

Public-Key Encryption with Quantum Keys

Khashayar Barooti¹, Alex B. Grilo², Loïs Huguenin-Dumittan¹, Giulio Malavolta³, Or Sattath⁴,
Quoc-Huy Vu², and Michael Walter⁵

¹ EPFL, Lausanne, Switzerland

² Sorbonne Université, CNRS, LIP6, France

³ Max-Planck Institute in Security and Privacy, Bochum, Germany

⁴ Computer Science Department, Ben-Gurion University of the Negev, Israel

⁵ Faculty of Computer Science, Ruhr University Bochum, Germany

Abstract. In the framework of Impagliazzo’s five worlds, a distinction is often made between two worlds, one where public-key encryption exists (Cryptomania), and one in which only one-way functions exist (MiniCrypt). However, the boundaries between these worlds can change when quantum information is taken into account. Recent work has shown that quantum variants of oblivious transfer and multi-party computation, both primitives that are classically in Cryptomania, can be constructed from one-way functions, placing them in the realm of quantum MiniCrypt (the so-called MiniQCrypt). This naturally raises the following question: *Is it possible to construct a quantum variant of public-key encryption, which is at the heart of Cryptomania, from one-way functions or potentially weaker assumptions?*

In this work, we initiate the formal study of the notion of quantum public-key encryption (qPKE), i.e., public-key encryption where keys are allowed to be quantum states. We propose new definitions of security and several constructions of qPKE based on the existence of one-way functions (OWF), or even weaker assumptions, such as pseudorandom function-like states (PRFS) and pseudorandom function-like states with proof of destruction (PRFSPD). Finally, to give a tight characterization of this primitive, we show that computational assumptions are necessary to build quantum public-key encryption. That is, we give a self-contained proof that no quantum public-key encryption scheme can provide information-theoretic security.

1 Introduction

The use of quantum resources to enable cryptographic tasks under weaker assumptions than classically needed (or even *unconditionally*) were actually the first concrete proposals of quantum computing, with the seminal quantum money protocol of Wiesner [Wie83] and the key-exchange protocol of Bennett and Brassard [BB84]. Ever since, the field of quantum cryptography has seen a surge of primitives that leverage quantum information to perform tasks that classically require stronger assumptions, or are downright impossible. Recent works [BCKM21, GLSV21] have shown that there exist quantum protocols for oblivious transfer, and therefore arbitrary multi-party computation (MPC), based solely on the existence of one-way functions (OWF) [BCKM21, GLSV21], or pseudorandom states (PRS) [JLS18], which potentially entail even weaker computational assumptions [Kre21, KQST22]. It is well-known that, classically, oblivious transfer and MPC are “Cryptomania” objects, i.e., they can only be constructed from more structured assumptions that imply public-key encryption (PKE). Thus, the above results seem to challenge the boundary between Cryptomania and MiniCrypt, in the presence of quantum information. Motivated by this state of affairs, in this work we investigate the notion of *PKE itself*, the heart of Cryptomania, through the lenses of quantum computing. That is, we ask the following question:

Does public-key encryption (PKE) belong to MiniQCrypt?

Known results around this question are mostly negative: It is known that PKE cannot be constructed in a black-box manner from OWFs [IR90], and this result has been recently re-proven in the more challenging setting where the encryption or decryption algorithms are quantum [ACC⁺22]. However, a tantalizing possibility left open by these works is to realize PKE schemes from OWFs (or weaker assumptions), where public-key or ciphertexts are quantum states.

1.1 Our results

In this work we initiate the systematic study of quantum public-key encryption (qPKE), i.e., public-key encryption where public-keys and ciphertexts are allowed to be quantum states. We break down our contributions as follows.

1. Definitions. We provide a general definitional framework for qPKE, where both the public-key and ciphertext might be general quantum states. In the classical setting, there is no need to provide oracle access to the encryption, since the public-key can be used to implement that. In contrast, if the public-key is a quantum state, it might be measured during the encryption procedure, and the ciphertexts might depend on the measurement outcome. In fact, this is the approach taken in some of our constructions. This motivates a stronger security definition, similar to the classical counterpart, in which the adversary gets additional access to an encryption oracle that uses the same quantum public-key that is used during the challenge phase. We define IND-CPA-EO (respectively, IND-CCA-EO) security by adding the encryption oracle (EO) to the standard IND-CPA (respectively, IND-CCA) security game.

2. Constructions. With our new security definition at hand, we propose three protocols for implementing qPKE from OWF and potentially weaker assumptions, each with its own different advantages and disadvantages. More concretely, we show the existence of:

1. A qPKE scheme with quantum public-keys and classical ciphertexts that is IND-CCA-EO⁶ secure, based on post-quantum OWF, in Section 4.1.
2. A qPKE scheme with quantum public-key and quantum ciphertext that is IND-CCA1 secure, based on pseudo-random function-like states (PRFS) with super-logarithmic input-size⁷, in Section 4.2. Since this scheme is not EO secure, each quantum public-key enables the encryption of a single message.
3. A qPKE scheme with quantum public-key and classical ciphertext that is IND-CPA-EO secure based on pseudo-random function-like states with proof of destruction (PRFSPDs), in Section 5.

We wish to remark that it has been recently shown that OWF imply PRFS with super-logarithmic input-size [AQY22] and PRFSPDs [BSS23]. Therefore, the security of the second and third protocols is based on a potentially weaker cryptographic assumption than the first one. Furthermore, PRFS with super-logarithmic input-size are *oracle separated* from one-way functions [Kre21]; therefore, our second result shows a black-box separation between a certain form of quantum public-key encryption and one-way functions. On the other hand, for the other two constructions, even if the

⁶ Throughout this paper, unless explicitly specified, by IND-CCA we refer to the notion of adaptive IND-CCA2 security.

⁷ Note that PRS implies PRFS with logarithmic size inputs, but no such implication is known for super-logarithmic inputs.

public-key is a quantum state, the ciphertexts are classical and, furthermore, one quantum public-key can be used to encrypt multiple messages. The first protocol is much simpler to describe and understand since it only uses standard (classical) cryptographic objects. Moreover, we show that this scheme guarantees the notion of adaptive CCA2 security and is the only scheme that achieves perfect correctness.

3. Lower Bounds. To complete the picture, we demonstrate that *information-theoretically secure* qPKE does not exist. Due to the public-keys being quantum states, this implication is much less obvious than for the classical case. In fact, some of the existing constructions of qPKE [Got05] have been conjectured to be unconditionally secure, a conjecture that we invalidate in this work. While this general statement follows by known implications in the literature (see Section 6 for more details), in this work we present a self-contained proof of this fact, borrowing techniques from shadow tomography, which we consider to be of independent interest.

1.2 Technical overview

In this section, we provide a technical overview of our results. In Section 1.2.1, we explain the challenges and choices to define qPKE and its security definition. In Section 1.2.2, we present 3 instantiations of qPKE, each based on a different assumption and with different security guarantees. Ultimately, Section 1.2.3 is dedicated to the impossibility of information-theoretically secure qPKE and a high-level overview of the proof technique.

1.2.1 Definitions of qPKE

In order to consider public-key encryption schemes with quantum public-keys, we need to revisit the traditional security definitions. In the case of quantum public-keys, there are several immediate issues that require revision.

The first issue is related to the access the adversary is given to the public-key. In the classical-key case (even with quantum ciphertexts), the adversary is given the classical public-key pk . Given a single quantum public-key, one cannot create arbitrary number of copies of the quantum public-key, due to no-cloning. Hence, to naturally extend notions such as IND-CPA security, we provide multiple copies of the quantum public-key to the adversary (via the mean of oracle access to the quantum public-key generation algorithm).

The second issue concerns the quantum public-key's *reusability*. Classically, one can use the public-key to encrypt multiple messages. With quantum public-keys, this might not be the case: the quantum public-key might be consumed during the encryption. In a non-reusable scheme, the user needs a fresh quantum public-key for every plaintext they wish to encrypt. This is not only a theoretical concern: in the PRFS-based construction (see Section 4.2), part of the quantum public-key is sent as the (quantum) ciphertext, so clearly, this construction is *not* reusable.

Thirdly, it could be the case that in a reusable scheme, each encryption call changes the public-key state ρ_{ppk} in an irreversible way. Hence, we make a syntactic change: $\text{Enc}(\rho_{\text{ppk}}, m)$ outputs (ρ'_{ppk}, c) , where c is used as the ciphertext and ρ'_{ppk} is used as the key to encrypt the next message. Note that in this scenario the updated public-key is not publicly available anymore and is only held by the party who performed the encryption.

Lastly, the syntactic change mentioned above also has security effects. Recall that classically, there is no need to give the adversary access to an encryption oracle, since the adversary can generate

encryption on their own. Alas, with quantum public-keys, the distribution of ciphers might depend on the changes that were made to the quantum public-key by the challenger whenever the key is used to encrypt several messages. Therefore, for reusable schemes, we define two new security notions, denoted CPA-EO and CCA-EO, that are similar to CPA and CCA but where the adversary is given access to an encryption oracle (EO). We note there are several works considering the notions of chosen-ciphertext security in the quantum setting, because it is not clear how to prevent the adversary from querying the challenge ciphertext, if it contains a quantum states. However, we only consider CCA-security for schemes with classical ciphertexts, and therefore this issue does not appear in this work.

Pure vs Mixed States. We mention explicitly that we require our public-keys to be *pure states*. This is motivated by the following concern: there is no general method to authenticate quantum states. One proposal to ensure that the certificate authority (CA) is sending the correct state is to distribute various copies of the keys to different CAs and test whether they are all sending the same state [Got05]. This ensures that, as long as at least one CA is honest, the user will reject a malformed key with some constant probability. However, this argument crucially relies on the public-key being a pure state (in which case comparison can be implemented with a SWAP-test). On the other hand, if the public-key was a mixed state, there would be no way to run the above test without false positives.

We also mention that, if mixed states are allowed, then there is a trivial construction of qPKE from any given symmetric encryption scheme (SKE.key-gen, SKE.Enc, SKE.Dec), as also observed in [MY22a, Theorem C.6], which we describe in the following. To generate the keys, we use the output of SKE.key-gen as the secret-key and use it to create the uniform mixture

$$\frac{1}{2^n} \sum_{x \in \{0,1\}^n} |x\rangle\langle x| \otimes |\text{Enc}_{sk}(x)\rangle\langle \text{Enc}_{sk}(x)| \quad (1)$$

as the public-key. The ciphertext corresponding to a message m is given by $(\text{Enc}_x(m), \text{Enc}_{sk}(x))$. To decrypt, the decryptor would first recover x by decrypting the second element in the ciphertext using sk , and then recover m by decrypting the first item using x as the secret key.

1.2.2 Constructions for qPKE

As previously mentioned, we propose in this work three schemes for qPKE, based on three different assumptions, each providing a different security guarantee.

qPKE from OWF. Our first instantiation of qPKE is based on the existence of post-quantum OWFs. For this construction, we aim for the strong security notion of indistinguishability against adaptive chosen ciphertext attacks with encryption oracle referred to as IND-CCA-EO. We start with a simple bit-encryption construction that provides IND-CCA security and we discuss how one can modify the scheme to encrypt multi-bit messages and also provide EO security.

