, contrary to good theory practice, (iii) checking whether a particular adversary satisfies the demands cannot be done efficiently.

**A Promising Approach.** We now describe a more promising solution, beginning with unforgeability. We will express security in terms of the performance

of adversaries  $\mathcal{A}$  in two games: (1) F-Real, where  $\mathcal{A}$  gets oracle access to  $\text{Enc}_k$  and wins if he outputs *any* valid ciphertext, and (2) F-Cheat, where we attempt to ascertain if  $\mathcal{A}$  is cheating by feeding us an output of the oracle. How do we detect this kind of cheating? Recall that, even in the one-time setting, quantum authentication implies indistinguishability of ciphertexts. A consequence of this is that, whenever  $\mathcal{A}$  performs an encryption query on a certain plaintext state, we are free to respond with an encryption of a *different state* – for example, half of a maximally-entangled state. This will be our approach: we prepare an entangled pair  $|\phi^+\rangle_{MM'}$ , apply  $\text{Enc}_k$  to register  $M$ , give the resulting ciphertext register to  $\mathcal{A}$ , and keep  $M'$ . When the game ends, we decrypt the output of  $\mathcal{A}$  into a register  $O$ , and then perform the measurement  $\{\Pi_{\phi^+}, \mathbb{1} - \Pi_{\phi^+}\}$  on  $OM'$ . We then declare that  $\mathcal{A}$  is cheating if and only if the first outcome is recorded.

This idea can also be applied to the multiple-query setting. There, we respond to the  $j$ th query with an encryption of register  $M$  of  $|\phi^+\rangle_{MM_j}$ , and save  $M_j$ ; at the end of the game, we perform the aforementioned measurement on  $OM_j$  for all  $j$  and declare that  $\mathcal{A}$  cheated if any of them return the first outcome.

To define a quantum analogue of IND-CCA2, we can try a similar strategy. We again compare the performance of  $\mathcal{A}$  in two games: (1) C-Real, which is just like the classical IND-CCA2 game, except with no restrictions on  $\mathcal{A}$ 's use of the  $\text{Dec}_k$  oracle, and (2) C-Cheat, where we again attempt to detect cheating. In C-Cheat, when the adversary sends us the challenge plaintext, we discard it and respond with the ciphertext register of  $(\text{Enc}_k \otimes \mathbb{1}_{M'})|\phi^+\rangle_{MM'}$  instead, while keeping  $M'$  to ourselves. Whenever  $\mathcal{A}$  queries the decryption oracle, we first apply  $\text{Dec}_k$  and place the resulting plaintext in a register  $O$ . Then we apply the measurement  $\{\Pi_{\phi^+}, \mathbb{1} - \Pi_{\phi^+}\}$  to  $OM'$  to see if the adversary is cheating. If we get the first outcome, we declare that  $\mathcal{A}$  cheated.

The above ideas do lead to reasonable security definitions, which (at least partly) fulfill our original goals. However, they suffer from a number of drawbacks. First, repeated measurement of the plaintext requires the use of a so-called “gentle measurement lemma” [29], and thus can only apply to large plaintext spaces (e.g.,  $n^c$  qubits for  $c > 0$ ). Second, they only offer *plaintext authentication* and a kind-of *plaintext CCA security*; modification of ciphertexts (that does not also modify the underlying plaintext) cannot be detected. Our classical experience tells us that this is insufficient, and that we should demand impossibility of any ciphertext manipulation whatsoever. Addressing these problems is where many of our new technical contributions (in addition to the above ideas) are needed. While our actual approach will be different, and more sophisticated techniques are required, we will still follow the spirit of the idea outlined above.

## 1.2 Summary of Results

Recall that, in the setting of quantum data, copying is impossible and authentication implies encryption [9]. In particular, there is no direct quantum analogue of a MAC. As a result, the central objects of study in our work will be symmetric-key quantum encryption schemes, or SKQES for short, but our results on quantum CCA2 security carry over to the public-key setting as well.

**Quantum Ciphertext Authentication.** All previous definitions of authentication for quantum data allow manipulation of the ciphertext (see Sect. 2), thus only authenticating the plaintext state. In our first main contribution, we solve this problem, laying the necessary groundwork for our remaining results.

- We give a new definition: information-theoretic *quantum one-time ciphertext authentication* (QCA), inspired by ideas of [5, 15].
- We prove that QCA is a strengthening of “DNS”-authentication [19].

**Theorem 1 (informal).** *If a SKQES authenticates ciphertexts (QCA), then it also authenticates plaintexts (DNS); in particular, it satisfies secrecy (QIND).*

- We define computational-security (one-time) analogues: cQCA and cDNS.

**Quantum Unforgeability.** In this setting, the adversary is granted access to an encryption oracle, and must generate a valid “fresh” ciphertext.

- We give a new definition: *quantum unforgeability* (QUF), combining ideas of Sect. 1.1 and [5]. We also define a bounded-query analogue ( $t$ -QUF).
- We show that UF, the classical analogue of QUF, is remarkably strong.

**Theorem 2 (informal).** *For classical schemes, UF  $\iff$  AE.*

**Quantum Chosen-Ciphertext Security.** We address the longstanding problem of defining quantum security under adaptive chosen-ciphertext attack [3, 12, 20]; the state of the art was previously the non-adaptive QIND-CCA1 [3].

- We give a new definition: *quantum indistinguishability under adaptive chosen-ciphertext attack* (QIND-CCA2), using all of the aforementioned ideas.
- We relate QIND-CCA2 to existing security notions.

**Theorem 3 (informal).**

1. *For quantum schemes, QIND-CCA2  $\implies$  QIND-CCA1.*
2. *The classical analogue of QIND-CCA2 is equivalent to classical IND-CCA2.*

**Quantum Authenticated Encryption.** In our main contribution, we define a natural quantum analogue of the classical concept of *authenticated encryption* (AE). All previous quantum security notions lacked both unforgeability and adaptive chosen-ciphertext security.

- We give a new definition: *quantum authenticated encryption* (QAE), combining the ideas of Sect. 1.1, the notion of QCA, and a real/ideal approach [28].
- We give evidence that QAE is indeed the correct quantum analogue of AE.

**Theorem 4 (informal).**

1. *Unforgeability and secure authentication: QAE  $\implies$  QUF  $\wedge$  cQCA.*
2. *Chosen-ciphertext security: QAE  $\implies$  QIND-CCA2.*
3. *The classical analogue of QAE is equivalent to classical AE.*

The new notions and connections we develop are summarized in Fig. 1.

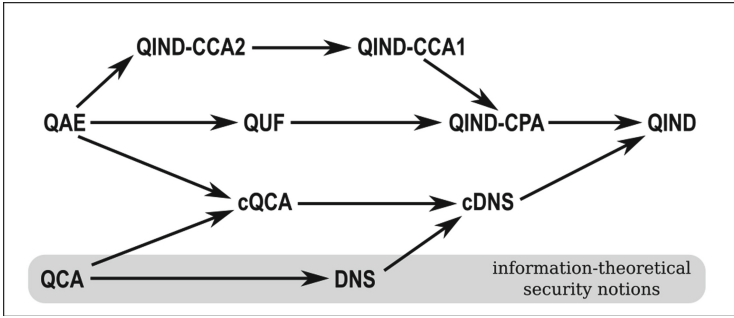


Fig. 1. Implications between quantum security notions

**Constructions and Separations.** Our new constructions combine a SKQES  $\Pi$  with a classical keyed function family  $f$  to build a new SKQES  $\Pi^f$ , as follows. In  $\Pi^f$ , key generation outputs a key for  $f$ ; to encrypt a state  $\rho$ , we generate a random  $r$  and output  $(r, \text{Enc}_{f_k(r)}^\Pi(\rho))$ . For example, if  $\Pi$  is the quantum one-time pad and  $f$  is a pqPRF (i.e., a post-quantum-secure pseudo-random function), then  $\Pi^f$  is the IND-CCA1-secure scheme from [3]. We will also need the standard one-time authentication scheme 2desTag, defined by  $\text{Enc}_k : \rho \mapsto C_k(\rho \otimes |0^n\rangle\langle 0^n|)C_k^\dagger$  where  $C$  is an (exact or approximate) unitary two-design.

**Theorem 5 (informal).** *Let  $\Pi$  be a 2desTag scheme, let  $f$  be a pqPRF, and let  $g$  be a  $t$ -wise independent classical function family. Then*

1.  $\Pi$  is one-time ciphertext authenticating (QCA).
2.  $\Pi^g$  is  $t$ -time quantum unforgeable ( $t$ -QUF).
3.  $\Pi^f$  satisfies quantum authenticated encryption (QAE); in particular, it is quantum unforgeable (QUF) and chosen-ciphertext secure (QIND-CCA2).

**Theorem 6 (informal).**

1. There exists an SKQES which is QIND-CCA1 but not QIND-CCA2.
2. There exists an SKQES which is QIND-CCA2 but not QAE.

**Our Choice of Primitives.** The reader may wonder why our constructions do not need “quantum-oracle-secure” primitives (e.g., QPRFs for unforgeability and  $2t$ -wise independence for  $t$ -time security, as in the quantum-secure classical setting of [11]). In our work, the classical portion of the ciphertext is generated by honest parties during encryption, and measured during decryption. As a result, oracle access to  $\text{Enc}_k$  and  $\text{Dec}_k$  (as CPTP maps) never grants quantum oracle access to the underlying classical primitive. Of course, one could grant the adversary more powerful oracles that do grant this kind of access, and then quantum-oracle-secure primitives (such as QPRFs) would indeed be required.

**A Remark on Applicability.** While all of our definitions apply to arbitrary quantum encryption schemes, security reductions sometimes require the following additional condition. As discussed in Sect. 3, all quantum encryption algorithms can be characterized as (1) drawing a random pure state from a probability distribution, (2) attaching it to the plaintext, and (3) applying a unitary operator. For the implication  $\text{QAE} \Rightarrow \text{cQCA}$  of Theorem 4 to hold, it is required that (1), (2) and (3) are efficiently implementable. This condition holds for all schemes known to us. However, it is *in principle* possible that there are schemes for which  $\text{Enc}_k$  is efficiently implementable, but the particular implementation “(1), then (2), then (3)” is not. We leave this as an open problem.

## 2 Preliminaries

**Basic Notation and Conventions.** In the rest of this work, we use “classical” to denote “non-quantum”, “iff” for “if and only if”, and  $n$  to denote the security parameter. A function  $\varepsilon(n)$  is negligible (denoted  $\varepsilon(n) \leq \text{negl}(n)$ ) if it is asymptotically smaller than  $1/p(n)$  for every polynomial function  $p$ . The notation  $x \stackrel{\$}{\leftarrow} X$  means that  $x$  is a sample from the uniform distribution over the set  $X$ . By “PPT” we mean a polynomial-time uniform family of probabilistic circuits, and by “QPT” we mean a polynomial-time uniform family of quantum circuits. We will frequently give such algorithms names like “adversary” or “challenger,” but this is only to help remember the role of the algorithm.

For notation and conventions regarding quantum information, we refer the reader to [26]. We recall a few basics here. We denote by  $\mathcal{H}_M$  a complex Hilbert space with label  $M$  and finite dimension  $\dim M$ . We use the standard bra-ket notation to work with pure states  $|\varphi\rangle \in \mathcal{H}_M$ . The class of positive, Hermitian, trace-one linear operators on  $\mathcal{H}_M$  is denoted by  $\mathfrak{D}(\mathcal{H}_M)$ . A *quantum register* is a physical system whose set of valid states is  $\mathfrak{D}(\mathcal{H}_M)$ ; in this case we label by  $M$  the register itself. We reserve the notation  $\tau_M$  for the maximally mixed state (i.e., uniform classical distribution)  $\mathbb{1}/\dim M$  on  $M$ .

In a typical cryptographic scenario, 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. This family is parameterized by  $n$ , which is typically also the security parameter. We will consider completely positive (CP), trace-preserving (TP) maps (i.e., quantum channels) when describing quantum algorithms. To indicate that  $\Phi$  is a channel from register  $A$  to  $B$ , we will write  $\Phi_{A \rightarrow B}$ . 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  $\Psi_{BC \rightarrow D}$ , we will write  $\Psi \circ \Phi$  to denote the map  $\Psi \circ (\Phi \otimes \mathbb{1}_C)$  from  $AC$  to  $D$ .