Our first scheme assumes the existence of a *quantum-secure pseudorandom function (PRF)*, which can be built from quantum-secure one-way functions [Zha12]. Given a PRF ensemble $\{f_k\}_k$, the public key consists of a pair of pure quantum states $qp\mathcal{K} = (|qp\mathcal{K}_0\rangle, |qp\mathcal{K}_1\rangle)$ and the secret key consists of a pair of bit-strings $dk = (dk_0, dk_1)$ such that, for all $b \in \{0, 1\}$,

$$|qp\mathcal{K}_b\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x, f_{dk_b}(x)\rangle,$$

where f_k denotes the quantum-secure PRF keyed by k . To encrypt a bit b , one simply measures all qubits of $|qp\kappa_b\rangle$ in the computational basis. The result takes the form $(x, f_{dk_b}(x))$ for some uniformly random $x \in \{0, 1\}^n$ and this is returned as the ciphertext, i.e., $(qc_0, qc_1) = (x, f_{dk_b}(x))$.

To decrypt a ciphertext (qc_0, qc_1) , we apply both f_{dk_0} and f_{dk_1} to qc_0 and return the value of $b \in \{0, 1\}$ such that $f_{dk_b}(qc_0) = qc_1$. In case this does happen for neither or both of the keys, the decryption aborts.

The IND-CCA security of the simple bit-encryption scheme can be proven with a hybrid argument (see Appendix A). However, there are a few caveats to the scheme that can be pointed out. First, the scheme is not reusable. It can be easily noticed that after using a public-key for an encryption, the public-key state collapses, meaning that all the subsequent encryption calls are derandomized. This would mean if the same public-key is reused, it can not even guarantee IND-CPA security as the encryption is deterministic.

The second issue is lifting this CCA-secure bit-encryption scheme to a many-bit CCA-secure encryption scheme. Note that although not trivial, as proven by Myers and Shelat [Ms09], classically it is possible to construct CCA-secure many-bit encryption from CCA-secure bit-encryption. However, the argument cannot be extended to qPKE in a generic way. The main issue is that the construction from [Ms09], similar to the Fujisaki-Okamoto transform, derandomizes the encryption procedure for some fixed random coins. Later these fixed random coins are encrypted and attached to the ciphertext, so that the decryptor can re-encrypt the plaintext to make sure they were handed the correct randomness. Looking at our construction, it is quite clear that it is not possible to derandomize the encryption procedure as the randomness is a consequence of the measurement.

Let us show how the same approach can be modified to circumvent the issues mentioned. Our main observation is that we can use public-keys of the form mentioned before for a key agreement stage and then use the agreed key to encrypt many-bit messages with a symmetric-key encryption scheme (SKE). Let us elaborate. Let $\{f_k\}_k$ be a PRF family and $(SE.Enc, SE.Dec)$ be a symmetric-key encryption scheme. Note that quantum-secure one-way functions imply a quantum-secure PRF [Zha12], and post-quantum IND-CCA symmetric encryption [BZ13a]⁸. Consider the following scheme: the secret key dk is a uniformly random key for the PRF, and for a fixed dk , the quantum public-key state is

$$|qp\kappa_{dk}\rangle = \frac{1}{\sqrt{2^\lambda}} \sum_{x \in \{0,1\}^\lambda} |x\rangle |f_{dk}(x)\rangle. \quad (2)$$

The encryption algorithm will then measure $|qp\kappa_{dk}\rangle$ in the computational basis leading to the outcome $(x^*, y^* = f_{dk}(x^*))$. The ciphertext of a message m is given by $(x^*, SE.Enc(y^*, m))$. To decrypt a ciphertext (\hat{x}, \hat{c}) , we first compute $\hat{y} = f_{dk}(\hat{x})$ and return $\hat{m} = SE.Dec(f_{dk}(\hat{x}), \hat{c})$.

We emphasize that this scheme is reusable since it allows the encryption of many messages using the same measurement outcome $(x^*, f_{dk}(x^*))$. Using a hybrid argument, it can be shown that if the underlying SKE guarantees IND-CCA security, this construction fulfills our strongest security notion, i.e. IND-CCA-EO security. A formal description of the scheme, along with a security proof can be found in Section 4.1.

QPKE from PRFS. The second construction we present in this paper is an IND-CCA1 secure public-key scheme based on the existence of pseudorandom function-like state generators. Our

⁸ IND-CCA SKE can be built from an IND-CPA SKE and a MAC using the encrypt-then-MAC paradigm.

approach is based on first showing bit-encryption, and the discussion regarding how to lift that restriction is discussed in Section 4.2. The ciphertexts generated by our scheme are quantum states, and as the public-keys of this construction are not reusable, we do not consider the notion of EO security. A family of states $\{|\psi_{k,x}\rangle\}_{k,x}$ is pseudo-random function-like [AQY22] if

1. There is a quantum polynomial-time algorithm Gen such that

$$\text{Gen}(k, \sum_x \alpha_x |x\rangle) = \sum_x \alpha_x |x\rangle |\psi_{k,x}\rangle, \text{ and}$$

2. No QPT adversary can distinguish $(|\psi_1\rangle, \dots, |\psi_\ell\rangle)$ from $(|\phi_1\rangle, \dots, |\phi_\ell\rangle)$, where $|\psi_i\rangle = \sum_x \alpha_x^i |x\rangle |\psi_{k,x}\rangle$, $|\phi_i\rangle = \sum_x \alpha_x^i |x\rangle |\phi_x\rangle$, and $\{|\phi_x\rangle\}_x$ are Haar random states and the states $|\sigma_i\rangle = \sum_x \alpha_x^i |x\rangle$ are chosen by the adversary.

We continue by providing a high-level description of the scheme. The key generation algorithm picks a uniform PRFS key dk and generates the corresponding public-keys as stated below:

$$\frac{1}{\sqrt{2^\lambda}} \sum_{x \in \{0,1\}^\lambda} |x\rangle |\psi_{\text{dk},x}\rangle^{\otimes n}, \quad (3)$$

where $\{|\psi_{k,x}\rangle\}_{k,x}$ is a PRFS family, the size of the input x is super-logarithmic in the security parameter and n is a polynomial in the security parameter.

To encrypt a bit m , the encryptor will then measure the first register of $|qp\kappa\rangle$ to obtain x^* and the residual state after this measurement will be of form $|x^*\rangle |\psi_{\text{dk},x^*}\rangle^{\otimes n}$. They also sample a uniform key dk_1 and compute the state $|\psi_{\text{dk}_1,x^*}\rangle$ then compute the ciphertext $c = (x^*, \rho)$ where

$$\rho = \begin{cases} |\psi_{\text{dk},x^*}\rangle^{\otimes n}, & \text{if } m = 0 \\ |\psi_{\text{dk}_1,x^*}\rangle^{\otimes n}, & \text{if } m = 1 \end{cases}. \quad (4)$$

To decrypt a ciphertext $(\hat{x}, \hat{\rho})$, we first compute n copies of the state $|\psi_{\text{dk},\hat{x}}\rangle$ and performs a SWAP tests between each copy and the subsystems of $\hat{\rho}$ with the same size as $|\psi_{\text{dk},\hat{x}}\rangle$. If all the SWAP tests return 0 the decryption algorithm returns $\hat{m} = 0$ and otherwise it returns $\hat{m} = 1$. For a large enough n , our scheme achieves statistical correctness.

We prove that this construction guarantees IND-CCA1 security by a hybrid argument in Section 4.2. We emphasize that as the ciphertexts of the scheme are quantum states it is challenging to define adaptive CCA2 security.

QPKE from PRFSPDs. Our third scheme is based on pseudo-random function-like states with proof of destruction (PRFSPDs), which was recently defined in [BBSS23]. The authors extended the notion of PRFS to pseudo-random function-like states with proof of destruction, where we have two algorithms $\mathcal{Destruct}$ and \mathcal{Ver} , which allows us to verify if a copy of the PRFS was destroyed.

We will discuss now how to provide non-reusable CPA security security⁹ of the encryption of a one-bit message and we discuss later how to use it to achieve reusable security, i.e., CPA-EO security. The quantum public-key in this simplified case is

$$\frac{1}{\sqrt{2^\lambda}} \sum_{x \in \{0,1\}^\lambda} |x\rangle |\psi_{\text{dk},x}\rangle. \quad (5)$$

⁹ Meaning that one can only encrypt once using a $|qp\kappa\rangle$.

The encryptor will then measure the first register of $|qp\kappa\rangle$ and the post-measurement state is $|x^*\rangle|\psi_{dk,x^*}\rangle$. The encryptor will then generate a (classical) proof of destruction $\pi = \mathcal{Destruct}(|\psi_{dk,x^*}\rangle)$. The encryption procedure also picks dk_1 uniformly at random, generated $|\psi_{dk_1,x^*}\rangle$ and generates the proof of destruction $\pi' = \mathcal{Destruct}(|\psi_{dk_1,x^*}\rangle)$. The corresponding ciphertext for a bit b is given by $c = (x^*, y)$, where

$$y = \begin{cases} \pi', & \text{if } b = 0 \\ \pi, & \text{if } b = 1 \end{cases}.$$

The decryptor will receive some value (\hat{x}, \hat{y}) and decrypt the message $\hat{b} = \mathcal{Ver}(dk, \hat{x}, \hat{y})$. The proof of the security of the aforementioned construction follows from a hybrid argument reminiscent of the security proof of the previous schemes (see Section 5). Notice that repeating such a process in parallel trivially gives a one-shot security of the encryption of a string m and moreover, such an encryption is classical. Therefore, in order to achieve IND-CPA-EO secure qPKE scheme, we can actually encrypt a secret key sk that is chosen by the encryptor, and send the message encrypted under sk . We leave the details of such a construction and its proof of security to Section 5.

1.2.3 Impossibility of Information-Theoretically Secure qPKE

So far, we have established that qPKE can be built from assumptions weaker than the ones required for the classical counterpart, and potentially even weaker than those needed to build secret-key encryption classically. This naturally leads to the question of whether it is possible to build an information-theoretically secure qPKE. In the following, we present a self-contained proof of this fact, using techniques from the literature on shadow tomography. Although proving the impossibility for classical PKE is immediate, there are a few challenges when trying to prove a result of a similar flavor for qPKE. Even when considering security against a computationally unbounded adversary, there is a limitation that such adversary has, namely, they are only provided with polynomially many copies of the public-key.

The first step of the proof is reducing winning the IND-CPA game to finding a secret-key/public-key pair $(dk, |qp\kappa_{dk}\rangle)$ such that

$$\langle qp\kappa^* | qp\kappa_{dk} \rangle \approx 1.$$

In other words, we show that if $|qp\kappa_{dk}\rangle$ is relatively close to $|qp\kappa^*\rangle$, there is a good chance that dk can decrypt ciphertexts encrypted by $|qp\kappa^*\rangle$ correctly. A formal statement and the proof of this argument can be found in Lemma 1.

Given this lemma, the second part of the proof consists in constructing an adversary that takes polynomially many copies of $|qp\kappa^*\rangle$ as input and outputs $(dk, |qp\kappa_{dk}\rangle)$ such that $|qp\kappa_{dk}\rangle$ is relatively close to $|qp\kappa^*\rangle$. The technique to realize this adversary is *shadow tomography*, which shows procedures to estimate the values $\langle qp\kappa_{dk} | qp\kappa^* \rangle$ for all $(|qp\kappa_{dk}\rangle, dk)$ pairs. Note that doing this naively, i.e. by SWAP-testing multiple copies of $|qp\kappa^*\rangle$ with each $|qp\kappa_{dk}\rangle$, would require exponentially many copies of the public-key $|qp\kappa^*\rangle$. The way we circumvent this problem is by using the a recent result by Huang, Kueng, and Preskill [HKP20]. Informally, this theorem states that for M rank 1 projective measurements O_1, \dots, O_M and an unknown n -qubit state ρ , it is possible to estimate $\text{Tr}(O_i\rho)$ for all i , up to precision ϵ , by only performing $T = O(\log(M)/\epsilon^2)$ single-copy random Clifford measurements on ρ .

Employing this theorem, we show that a computationally unbounded adversary can estimate all the values $\langle qp\kappa_{dk} | qp\kappa^* \rangle$ from random Clifford measurements on polynomially many copies of $|qp\kappa^*\rangle$.

Having the estimated values of $\langle qp\kappa_{dk} | qp\kappa^* \rangle$ the adversary picks a dk such that the estimated value is relatively large and uses this key to decrypt the challenge ciphertext. Now invoking Lemma 1 we conclude that the probability of this adversary winning the IND-CPA game is significantly more than $1/2$.

1.3 Related works

The notion of qPKE was already considered in the literature, although without introducing formal security definitions. For instance, Gottesman [Got05] proposed a candidate construction in an oral presentation, without a formal security analysis. The scheme has quantum public-keys and quantum ciphers, which consumes the public-key for encryption. Kawachi et al. [KKNY05] proposed a construction of qPKE (with quantum keys and ciphertexts) from a newly introduced hardness assumption, related to the graph automorphism problem. [OTU00] defines and constructs a public-key encryption where the keys, plaintexts and ciphers are classical, but the algorithms are quantum (the key-generation uses Shor’s algorithm). One of the contributions of this work, is to provide a unifying framework for these results, as well as improve in terms of computational assumptions and security guarantees.

In [NI09], the authors define and provide impossibility results regarding encryption with quantum public-keys. Classically, it is easy to show that a (public) encryption scheme cannot have deterministic ciphers; in other words, encryption must use randomness. They show that this is also true for a quantum encryption scheme with quantum public-keys. In [Dol20], a secure encryption scheme with quantum public keys based on the LWE assumption is introduced. That work shows (passive) indistinguishable security, and is not IND-CPA secure.

In [MY22b,MY22a], the authors study digital signatures with quantum signatures, and more importantly in the context of this work, quantum public-keys.

The study of quantum pseudorandomness and its applications has recently experienced rapid advancements. One of the most astonishing aspects is that PRS (Pseudorandom states) and some of its variations are considered weaker than one-way functions. In other words, they are implied by one-way functions, and there exists a black-box separation between them. However, it has been demonstrated that these primitives are sufficient for many applications in Minicrypt and even extend beyond it. A graph presenting the various notions of quantum pseudorandomness and its application is available at <https://sattath.github.io/qcrypto-graph/>.

1.4 Concurrent and subsequent work

This work is a merge of two concurrent and independent works [BMW23,GSV23], with a unified presentation and more results.

In a concurrent and independent work, Coladangelo [Col23] proposes a qPKE scheme with a construction that is very different from ours, and uses a quantum trapdoor function, which is a new notion first introduced in their work. Their construction is based on the existence of quantum-secure OWF. However, in their construction, each quantum public-key can be used to encrypt a single message (compared to our construction from OWF, where the public-key can be used to encrypt multiple messages), and the ciphertexts are quantum (whereas our construction from OWF has classical ciphertexts). They do not consider the stronger notion of IND-CCA security.

Our paper has already generated interest in the community: Two follow-up works [KMNY23,MW23] consider a *stronger* notion of qPKE where the public-key consists of a classical and a quantum part,

and the adversary is allowed to tamper arbitrarily with the quantum part (but not with the classical component).¹⁰ The authors provide constructions assuming quantum-secure OWF. While their security definition is stronger, we remark that our approach is more general, as exemplified by the fact that we propose constructions from potentially weaker computational assumptions. In [BS23], the authors give another solution for the quantum public-key distribution problem using time-dependent signatures, which can be constructed from quantum-secure OWF, but the (classical) verification key needs to be continually updated.

2 Preliminaries

2.1 Notation

Throughout this paper, λ denotes the security parameter. The notation $\text{negl}(\lambda)$ denotes any function f such that $f(\lambda) = \lambda^{-\omega(1)}$, and $\text{poly}(\lambda)$ denotes any function f such that $f(\lambda) = \lambda^{\mathcal{O}(1)}$. When sampling uniformly at random a value a from a set \mathcal{U} , we employ the notation $a \leftarrow_{\$} \mathcal{U}$. When sampling a value a from a probabilistic algorithm \mathcal{A} , we employ the notation $a \leftarrow \mathcal{A}$. Let $|\cdot|$ denote either the length of a string, or the cardinal of a finite set, or the absolute value. By PPT we mean a polynomial-time non-uniform family of probabilistic circuits, and by QPT we mean a polynomial-time family of quantum circuits.

2.2 Quantum Information

For a more in-depth introduction to quantum information, we refer the reader to [NC16]. 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 $|\psi\rangle \in \mathcal{H}_M$. The class of positive, Hermitian, trace-one linear operators on \mathcal{H}_M is denoted by $\mathcal{D}(\mathcal{H}_M)$. A quantum register is a physical system whose set of valid states is $\mathcal{D}(\mathcal{H}_M)$; in this case we label by M the register itself. The maximally mixed state (i.e., uniform classical distribution) is written as $\mathcal{I}/\dim M$ on M .

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

A quantum t -design (for a fixed t) is a probability distribution over pure quantum states which can duplicate properties of the probability distribution over the Haar measure for polynomials of degree t or less. A quantum t -design with n -qubit output can be efficiently implemented with a random $\text{poly}(t, n)$ -size quantum circuits.

We recall the SWAP test on two quantum states $|\psi\rangle, |\phi\rangle$ which is an efficient algorithm that outputs 0 with probability $\frac{1}{2} + \frac{1}{2}|\langle\psi|\phi\rangle|^2$. In particular, if the states are equal, the output of SWAP test is always 0.

Next, we state a well-known fact about the quantum evaluation of classical circuits.

¹⁰ Because of this stronger security definition, here the notion of public-keys with mixed states is meaningful since there is an alternative procedure to ensure that the key is well-formed (e.g., signing the classical component).

Fact 1. Let $f: \{0, 1\}^n \rightarrow \{0, 1\}^m$ be a function which is efficiently computable by a classical circuit. Then there exists a unitary U_f on $(\mathbb{C}^2)^{\otimes n+m}$ which is efficiently computable by a quantum circuit (possibly using ancillas) such that, for all $x \in \{0, 1\}^n$ and $y \in \{0, 1\}^m$,

$$U_f: |x\rangle|y\rangle \mapsto |x\rangle|y \oplus f(x)\rangle.$$

2.3 Quantum-Secure Pseudorandom Functions

Throughout this work, we often refer to a *pseudorandom function* (PRF) first introduced in [GGM86]. This is a keyed function, denoted PRF, that can be evaluated in polynomial time satisfying a certain security property. In this work, we require PRF to be *quantum-secure*, which, loosely speaking, says that an adversary with oracle access to PRF cannot distinguish it from a truly random function, even given superposition queries. It is known that quantum-secure PRF can be constructed from any quantum-secure one-way function [Zha12].

Definition 1 (Quantum-secure PRF). We say that a keyed family of functions $\{f_k\}_k$ is a quantum-secure pseudorandom function (PRF) ensemble if, for any QPT adversary \mathcal{A} , we have

$$\left| \Pr \left[1 \leftarrow \mathcal{A}(1^\lambda)^{f_k} \right] - \Pr \left[1 \leftarrow \mathcal{A}(1^\lambda)^f \right] \right| \leq \mu(\lambda),$$

where $k \xleftarrow{\$} \{0, 1\}^\lambda$, f is a truly random function, and the oracles can be accessed in superposition, that is, they implement the following unitaries

$$|x\rangle|z\rangle \xrightarrow{U_{f_k}} |x\rangle|z \oplus f_k(x)\rangle \quad \text{and} \quad |x\rangle|z\rangle \xrightarrow{U_f} |x\rangle|z \oplus f(x)\rangle,$$

respectively.

2.4 Post-Quantum IND-CCA Symmetric-Key Encryption

We briefly recall the definition of a symmetric-key encryption scheme (SKE).

Definition 2. An SKE consists of 2 algorithms with the following syntax:

1. $\text{Enc}(\text{sk}, \text{pt})$: a PPT algorithm, which receives a symmetric-key $\text{sk} \in \{0, 1\}^\lambda$ and a plaintext pt , and outputs a ciphertext ct .
2. $\text{Dec}(\text{sk}, \text{ct})$: a deterministic polynomial-time algorithm, which takes a symmetric-key sk and a ciphertext ct , and outputs a plaintext pt .

We say that a SKE scheme is perfectly *correct* if for every plaintext $\text{pt} \in \{0, 1\}^*$ and symmetric-key $\text{sk} \in \{0, 1\}^\lambda$, $\text{Dec}(\text{sk}, \text{Enc}(\text{sk}, \text{pt})) = \text{pt}$.

Definition 3. An SKE is post-quantum IND-CCA secure if for every QPT adversary $\mathcal{A} := (\mathcal{A}_1, \mathcal{A}_2)$, there exists a negligible function ϵ such that the following holds for all λ :

$$\Pr \left[\begin{array}{l} \text{sk} \xleftarrow{\$} \{0, 1\}^\lambda \\ \text{pt}_0, \text{pt}_1 \leftarrow \mathcal{A}_1^{\text{Enc}(\text{sk}, \cdot), \text{Dec}(\text{sk}, \cdot)}(1^\lambda) \\ \tilde{b} = b \\ b \xleftarrow{\$} \{0, 1\} \\ \text{ct}^* \leftarrow \text{Enc}(\text{sk}, \text{pt}_b) \\ \tilde{b} \leftarrow \mathcal{A}_2^{\text{Enc}(\text{sk}, \cdot), \text{Dec}^*(\text{sk}, \cdot)}(\text{ct}^*, 1^\lambda) \end{array} \right] \leq 1/2 + \epsilon(\lambda),$$

Where $\text{Dec}^*(\text{sk}, \cdot)$ is the same as $\text{Dec}(\text{sk}, \cdot)$ but returns \perp on input the challenge ciphertext ct^* .