The support of a quantum state  $\varrho$  is its cokernel (as a linear operator). Equivalently, this is the span of the pure states making up any decomposition of  $\varrho$  as a convex combination of pure states. We will denote the orthogonal projection operator onto this subspace by  $P_\varrho$ . The two-outcome projective measurement (to test if a state has the same or different support as  $\varrho$ ) is then  $\{P_\varrho, \mathbb{1} - P_\varrho\}$ .

Next, we single out some unitary operators that will appear frequently. First, the group of  $n$ -qubit operators generated by Paulis  $I, X, Y, Z$  (applied to individual qubits) is a well-known *unitary one-design*. The Clifford group on  $n$  qubits is defined to be the normalizer of the Pauli group inside the unitary group. It can also be seen as the group generated by the gate set  $(H, P, CNOT)$  [22]; it is also a *unitary two-design* [16].

A *unitary  $t$ -design* (for a fixed  $t$ ) is an infinite collection  $\mathcal{U} = \{\mathcal{U}^{(n)} : n \in \mathbb{N}\}$ , where  $\mathcal{U}^{(n)}$  forms an  $n$ -qubit unitary  $t$ -design in the standard sense, i.e.,

$$\frac{1}{|\mathcal{U}^{(n)}|} \sum_{U \in \mathcal{U}^{(n)}} U^{\otimes t} X (U^\dagger)^{\otimes t} = \int U^{\otimes t} X (U^\dagger)^{\otimes t} dU. \tag{1}$$

In the above, the integral is taken over the  $n$ -qubit unitary group according to the Haar measure. We assume that there is an explicit polynomial function  $m(n)$  and a deterministic polynomial-time algorithm which, given  $1^n$  and  $k \xleftarrow{\$} \{0, 1\}^{m(n)}$ , produces a circuit for a unitary operator  $U_{k,n}$  which is distributed uniformly at random in  $\mathcal{U}^{(n)}$ . We will not refer to this algorithm explicitly and will simply write  $\{U_{k,n} : k \in \{0, 1\}^{m(n)}\}$  for the resulting distribution on unitary operators; we will also frequently suppress one index and write  $U_k$  when  $n$  is clear from context. We refer to the polynomial  $m$  as the *key length* of the  $t$ -design. Standard examples are: (i) the Pauli one-design (where we apply  $X^a Z^b$  to each qubit for random  $a, b \in \{0, 1\}$ ) is a unitary one-design on  $n$  qubits with key length  $2n$ ; (ii) the Clifford group (where we apply a uniformly random element of the  $n$ -qubit Clifford group, efficiently generated via the Gottesman-Knill theorem [1]) is a unitary 3-design, and therefore in particular a unitary 2-design, on  $n$  qubits with key length  $O(n^2)$ ; (iii) random  $\text{poly}(t, n)$ -size quantum circuits, randomly generated from a universal gate set, are approximate  $t$ -designs on  $n$  qubits [13].

In this work, we will only require one-designs and two-designs, and we will assume for simplicity that the designs are exact. While approximate designs would also suffice, some additional (but straightforward) analysis would be required.

**Quantum Encryption.** We will follow the conventions set in [3]; the exception is that decryption can reject by outputting a special symbol  $\perp$ .

**Definition 1.** A symmetric-key quantum encryption scheme (or SKQES) is a triple of QPT algorithms:

1. (key generation)<sup>1</sup>  $\text{KeyGen} : \text{on input } 1^n, \text{ outputs } k \xleftarrow{\$} \mathcal{K}$
2. (encryption)  $\text{Enc} : \mathcal{K} \times \mathcal{D}(\mathcal{H}_M) \rightarrow \mathcal{D}(\mathcal{H}_C)$
3. (decryption)  $\text{Dec} : \mathcal{K} \times \mathcal{D}(\mathcal{H}_C) \rightarrow \mathcal{D}(\mathcal{H}_M \oplus |\perp\rangle\langle\perp|)$

such that  $\|\text{Dec}_k \circ \text{Enc}_k - \mathbf{1}_M \oplus 0_\perp\|_\diamond \leq \text{negl}(n)$  for all  $k \in \text{supp KeyGen}(1^n)$ .

---

<sup>1</sup> A more general definition uses arbitrary key generation algorithms. We assume a uniform key in this paper for technical and notational convenience.



It is implicit that the key space  $\mathcal{K}$  is classical and of size  $\text{poly}(n)$ ; likewise, the registers  $C$  and  $M$  are quantum registers of at most  $\text{poly}(n)$  qubits. We will only consider SKQES of *fixed-length*, meaning that the number of qubits in  $M$  is a fixed function of the security parameter  $n$ . We assume that honest parties will apply the measurement  $\{II_{\perp}, \mathbb{1} - II_{\perp}\}$  (where  $II_{\perp} = |\perp\rangle\langle\perp|$ ) immediately after decryption. This allows us to write, e.g.,  $\text{Dec}_k(\varrho) \neq \perp$  to mean that decryption (followed by this measurement) successfully produced a valid plaintext.

We will often combine quantum schemes with classical (keyed) function families. A keyed function family consists of functions  $f : \{0, 1\}^{p(n)} \times \{0, 1\}^{q(n)} \rightarrow \{0, 1\}^{s(n)}$  where  $p, q, s$  are polynomials in  $n$ . In typical usage, we sample a key  $k \xleftarrow{\$} \{0, 1\}^{p(n)}$  and then consider the restricted function  $f_k : \{0, 1\}^{q(n)} \rightarrow \{0, 1\}^{s(n)}$  defined by  $f_k(x) = f(k, x)$ . All keyed function families are assumed to be computable by a deterministic polynomial-time uniform classical algorithm.

**Definition 2.** Let  $\Pi = (\text{KeyGen}^{\Pi}, \text{Enc}^{\Pi}, \text{Dec}^{\Pi})$  be a SKQES, and  $f : \{0, 1\}^{p(n)} \times \{0, 1\}^{q(n)} \rightarrow \{0, 1\}^{s(n)}$  a classical keyed function family. Define a new SKQES  $\Pi^f = (\text{KeyGen}, \text{Enc}, \text{Dec})$  as follows:

1.  $\text{KeyGen} : \text{on input } 1^n, \text{ outputs } k \xleftarrow{\$} \{0, 1\}^{p(n)}$ ;
2.  $\text{Enc}_k : \text{on input } \varrho, \text{ outputs } |r\rangle\langle r| \otimes \text{Enc}_{f_k(r)}^{\Pi}(\varrho), \text{ where } r \xleftarrow{\$} \{0, 1\}^{q(n)}$ ;
3.  $\text{Dec}_k : |s\rangle\langle s| \otimes \sigma \mapsto \text{Dec}_{f_k(s)}^{\Pi}(\sigma)$ .

We extend  $\text{Dec}_k$  to arbitrary inputs by postulating that it begins by measuring the first register in the computational basis. Note that  $\Pi^f$  has plaintext length  $t(s(n))$  where  $t(\cdot)$  is the plaintext length of  $\Pi$  as a function of  $\Pi$ 's key length. This construction can be extended to schemes  $\Pi$  with a non-uniform key by using the output of the keyed function family as a random tape for  $\text{KeyGen}^{\Pi}$ .

**Quantum Secrecy.** The literature contains a number of information-theoretic definitions of quantum secrecy (see, e.g., [3, 6, 7, 14]). It is well-known that a unitary one-design (e.g., the Pauli group) is an information-theoretically secret scheme. In this work, however, we focus on the computational setting [3, 14].

**Definition 3 (QIND).** A SKQES  $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$  has *indistinguishable encryptions* (or is QIND) if for every QPT adversary  $\mathcal{A} = (\mathcal{M}, \mathcal{D})$  we have:

$$\left| \Pr[\mathcal{D}\{(\text{Enc}_k \otimes \mathbb{1}_E)\varrho_{ME}\} = 1] - \Pr[\mathcal{D}\{(\text{Enc}_k \otimes \mathbb{1}_E)(|0\rangle\langle 0|_M \otimes \varrho_E)\} = 1] \right| \leq \text{negl}(n),$$

where  $\varrho_{ME} \leftarrow \mathcal{M}(1^n)$ ,  $\varrho_E = \text{Tr}_M(\varrho_{ME})$ , and the probabilities are taken over  $k \leftarrow \text{KeyGen}(1^n)$  and the coins and measurements of  $\text{Enc}, \mathcal{M}, \mathcal{D}$ . We also define:

- QIND-CPA: In addition to the above,  $\mathcal{M}$  and  $\mathcal{D}$  have oracle access to  $\text{Enc}_k$ .
- QIND-CCA1: In addition to QIND-CPA,  $\mathcal{M}$  has oracle access to  $\text{Dec}_k$ .

Recall that a pqPRF (post-quantum pseudorandom function) is a classical, deterministic, efficiently computable keyed function family  $\{f_k\}_k$  which appears random to QPT algorithms with classical oracle access to  $f_k$  for uniformly random  $k$ . The strongest notion (QIND-CCA1) is satisfied by  $\Pi^f$  where  $\Pi$  is a one-design and  $f$  is a pqPRF [3]. We let  $\text{1des}^{\text{PRF}}$  denote such schemes.

**One-Time Authentication.** We recall quantum authentication as defined by Dupuis et al. [19], and adapt it to our conventions. Given an attack map  $A_{CB \rightarrow C\bar{B}}$  on a scheme  $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$  (where the adversary holds  $B$  and  $\bar{B}$ ), we define the “averaged effective plaintext map” (or just “effective map”) as follows.

$$A_{MB \rightarrow M\bar{B}}^{\Pi} := \mathbb{E}_{k \leftarrow \text{KeyGen}(1^n)} [\text{Dec}_k \circ A \circ \text{Enc}_k] .$$

We then require that, conditioned on acceptance, this map is the identity on  $M$ .

**Definition 4 ([19]).** A SKQES  $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$  is DNS-authenticating if, for any CP-map  $A_{CB \rightarrow C\bar{B}}$ , there exist CP-maps  $A_{B \rightarrow \bar{B}}^{\text{acc}}$  and  $A_{B \rightarrow \bar{B}}^{\text{rej}}$  that sum to a TP map, such that

$$\left\| A_{MB \rightarrow M\bar{B}}^{\Pi} - \left( \text{id}_M \otimes A_{B \rightarrow \bar{B}}^{\text{acc}} + |\perp\rangle\langle\perp|_M \otimes A_{B \rightarrow \bar{B}}^{\text{rej}} \right) \right\|_{\diamond} \leq \text{negl}(n) . \quad (2)$$

An important observation is that this definition only provides for authentication of the plaintext state. To see that this cannot be “ciphertext authentication,” simply take a scheme which is DNS and change it so that (i) an extra bit is added to the ciphertext during encryption, and (ii) that same bit is ignored during decryption. The resulting scheme still satisfies DNS, but the adversary can clearly forge ciphertexts by flipping the extra bit. A perhaps more compelling example just adds encoding (in some QEC code) after encryption, and decoding prior to decryption. The adversary is then free to modify ciphertexts with correctable errors without violating DNS. We remark that, in this respect, the recent strengthening of DNS due to Garg et al. [21] is no different: a scheme secure according to this stronger notion of authentication can be modified in the same way without losing security.

Next, we recall a standard one-time authentication scheme. We encrypt by appending  $n$  “tag” qubits in the fixed state  $|0\rangle$  and then applying a random element of a 2-design. Decryption first undoes the 2-design, then outputs the plaintext iff all tag qubits measure to 0; otherwise it outputs  $\perp$ .

**Scheme 1.** The scheme family 2desTag is defined as follows. Select a unitary 2-design  $\mathcal{U}$  with key length  $m(\cdot)$ , and define algorithms:

1. **KeyGen:** on input  $1^n$ , output  $k \xleftarrow{\$} \{0, 1\}^{m(2n)}$ ;
2. **Enc<sub>k</sub>:** on input  $\varrho_M$ , output  $U_k(\varrho_M \otimes |0^n\rangle\langle 0^n|_T)U_k^\dagger$
3. **Dec<sub>k</sub>:** on input  $\sigma_{MT}$ , output

$$\langle 0^n|_T U_k^\dagger \sigma_{MT} U_k |0^n\rangle_T + \text{Tr}[(\mathbb{1} - |0^n\rangle\langle 0^n|_T)U_k^\dagger \sigma_{MT} U_k] |\perp\rangle\langle\perp|_M .$$

We chose 2desTag to have plaintext and tag length  $n$ . It is well-known that, for plaintexts of at most polynomial length and tags of length at least  $n^c$ , these schemes are DNS-authenticating [2, 19].