Note that as the adversary is not given superposition access to the Enc, Dec oracles one can build post-quantum IND-CCA SKE from quantum-secure OWF the same way as it is done classically with message authentication codes.

2.5 Pseudorandom Function-Like State (PRFS) Generators

The notion of pseudorandom function like states was first introduced by Ananth, Qian and Yuen in [AQY22]. A stronger definition where the adversary is allowed to make superposition queries to the challenge oracles was introduced in the follow-up work [AGQY22]. We reproduce their definition here:

Definition 4 (Quantum-accessible PRFS generator). *We say that a QPT algorithm G is a quantum-accessible secure pseudorandom function-like state generator if for all QPT (non-uniform) distinguishers A if there exists a negligible function ϵ , such that for all λ , the following holds:*

$$\left| \Pr_{k \leftarrow \{0,1\}^{1^\lambda}} \left[A_\lambda^{\mathcal{O}_{\text{PRFS}}(k, \cdot)}(\rho_\lambda) = 1 \right] - \Pr_{\mathcal{O}_{\text{Haar}}} \left[A_\lambda^{\mathcal{O}_{\text{Haar}}(\cdot)}(\rho_\lambda) = 1 \right] \right| \leq \epsilon(\lambda),$$

where:

- $\mathcal{O}_{\text{PRFS}}(k, \cdot)$, on input a d -qubit register \mathbf{X} , does the following: it applies an isometry channel that is controlled on the register \mathbf{X} containing x , it creates and stores $G_{1^\lambda}(k, x)$ in a new register \mathbf{Y} . It outputs the state on the registers \mathbf{X} and \mathbf{Y} .
- $\mathcal{O}_{\text{Haar}}(\cdot)$, modeled as a channel, on input a d -qubit register \mathbf{X} , does the following: it applies a channel that controlled on the register \mathbf{X} containing x , stores $|\vartheta_x\rangle\langle\vartheta_x|$ in a new register \mathbf{Y} , where $|\vartheta_x\rangle$ is sampled from the Haar distribution. It outputs the state on the registers \mathbf{X} and \mathbf{Y} .

Moreover, A_{1^λ} has superposition access to $\mathcal{O}_{\text{PRFS}}(k, \cdot)$ and $\mathcal{O}_{\text{Haar}}(\cdot)$ (denoted using the ket notation).

We say that G is a $(d(\lambda), n(\lambda))$ -QAPRFS generator to succinctly indicate that its input length is $d(\lambda)$ and its output length is $n(\lambda)$.

2.6 Quantum Pseudorandomness with Proofs of Destruction

We import the definition of pseudorandom function-like states with proofs of destruction (PRFSPD) from [BSS23].

Definition 5 (PRFS generator with proof of destruction). *A PRFSPD scheme with key-length $w(\lambda)$, input-length $d(\lambda)$, output length $n(\lambda)$ and proof length $c(\lambda)$ is a tuple of QPT algorithms Gen , Destruct , Ver with the following syntax:*

1. $|\psi_k^x\rangle \leftarrow \text{Gen}(k, x)$: takes a key $k \in \{0, 1\}^w$, an input string $x \in \{0, 1\}^{d(\lambda)}$, and outputs an n -qubit pure state $|\psi_k^x\rangle$.
2. $p \leftarrow \text{Destruct}(|\phi\rangle)$: takes an n -qubit quantum state $|\phi\rangle$ as input, and outputs a c -bit classical string, p .
3. $b \leftarrow \text{Ver}(k, x, p)$: takes a key $k \in \{0, 1\}^w$, a d -bit input string x , a c -bit classical string p and outputs a boolean output b .

Correctness. A PRFSPD scheme is said to be correct if for every $x \in \{0, 1\}^d$,

$$\Pr_{k \xleftarrow{u} \{0,1\}^w} [1 \leftarrow \text{Ver}(k, x, p) \mid p \leftarrow \text{Destruct}(|\psi_k^x\rangle); |\psi_k^x\rangle \leftarrow \text{Gen}(k, x)] = 1$$

Security.

1. **Pseudorandomness:** A PRFSPD scheme is said to be (adaptively) pseudorandom if for any QPT adversary \mathcal{A} , and any polynomial $m(\lambda)$, there exists a negligible function $\text{negl}(\lambda)$, such that

$$\left| \Pr_{k \leftarrow \{0,1\}^w} [\mathcal{A}^{\text{Gen}(k,\cdot)}(1^\lambda) = 1] - \Pr_{\forall x \in \{0,1\}^d, |\phi^x\rangle \leftarrow \mu_{(\mathbb{C}^2)^{\otimes n}} [\mathcal{A}^{\mathcal{H}\text{aar}^{\{|\phi^x\rangle\}_{x \in \{0,1\}^d}}(\cdot)}(1^\lambda) = 1] \right| = \text{negl}(\lambda)$$

where $\forall x \in \{0,1\}^d$, $\mathcal{H}\text{aar}^{\{|\phi^x\rangle\}_{x \in \{0,1\}^d}}(x)$ outputs $|\phi^x\rangle$. Here we note that \mathcal{A} gets quantum access to the oracles.

2. **Unclonability-of-proofs:** A PRFSPD scheme satisfies Unclonability-of-proofs if for any QPT adversary \mathcal{A} in cloning game (see Game 1), there exists a negligible function $\text{negl}(\lambda)$ such that

$$\Pr[\text{Cloning-Exp}_\lambda^{\mathcal{A}, \text{PRFSPD}} = 1] = \text{negl}(\lambda).$$

Game 1 Cloning-Exp $_\lambda^{\mathcal{A}, \text{PRFSPD}}$

- 1: Given input 1^λ , Challenger samples $k \leftarrow \{0,1\}^{w(\lambda)}$ uniformly at random.
 - 2: Initialize an empty set of variables, S .
 - 3: \mathcal{A} gets oracle access to $\text{Gen}(k, \cdot)$, $\text{Ver}(k, \cdot, \cdot)$ as oracle.
 - 4: **for** Gen query x made by \mathcal{A} **do**
 - 5: **if** \exists variable $t_x \in S$ **then** $t_x = t_x + 1$.
 - 6: **else** Create a variable t_x in S , initialized to 1.
 - 7: **end if**
 - 8: **end for**
 - 9: \mathcal{A} outputs $x, c_1, c_2, \dots, c_{t_x+1}$ to the challenger.
 - 10: Challenger rejects if c_i 's are not distinct.
 - 11: **for** $i \in [m+1]$ **do** $b_i \leftarrow \text{Ver}(k, x, c_i)$
 - 12: **end for**
 - 13: Return $\bigwedge_{i=1}^{m+1} b_i$.
-

3 Definitions of qPKE

In this section, we introduce the new notion of encryption with quantum public keys (Definition 6). The indistinguishability security notions are defined in Section 3.1 and Section 3.2.

Definition 6 (Encryption with quantum public keys). *Encryption with quantum public keys (qPKE) consists of 4 algorithms with the following syntax:*

1. $\text{dk} \leftarrow \text{Gen}(1^\lambda)$: a QPT algorithm, which receives the security parameter and outputs a classical decryption key.

2. $|\mathit{qp}\kappa\rangle \leftarrow \mathit{QP}\mathcal{K}\mathit{Gen}(\mathit{dk})$: a QPT algorithm, which receives a classical decryption key dk , and outputs a quantum public key $|\mathit{qp}\kappa\rangle$. In this work, we require that the output is a pure state, and that t calls to $\mathit{QP}\mathcal{K}\mathit{Gen}(\mathit{dk})$ should yield the same state, that is, $|\mathit{qp}\kappa\rangle^{\otimes t}$.
3. $(\mathit{qp}\kappa', \mathit{qc}) \leftarrow \mathit{Enc}(\mathit{qp}\kappa, m)$: a QPT algorithm, which receives a quantum public key $\mathit{qp}\kappa$ and a plaintext m , and outputs a (possibly classical) ciphertext qc and a recycled public key $\mathit{qp}\kappa'$.
4. $m \leftarrow \mathit{Dec}(\mathit{dk}, \mathit{qc})$: a QPT algorithm, which uses a decryption key dk and a ciphertext qc , and outputs a classical plaintext m .

We say that a qPKE scheme is *correct* if for every message $m \in \{0, 1\}^*$ and any security parameter $\lambda \in \mathbb{N}$, the following holds:

$$\Pr \left[\mathit{Dec}(\mathit{dk}, \mathit{qc}) = m \left| \begin{array}{l} \mathit{dk} \leftarrow \mathit{Gen}(1^\lambda) \\ |\mathit{qp}\kappa\rangle \leftarrow \mathit{QP}\mathcal{K}\mathit{Gen}(\mathit{dk}) \\ (\mathit{qp}\kappa', \mathit{qc}) \leftarrow \mathit{Enc}(|\mathit{qp}\kappa\rangle, m) \end{array} \right. \right] \geq 1 - \mathit{negl}(\lambda),$$

where the probability is taken over the randomness of Gen , $\mathit{QP}\mathcal{K}\mathit{Gen}$, Enc and Dec . We say that the scheme is reusable if completeness holds to polynomially many messages using a single quantum public key. More precisely, we say that a qPKE scheme is *reusable* if for every security parameter $\lambda \in \mathbb{N}$, polynomial number of messages $m_1, \dots, m_{n(\lambda)} \in \{0, 1\}^*$, the following holds:

$$\Pr \left[\forall i \in [n(\lambda)], \mathit{Dec}(\mathit{dk}, \mathit{qc}_i) = m_i \left| \begin{array}{l} \mathit{dk} \leftarrow \mathit{Gen}(1^\lambda) \\ |\mathit{qp}\kappa_1\rangle \leftarrow \mathit{QP}\mathcal{K}\mathit{Gen}(\mathit{dk}) \\ (\mathit{qp}\kappa_2, \mathit{qc}_2) \leftarrow \mathit{Enc}(|\mathit{qp}\kappa_1\rangle, m_1) \\ \vdots \\ (\mathit{qp}\kappa_{n+1}, \mathit{qc}_n) \leftarrow \mathit{Enc}(|\mathit{qp}\kappa_i\rangle, m_{n(\lambda)}) \end{array} \right. \right] \geq 1 - \mathit{negl}(\lambda).$$

3.1 Security Definitions for qPKE with Classical Ciphertexts

In this section, we present a quantum analogue of classical indistinguishability security for qPKE with classical ciphertexts. We note that there are few differences. Firstly, since in general the public keys are quantum states and unclonable, in the security games, we allow the adversary to receive polynomially many copies of $|\mathit{qp}\kappa\rangle$, by making several calls to the $\mathit{QP}\mathcal{K}\mathit{Gen}(\mathit{dk})$ oracle. Secondly, in the classical setting, there is no need to provide access to an encryption oracle since the adversary can use the public key to apply the encryption themselves. In the quantum setting, this is not the case: as we will see, the quantum public key might be measured, and the ciphertexts might depend on the measurement outcome. Furthermore, the quantum public key can be reused to encrypt multiple different messages. This motivates a stronger definition of indistinguishability with encryption oracle, in which the adversary gets oracle access to the encryption, denoted as IND-ATK-EO security, where ATK can be either chosen-plaintext attacks (CPA), (adaptive or non-adaptive) chosen-ciphertext attacks (CCA1 and CCA2).

We define the oracles $\mathcal{O}_1, \mathcal{O}_2$ depending on the level of security as follows.

ATK	Oracle \mathcal{O}_1	Oracle \mathcal{O}_2
CPA	\emptyset	\emptyset
CCA1	$\mathit{Dec}(\mathit{dk}, \cdot)$	\emptyset
CCA2	$\mathit{Dec}(\mathit{dk}, \cdot)$	$\mathit{Dec}^*(\mathit{dk}, \cdot)$

Game 2 Indistinguishability security with an encryption oracle (IND-ATK-EO) for encryption with quantum public key and classical ciphertext schemes.

- 1: The challenger generates $\text{dk} \leftarrow \text{Gen}(1^\lambda)$.
- 2: The adversary gets 1^λ as an input, and oracle access to $\text{QPKEGen}(\text{dk})$.
- 3: The challenger generates $|\text{qpk}\rangle \leftarrow \text{QPKEGen}(\text{dk})$. Let $\text{qpk}_1 := |\text{qpk}\rangle$.
- 4: For $i = 1, \dots, \ell$, the adversary creates a classical message m_i and send it to the challenger.
- 5: The challenger computes $(q_{c_i}, \text{qpk}_{i+1}) \leftarrow \text{Enc}(\text{qpk}_i, m_i)$ and send q_{c_i} to the adversary.
- 6: During step (2) to step (5), the adversary also gets classical oracle access to an oracle \mathcal{O}_1 .
- 7: The adversary sends two messages m'_0, m'_1 of the same length to the challenger.
- 8: The challenger samples $b \in_R \{0, 1\}$, computes $(q_{c^*}, \text{qpk}_{\ell+2}) \leftarrow \text{Enc}(\text{qpk}_{\ell+1}, m'_b)$ and sends q_{c^*} to the adversary.
- 9: For $i = \ell + 2, \dots, \ell'$, the adversary creates a classical message m_i and send it to the challenger.
- 10: The challenger computes $(q_{c_i}, \text{qpk}_{i+1}) \leftarrow \text{Enc}(\text{qpk}_i, m_i)$ and send q_{c_i} to the adversary.
- 11: During step (9) to step (10), the adversary also gets classical oracle access to an oracle \mathcal{O}_2 . Note that after step (7), the adversary no longer gets access to oracle \mathcal{O}_1 .
- 12: The adversary outputs a bit b' .

We say that the adversary wins the game (or alternatively, that the outcome of the game is 1) iff $b = b'$.

Here $\text{Dec}^*(\text{dk}, \cdot)$ is defined as $\text{Dec}(\text{dk}, \cdot)$, except that it return \perp on input the challenge ciphertext q_{c^*} .

Definition 7. A qPKE scheme is IND-ATK-EO secure if for every QPT adversary, there exists a negligible function ϵ such that the probability of winning the IND-ATK-EO security game (Game 2) is at most $\frac{1}{2} + \epsilon(\lambda)$.

Remark 1. The definition presented in Definition 7 is stated for the single challenge query setting. Using the standard hybrid argument, it is straightforward to show that single-challenge definitions also imply many-challenge definitions where the adversary can make many challenge queries.

Remark 2. Note that the IND-CCA2-EO definition is only well-defined for schemes with classical ciphertexts. The other two notions are well-defined even for quantum ciphertexts, though we do not use those.

3.2 Security Definitions for qPKE with Quantum Ciphertexts

We now give a definition for qPKE with quantum ciphertexts. In the case of adaptive chosen ciphertext security, the definition is non-trivial due to the no-cloning and the destructiveness of quantum measurements. We note there are indeed several works considering the notions of chosen-ciphertext security in the quantum setting: [AGM18] defines chosen-ciphertext security for quantum symmetric-key encryption (when the message is a quantum state), and [BZ13b,CEV22] defines chosen-ciphertext security for classical encryption under superposition attacks. However, extending the technique from [AGM18] to the public-key setting is non-trivial, and we leave this open problem for future work. In this section, we only consider security notions under chosen-plaintext attacks and non-adaptive chosen-ciphertext attacks.

Even though one can similarly define security notions with encryption oracle for schemes with quantum ciphertexts as in Section 3.1, we note that in all constructions of qPKE with quantum ciphertexts present in this work are not reusable, and thus we do not present the definition in which the adversary has oracle access to the encryption oracle for the sake of simplicity. We denote these notions as IND-ATK, where ATK is either chosen-plaintext attacks (CPA) or non-adaptive chosen-ciphertext attacks (CCA1).

Game 3 IND-ATK security game for encryption with quantum public key and quantum ciphertexts schemes.

- 1: The challenger generates $\text{dk} \leftarrow \text{Gen}(1^\lambda)$.
- 2: The adversary \mathcal{A}_1 gets 1^λ as an input, and oracle access to $\text{QPKGen}(\text{dk})$, $\text{Enc}(\text{qpk}, \cdot)$ and \mathcal{O}_1 , and sends m_0, m_1 of the same length to the challenger. \mathcal{A}_1 also output a state $|\text{st}\rangle$ and sends it to \mathcal{A}_2 .
- 3: The challenger samples $b \in_R \{0, 1\}$, generates $|\text{qpk}\rangle \leftarrow \text{QPKGen}(\text{dk})$ and sends $c^* \leftarrow \text{Enc}(|\text{qpk}\rangle, m_b)$ to the adversary \mathcal{A}_2 .
- 4: \mathcal{A}_2 gets oracle access to $\text{QPKGen}(\text{dk})$, $\text{Enc}(\text{qpk}, \cdot)$.
- 5: The adversary \mathcal{A}_2 outputs a bit b' .

We say that the adversary wins the game (or alternatively, that the outcome of the game is 1) iff $b = b'$.

The oracles \mathcal{O}_1 is defined depending on the level of security as follows.

ATK	Oracle \mathcal{O}_1
CPA	\emptyset
CCA1	$\text{Dec}(\text{dk}, \cdot)$

Definition 8. A qPKE scheme with quantum ciphertexts is IND-ATK secure if for every QPT adversary $\mathcal{A} := (\mathcal{A}_1, \mathcal{A}_2)$, there exists a negligible function ϵ such that the probability of winning the IND-ATK security game (Game 3) is at most $\frac{1}{2} + \epsilon(\lambda)$.

4 Constructions of CCA-Secure qPKE

In this section, we present our qPKE constructions from OWF and PRFS and prove that their CCA security. The former (given in Section 4.1) has classical ciphertexts, and allows to encrypt arbitrary long messages. The latter (given in Section 4.2) has quantum ciphertexts, and only allows to encrypt a single-bit message. However, we note that the latter is based on a weaker assumption than the former. Finally, in Section 4.3, we give a remark on the black-box construction of non-malleable qPKE from CPA-secure qPKE using the same classical approach.

4.1 CCA-Secure Many-Bit Encryption from OWF

We start by presenting a simple qPKE construction from OWF which prove that it provides our strongest notion of security, i.e. IND-CCA-EO security. The scheme is formally presented in Construction 1. The ciphertexts produced by the scheme are classical, and the public-keys are reusable. The cryptographic components of our construction are a quantum secure PRF family $\{f_k\}$ and a post-quantum IND-CCA secure symmetric-key encryption scheme (SE.Enc, SE.Dec) which can both be built from a quantum-secure OWF [Zha12,BZ13a].

Construction 1 (IND-CCA-EO secure qPKE from OWF).

- **Assumptions:** A family of quantum-secure pseudorandom functions $\{f_k\}_k$, and post-quantum IND-CCA SKE (SE.Enc, SE.Dec).
- $\text{Gen}(1^\lambda)$
 1. $\text{dk} \xleftarrow{\$} \{0, 1\}^\lambda$
 2. $|\text{qpk}\rangle \leftarrow \sum_{x \in \{0, 1\}^\lambda} |x, f_{\text{dk}}(x)\rangle$

- $\text{Enc}(|qp\kappa\rangle, m)$
 1. Measure $|qp\kappa\rangle$ to obtain classical strings x, y .
 2. Let $c_0 \leftarrow x$ and $c_1 \leftarrow \text{SE.Enc}(y, m)$.
 3. Output (c_0, c_1)
- $\text{Dec}(\text{dk}, (c_0, c_1))$
 1. Compute $y \leftarrow f_{\text{dk}}(c_0)$.
 2. Compute $m \leftarrow \text{SE.Dec}(y, c_1)$ and return m .

It can be trivially shown that the scheme achieves perfect correctness if the underlying SKE provides the perfect correctness property.

Theorem 1. *Let $\{f_k\}_k$ be a quantum secure PRF and $(\text{SE.Enc}, \text{SE.Dec})$ be a post-quantum IND-CCA secure SKE. Then, the quantum qPKE given in Construction 1 is IND-CCA-EO secure.*

Proof. We proceed with a sequence of hybrid games detailed in

- **Hybrid H_0 :** This is the IND-CCA game with Π with the challenge ciphertext fixed to $(x^*, c^*) = \text{Enc}(|p\kappa\rangle, m'_0)$.
- **Hybrid H_1 :** This is identical to H_0 except instead of measuring $|qp\kappa\rangle$ when the adversary queries the encryption oracle, the challenger measures a copy of $|qp\kappa\rangle$ in advance to obtain $(x^*, y^* = f_{\text{dk}}(x^*))$ and answers queries to the encryption oracle using (x^*, y^*) instead. The decryption oracle still returns \perp when queried (x^*, c^*) . This change is only syntactical so the two hybrids are the same from the adversary's view.

The hybrids H_2 to H_5 have 2 main goals: (i) to decorrelate the encryption/decryption oracles Dec^*, Enc from the public-keys handed to the adversary and (ii) to remove the oracles' dependency on dk .

- **Hybrid H_2 :** This is identical to H_1 , except (x^*, y^*) is removed from the copies of $|qp\kappa\rangle$ handed to the adversary. More precisely, the adversary is handed $|qp\kappa'\rangle$ of the following form:

$$|qp\kappa'\rangle = \frac{1}{\sqrt{2^{|x^*|} - 1}} \sum_{x: x \neq x'} |x\rangle |f_{\text{dk}}(x)\rangle \quad (6)$$

The decryption oracle still returns \perp when queried on the challenge ciphertext. Note that $|qp\kappa\rangle$ and $|qp\kappa'\rangle$ have $\text{negl}(\lambda)$ trace distance so the advantage of distinguishing H_1 and H_2 is $\text{negl}(\lambda)$.