### 3 One-Time Ciphertext Authentication

One-time quantum authentication has been extensively studied [5, 9, 15, 18, 19, 21]. As we observed above, all of these works concern *plaintext authentication*, which ensures that manipulated ciphertexts decrypt to either the original plaintext or the reject symbol. Classical MACs, on the other hand, provide *ciphertext authentication*, which ensures that any ciphertext manipulation whatsoever will result in rejection. This distinction is important; for instance, in classical IND-CCA2, the adversary can defeat plaintext-authenticating schemes by invoking the decryption oracle on a modified challenge ciphertext.

In this section we show how to define and construct ciphertext authentication in the quantum setting. These ideas will be crucial to defining more advanced notions (such as ciphertext unforgeability and adaptive chosen-ciphertext security) later in the paper. We start with the information-theoretical security setting, and then we discuss how to apply these notions to the computational setting.

**A Characterization of Encryption Schemes.** We recall a lemma from [5] stating that all SKQES encrypt by (i) attaching some (possibly key-dependent) auxiliary state, and (ii) applying a unitary<sup>2</sup> operator. Decryption undoes the unitary, and then checks if the support of the state in the auxiliary register has changed. We emphasize that this characterization follows from correctness only, and thus applies to all schemes.

**Lemma 1 (Lemma B.9 in [5], restated).** *Let  $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$  be a SKQES. Then Enc and Dec have the following form:*

$$\begin{aligned} \text{Enc}_k(X_M) &= V_k (X_M \otimes (\sigma_k)_T) V_k^\dagger \\ \text{Dec}_k(Y_C) &= \text{Tr}_T \left[ P_T^{\sigma_k} \left( V_k^\dagger Y_C V_k \right) P_T^{\sigma_k} \right] + \hat{D}_k \left[ \bar{P}_T^{\sigma_k} \left( V_k^\dagger Y_C V_k \right) \bar{P}_T^{\sigma_k} \right]. \end{aligned}$$

Here,  $\sigma_k$  is a state on register  $T$ ,  $P_T^{\sigma_k}$  and  $\bar{P}_T^{\sigma_k}$  are the orthogonal projectors onto the support of  $\sigma^{(k)}$  (see Sect. 2) and its complement (respectively),  $V_k$  is a unitary operator, and  $\hat{D}_k$  is a channel.

In practice,  $\hat{D}_k$  (i.e., the map that is applied to any ciphertext outside of the range of  $\text{Enc}_k$ ) will just discard the state and replace it with  $\perp$ . Let us explain how the schemes we have seen so far fit into this characterization. For  $2\text{desTag}$ ,  $\sigma_k$  is simply the (key-independent) pure state  $|0^n\rangle\langle 0^n|_T$ ,  $V_k$  is the unitary operator of the two-design corresponding to key  $k$ ,  $P^{\sigma_k} = |0^n\rangle\langle 0^n|$ , and  $\hat{D}_k$  replaces the state with  $\perp$ . For  $1\text{des}^{\text{PRF}}$ ,  $\sigma_k$  is the maximally mixed state  $\tau$  (i.e., the classical randomness  $r$  from Definition 2), and  $V_k$  is the controlled-unitary which applies a quantum one-time pad on the first register, controlled on the contents of the second register (using the pqPRF  $f$ ), i.e.,  $|x\rangle|r\rangle \mapsto P_{f_k(r)}|x\rangle|r\rangle$ . Decryption

<sup>2</sup> If the dimension of the plaintext space does not divide the dimension of the ciphertext space, then we may need an isometry. In our case, all spaces are made up of qubits.

undoes the controlled unitary and never rejects, i.e.,  $P^{\sigma_k} = 1$ . This corresponds to the fact that  $\tau$  has full support.

By considering the spectral decomposition of the state  $\sigma_k$  from Lemma 1, it is straightforward to show that encryption can always be implemented using unitary operators and only classical randomness. We state this fact as follows.

**Corollary 1.** *Let  $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$  be a SKQES. Then for every  $k$ , there exists a probability distribution  $p_k : \{0, 1\}^t \rightarrow [0, 1]$  and a family of quantum states  $|\psi^{(k,r)}\rangle_T$  such that  $\text{Enc}_k$  is equivalent to the following algorithm:*

1. *sample  $r \in \{0, 1\}^t$  according to  $p_k$ ;*
2. *apply the following map:  $\text{Enc}_{k;r}(X_M) = V_k (X_M \otimes |\psi^{(k,r)}\rangle\langle\psi^{(k,r)}|_T) V_k^\dagger$ .*

*Here  $V_k$  and  $T$  are defined as in Lemma 1, and  $t$  is the number of qubits in  $T$ .*

For example, in the case of `2desTag`, the distribution is a point distribution and  $|\psi^{(k,r)}\rangle = |0^t\rangle$ . In `1desPRF`, the distribution is uniform and  $|\psi^{(k,r)}\rangle = |r\rangle$ .

It is important to remark here that, even if  $\text{Enc}_k$  is a polynomial-time algorithm, the functionally-equivalent algorithm provided by Corollary 1 may not be. We thus define the following.

**Condition 1.** *Let  $\Pi$  be a SKQES, and let  $p_k$ ,  $|\psi^{(k,r)}\rangle$  and  $V_k$  be as given in Corollary 1. We say that  $\Pi$  satisfies Condition 1 if there exist efficient quantum algorithms for (i) sampling from  $p_k$ , (ii) preparing  $|\psi^{(k,r)}\rangle$ , and (iii) implementing  $V_k$ , and this holds for all but a negligible fraction of  $k$  and  $r$ .*

We are not aware of any examples of SKQES that violate Condition 1. In fact, in all schemes we will consider (including all schemes constructed via Definition 2), the distribution  $p_k$  and the states  $|\psi^{(k,r)}\rangle$  are trivial to prepare, and the unitaries  $V_k$  are implementable by poly-size quantum circuits. In any case, when Condition 1 is required for a particular result, we will state this explicitly.

**Defining Ciphertext Authentication.** We begin by outlining our approach. Fix an encryption scheme  $\Pi$  with plaintext register  $M$  and ciphertext register  $C$ . Let  $A_{CB \rightarrow C\bar{B}}$  be an attack map. Intuitively, we would like to decide whether to accept or reject conditioned on whether  $A$  has changed the ciphertext. A possible approach would be to use the simulator from Theorem 5.1 in [15]: in the case of acceptance, this simulator<sup>3</sup> ensures that  $A$  is equivalent to  $\mathbb{1}_C \otimes \Phi$  for some side-information map  $\Phi_{B \rightarrow \bar{B}}$ . While this approach is on the right track, it is unnecessarily strong as a definition of security: it prevents the adversary from even looking at (or copying) classical parts of the ciphertext! This would place strange requirements on encryption. It would disallow constant classical messages (e.g., “begin PGP message”) accompanying ciphertexts. It would also disallow a large class of natural schemes, including all schemes  $\Pi^f$  from Sect. 2.

<sup>3</sup> In [15], this simulator was used to prove DNS security of the `2desTag` scheme. Here, we consider whether that simulator can be used to *define* secure authentication.

This class has many schemes that (intuitively speaking) should be adequate for authenticating poly-many quantum ciphertexts, such as the case where  $\Pi$  applies a random unitary and  $f$  is a random function.

The key to finding the middle ground lies in Corollary 1: any scheme can be decomposed in a way that enables us to check separately whether the identity has been applied to the quantum part, and whether the classical register has changed. In effect, this will amount to an additional constraint over DNS-authentication<sup>4</sup> (Definition 4), demanding extra structure from the simulator.

Recall that an attack  $\Lambda_{CB \rightarrow C\tilde{B}}$  on the scheme  $\Pi$  defines the averaged effective plaintext map  $\Lambda_{MB \rightarrow M\tilde{B}}^\Pi = \mathbb{E}_k[\text{Dec}_k \circ \Lambda \circ \text{Enc}_k]$ . We define ciphertext authentication as follows, using notation from Lemma 1 and Corollary 1.

**Definition 5.** A SKQES  $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$  is ciphertext authenticating, or QCA, if for all CP-maps  $\Lambda_{CB \rightarrow C\tilde{B}}$ , there exists a CP-map  $\Lambda_{B \rightarrow \tilde{B}}^{\text{rej}}$  such that:

$$\left\| \Lambda_{MB \rightarrow M\tilde{B}}^\Pi - \left( \text{id}_M \otimes \Lambda_{B \rightarrow \tilde{B}}^{\text{acc}} + |\perp\rangle\langle\perp|_M \otimes \Lambda_{B \rightarrow \tilde{B}}^{\text{rej}} \right) \right\|_\diamond \leq \text{negl}(n), \quad (3)$$

and  $\Lambda_{B \rightarrow \tilde{B}}^{\text{acc}} + \Lambda_{B \rightarrow \tilde{B}}^{\text{rej}}$  is TP. Here  $\Lambda_{B \rightarrow \tilde{B}}^{\text{acc}}$  is given by:

$$\Lambda_{B \rightarrow \tilde{B}}^{\text{acc}}(Z_B) = \mathbb{E}_{k,r} \left[ \langle \Phi_{k,r} | V_k^\dagger \Lambda(\text{Enc}_{k;r}(\phi_{MM'}^+ \otimes Z_B)) V_k | \Phi_{k,r} \rangle \right] \quad (4)$$

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

Condition (3) is simply DNS. It ensures that, in the accept case, the adversary performs the identity on the plaintext. Condition (4) demands that the rest of the action (i.e., on the side-information) is well-simulated by the following:

1. prepare a maximally entangled state  $\phi_{MM'}^+$  and attach it to the input  $B$ ;
2. run encryption, saving the classical randomness  $r$  used (meaning that the tag register  $T$  was prepared in the state  $|\psi^{(k,r)}\rangle$ );
3. apply decryption while conditioning on (i) the plaintext still being maximally entangled with  $M'$ , and (ii) register  $T$  still containing  $|\psi^{(k,r)}\rangle$ ;
4. output the contents of  $\tilde{B}$ .

Note that this definition only adds further constraints to DNS. Recalling that DNS implies QIND [9,21], we thus have the following.

**Theorem 7.** If a SKQES is QCA, then it is also DNS; in particular, it is QIND.

It is not difficult to see that the security proof in Theorem 5.1 of [15] (for establishing DNS of the Clifford scheme) actually applies to arbitrary 2-designs, and in fact proves QCA and not only DNS. We thus have that the scheme 2desTag fulfills ciphertext authentication. For details on the separation between QCA and DNS, see the appendix of the full version of this paper [4].

<sup>4</sup> One might also start from the authentication definitions of [21,27] rather than DNS. However, this is not necessary: these definitions' advantage over DNS is in key recycling; our setting is non-interactive and has no back-channel for key recycling.

**Computational-Security Variant.** We now briefly record a computational-security variant of one-time ciphertext authentication, which simply requires that all elements in Definition 5 are efficient.

**Definition 6.** A SKQES  $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$  is computationally ciphertext authenticating (cQCA) if, for any efficiently implementable attack map  $\Lambda_{CB \rightarrow C\tilde{B}}$ , the effective attack  $\Lambda_{MB \rightarrow M\tilde{B}}$  is computationally indistinguishable from the simulator:

$$\Lambda_{MB \rightarrow M\tilde{B}}^{\text{sim}} = \text{id}_M \otimes \Lambda_{B \rightarrow \tilde{B}}^{\text{acc}} + |\perp\rangle\langle\perp|_M \otimes \Lambda_{B \rightarrow \tilde{B}}^{\text{reject}}. \quad (5)$$

Here the simulator is given by:

$$\begin{aligned} \Lambda_{B \rightarrow \tilde{B}}^{\text{acc}} &= \mathbb{E}_{k,r} \left[ \langle \Phi_{k,r} | V_k^\dagger \Lambda(\text{Enc}_{k;r}(\phi_{MM'}^+ \otimes (\cdot)_B)) V_k | \Phi_{k,r} \rangle \right] \text{ and} \\ \Lambda_{B \rightarrow \tilde{B}}^{\text{reject}} &= \mathbb{E}_{k,r} \left[ \text{Tr}(\mathbf{1} - |\Phi_{k,r}\rangle\langle\Phi_{k,r}|) V_k^\dagger \Lambda(\text{Enc}_{k;r}(\phi_{MM'}^+ \otimes (\cdot)_B)) V_k \right], \end{aligned} \quad (6)$$

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

Because we fix the form of the simulator in the reject case, the simulator is efficiently implementable just as in [15] for schemes that satisfy Condition 1. It is straightforward to define a computational variant of DNS [15], which we denote by cDNS. Given that Theorem 7 only talks about computationally bounded quantum adversaries, it also applies to cDNS. In particular we have the following.

**Proposition 1.** *If a SKQES is cQCA, then it is also cDNS; in particular, it satisfies QIND.*

## 4 Quantum Unforgeability

Translating the standard classical intuition of ciphertext unforgeability to the quantum setting appears nontrivial. As we develop our approach, it will be useful to keep in mind a “prototype” scheme that should (intuitively) satisfy quantum unforgeability against a polynomial-time adversary making an arbitrary number of queries. This is the scheme  $2\text{desTag}^{\text{PRF}}$ , which encrypts via:

$$\text{Enc}_k(\rho) = U_{f_k(r)}(\rho \otimes |0^n\rangle\langle 0^n|) U_{f_k(r)}^\dagger \otimes |r\rangle\langle r|$$

where  $k$  is a key for the pqPRF  $f$  and  $r$  is randomness selected freshly for each encryption. This scheme is characterized (via Lemma 1) by the key-independent “tag state”  $|0^n\rangle\langle 0^n| \otimes \tau$  (where  $\tau$  is the maximally mixed state) and the unitary  $V_k$  which applies  $U_{f_k(\cdot)}$  on the first two registers, controlled on the third register (i.e., the randomness  $r$ ).

To see why this scheme should be unforgeable, assume for the moment that  $U_s$  is a Haar-random unitary and  $f_k$  is a perfectly random function. Intuitively, from the point of view of the adversary, each plaintext is mapped into a subspace which is fresh, independent, random, and exponentially-small as a fraction of the total dimension (of the ciphertext space). Security should then reduce to the security of multiple uses of a QCA one-time scheme, each time with a freshly generated key. We will carefully formalize this intuition in a later section.

**Formal Definitions.** Our definition will compare the performance of an adversary in two games: an unrestricted forgery game, and a cheat-detecting game. Fix an SKQES  $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$  and let  $\mathcal{A}$  be an adversary in the following.

**Experiment 1.** The QUF-Forge( $\Pi, \mathcal{A}, n$ ) experiment:

- 1:  $k \leftarrow \text{KeyGen}(1^n)$ ;
- 2: **if**  $\text{Dec}_k(\mathcal{A}^{\text{Enc}_k}(1^n)) \neq \perp$ , **output** win; otherwise **output** reject.

We will think about this experiment as taking place between the adversary  $\mathcal{A}$  and a challenger  $\mathcal{C}$ , who generates the key  $k$ , answers the queries of  $\mathcal{A}$ , and then decrypts to see the outcome of the game.

We now consider a different experiment where  $\mathcal{C}$  attempts to check  $\mathcal{A}$  for cheating. We will make use of the maximally entangled state  $|\phi^+\rangle_{M'M''}$  on two copies ( $M'$  and  $M''$ ) of the plaintext register, and the corresponding measurement  $\{\Pi_{M'M''}^+, \mathbb{1} - \Pi_{M'M''}^+\}$ . We will also need a measurement that will help  $\mathcal{C}$  identify previously generated ciphertexts. Recall from Sect. 3 that correctness implies that  $\text{Enc}$  can be written in the form  $\text{Enc}_k(X) = V_k(X_M \otimes \sigma_k)V_k^\dagger$  where  $\sigma_T^{(k)} = \sum_r p_k(r)\Pi_{k,r}$  and  $\Pi_{k,r} = |\psi^{(k,r)}\rangle\langle\psi^{(k,r)}|_T$ . This also defines, for each  $(k, r)$ , the two-outcome measurement  $\{\Pi_{k,r}, \mathbb{1} - \Pi_{k,r}\}$ . In all these two-outcome measurements, we denote the first outcome by 0 and the second outcome by 1. Notice that these projectors commute, as  $|\psi^{(k,r)}\rangle_T$  are elements of an orthonormal basis of eigenvectors.

**Experiment 2.** The QUF-Cheat( $\Pi, \mathcal{A}, n$ ) experiment:

- 1:  $\mathcal{C}$  runs  $k \leftarrow \text{KeyGen}(1^n)$ ;
- 2:  $\mathcal{A}$  receives  $1^n$  and oracle access to  $E_k$  (controlled by  $\mathcal{C}$ ), defined as follows:
  - (1)  $\mathcal{A}$  sends plaintext register  $M$  to  $\mathcal{C}$ ;
  - (2)  $\mathcal{C}$  discards  $M$  and prepares  $|\phi^+\rangle_{M'M''}$ ;
  - (3)  $\mathcal{C}$  applies  $\text{Enc}_k$  to  $M'$  using fresh randomness  $r$ , sends result  $C$  to  $\mathcal{A}$ ;
  - (4)  $\mathcal{C}$  stores  $(M'', r)$  in a set  $\mathcal{M}$ .
- 3:  $\mathcal{A}$  sends final output register  $C_{\text{out}}$  to  $\mathcal{C}$ ;
- 4:  $\mathcal{C}$  applies  $V_k^\dagger$  to  $C_{\text{out}}$ , places results in  $MT$ ;
- 5: **for each**  $(M'', r) \in \mathcal{M}$  **do**
- 6:      $\mathcal{C}$  applies  $\{\Pi_{k,r}, \mathbb{1} - \Pi_{k,r}\}$  to  $T$ ;
- 7:     **if** outcome is 0 **then**:
- 8:          $\mathcal{C}$  applies  $\{\Pi^+, \mathbb{1} - \Pi^+\}$  to  $MM''$ ;
- 9:         **if** outcome is 0: **output** cheat; **end if**
- 10:    **end if**
- 11: **end for**
- 12: **output** reject.

Note that the experiment always outputs **reject** if  $\mathcal{A}$  makes no queries. We emphasize that  $\mathcal{C}$  is a fixed algorithm defined by the security game and the properties of  $\Pi$ . The challenger is efficient if the states  $|\psi^{(k,r)}\rangle\langle\psi^{(k,r)}|$  and the unitary  $V_k$  are efficiently implementable and the probability distribution  $p_k$  is efficiently sampleable. We believe this is not a significant constraint. It is easily

satisfied in all schemes we are aware of. Moreover, in light of Lemma 1, it seems unlikely that any reasonable form of ciphertext unforgeability can be defined without this requirement. We are now ready to define security.

**Definition 7.** *A SKQES  $\Pi$  has unforgeable ciphertexts (or is QUF) if, for all QPT adversaries  $\mathcal{A}$ , it holds:*

$$|\Pr[\text{QUF-Forge}(\Pi, \mathcal{A}, n) \rightarrow \text{win}] - \Pr[\text{QUF-Cheat}(\Pi, \mathcal{A}, n) \rightarrow \text{cheat}]| \leq \text{negl}(n).$$

It is straightforward to adapt the above definition to the bounded-query setting, where we fix some positive integer  $t$  (at scheme design time) and demand that adversaries can make no more than  $t$  queries. We call the resulting notion  $\text{QUF}_t$ . One then has the obvious implications  $\text{QUF} \Rightarrow \text{QUF}_t \Rightarrow \text{QUF}_{t-1} \forall t \in \mathbb{N}$ .

Let us briefly discuss a potential concern with these definitions. Consider the repeated measurements applied to the adversary's final output  $C_{\text{out}}$  (Line 6 and Line 8) in QUF-Cheat. The first measurement simply compares the randomness of  $C_{\text{out}}$  to that of previously generated ciphertexts. Such measurements will not disturb properly-formed ciphertexts at all, and malformed ones will not affect our security definition. The second measurement actually measures the plaintext register  $M$ , and thus might (a priori) appear to be concerning. Indeed, if multiple such measurements are applied to  $M$ , this might open up a vulnerability to attacks. As it turns out, this is not a problem. We will shortly show (see Theorem 8 below) that QUF implies QIND-CPA. For QIND-CPA schemes, any given random string  $r$  is only chosen with negligible probability at encryption time (if not, querying the encryption oracle a polynomial number of times with the challenge plaintext would be enough to compromise security). It follows that, with overwhelming probability, the random strings chosen in the different oracle calls in QUF-Cheat are pairwise distinct. This, in turn, implies that the measurement in Line 8 is applied at most once in a given run of the experiment.

**Relationship to Other Security Notions.** It is well-known that even one-time quantum authentication implies QIND secrecy [9]. As we now show, QUF implies an even stronger notion of secrecy, QIND-CPA. This is a significant departure from classical unforgeability, which is completely independent of secrecy.

**Theorem 8.** *If a SKQES satisfies QUF, then it also satisfies QIND-CPA.*

*Proof.* Let  $\Pi$  be a SKQES, and let  $\mathcal{A}$  be an adversary winning QIND-CPA with non-negligible advantage  $\nu$  over guessing, with pre-challenge algorithm  $\mathcal{A}_1$  and post-challenge algorithm  $\mathcal{A}_2$ . We will build an adversary  $\mathcal{B}$  with black-box oracle access to  $\mathcal{A}$ , able to distinguish between the QUF-Forge game and the QUF-Cheat game with non-negligible advantage over guessing, as follows:

1.  $\mathcal{B}$  runs  $\mathcal{A}_1(1^n)$ , answering its queries using his own oracle  $\mathcal{O}$ ;
2. get registers  $M$  (challenge plaintext) and  $B$  (side information) from  $\mathcal{A}_1$ ;
3. choose a random bit  $b \xleftarrow{\$} \{0, 1\}$ ; if  $b = 1$ , then replace contents of  $M$  with a maximally-mixed state;



4. invoke oracle  $\mathcal{O}$  on  $M$  and place result in register  $C$ ;
5. run  $\mathcal{A}_2$  on registers  $C$  and  $B$ , receiving output  $b' \in \{0, 1\}$ ;
6. if  $b = b'$ , then output *real*; else output *real* or *ideal* with equal probability.

Note that, if  $\mathcal{B}$  is playing QUF-Forge, then  $\mathcal{O} = \text{Enc}_k$  and we are faithfully simulating the QIND-CPA game for  $\mathcal{A}$ . It follows that  $b = b'$  with probability at least  $1/2 + \nu$ . If  $\mathcal{B}$  is playing QUF-Cheat instead,  $\mathcal{O}$  discards its input (and replaces it with half of a maximally-entangled state) on every call. In that case, all inputs to  $\mathcal{A}_1$  and  $\mathcal{A}_2$  are completely uncorrelated with  $b$ , so that  $b' = b$  with probability  $1/2$ . Therefore,  $\mathcal{A}'$  will correctly guess the game it is playing in with non-negligible advantage.

Now it is easy to see how to use  $\mathcal{B}$  to violate the main condition in the definition of QUF with the same distinguishing advantage. First, query the oracle once and store the output in register  $C$ . Next, run  $\mathcal{B}$ . If  $\mathcal{B}$  outputs *real*, then output the contents of  $C$  (achieving win in QUF-Forge). Otherwise, output a random state in the ciphertext register (achieving reject in QUF-Cheat).  $\square$

We also study the restriction of the quantum notion QUF to the classical case, i.e., classical symmetric-key encryption schemes (SKES) vs classical adversaries. We denote this classical restriction by UF. In this notion, the classical unrestricted forgery game UF-Forge is defined precisely as in Experiment 1. Regarding the quantum game QUF-Cheat, notice that, in any classical scheme, one can apply ciphertext verification to a string  $c$  as follows: (i) make a copy  $c'$  of  $c$ , (ii) decrypt  $c$ , (iii) if decryption rejected, output *reject*, and otherwise output  $c'$ . In other words, all classical encryption schemes automatically satisfy Condition 1. The appropriate classical restriction UF-Cheat of this game thus proceeds as Experiment 2, with two modifications: (i) in step 2:,  $\mathcal{C}$  replaces the plaintext in register  $M_j$  by a random plaintext, encrypts it, and stores a copy of the resulting ciphertext in  $C_j$ ; and (ii) in step 4:, without decrypting, the game outputs *cheat* if the challenge ciphertext  $C$  equals any one of the saved  $C_j$ 's. We then have the following.