- **Hybrid H_3 :** This (inefficient) hybrid is identical to H_2 other than f_{dk} being replaced with a truly random function f , i.e. the public-keys are change to:

$$|qp\kappa'\rangle = \frac{1}{\sqrt{2^{|x^*|} - 1}} \sum_{x: x \neq x'} |x\rangle |f(x)\rangle \quad (7)$$

The encryption and decryption oracle can be simulated by oracle access to f . The decryption oracle returns \perp when queried (x^*, c^*) . The indistinguishability of H_3 and H_2 follows directly from pseudorandomness property of $\{f_k\}_k$.

- **Hybrid H_4 :** This hybrid is identical to H_3 other than y^* being sampled uniformly at random. Upon quering (c_0, c_1) to the decryption oracle if $c_0 \neq x^*$, the oracle computes $y = f(c_0)$ and returns $m = \text{SE.Dec}(y, c_1)$. In case $c_0 = x^*$ and $c_1 \neq c^*$, the decryption oracle returns $m = \text{SE.Dec}(y^*, c_1)$. On (x^*, c^*) the oracle returns \perp . The encryption oracle returns $(x^*, \text{SE.Enc}(y^*, m))$ when queried on m . As x^* does not appear in any of the public-keys this change is only syntactical.

- **Hybrid H_5 :** This hybrid reverts the changes of H_3 , i.e. dk' is sampled uniformly at random and the public-keys are changed as follows:

$$|\text{qp}\mathcal{K}'\rangle = \frac{1}{\sqrt{2^{|x^*|} - 1}} \sum_{x:x \neq x'} |x\rangle |f_{\text{dk}'}(x)\rangle \quad (8)$$

With this change, on query (c_0, c_1) if $c_0 \neq x^*$, the decryption oracle computes $y = f_{\text{dk}'}(c_0)$ and returns $m = \text{SE.Dec}(y, c_1)$. In case $c_0 = x^*$, the decryption oracle simply returns $m = \text{SE.Dec}(y^*, c_1)$ when $c_1 \neq c^*$ and \perp otherwise. The encryption oracle is unchanged from H_4 . The indistinguishability of H_4 and H_5 follows from the pseudorandomness of f and the fact that $|\text{qp}\mathcal{K}'\rangle$ and (x^*, y^*) are decorrelated. The hybrid is efficient again.

The next step is to remove the dependency of the encryption and decryption oracles on y^* . This is done by querying the encryption and decryption oracles of the SKE.

- **Hybrid H_6 :** Let SE.OEnc and SE.ODec^* be two oracles implementing the encryption and decryption procedures of SE with the key y^* . SE.ODec^* returns \perp when queried y^* . In this hybrid, we syntactically change the encryption and decryption oracle using these two oracles. To implement the encryption oracle, on query m we simply query SE.OEnc on message m and return $(x^*, \text{SE.OEnc}(m))$. To simulate the decryption oracle, on query (c_0, c_1) we act the same as in H_5 when $c_0 \neq x^*$, but on queries of form (x^*, c) we query SE.ODec^* on c and return $\text{SE.ODec}^*(c)$. Due to the definition of OEnc and ODec^* these changes are also just syntactical. Note that although SE.ODec^* always returns \perp on y^* , it is only queried when $c_0 = x^*$, i.e. to cause this event the decryption oracle should be queried on the challenge ciphertext (x^*, c^*) .
 - **Hybrid H_7 :** We provide the adversary with $x^*, \text{SE.OEnc}, \text{SE.ODec}^*$, instead of access to the encryption and decryption oracle. Note that the adversary can implement the encryption and decryption oracles themselves by having access to $x^*, \text{SE.OEnc}, \text{SE.ODec}^*$ and sampling a uniform dk' themselves and vice versa (SE.ODec^* can be queried on c by querying the decryption oracle (x^*, c) and SE.OEnc can be queried on m by querying the encryption oracle on m). This demonstrates that the hybrids are only syntactically different and hence are indistinguishable.
 - **Hybrid H_8 :** This hybrid is identical to H_7 with the only difference that the challenge ciphertext is swapped with $(x^*, \text{SE.OEnc}(0))$. Now notice that any adversary that can distinguish H_8 from H_7 can effectively break the IND-CCA security of SE. Hence, the indistinguishability of the two hybrids follows directly from the IND-CCA security of SE.
- Following the same exact hybrids for challenge ciphertext $\mathcal{E}nc(|\text{qp}\mathcal{K}\rangle, m'_1)$ we can deduce that the scheme is IND-CCA-EO secure.

□

4.2 CCA1-Secure Many-Bit Encryption from PRFS

We continue by presenting a CCA1-secure bit-encryption from PRFS. Extending this scheme to polynomially many bits is discussed at the end of this section, see Remark 3. The description of the scheme is given below in Construction 2.

Construction 2 (IND-CCA1 secure qPKE from PRFS).

- **Assumptions:** A PRFS family $\{|\psi_{\text{dk},x}\rangle\}_{\text{dk},x}$ with super-logarithmic input-size. Let $n := n(\lambda)$.

- Gen(1^λ)
 1. Output $\text{dk} \leftarrow_R \{0, 1\}^\lambda$.
- QPKGen(dk)
 1. Output $|\text{qpk}\rangle \leftarrow \sum_x |x\rangle_R |\psi_{\text{dk},x}\rangle_S^{\otimes n}$, where $x \in \{0, 1\}^{\omega(\log \lambda)}$.
- Enc($|\text{qpk}\rangle, m$) for $m \in \{0, 1\}$
 1. Measure the R registers of $|\text{qpk}\rangle$ to obtain a classical string x . Let $|\phi\rangle := |\psi_{\text{dk},x}\rangle^{\otimes n}$ denote the residual state.
 2. If $m = 0$, output the ciphertext as $(x, |\phi\rangle)$.
 3. Else, sample a uniformly random key dk_1 , and output the ciphertext as $(x, |\psi_{\text{dk}_1,x}\rangle^{\otimes n})$.
- Dec($\text{dk}, (x, \Psi)$)
 1. Compute $|\psi_{\text{dk},x}\rangle^{\otimes n}$ and perform n SWAP tests for each subsystem of Ψ of the same size as $|\psi_{\text{dk},x}\rangle$ with $|\psi_{\text{dk},x}\rangle$.
 2. If the outcome of the SWAP tests is 0 all the time, output 0, otherwise output 1.

The correctness of the scheme follows from the fact that the states $|\psi_{\text{dk}_1,x}\rangle$ are relatively well spread out for a random choice of dk . This is due to the pseudorandomness of the state generator. Note that if in step 3 instead of picking dk_1 randomly and computing $|\psi_{\text{dk}_1,x}\rangle$, the encryption algorithm sampled $|\vartheta\rangle^{\otimes n}$, from the Haar measure, the expected probability of n SWAP tests between $|\psi_{x,\text{dk}}\rangle$ and $|\vartheta\rangle$ all returning 0 would be 2^{-n} . Hence, if the probability is more than negligibly apart for n SWAP tests between $|\psi_{x,\text{dk}_1}\rangle$ and $|\psi_{x,\text{dk}}\rangle$ for a random choice of dk_1 , with a Chernoff bound argument one can show that this would lead to a distinguisher for the PRFS. Hence, for n polynomial in λ the scheme has negligible correctness error.

Theorem 2. *The construction in Construction 2 is IND-CCA1 secure (see Definition 8), assuming $\{|\psi_{\text{dk},x}\rangle\}_{\text{dk},x}$ is a PRFS with super-logarithmic input-size.*

Proof. We prove the theorem via a series of hybrids.

- **Hybrid H_0 .** The original security game as defined in Definition 8.
- **Hybrid H_1 .** This is identical to hybrid H_0 , except that the challenger, instead of measuring $|\text{qpk}\rangle$ when the adversary queries the encryption oracle for the first time, the challenger measures (the R registers of) this state before providing the copies of $|\text{qpk}\rangle$ to the adversary. Note that by measuring $|\text{qpk}\rangle$ in the computational basis, the challenger would obtain a classical uniformly random string x^* , let the residual state be $|\phi^*\rangle := |\psi_{\text{dk},x^*}\rangle^{\otimes n}$. Note that the two operations corresponding to the challenger's measurement of $|\text{qpk}\rangle$ and the creation of the copies of $|\text{qpk}\rangle$ given to the adversary commute. Thus, the distribution of the two hybrids are identical and no adversary can distinguish H_0 from H_1 with non-zero advantage.
- **Hybrid H_2 .** This is identical to hybrid H_1 , except that the challenger samples x^* as in the previous hybrid, and instead of providing $|\text{qpk}\rangle$ to the adversary, it provides

$$|\text{qpk}'\rangle := \frac{1}{\sqrt{2^{|x^*|} - 1}} \sum_{x:x \neq x^*} |x\rangle |\psi_{\text{dk},x}\rangle^{\otimes n}.$$

Moreover, in the challenge query, the challenger uses $(x^*, |\phi^*\rangle)$ for the encryption of the chosen message m , without measuring a fresh copy of $|\text{qpk}\rangle$ (that is, it skips the first step of the encryption algorithm). We note that this state $|\text{qpk}'\rangle$ can be efficiently prepared.

The distinguishing probability of the two hybrids H_1 and H_2 implies that we can distinguish the following quantum states $|qp\kappa\rangle^{\otimes p} \otimes |x^*\rangle$ and $|qp\kappa'\rangle^{\otimes p} \otimes |x^*\rangle$ with the same probability, but these two quantum states have $\text{negl}(\lambda)$ trace-distance for any polynomial p . Therefore, any adversary can only distinguish H_1 and H_2 with success probability at most $\text{negl}(\lambda)$.

- **Hybrid H_3 .** This (inefficient) hybrid is identical to H_2 , except that the challenger uses a Haar oracle $\mathcal{O}_{\text{Haar}}$ to generate $|qp\kappa'\rangle$ in place of $|\psi_{\text{dk},\cdot}\rangle$. In particular, the quantum public key in the hybrid H_3 is computed as:

$$|qp\kappa'\rangle \leftarrow \sum_{x:x \neq x^*} |x\rangle \otimes |\vartheta_x\rangle^{\otimes n},$$

where each $|\vartheta_x\rangle$ is an output of $\mathcal{O}_{\text{Haar}}$ on input x . The decryption oracle is the same as the decryption algorithm with the difference that $\mathcal{O}_{\text{PRFS}}$ (the algorithm generating the PRFS) is swapped with $\mathcal{O}_{\text{Haar}}$. The crucial point here is that the decryption oracle only uses the PRFS in a black-box way (in particular, it only uses $\mathcal{O}_{\text{PRFS}}$ and does not use $\mathcal{O}_{\text{PRFS}}^\dagger$).

Note that the decryption oracle can return \perp on query (x^*, \cdot) . This can not be used to distinguish the two hybrids as the adversary has a negligible chance of querying x^* as x^* is picked uniformly at random. The adversary is only provided with the value of x^* when given the challenge ciphertext, at which point they do not have access to the decryption oracle anymore.