**Definition 8.** *A SKES  $\Pi$  has unforgeable ciphertexts (or is UF) if, for all PPT adversaries  $\mathcal{A}$ ,*

$$|\Pr[\text{UF-Forge}(\Pi, \mathcal{A}, n) \rightarrow \text{win}] - \Pr[\text{UF-Cheat}(\Pi, \mathcal{A}, n) \rightarrow \text{cheat}]| \leq \text{negl}(n).$$

The proof of Theorem 8 carries over easily to the classical case. Moreover, one can show how UF implies the classical security notion of *integrity of ciphertexts* INT-CTXT [10], which states that no bounded adversary with oracle access to an encryption oracle can produce a ciphertext which is at the same time (i) valid, and (ii) fresh, i.e., never output by the oracle. Recall that, classically, it is known [10] that INT-CTXT plus IND-CPA defines authenticated encryption AE. Therefore, the notion of unforgeability of ciphertexts, when restricted to the classical case, is at least as strong as authenticated encryption. However, one can also show the converse, i.e., AE implies UF.

**Theorem 9.**  $UF \iff AE$ .

*Proof.* The first non-trivial part to prove is  $UF \implies INT\text{-}CTXT$ . Let  $\Pi$  be an INT-CTXT insecure SKES. Then there exists an adversary  $\mathcal{A}$  with oracle access to  $\text{Enc}_k$  which, with non-negligible probability  $\nu$ , outputs a ciphertext  $c$  which was never output by the encryption oracle. Define a PPT algorithm  $\mathcal{B}$  with oracle access to  $\text{Enc}_k$ , as follows. First,  $\mathcal{B}$  executes  $\mathcal{A}$  and records a list  $L$  of all  $\text{Enc}_k$ 's answers  $c_j$  output to  $\mathcal{A}$ . When  $\mathcal{A}$  outputs a ciphertext  $c$ , if  $c \in L$ ,  $\mathcal{B}$  outputs a random ciphertext  $c'$ ; else it outputs  $c$ . For  $\mathcal{B}$ , the success probabilities in the games defining UF are as follows:

- in the UF-Forge experiment, since  $c$  is a fresh ciphertext with non-negligible probability  $\nu$ ,  $\mathcal{B}$  wins UF-Forge with probability at least  $\nu$ .
- In UF-Cheat instead, whenever the ciphertext is not fresh,  $\mathcal{B}$  replaces it with a random one, and hence only wins UF-Cheat with negligible probability.

The fact that a random ciphertext is invalid with overwhelming probability follows by considering an adversary that does not make any queries. So we have:

$$|\Pr[\text{UF-Forge}(\Pi, \mathcal{A}', n) \rightarrow \text{win}] - \Pr[\text{UF-Cheat}(\Pi, \mathcal{A}', n) \rightarrow \text{cheat}]| \geq \nu,$$

and hence  $\Pi$  cannot be UF.

The other direction to prove is  $AE \implies UF$ . For this, we will use an equivalent characterization of AE, also known in the literature as IND-CCA3 [28]. In this definition, the adversary's goal is to distinguish whether he's playing in the AE-Real world, or in the AE-Ideal world. In the AE-Real world, the adversary can interact freely with an encryption oracle  $\text{Enc}_k$ , and with a restricted decryption oracle  $\text{Dec}_k$  which always rejects ( $\perp$ ) decryption queries over any ciphertext which was output by  $\text{Enc}_k$ . In the AE-Ideal world, instead, the adversary is interacting with an oracle  $\text{Enc}_k(\$)$  (which ignores the input query, and always returns the encryption of a fresh random plaintext), and a constant  $\perp$  oracle (which simulates the decryption oracle but always rejects any query). A scheme  $\Pi$  is AE secure iff, for any adversary  $\mathcal{A}$  it holds:

$$|\Pr[\text{AE-Real}(\Pi, \mathcal{A}, n) \rightarrow 1] - \Pr[\text{AE-Ideal}(\Pi, \mathcal{A}, n) \rightarrow 1]| \leq \text{negl}(n).$$

Now, let  $\mathcal{A}$  be a PPT adversary breaking UF for a scheme  $\Pi$ . This means that there exists a non-negligible function  $\nu$  such that:

$$|\Pr[\text{UF-Forge}(\Pi, \mathcal{A}, n) \rightarrow \text{win}] - \Pr[\text{UF-Cheat}(\Pi, \mathcal{A}, n) \rightarrow \text{cheat}]| \geq \nu(n).$$

We use  $\mathcal{A}$  to build an adversary  $\mathcal{B}$  able to distinguish AE-Real from AE-Ideal. The new adversary  $\mathcal{B}$  runs  $\mathcal{A}$  and forwards all of  $\mathcal{A}$ 's encryption queries to his own encryption oracle. Finally, when  $\mathcal{A}$  outputs a ciphertext  $c$ ,  $\mathcal{B}$  queries his own decryption oracle on  $c$ , and looks at the oracle's response. If the response is *not*  $\perp$ , then  $\mathcal{B}$  returns *real*, otherwise returns *real* or *ideal* with equal chance.

It is easy to see that  $\mathcal{B}$  distinguishes AE-Ideal from AE-Real with non-negligible advantage at least  $\nu/2$  over guessing. The reason is as follows. If  $\mathcal{B}$  is

in the AE-Real world (probability  $1/2$ ), then he is correctly simulating for  $\mathcal{A}$  the UF-Forge game. Since  $\mathcal{A}$  breaks UF by assumption, it means that, with probability at least  $\nu$ , his output  $c$  will be a fresh valid ciphertext; in that case, also  $\mathcal{B}$  wins. On the other hand, if the world is AE-Ideal,  $\mathcal{B}$  still wins with probability  $1/2$ .  $\square$

This means that UF is actually *another characterization of authenticated encryption*. This is an interesting observation, given that UF comes from the classical restriction of a quantum notion “merely” concerning the unforgeability of ciphertexts. However, we stress that this equivalence only holds at the classical level, and that this is insufficient evidence to declare that UF serves the same purpose quantumly as AE does classically. In fact, in Sect. 6 we introduce a quantum analogue of AE which we call QAE, and provide stronger evidence that the latter is in fact the correct analogue.

## 5 Quantum IND-CCA2

Next, we move to the problem of defining adaptive chosen-ciphertext security for quantum encryption. In the usual classical formulation (IND-CCA2), the adversary  $\mathcal{A}$  receives both an encryption oracle and a decryption oracle for the entire duration of the indistinguishability game. To eliminate the trivial strategy, we do not permit  $\mathcal{A}$  to query the decryption oracle on the challenge ciphertext. This last condition does not make sense in the quantum setting, for a number of reasons we’ve seen before: no-cloning prevents us from storing a copy of the challenge, measurement may destroy the states involved, and so on. However, our approach to defining unforgeability can be adapted to this case. The resulting notion of *quantum indistinguishability under adaptive chosen-ciphertext attacks* (QIND-CCA2) can also be recast in the public-key quantum encryption setting.

**Formal Definition.** As before, we will compare the performance of the adversary in two games. In each case, the adversary  $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$  consists of two parts (pre-challenge and post-challenge), and is playing against the challenger  $\mathcal{C}$ , which is a fixed algorithm determined only by the security game and the scheme.

**Experiment 3.** The QCCA2-Test( $\Pi, \mathcal{A}, n$ ) experiment:

- 1:  $\mathcal{C}$  runs  $k \leftarrow \text{KeyGen}(1^n)$  and flips a coin  $b \xleftarrow{\$} \{0, 1\}$ ;
- 2:  $\mathcal{A}_1$  receives  $1^n$  and access to oracles  $\text{Enc}_k$  and  $\text{Dec}_k$ ;
- 3:  $\mathcal{A}_1$  prepares a side register  $S$ , and sends  $\mathcal{C}$  a challenge register  $M$ ;
- 4:  $\mathcal{C}$  puts into  $C$  either  $\text{Enc}_k(M)$  (if  $b = 0$ ) or  $\text{Enc}_k(\tau_M)$  (if  $b = 1$ );
- 5:  $\mathcal{A}_2$  receives registers  $C$  and  $S$  and oracles  $\text{Enc}_k$  and  $\text{Dec}_k$ ;
- 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  $\text{Dec}_k$  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 4.** The QCCA2-Fake( $\Pi, \mathcal{A}, n$ ) experiment:

- 1:  $\mathcal{C}$  runs  $k \leftarrow \text{KeyGen}(1^n)$ ;
- 2:  $\mathcal{A}_1$  receives  $1^n$  and access to oracles  $\text{Enc}_k$  and  $\text{Dec}_k$ ;
- 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  $\text{Enc}_k$  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  $\{\Pi_{k,r}, \mathbb{1} - \Pi_{k,r}\}$  to  $T$ ;
  - (3) **if** outcome is 0 **then**:
  - (4)      $\mathcal{C}$  applies  $\{\Pi^+, \mathbb{1} - \Pi^+\}$  to  $MM''$ ;
  - (5)     **if** outcome is 0: **output** cheat;
  - (6)     **end if**
  - (7) **return**  $M$ ;
- 6:  $\mathcal{C}$  draws a bit  $b$  at random. **If**  $b = 1$ , **output** cheat; **if**  $b = 0$  **output** reject.

We now define quantum IND-CCA2 in terms of the advantage gap of adversaries between the above two games.<sup>5</sup>

**Definition 9.** A SKQES  $\Pi$  is QIND-CCA2 if, for all QPT adversaries  $\mathcal{A}$ ,

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

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. We note that QIND-CCA2 clearly implies QIND-CCA1.

**Proposition 2.** QIND-CCA2  $\implies$  QIND-CCA1.

*Proof.* Suppose we have a scheme  $\Pi$  which is not QIND-CCA1, i.e., there exists an adversary  $\mathcal{A}$  which wins the usual QIND-CCA1 game with non-negligible advantage  $\nu$  over guessing. Clearly  $\mathcal{A}$  can also play the games QCCA2-Test and QCCA2-Fake, but will not query the decryption oracle post-challenge. Note that  $\mathcal{A}$  wins QCCA2-Test with probability  $1/2 + \nu$ , but is declared as cheating in QCCA2-Fake with probability exactly  $1/2$ . Hence  $\Pi$  is not QIND-CCA2.  $\square$

Next, we show that the classical restriction of QIND-CCA2 is equivalent to the classical security notion IND-CCA2. We denote the classical restriction of QIND-CCA2 by IND-CCA2'. This is defined by adapting the replacement and verification procedure of the challenger in QCCA2-Test in the same way as when defining UF. We denote the classical versions of the games QCCA2-Test and QCCA2-Fake by CCA2-Test and CCA2-Fake, respectively.

<sup>5</sup> 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.

**Theorem 10.** *A SKES  $\Pi$  is IND-CCA2' iff it is IND-CCA2.*

*Proof.* Suppose first that  $\mathcal{A}$  is an adversary breaking IND-CCA2', i.e., winning CCA2-Test with non-negligible advantage  $\nu$  over the probability of winning CCA2-Fake. We construct an adversary  $\mathcal{A}'$ , that runs  $\mathcal{A}$ , keeps a copy of the challenge ciphertext and aborts by giving a random answer whenever  $\mathcal{A}$  is about to query the decryption oracle with the challenge ciphertext. Note that  $\mathcal{A}'$  wins CCA2-Fake with probability exactly  $1/2$ . We call  $\mathcal{A}'$  the *self-checking version* of  $\mathcal{A}$ . It is easy to show that  $\mathcal{A}'$  wins the CCA2-Test game with probability at least  $1/2 + \nu$ :

$$\begin{aligned} & \Pr[\mathcal{A}' \text{ wins CCA2-Test}] \\ &= \Pr[\mathcal{A} \text{ wins CCA2-Test} \wedge \mathcal{A} \text{ does not cheat}] + \frac{1}{2} \Pr[\mathcal{A} \text{ cheats}] \\ &\geq \Pr[\mathcal{A} \text{ wins CCA2-Test}] - \frac{1}{2} \Pr[\mathcal{A} \text{ cheats}] \\ &\geq \Pr[\mathcal{A} \text{ wins CCA2-Fake}] + \frac{1}{\nu} - \frac{1}{2} \Pr[\mathcal{A} \text{ cheats}] = \frac{1}{2} + \nu. \end{aligned}$$

The first inequality is  $\Pr[A \wedge B] \geq \Pr[A] - \Pr[\neg B]$  and the second inequality is the assumption. But the CCA2-Test and IND-CCA2 games are identical for adversaries that do not query the challenge, and  $\mathcal{A}'$  has been constructed not to, i.e.,  $\mathcal{A}'$  wins the IND-CCA2 game with probability  $1/2 + \nu$ .