We note that the adversary does not have direct access to this $\mathcal{O}_{\text{Haar}}$, but only via the decryption oracle. By pseudorandomness property of $|\psi_{\text{dk},\cdot}\rangle$, we have that H_2 and H_3 are computationally indistinguishable.

- **Hybrid H_4 .** In this hybrid, we revert the changes in H_3 , except that the challenger samples a uniformly random key dk' to compute all states in $|qp\kappa'\rangle$, except for the one used to encrypt the challenge query. In particular, the public key $|qp\kappa'\rangle$ is now generated using the PRFS generator with the key dk' , and the secret key dk and its public counterpart $(x^*, |\psi_{\text{dk},x^*}\rangle^{\otimes n})$ are used for the challenge encryption. We note that the hybrid is now efficient again. Similar to the previous argument, H_3 and H_4 are also computationally indistinguishable due to pseudorandomness property of $|\psi_{\text{dk}',\cdot}\rangle$.
- **Hybrid H_5 .** This hybrid is identical to H_4 , except that in the challenge query, instead of encrypting 0 as $(x^*, |\psi_{\text{dk},x^*}\rangle^{\otimes n})$, the challenger encrypts 0 as $(x^*, |\vartheta_{x^*}\rangle^{\otimes n})$, where each $|\vartheta_x\rangle$ is an output of $\mathcal{O}_{\text{Haar}}$ on input x .

Notice that in this hybrid, the secret key dk and its public counterpart $(x^*, |\psi_{\text{dk},x^*}\rangle^{\otimes n})$ are not correlated with any of other variables in the hybrid. Furthermore, after receiving the challenge ciphertext, the adversary no longer gets access to the decryption oracle. By the pseudorandomness property of $|\psi_{\text{dk},x^*}\rangle$, we have that H_4 and H_5 are computationally indistinguishable.

Furthermore, in this final hybrid, the adversary needs to distinguish the output of PRFS with a uniformly random key dk_1 (for encryption of 1) and the output of a Haar random oracle (for encryption of 0). By the same argument as above, the winning advantage of the adversary is also negligible.

Overall, since all hybrids are negligibly close and the winning advantage of the adversary in the last hybrid is negligible, we conclude the proof. \square

Remark 3. We sketch here how to achieve many-bit encryption (i.e., non-restricted length encryption) from our scheme present above. We do this through several steps.

- The scheme stated in Construction 2 can easily be extended to a length-restricted scheme, by applying bit-by-bit encryption.

- Given a qPKE length-restricted CCA1 encryption, and a (non-restricted length) symmetric key encryption, we can define a hybrid encryption scheme, where the qPKE scheme is used first to encrypt a random (fixed length) secret key, which is later used to encrypt an arbitrarily long message. The entire scheme is CPA- (respectively, CCA1-) secure if the symmetric key encryption has CPA- (respectively, CCA1-) security.
- Finally, we note that the following many-bit symmetric key encryption scheme can be proven to be CCA1 secure, using the same proof strategy as in Theorem 2, based on the existence of PRFS alone. Given a secret key dk , to encrypt a message $m \in \{0, 1\}^\ell$, we sample ℓ distinct uniformly random strings x_i , and compute $|\psi_{\text{dk}, x_i}\rangle^{\otimes n}$. Then each bit m_i will be encrypted using $(x_i, |\psi_{\text{dk}, x_i}\rangle^{\otimes n})$ if $m_i = 0$, or $(x_i, |\psi_{\text{dk}', x_i}\rangle^{\otimes n})$ if $m_i = 1$ for a fresh key dk' .

4.3 Generic Construction of Non-Malleable qPKE

We remark that known implications from the literature can be used to show that IND-CPA secure qPKE *with classical ciphertexts* implies non-malleable qPKE: The work of [CDMW18] shows a black-box compiler from IND-CPA encryption to non-malleable encryption, which also applies to the settings of quantum public-keys. The only subtlety is that the compiler assumes the existence of a one-time signature scheme to sign the ciphertext. In [MY22b, MY22a] it is shown that one-time signatures (with quantum verification keys) exist assuming one-way state generators, which in turn are implied by qPKE. Combining the implications of these two works, we obtain a generic construction of non-malleable qPKE from any IND-CPA secure one.

5 IND-CPA-EO secure qPKE from PRFSPD

In this section, we propose a construction for qPKE from pseudo-random function-like states with proof of destruction. The construction is reusable, has classical ciphers, and is CPA-EO secure.

We first import the following result that builds *symmetric*-key encryption from PRFSPD.

Proposition 1 ([BBS23]). *If quantum-secure PRFSPD exists, then there exists a quantum CPA symmetric encryption with classical ciphertexts.*

We give the formal construction for many-bit reusable encryption scheme from PRFSPD in Construction 3.

Construction 3 (IND-CPA-EO secure qPKE from PRFSPD).

- **Assumptions:** A PRFSPD family $\{|\psi_{\text{dk}, x}\rangle\}_{\text{dk}, x}$ and a quantum symmetric encryption scheme with classical ciphers $\{\text{Enc}, \text{Dec}\}$.
- $\text{Gen}(1^\lambda)$
 1. Let $\text{dk}_{0,i} \leftarrow_R \{0, 1\}^\lambda$ for all $i \in [1, \lambda]$.
 2. Output $\text{dk} \leftarrow \{\text{dk}_{0,i}\}_{i \in [1, \lambda]}$.
- $\text{QPKEGen}(\text{dk})$
 1. Output $|\text{qp}\mathcal{K}\rangle = \bigotimes_{i \in [\lambda]} \frac{1}{\sqrt{2^\lambda}} \sum_{x_i \in \{0, 1\}^\lambda} |x_i\rangle |\psi_{\text{dk}_{0,i}, x_i}\rangle$.
- $\text{Enc}(|\text{qp}\mathcal{K}\rangle, m)$ for $m \in \{0, 1\}^*$
 1. Let $|\text{qp}\mathcal{K}_i\rangle := \frac{1}{\sqrt{2^\lambda}} \sum_{x_i \in \{0, 1\}^\lambda} |x_i\rangle |\psi_{\text{dk}_{0,i}, x_i}\rangle$, and write $|\text{qp}\mathcal{K}\rangle$ as $|\text{qp}\mathcal{K}\rangle = \bigotimes_{i \in [\lambda]} |\text{qp}\mathcal{K}_i\rangle$.

2. Measure the left registers of $|\mathit{qp}\mathit{k}_i\rangle$ to obtain classical strings x_i . Denote the post-measurement states as $|\psi'_i\rangle$.
 3. Set $y_i \leftarrow \mathit{Destruct}(|\psi'_i\rangle)$.
 4. For $i \in [1, \lambda]$, pick $\mathit{dk}_{1,i} \leftarrow \{0, 1\}^\lambda$ and compute $|\psi_{\mathit{dk}_{1,i}, x_i}\rangle$.
 5. Set $y'_i \leftarrow \mathit{Destruct}(|\psi_{\mathit{dk}_{1,i}, x_i}\rangle)$ for all $i \in [\lambda]$.
 6. Pick a uniformly random key $k \leftarrow \{0, 1\}^\lambda$.
 7. Set $\tilde{y}_i = \begin{cases} y'_i & , \text{if } k_i = 0 \\ y_i & , \text{if } k_i = 1 \end{cases}$.
 8. Output $(\mathit{Enc}(k, m), ((x_i, \tilde{y}_i))_i)$ as ciphertext and $(k, ((x_i, \tilde{y}_i))_i)$ as the recycled public-key.
- $\mathit{Dec}(\mathit{dk}, c)$
 1. Interpret c as $(c', ((x_i, \tilde{y}_i))_i)$.
 2. Let $k'_i = \mathit{Ver}(\mathit{dk}_{0,i}, x_i, \tilde{y}_i)$ and let $k' = k'_0 \dots k'_\lambda$.
 3. Output $\mathit{Dec}(k', c')$.

The correctness of our scheme relies on the existence of PRFSPD with pseudorandomness and unclonability of proofs properties. The proof of correctness can be shown similarly to that of Construction 2. Next, we show that this construction achieves IND-CPA-EO security in Theorem 3.

Theorem 3. *If quantum-secure PRFSPD with super-logarithmic input size exists, then there exists a public-key encryption with classical ciphertexts which is IND-CPA-EO secure.*

Proof. Our construction is given in Construction 3. It uses a PRFSPD family $\{|\psi_{\mathit{dk}, x}\rangle\}_{\mathit{dk}, x}$ and a quantum symmetric encryption scheme with classical ciphers $\{\mathit{Enc}, \mathit{Dec}\}$. We prove the security of our scheme through a series of hybrids.

- **Hybrid H_0 .** The original security game as defined in Definition 7.
- **Hybrid H_1 .** This is identical to hybrid H_0 , except that the challenger, instead of measuring $|\mathit{qp}\mathit{k}_i\rangle$ (for all $i \in [\lambda]$) when the adversary queries the encryption oracle for the first time, the challenger measures the left register of each $|\mathit{qp}\mathit{k}_i\rangle$ before providing the copies of $|\mathit{qp}\mathit{k}\rangle$ to the adversary. Note that by measuring $|\mathit{qp}\mathit{k}_i\rangle$ in the computational basis, the challenger would obtain a classical uniformly random string x_i^* .
Note that the two operations corresponding to the challenger's measurement of $|\mathit{qp}\mathit{k}\rangle$ and the creation of the copies of $|\mathit{qp}\mathit{k}\rangle$ given to the adversary commute. Thus, the distribution of the two hybrids are identical and no adversary can distinguish H_0 from H_1 with non-zero advantage.
- **Hybrid H_2 .** This is identical to hybrid H_1 , except that the challenger samples x_i^* as in the previous hybrid, and instead of providing $|\mathit{qp}\mathit{k}\rangle$ to the adversary, it provides

$$|\mathit{qp}\mathit{k}'\rangle := \bigotimes_{i \in [\lambda]} \frac{1}{\sqrt{2^{|x_i^*|} - 1}} \sum_{x_i: x_i \neq x_i^*} |x_i\rangle |\psi_{\mathit{dk}_{0,i}, x_i}\rangle.$$

Moreover, in the challenge query, the challenger uses $(x_i^*, |\psi_{\mathit{dk}_{0,i}, x_i^*}\rangle)$ for all $i \in [\lambda]$ for the encryption of the chosen message m , without measuring a fresh copy of $|\mathit{qp}\mathit{k}\rangle$ (that is, it skips the first step of the encryption algorithm). We note that this state $|\mathit{qp}\mathit{k}'\rangle$ can be efficiently prepared.

The distinguishing probability of the two hybrids H_1 and H_2 implies that we can distinguish the following quantum states $|\mathit{qp}\mathit{k}\rangle^{\otimes p} \otimes \bigotimes_{i \in [\lambda]} |x_i^*\rangle$ and $|\mathit{qp}\mathit{k}'\rangle^{\otimes p} \otimes \bigotimes_{i \in [\lambda]} |x_i^*\rangle$ with the same

probability, but these two quantum states have $\text{negl}(\lambda)$ trace-distance for any polynomial p . Therefore, any adversary can only distinguish H_1 and H_2 with success probability at most $\text{negl}(\lambda)$.