For the other direction, let  $\mathcal{A}$  be an adversary that wins the IND-CCA2 game with non-negligible advantage. Let  $\mathcal{A}'$  be the self-checking version of  $\mathcal{A}$ . Note that  $\mathcal{A}$  and  $\mathcal{A}'$  behave the same in both the IND-CCA2 and CCA2-Test games, as  $\mathcal{A}$  never submits the challenge ciphertext there by assumption. In the CCA2-Fake game, however,  $\mathcal{A}$  could, in principle, query the oracle with the challenge ciphertext, which is why we have to resort to the use of  $\mathcal{A}'$ . The latter is a successful adversary for IND-CCA2': It wins the CCA2-Test game with non-negligible advantage over random guessing by assumption, but it wins the CCA2-Fake game with probability exactly  $\frac{1}{2}$ . □

## 6 Quantum Authenticated Encryption

In the classical setting, authenticated encryption (AE) is defined as IND-CCA2 and unforgeability of ciphertexts (see Definition 4.17 in [25]) or, equivalently, IND-CPA and unforgeability of ciphertexts [10]. A third equivalent formulation due to Shrimpton [28] defines AE in terms of a real vs ideal scenario. According to this definition, a classical scheme  $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$  is AE if no adversary, given oracles  $E$  and  $D$ , can distinguish these two scenarios:

- AE-Real:  $(E, D)$  is  $(\text{Enc}_k, \text{Dec}_k)$  with  $k \leftarrow \text{KeyGen}$ ;
- AE-Ideal:  $E$  discards the input and returns  $\text{Enc}_k(m)$  for random  $m$ , and  $D$  always rejects; here again  $k \leftarrow \text{KeyGen}$ ;

This is not yet enough, because the adversary  $\mathcal{A}$  can always distinguish real from ideal by composing  $E$  with  $D$ . To patch this problem, we can (i) demand that

$\mathcal{A}$  cannot do that, as in [28], or (ii) add the condition  $D \circ E = \mathbb{1}$  to the ideal case<sup>6</sup>. We will take the latter approach.

Motivated by this formulation of AE and our general strategy so far, we will define quantum authenticated encryption by comparing the performance of the adversary in a real world and an ideal world. In the real world, the adversary gets unrestricted access to  $\text{Enc}_k$  and  $\text{Dec}_k$ . In the ideal world, the challenger  $\mathcal{C}$  stores the  $\text{Enc}_k$  queries, replacing them with halves of maximally-entangled states; when a  $\text{Dec}_k$  query is detected as corresponding to a particular earlier  $\text{Enc}_k$  query,  $\mathcal{C}$  replies with the contents of the stored register; otherwise  $\text{Dec}_k$  rejects. Cheat detection is performed just as in the unforgeability game QUF-Cheat.

**Formal Definition.** We now formally define the two worlds: the real world QAE-Real, and the ideal (or cheat-detecting) world QAE-Ideal. In both cases, the adversary  $\mathcal{A}$  receives two oracles and then outputs a single bit.

**Experiment 5.** The QAE-Real( $\Pi, \mathcal{A}, n$ ) experiment:

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

In the ideal setting, it will be convenient to describe the experiment in terms of an interaction between  $\mathcal{A}$  and the challenger  $\mathcal{C}$ , a fixed algorithm determined only by the security game and the properties of  $\Pi$ .

**Experiment 6.** The QAE-Ideal( $\Pi, \mathcal{A}, n$ ) experiment:

- 1:  $\mathcal{C}$  runs  $k \leftarrow \text{KeyGen}(1^n)$ ;
- 2: initialize oracles  $E_{M \rightarrow C}$  and  $D_{C \rightarrow M}$ :
  - $E$  is defined as follows. On input a register  $M$ :
    - (1)  $\mathcal{C}$  prepares  $|\phi^+\rangle_{M'M''}$ , and generates fresh randomness  $r$ ;
    - (2)  $\mathcal{C}$  stores  $(r, M'', M)$  in a set  $\mathcal{M}$ ;
    - (3)  $\mathcal{C}$  applies  $\text{Enc}_k$  to  $M'$  using randomness  $r$ ; **return** result to  $\mathcal{A}$ .
  - $D$  is defined as follows. On input a register  $C$ :
    - (1)  $\mathcal{C}$  applies  $V_k^\dagger$  to  $C$ , places results in  $M'T$ ;
    - (2) **for each**  $(r, M'', M) \in \mathcal{M}$  **do**:
    - (3)  $\mathcal{C}$  applies  $\{\Pi_{k,r}, \mathbb{1} - \Pi_{k,r}\}$  to  $T$ ;
    - (4) **if** outcome is 0 **then**:
    - (5)  $\mathcal{C}$  applies  $\{\Pi^+, \mathbb{1} - \Pi^+\}$  to  $M'M''$ ;
    - (6) **if** outcome is 0: **return**  $M$ ;
    - (7) **end if**
    - (8) **end for**
    - (9) **return**  $|\perp\rangle\langle\perp|$ ;
- 3: **output**  $\mathcal{A}^{E,D}(1^n)$ .

Note that, as before, we number the measurement outcomes by 0 (the first outcome) and 1 (the second outcome). With the above games defined, we can now set down our definition of quantum authenticated encryption.

<sup>6</sup> More precisely, the ideal world maintains a list of all queries that  $\mathcal{A}$  makes to  $E$ , and ensures that  $D$  will respond correctly if queried on an output of  $E$ .

**Definition 10.** A SKQES  $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$  is an authenticated quantum encryption scheme (or is QAE) if, for all QPT adversaries  $\mathcal{A}$ :

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

**Relationship to Other Security Notions.** Next, we give evidence that QAE is indeed the correct formalization of a quantum analogue of AE, by showing that it implies all of the quantum security notions defined thus far. We begin with adaptive chosen-ciphertext security.

**Theorem 11.** QAE  $\implies$  QIND-CCA2.

*Proof.* The proof is similar to that of Theorem 8. For a scheme  $\Pi$ , let  $\mathcal{A}$  be an adversary against QIND-CCA2, e.g., let us say that:

$$\Pr[\text{QCCA2-Test}(\Pi, \mathcal{A}, n) \rightarrow \text{win}] = \Pr[\text{QCCA2-Fake}(\Pi, \mathcal{A}, n) \rightarrow \text{cheat}] + \nu(n),$$

for non-negligible  $\nu$ . We then show how to build another adversary  $\mathcal{B}$  with black-box access to  $\mathcal{A}$ , able to distinguish QAE-Real from QAE-Ideal.

$\mathcal{B}$  runs  $\mathcal{A}$ , and forwards all of  $\mathcal{A}$ 's queries to his own oracles. When eventually  $\mathcal{A}$  outputs a challenge plaintext state,  $\mathcal{B}$  flips a random bit  $b$ . If  $b = 0$ , then  $\mathcal{B}$  forwards the challenge plaintext to his encryption oracle as usual. Otherwise, if  $b = 1$ ,  $\mathcal{B}$  replaces the challenge with a totally mixed plaintext state before relaying it to the oracle. After that,  $\mathcal{B}$  continues to answer  $\mathcal{A}$ 's queries during the second quantum CCA phase as before, by forwarding all the queries to his oracles, until  $\mathcal{A}$  produces an output bit  $b'$ . Finally, if  $b = b'$ , then  $\mathcal{B}$  outputs real, otherwise he outputs ideal.

Now notice the following: If we are in the QAE-Real environment (that is,  $\mathcal{B}$  has unrestricted Enc and Dec oracles), then  $\mathcal{B}$  is faithfully simulating for  $\mathcal{A}$  the QCCA2-Test game, which means that the probability of  $\mathcal{B}$  correctly outputting real is exactly the same probability of  $\mathcal{A}$  of winning QCCA2-Test.

If we are in the QAE-Ideal world, instead,  $\mathcal{B}$  is playing in a “malformed” game, where all his encryption queries are replaced by random plaintexts before encryption. This means that the best  $\mathcal{A}$  could do in order to guess the secret bit  $b$  is guessing at random, *unless*  $\mathcal{A}$  uses a “cheating decryption query” on the challenge ciphertext (in this case the modified decryption oracle of the game QAE-Ideal would actually return the encrypted plaintext). Looking at the description of the QCCA2-Fake game, it is clear that this is exactly the same as  $\Pr[\text{QCCA2-Fake}(\Pi, \mathcal{A}, n) \rightarrow \text{cheat}]$ . So, summing up:

$$\begin{aligned} & \left| \Pr[\text{QAE-Real}(\Pi, \mathcal{B}, n) \rightarrow \text{Real}] - \Pr[\text{QAE-Ideal}(\Pi, \mathcal{B}, n) \rightarrow \text{Real}] \right| \\ &= \left| \Pr[\text{QCCA2-Test}(\Pi, \mathcal{A}, n) \rightarrow \text{win}] - \Pr[\text{QCCA2-Fake}(\Pi, \mathcal{A}, n) \rightarrow \text{cheat}] \right| = \nu, \end{aligned}$$

which concludes the proof.  $\square$

In terms of authentication security, we can show that QAE implies cQCA (computational one-time ciphertext authentication), and hence also cDNS.

**Theorem 12.** *Let  $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$  be a SKQES that is QAE secure and satisfies Condition 1. Then it is cQCA.*

*Proof.* Assume  $\Pi$  is not cQCA. Then there exists an algorithm  $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$  that achieves the following.  $\mathcal{A}_1$  gets an input  $1^n$  and outputs registers  $M$  (the plaintext register) and  $B$ .  $\mathcal{A}_2$  implements a map  $A_{CB \rightarrow C\tilde{B}}$  on two registers  $C$  (the ciphertext register) and  $B$ .  $\mathcal{A}_3$  is a distinguisher between the two states resulting from applying  $\tilde{A}_{CB \rightarrow C\tilde{B}}$  or the corresponding simulator according to Eqs. (5) and (6) to the output of  $\mathcal{A}_1$ .

The crucial observation is, that the map on registers  $MB$  resulting from sending  $M$  to the challenger  $C'_{\text{ideal}}$  as an encryption query in the ideal QAE game, applying  $A_{CB \rightarrow C\tilde{B}}$  to the output and sending the resulting  $C$ -register to  $C'_{\text{ideal}}$  as a decryption query, is exactly the simulator defined in Eqs. (5) and (6). Thus, the adversary that runs  $\mathcal{A}_1$ , queries the encryption oracle, runs  $\mathcal{A}_2$ , queries the decryption oracle and finally runs  $\mathcal{A}_3$  is a successful QAE adversary.  $\square$

In addition, QAE implies quantum unforgeability.

**Theorem 13.** QAE  $\implies$  QUF.

*Proof.* For a scheme  $\Pi$ , let  $\mathcal{A}$  be an adversary against QUF, e.g., let us say that:

$$\Pr[\text{QUF-Forge}(\Pi, \mathcal{A}, n) \rightarrow \text{win}] = \Pr[\text{QUF-Cheat}(\Pi, \mathcal{A}, n) \rightarrow \text{cheat}] + \nu,$$

where  $\nu$  is non-negligible. We then build another adversary  $\mathcal{B}$  with black-box access to  $\mathcal{A}$ , able to distinguish QAE-Real from QAE-Ideal with non-negligible advantage.  $\mathcal{B}$  runs  $\mathcal{A}$ , and forwards all of  $\mathcal{A}$ 's queries to his own encryption oracle. When eventually  $\mathcal{A}$  outputs a forgery,  $\mathcal{B}$  sends it for decryption to his own decryption oracle. If the decryption succeeds (that is, the oracle does not return  $|\perp\rangle\langle\perp|$ ), then  $\mathcal{B}$  outputs real, otherwise he outputs ideal.

The idea is the following: suppose the decryption of the forgery state succeeds (i.e., it does not decrypt to  $|\perp\rangle\langle\perp|$ ). This can happen in two cases:

1. we are in the QAE-Real game, and  $\mathcal{A}$  produced a valid forgery (i.e., he won the QUF-Forge game); or
2. we are in the QAE-Ideal game, and  $\mathcal{A}$  cheated by replaying an output of the encryption oracle (i.e., he won the QUF-Cheat game).

Recall that, by assumption,  $\mathcal{A}$  produces a valid forgery with probability at least  $\nu$  over cheating. Therefore the case 2. above happens with noticeable less probability than case 1., which is in fact the one  $\mathcal{B}$  “bets” on. Analogously, suppose the decryption fails. This can happen in two cases:

1. we are in the QAE-Real game, but  $\mathcal{A}$  produced an invalid forgery (i.e., he lost the QUF-Forge game); or
2. we are in the QAE-Ideal game, and  $\mathcal{A}$  did not cheat (i.e., he lost QUF-Cheat).



For the same reasoning as above, 2. is noticeably more likely than 1., which is in fact  $\mathcal{B}$ 's bet. More in detail, we have:

$$\begin{aligned} & \left| \Pr[\mathcal{B}(\text{QAE-Real}) \rightarrow \text{Real}] - \Pr[\mathcal{B}(\text{QAE-Ideal}) \rightarrow \text{Real}] \right| \\ &= \left| \Pr[\text{QAE-Real}] \cdot \Pr[\mathcal{A}(\text{QUF-Forge}) \rightarrow \text{win}] \right. \\ & \quad \left. - \Pr[\text{QAE-Ideal}] \cdot \Pr[\mathcal{A}(\text{QUF-Cheat}) \rightarrow \text{cheat}] \right| \\ &= \frac{1}{2} \left| \Pr[\mathcal{A}(\text{QUF-Forge}) \rightarrow \text{win}] - (\Pr[\mathcal{A}(\text{QUF-Forge}) \rightarrow \text{win}] - \nu) \right| = \frac{\nu}{2}, \end{aligned}$$

which is non-negligible.  $\square$

Finally, we consider the classical restriction  $\text{AE}'$  of QAE.

**Proposition 3.**  $\text{AE}' \iff \text{AE}$ .

*Proof.* The security notion  $\text{AE}'$  is given in terms of two experiments which are like the  $\text{AE-Real}$  and  $\text{AE-Ideal}$  experiments in Shrimpton's formulation of  $\text{AE}$  security, with the following difference:

1. in the modified  $\text{AE-Real}$  experiment, the decryption oracle does not reject non-fresh ciphertexts, i.e. it is unrestricted; and
2. in the modified  $\text{AE-Ideal}$  experiment, the decryption oracle does not always return  $\perp$ : in case it is queried on a non-fresh ciphertext, it decrypts correctly.

Since classically we can store and compare plaintexts and ciphertexts, it is easy to construct an efficient simulator able to switch between the experiments of  $\text{AE}$  and  $\text{AE}'$ , by inspecting  $\mathcal{A}$ 's decryption queries and reacting accordingly. Namely:

1. to switch from  $\text{AE}$  to  $\text{AE}'$ , record  $\mathcal{A}$ 's plaintexts and ciphertexts during encryption queries, and reply with the right plaintext whenever  $\mathcal{A}$  asks to decrypt a non-fresh ciphertext (otherwise, just send the query to the decryption oracle); and
2. to switch from  $\text{AE}'$  to  $\text{AE}$ , record  $\mathcal{A}$ 's received ciphertexts during encryption queries, and reply with  $\perp$  whenever  $\mathcal{A}$  asks to decrypt a non-fresh ciphertext (otherwise, just send the query to the decryption oracle).

This concludes the proof, as it shows the two cases to be equivalent.  $\square$

In particular,  $\text{AE}'$  is equivalent to  $\text{UF}$ . We provide evidence that a quantum analogue of this statement does not hold in the next section.

## 7 Constructions and Separations

In this section we exhibit constructions of  $\text{SKQES}$  that fulfill and separate the different security notions presented in the preceding sections. We begin by showing that augmenting a one-time scheme by a (perfectly) random function family using the construction in Definition 2 turns a  $\text{QCA}$  secure scheme into a  $\text{QAE}$  secure scheme. Then we will move on to show how to satisfy  $\text{QAE}$  with an efficiently implementable scheme. Recall that efficient  $\text{QCA}$ -secure  $\text{SKQES}$  can be constructed, e.g., from unitary two-designs like the Clifford group.

**Theorem 14.** *Let  $\Pi$  be a QCA-secure SKQES, and let  $f : \mathcal{K} \times \{0, 1\}^n \rightarrow \{0, 1\}^m$  be a random function family. Then the scheme  $\Pi^f$  in Definition 2 is QAE secure.*

*Proof.* We let  $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec})$  and  $\Pi^f = (\text{KeyGen}', \text{Enc}', \text{Dec}')$  where

1.  $\text{KeyGen}'(1^n)$  outputs a random function  $F$  from  $\{0, 1\}^n$  to  $\{0, 1\}^m$ ;
2.  $\text{Enc}'_F(X_M)$  outputs  $|s\rangle\langle s|_R \otimes \text{Enc}_{F(s)}(X)_C$ , where  $s \xleftarrow{\$} \{0, 1\}^n$ ;
3.  $\text{Dec}'_F(Y_{RC})$  first measures the  $R$  register to get outcome  $s'$ ; then it runs  $\text{Dec}_{F(s')}$  on register  $C$  and outputs the result.

Suppose  $\mathcal{A}$  is a QAE adversary against  $\Pi^f$ , i.e., a QPT algorithm with oracle access to  $\text{Enc}'_k$  and  $\text{Dec}'_k$ . Suppose  $\mathcal{A}$  makes  $\ell(n)$  queries to the oracle, where  $\ell$  is some polynomial function of  $n$ . We assume that the randomnesses  $s_i$  and the keys  $F(s_i)$  used for the scheme  $\Pi$  in the different encryption queries (for  $i = 1, \dots, \ell(n)$ ) are all distinct; this is true except with negligible probability.

Let us first analyze what happens in the QAE-Real experiment. Consider the  $i$ -th decryption oracle call. The decryption begins with a measurement of the  $R$  register, yielding some outcome  $s$  and thereby a key  $\bar{k} = F(s)$ . We can analyze the situation for each outcome  $s$  that occurs with non-negligible probability, separately. This is because if an adversary is successful, it is easy to see that there is also a modified successful adversary, that submits only decryption queries with a fixed string  $s$  in the randomness register.

Suppose first that  $\bar{k} = F(s) \neq F(s_i)$  for all  $i$ . In this case, the  $\Pi$ -encrypted part of the forgery candidate gets decrypted with a key different from all the ones used for encryption. We analyze the attack map  $\Lambda = \tilde{\mathcal{A}}(1^n)\text{Tr}_C$  against the QCA scheme  $\Pi$ , where  $\tilde{\mathcal{A}}$  is defined to first run  $\mathcal{A}$  until the  $i$ th decryption query, while answering each encryption query by sampling a fresh key for the scheme  $\Pi$ . Note that  $\Lambda$  does not use initial side information, therefore  $\sigma^{\text{acc}} := \Lambda^{\text{acc}}$  and  $\sigma^{\text{rej}} := \Lambda^{\text{rej}}$  are just positive semidefinite matrices whose trace sums to one.

According to Eq. (4) in the definition of QCA, the trace of  $\sigma^{\text{acc}}$  is the probability that the simulator applies the identity to the plaintext. The output of the attack map  $\Lambda$  does not depend on its input, i.e. the same holds for the effective map  $\Lambda^{\Pi}$  and hence for  $(\mathbb{1} - |\perp\rangle\langle\perp|)\Lambda^{\Pi}(\cdot)(\mathbb{1} - |\perp\rangle\langle\perp|)$ . Any such map is far from any non-negligible multiple of the identity channel so the trace of  $\sigma^{\text{acc}}$  is negligible according to Eq. 3. We have hence shown that the decryption oracle returns  $\perp$  with overwhelming probability, so we can take  $\sigma^{\text{crej}} = \text{Tr}_C \tilde{\mathcal{A}}(1^n)$ .

Let now  $s' = r_j$ , and write  $\mathcal{A} = \mathcal{A}_1 \text{Enc}_{\bar{k}} \mathcal{A}_0$ , splitting the adversary into two parts before and after the  $j$ -th encryption query. Let  $(\tilde{\mathcal{A}}_1)_{CE_1 \rightarrow CE_2}$  be defined analogous to  $\tilde{\mathcal{A}}$ .  $E_1$  and  $E_2$  are the internal memory registers of  $\mathcal{A}$  at the time of the  $j$ -th encryption query and the  $i$ -th decryption query, respectively.  $\Pi$  is QCA secure, implying that  $\tilde{\mathcal{A}}_1^{\Pi} = \mathbb{E}_{\bar{k}} \left[ \text{Dec}_{\bar{k}} \circ \tilde{\mathcal{A}}_1 \circ \text{Enc}_{\bar{k}} \right]$  fulfills:

$$\|(\tilde{\mathcal{A}}_1^{\Pi})_{ME_1 \rightarrow ME_2} - \text{id}_M \otimes (\tilde{\mathcal{A}}_1^{\text{acc}})_{E_1 \rightarrow E_2} - \perp \otimes (\tilde{\mathcal{A}}_1^{\text{rej}})_{E_1 \rightarrow E_2}\|_{\diamond} \leq \text{negl}(n), \quad (7)$$

where (using  $P_{\text{inv}} = \mathbb{1} - |\Phi_{\bar{k}, \bar{r}}\rangle\langle\Phi_{\bar{k}, \bar{r}}|$ ):

$$\begin{aligned}\tilde{\mathcal{A}}_1^{\text{acc}} &= \mathbb{E}_{\bar{k}, \bar{r}} \left[ \langle \Phi_{\bar{k}, \bar{r}} | V_{\bar{k}}^\dagger \tilde{\mathcal{A}}_1^{\text{acc}} (\text{Enc}_{\bar{k}; \bar{r}}^+(\phi_{MM'}^+) \otimes (\cdot)_{E_1}) V_{\bar{k}} | \Phi_{\bar{k}, \bar{r}} \rangle \right] \text{ and} \\ \tilde{\mathcal{A}}_1^{\text{rej}} &= \mathbb{E}_{\bar{k}, \bar{r}} \left[ \text{Tr}_{MM'T} P_{\text{inv}} V_{\bar{k}}^\dagger \tilde{\mathcal{A}}_1^{\text{acc}} (\text{Enc}_{\bar{k}; \bar{r}}^+(\phi_{MM'}^+) \otimes (\cdot)_{E_1}) V_{\bar{k}} \right].\end{aligned}\quad (8)$$

The form of the simulator in the reject case follows by using that the maximally entangled state is a point in the optimization defining the diamond norm in (3) and using the monotonicity of the trace norm under partial trace.

We now show indistinguishability of the real and ideal experiments by induction over the decryption queries. Since QCA implies IND, the two are indistinguishable before the first decryption query. Assume now that the two experiments cannot be distinguished using an algorithm that makes at most  $i - 1$  decryption queries. Consider  $\mathcal{A}$  running in the ideal experiment until right before the  $(i + 1)$ -th decryption query (or until the end, if  $i = \ell$ ). We make the same case distinction as before. In the first case the measurement in line (3) in the ideal decryption oracle in Experiment 5 never returns 0, i.e. the output is always reject. Therefore we can replace the  $i$ -th decryption oracle by the constant reject function, thereby reducing the number of decryption oracle calls of to  $i - 1$ . By the induction hypothesis, the contents of the internal register are therefore indistinguishable whether in the QAE-Real or in the QAE-Ideal experiment.

Turning to the second case, we make a very similar argument. We have  $s = s_j$ , i.e. the only encryption query where the measurement from line (3) in the definition of the ideal decryption oracle in Experiment 5 can possibly return 0 is the  $j$ -th. Here it is left to observe that the rest of the ideal decryption oracle implements exactly the same map as in the ideal world, i.e. the ones from Eqs. (7) and (8). Replacing the  $j$ -th encryption and the  $i$ -th decryption oracle call by this map, and using the induction hypothesis, we get that  $\mathcal{A}$  run until before the  $i + 1$ -th decryption oracle call cannot distinguish QAE-Real from QAE-Ideal. This ends the proof by induction.  $\square$

We now show how to satisfy QAE efficiently, by means of a post-quantum-secure pseudorandom function.

**Corollary 2.** *Let  $\Pi$  be a QCA-secure SKQES that satisfies Condition 1, and let  $f$  be a pqPRF. Then the scheme  $\Pi^f$  (from Definition 2) satisfies QAE.*

*Proof.* As a contradiction, suppose there exists a QPT algorithm  $\mathcal{A}$  that distinguishes QAE-Real from QAE-Ideal. We claim that this also holds if  $f$  is replaced with a completely random function family  $\mathcal{F}$ . If  $\mathcal{A}$  cannot break the random scheme  $\Pi^{\mathcal{F}}$ , then we can build a distinguisher for  $f$  versus  $\mathcal{F}$ , as follows. What we would like to do is the following. Given an oracle  $\mathcal{O}$ , we:

1. choose a random bit  $b \xleftarrow{\$} \{0, 1\}$ ;
2. if  $b = 0$ , we simulate the  $\text{QAE-Real}(\Pi^{\mathcal{O}}, \mathcal{A}, n)$  experiment using our oracle;
3. if  $b = 1$ , we simulate the  $\text{QAE-Ideal}(\Pi^{\mathcal{O}}, \mathcal{A}, n)$  experiment using our oracle;
4. output  $b \oplus s$  where  $s$  is the output of  $\mathcal{A}$ .

This may at first not seem possible using the classical oracle we are provided with, as the ideal decryption oracle has to implement the unitary  $V_{\bar{k}}^\dagger$ , which

seems to require superposition access to the random/pseudorandom function. However, observe that steps 5–11 of Experiment 2 commute with a measurement of the randomness register  $R$  in the computational basis, and afterwards this register is discarded. Therefore the outcome of the experiment is not changed by first measuring the register  $R$ , which yields an outcome  $r$ . Then the modified challenger can use classical oracle access to the random/pseudorandom function to implement  $V_k^\dagger$  on the measured input state.

Note that, if  $\Pi^{\mathcal{O}}$  is secure, then  $b$  and  $s$  are independent (up to negligible terms) and  $b \oplus s$  is a fair coin. If  $\Pi^{\mathcal{O}}$  is insecure, then it deviates from uniform by the QUF distinguishing advantage of  $\mathcal{A}$ . This yields a distinguisher between the case  $\mathcal{O} = f$  and  $\mathcal{O} = \mathcal{F}$ . The claim then follows from Theorem 14.  $\square$

In particular, the scheme family  $2\text{desTag}^{\text{pqPRF}}$  is sufficient for QAE. We remark that the proof uses the fact that, given classical oracle access to  $f$ , the scheme  $\Pi^f$  is efficiently implementable in the sense of Condition 1 – regardless of the nature of the family  $f$ . Of course, in the special case where  $f$  is a pqPRF, then  $\Pi^f$  simply satisfies Condition 1 without any need for oracles.

As QAE implies both QUF and QIND-CCA2 (see Theorems 11 and 13), we have the following corollary.

**Corollary 3.** *Let  $\Pi$  be a QCA-secure SKQES that satisfies Condition 1, and let  $f$  be a pqPRF. Then the scheme  $\Pi^f$  (from Definition 2) satisfies QUF and QIND-CCA2.*

We can also show how to satisfy bounded-query unforgeability, i.e.,  $\text{QUF}_t$ . Recall that a  $t$ -wise independent function is a deterministic, efficiently computable keyed function family  $\{f_k\}_k$  which appears random to any algorithm (of unbounded computational power) which gets classical oracle access to  $f_k$  for uniformly random  $k$ , and can make at most  $t$  queries. One can apply the proof technique of Corollary 2 and Theorem 14 to obtain the following.

**Corollary 4.** *Let  $\Pi$  be a QCA-secure SKQES, and let  $f$  be a  $t$ -wise independent function family. Then the scheme  $\Pi^f$  (as defined in Definition 2) satisfies  $\text{QUF}_t$ .*

*Proof (Sketch).* If there exists a QPT  $\mathcal{A}$  which can break  $\text{QUF}_t$  for  $\Pi^f$  using  $t$ -many queries, then it also breaks  $\Pi^{\mathcal{F}}$  where  $\mathcal{F}$  is a random function. If not, we construct an oracle distinguisher for  $\mathcal{O} = f$  versus  $\mathcal{O} = \mathcal{F}$  which simulates  $\mathcal{A}$  in one of the two games (each with probability  $1/2$ ) and outputs a bit which is biased depending on  $\mathcal{O}$ . Note that we only need  $t$  queries to do this, since we only run one of the games (and not both). It then remains to invoke Theorem 14, and observe that Theorem 13 holds in the case of a bounded number of queries.  $\square$

**Separations.** While QAE implies QIND-CCA2 according to Theorem 11, the converse does not hold. In fact, consider any QAE secure scheme and modify the decryption function by replacing the reject symbol by a fixed plaintext, e.g. the all zero state. Such a scheme is certainly still QIND-CCA2 secure, as any adversary against it can be used against the original scheme by simulating the

modified one. The modified scheme is, however, manifestly not QAE as it never outputs  $\perp$ . The same reasoning works for QUF in place of QAE.

**Proposition 4.** QIND-CCA2  $\not\approx$  QUF, and therefore QIND-CCA2  $\not\approx$  QAE.

Finally, we turn to the relationship of QAE and QUF, and propose a separation as follows. Let  $\Pi$  be a scheme that fulfills cQCA (Definition 6) for trivial register  $\tilde{B}$ , but can be broken using an efficient attack with nontrivial  $\tilde{B}$ . For any PRF  $f$ ,  $\Pi^f$  is clearly QUF, as the security notion ignores side information. It can however not be QAE, as QAE implies cQCA.

## 8 Discussion

In this work, we presented four new security notions for symmetric key quantum encryption: QCA, QUF, QIND-CCA2 and QAE. While we have already made significant progress on understanding these notions, a number of open questions remain. A few are as follows. Does an encryption scheme as discussed below Proposition 4 exist, proving QUF  $\not\approx$  QAE? If so, does QUF imply QIND-CCA2 or QIND-CCA1? Classically, unforgeability and IND-CCA2 imply AE; does this hold quantumly as well? Finally, is there a scheme that satisfies QIND-CCA2 but cannot be upgraded to QAE by simply modifying the decryption function?

**Acknowledgements.** The authors would like to thank Anne Broadbent, Frédéric Dupuis, Yfke Dulek, Alex Russell, Christian Schaffner, and Fang Song for insightful discussions about the problems solved in this work. Part of this work was done while T.G. was supported by the TU Darmstadt. Part of this work was done while G.A. and C.M. were at QMATH, University of Copenhagen. Part of this work was sponsored by the COST CryptoAction IC1306. T.G. acknowledges financial support from the European Commissions PERCY grant (agreement 321310). G.A. and C.M. acknowledge financial support from the European Research Council (ERC Grant Agreement no 337603), the Danish Council for Independent Research (Sapere Aude) and VILLUM FONDEN via the QMATH Centre of Excellence (Grant No. 10059). This work is part of the research programme “Cryptography in the Quantum Age” with project number 639.022.519, which is financed by the Netherlands Organisation for Scientific Research (NWO).

## References

1. Aaronson, S., Gottesman, D.: Improved simulation of stabilizer circuits. CoRR, quant-ph/0406196 (2004)
2. Aharonov, D., Ben-Or, M., Eban, E.: Interactive proofs for quantum computations. In: Proceedings of the Innovations in Computer Science - ICS 2010, Tsinghua University, Beijing, China, 5–7 January 2010, pp. 453–469 (2010)
3. Alagic, G., Broadbent, A., Fefferman, B., Gagliardini, T., Schaffner, C., St. Jules, M.: Computational security of quantum encryption. In: Nascimento, A.C.A., Barreto, P. (eds.) ICITS 2016. LNCS, vol. 10015, pp. 47–71. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-49175-2\\_3](https://doi.org/10.1007/978-3-319-49175-2_3)
4. Alagic, G., Gagliardini, T., Majenz, C.: Unforgeable quantum encryption. Cryptology ePrint Archive, Report 2017/960 (2017). <https://eprint.iacr.org/2017/960>

5. Alagic, G., Majenz, C.: Quantum non-malleability and authentication. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017. LNCS, vol. 10402, pp. 310–341. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-63715-0\\_11](https://doi.org/10.1007/978-3-319-63715-0_11)
6. Ambainis, A., Bouda, J., Winter, A.: Non-malleable encryption of quantum information. *J. Math. Phys.* **50**(4), 042106 (2009)
7. Ambainis, A., Mosca, M., Tapp, A., de Wolf, R.: Private quantum channels. In: 41st Annual Symposium on Foundations of Computer Science, FOCS 2000, Redondo Beach, California, USA, 12–14 November 2000, pp. 547–553 (2000)
8. Barak, B.: Cs127 course notes, Chap. 6. <http://www.boazbarak.org/cs127/chap06-CCA.pdf>. Accessed 7 Sept 2017
9. Barnum, H., Crépeau, C., Gottesman, D., Smith, A.D., Tapp, A.: Authentication of quantum messages. In: Proceedings of the 43rd Symposium on Foundations of Computer Science (FOCS 2002), Vancouver, BC, Canada, 16–19 November 2002, pp. 449–458 (2002)
10. Bellare, M., Namprepmpre, C.: Authenticated encryption: relations among notions and analysis of the generic composition paradigm. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 531–545. Springer, Heidelberg (2000). [https://doi.org/10.1007/3-540-44448-3\\_41](https://doi.org/10.1007/3-540-44448-3_41)
11. Boneh, D., Zhandry, M.: Quantum-secure message authentication codes. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 592–608. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-38348-9\\_35](https://doi.org/10.1007/978-3-642-38348-9_35)
12. Boneh, D., Zhandry, M.: Secure signatures and chosen ciphertext security in a quantum computing world. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8043, pp. 361–379. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-40084-1\\_21](https://doi.org/10.1007/978-3-642-40084-1_21)
13. Brandão, F.G.S.L., Harrow, A.W., Horodecki, M.: Local random quantum circuits are approximate polynomial-designs. *Commun. Math. Phys.* **346**(2), 397–434 (2016)
14. Broadbent, A., Jeffery, S.: Quantum homomorphic encryption for circuits of low T-gate complexity. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015. LNCS, vol. 9216, pp. 609–629. Springer, Heidelberg (2015). [https://doi.org/10.1007/978-3-662-48000-7\\_30](https://doi.org/10.1007/978-3-662-48000-7_30)
15. Broadbent, A., Wainwright, E.: Efficient simulation for quantum message authentication. In: Nascimento, A.C.A., Barreto, P. (eds.) ICITS 2016. LNCS, vol. 10015, pp. 72–91. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-49175-2\\_4](https://doi.org/10.1007/978-3-319-49175-2_4)
16. DiVincenzo, D.P., Leung, D.W., Terhal, B.M.: Quantum data hiding. *IEEE Trans. Inf. Theory* **48**(3), 580–598 (2002)
17. Dulek, Y., Schaffner, C., Speelman, F.: Quantum homomorphic encryption for polynomial-sized circuits. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9816, pp. 3–32. Springer, Heidelberg (2016). [https://doi.org/10.1007/978-3-662-53015-3\\_1](https://doi.org/10.1007/978-3-662-53015-3_1)
18. Dupuis, F., Nielsen, J.B., Salvail, L.: Secure two-party quantum evaluation of unitaries against specious adversaries. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 685–706. Springer, Heidelberg (2010). [https://doi.org/10.1007/978-3-642-14623-7\\_37](https://doi.org/10.1007/978-3-642-14623-7_37)
19. Dupuis, F., Nielsen, J.B., Salvail, L.: Actively secure two-party evaluation of any quantum operation. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 794–811. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-32009-5\\_46](https://doi.org/10.1007/978-3-642-32009-5_46)

20. Gagliardoni, T., Hülsing, A., Schaffner, C.: Semantic security and indistinguishability in the quantum world. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016. LNCS, vol. 9816, pp. 60–89. Springer, Heidelberg (2016). [https://doi.org/10.1007/978-3-662-53015-3\\_3](https://doi.org/10.1007/978-3-662-53015-3_3)
21. Garg, S., Yuen, H., Zhandry, M.: New security notions and feasibility results for authentication of quantum data. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017. LNCS, vol. 10402, pp. 342–371. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-63715-0\\_12](https://doi.org/10.1007/978-3-319-63715-0_12)
22. Gottesman, D.: The Heisenberg representation of quantum computers. arXiv quant-ph/9807006 (1998)
23. Gottesman, D.: Uncloneable encryption. *Quantum Inf. Comput.* **3**(6), 581–602 (2003)
24. Hayden, P., Leung, D.W., Mayers, D.W.: The universal composable security of quantum message authentication with key recycling. arXiv quant-ph/1610.09434 (2016)
25. Katz, J., Lindell, Y.: *Introduction to Modern Cryptography*, 2nd edn. CRC Press, Boca Raton (2014)
26. Nielsen, M.A., Chuang, I.L.: *Quantum Computation and Quantum Information: 10th Anniversary Edition*, 10th edn. Cambridge University Press, New York (2011)
27. Portmann, C.: Quantum authentication with key recycling. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10212, pp. 339–368. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-56617-7\\_12](https://doi.org/10.1007/978-3-319-56617-7_12)
28. Shrimpton, T.: A characterization of authenticated-encryption as a form of chosen-ciphertext security. *IACR Cryptology ePrint Archive* 2004:272 (2004)
29. Winter, A.J.: Coding theorem and strong converse for quantum channels. *IEEE Trans. Inf. Theory* **45**(7), 2481–2485 (1999)