- **Hybrid $H_{2,i}$ for $i \in [0, \lambda]$.** We define a series of (inefficient) hybrids $H_{2,i}$, in which $H_{2,0} := H_2$, and we denote $H_{2,\lambda} := H_3$. Each $H_{2,i+1}$ is identical as $H_{2,i}$, except that the challenger uses a Haar oracle $\mathcal{O}_{\text{Haar}_i}$ in place of $|\psi_{\text{dk}_{0,i},\cdot}\rangle$. In particular, the quantum public key in the hybrid $H_{2,i}$ is computed as:

$$|qp\kappa'\rangle \leftarrow \bigotimes_{j=1}^i \sum_{x_j: x_j \neq x_j^*} |x_j\rangle \otimes |\vartheta_{x_j}\rangle \otimes \bigotimes_{j=i+1}^{\lambda} \sum_{x_j: x_j \neq x_j^*} |x_j\rangle |\psi_{\text{dk}_{0,j},x_j}\rangle,$$

where each $|\vartheta_{x_j}\rangle$ is an output of $\mathcal{O}_{\text{Haar}_j}$ on input x_j . For the challenge encryption query, the challenger uses $(x_j^*, |\vartheta_{x_j^*}\rangle)$ for all $j \in [1, i]$, and $(x_j^*, |\psi_{\text{dk}_{0,j},x_j^*}\rangle)$ for all $j \in [i+1, \lambda]$.

By pseudorandomness property of $|\psi_{\text{dk}_{0,i},\cdot}\rangle$, we have that $H_{2,i}$ and $H_{2,i+1}$ are computationally indistinguishable.

- **Hybrid $H_{3,i}$ for $i \in [0, \lambda]$.** We define a series of (inefficient) hybrids $H_{3,i}$, in which $H_{3,0} := H_3$, and we denote $H_{3,\lambda} := H_4$. In each $H_{3,i+1}$, we revert the changes in $H_{3,i}$, except that the challenger samples uniformly random keys dk'_i to compute the i -th component in $|qp\kappa'\rangle$, except for the one used to encrypt the challenge query.

Similar to the previous argument, $H_{3,i+1}$ and $H_{3,i}$ are also computationally indistinguishable due to pseudorandomness property of $|\psi_{\text{dk}'_i,\cdot}\rangle$.

- **Hybrid $H_{4,i}$ for $i \in [0, \lambda]$.** We define a series of (inefficient) hybrids $H_{4,i}$, in which $H_{4,0} := H_4$, and we denote $H_{4,\lambda} := H_5$.

Each hybrid $H_{4,i}$ is identical to $H_{4,i+1}$, except that for the challenge encryption, the challenger does not sample $\text{dk}_{1,i}$ and compute $|\psi_{\text{dk}_{1,i},x_i^*}\rangle$. Instead, the challenger generates $|\vartheta_{x_i^*}\rangle$ using a Haar random oracle $\mathcal{O}_{\text{Haar}_i}$ and uses this state to compute y'_i (by applying $\mathcal{Destruct}$ to $|\vartheta_{x_i^*}\rangle$).

By the pseudorandomness of $|\psi_{\text{dk}_{1,i},\cdot}\rangle$, $H_{4,i}$ and $H_{4,i+1}$ are computationally indistinguishable.

- **Hybrid H_6 .** This hybrid is identical to H_5 , except that now the challenger sets $\tilde{y}_i = y_i$ for all i for the challenge encryption query.

Note that in this hybrid, both y_i and y'_i are computed by applying $\mathcal{Destruct}$ to a Haar random state, thus they are output of the same distribution. Therefore, H_5 and H_6 are identical.

- **Hybrid $H_{6,i}$ for $i \in [0, \lambda]$.** We define a series of hybrids $H_{6,i}$, in which $H_{6,0} := H_6$, and we denote $H_{6,\lambda} := H_7$.

Each hybrid $H_{6,i+1}$ is identical to $H_{6,i}$, except now instead of using a Haar random oracle in encryption of the challenge query, the challenger samples a fresh key dk_i and uses this key to compute \tilde{y}_i which is a proof of destruction of the state $|\psi_{\text{dk}_i,x_i^*}\rangle$.

By pseudorandomness of $|\psi_{\text{dk}_i,\cdot}\rangle$, $H_{6,i+1}$ and $H_{6,i}$ are computationally indistinguishable.

We also note that the hybrid H_7 is now efficient again. In this final hybrid, we note that the secret key k of the symmetric key encryption scheme is uniformly random and independent from all the other variables in the hybrid. Thus, we can easily reduce the winning probability of the adversary in this hybrid to the security of the symmetric key encryption scheme, which is negligible.

Overall, we obtain the winning probability of the adversary in the first hybrid H_0 is negligible, and conclude the proof. □

6 Impossibility of Unconditionally Secure qPKE

In the following, we investigate the question on whether qPKE is possible to construct with information-theoretic security, and we give strong bounds against this. First, let us mention that a recent work by Morimae et al. [MY22a] shows that an object called quantum pseudo-one-time pad (QPOTP) implies the existence of efficiently samplable, statistically far but computationally indistinguishable pairs of (mixed) quantum states (EFI pairs). QPOTP is a one-time symmetric encryption with quantum ciphertexts and classical keys, whose key length is shorter than the message length. qPKE immediately implies the existence of QPOTP, by increasing the message length, using bit-by-bit encryption. Since EFI pairs cannot exist information-theoretically, this chain of implications rules out the existence of unconditionally secure qPKE.¹¹

For the sake of completeness, we provide a new and direct proof of the impossibility statement using a shadow tomography argument.

A Proof from Shadow Tomography. In order to prove our impossibility result, we first show that if two public-keys $|qpk\rangle$ and $|qpk^*\rangle$ are close, if we encrypt a random bit using $|qpk^*\rangle$, the probability of decrypting correctly with dk is high, where dk is the corresponding secret-key of $|qpk\rangle$.

Lemma 1. *Let λ be the security parameter and $\Gamma = (\text{Gen}, \text{QPXGen}, \text{Enc}, \text{Dec})$ be a qPKE. Let $dk^*, |qpk^*\rangle$ be a fixed pair of honestly generated keys and for all decryption keys dk define p_{dk} to be:*

$$p_{dk} = \Pr \left[\text{Dec}(dk, qc) = pt \mid \begin{array}{l} pt \xleftarrow{\$} \{0, 1\} \\ (qc, \cdot) \leftarrow \text{Enc}(qpk^*, pt) \end{array} \right]$$

and let $|qpk_{dk}\rangle \leftarrow \text{QPXGen}(dk)$. For all dk , if $|\langle qpk^* | qpk_{dk} \rangle| \geq 1 - \epsilon$, then $p_{dk} \geq 1 - \sqrt{3\epsilon}$.

Proof. Let U_{Enc} be the purified implementation of the encryption procedures, i.e. given the state $|qpk^*\rangle|b\rangle|0\rangle$, U_{Enc} computes the state computed by Enc prior to the measurement. We argue that for any $|qpk_{dk}\rangle$ which is close to $|qpk^*\rangle$, the purified ciphertexts generated by the two keys are also close. For any bit b , the purified ciphertext are defined as $\tilde{qc}_b = U_{\text{Enc}}|qpk^*\rangle|b\rangle|0\rangle\langle 0| \langle b| \langle qpk^* | U_{\text{Enc}}^\dagger$ and $\tilde{qc}'_b = U_{\text{Enc}}|qpk_{dk}\rangle|b\rangle|0\rangle\langle 0| \langle b| \langle qpk_{dk} | U_{\text{Enc}}^\dagger$. We refer to these as purified ciphertexts. Now we can show,

$$\text{Tr}(\tilde{qc}_b \tilde{qc}'_b{}^\dagger) = \text{Tr}(U_{\text{Enc}} \langle qpk^* | qpk_{dk} \rangle |qpk^*\rangle \langle qpk_{dk} | U_{\text{Enc}}^\dagger) \quad (9)$$

$$= |\langle qpk^* | qpk_{dk} \rangle|^2 \geq (1 - \epsilon)^2 \quad (10)$$

The transition from Equation (9) to Equation (10) follows from the trace-preserving property of unitaries. Let $\{II_{dk}^b\}_{dk}$ be the POVM corresponding to decrypting a purified ciphertext with key dk , i.e. the probability of a purified ciphertext qc being decrypted to b by dk is given by $\text{Tr}(II_{dk}^b qc)$. Now the term p_{dk} can be rewritten as follows:

$$p_{dk} = \frac{1}{2} [\text{Tr}(II_{dk}^0 \tilde{qc}_0) + \text{Tr}(II_{dk}^1 \tilde{qc}_1)] \quad (11)$$

¹¹ This observation was pointed out to us by Takashi Yamakawa.

Now note that, $\text{Tr}(\Pi_{\text{dk}}^0 q c'_0) = \text{Tr}(\Pi_{\text{dk}}^1 q c'_1) = 1 - \text{negl}(\lambda)$ as we assumed Γ has negligible correctness error. Now we can bound p_{dk} as follows,

$$p_{\text{dk}} = \frac{1}{2} [\text{Tr}(\Pi_{\text{dk}}^0 \tilde{q} c_0) + \text{Tr}(\Pi_{\text{dk}}^1 \tilde{q} c_1)] \quad (12)$$

$$\geq 1 - \text{negl}(\lambda) - \frac{1}{2} [\text{Tr}(|\Pi_{\text{dk}}^0(\tilde{q} c_0 - \tilde{q} c'_0)|) + \text{Tr}(|\Pi_{\text{dk}}^1(\tilde{q} c_1 - \tilde{q} c'_1)|)] \quad (13)$$

$$\geq 1 - \text{negl}(\lambda) - \frac{1}{2} [\text{Tr}(|\tilde{q} c_0 - \tilde{q} c'_0|) + \text{Tr}(|\tilde{q} c_1 - \tilde{q} c'_1|)] \quad (14)$$

$$= 1 - \text{negl}(\lambda) - \frac{1}{2} [\sqrt{1 - \text{Tr}(\tilde{q} c_0 \tilde{q} c_0'^{\dagger})} + \sqrt{1 - \text{Tr}(\tilde{q} c_1 \tilde{q} c_1'^{\dagger})}] \quad (15)$$

$$\geq 1 - \text{negl}(\lambda) - \sqrt{2\epsilon} \geq 1 - \sqrt{3\epsilon} \quad (16)$$

The transition from Equation (14) to Equation (15) is due to $\tilde{q} c_b$ and $\tilde{q} c'_b$ being pure states. This concludes the proof of the lemma. \square

Given Lemma 1 one can reduce the adversary's task in the IND-CPA game to finding a decryption key dk such that the state $|qp\kappa_{\text{dk}}\rangle \leftarrow \text{QPKGen}(\text{dk})$ is close to $|qp\kappa^*\rangle$ in inner product distance. The main technique we use to realize this subroutine of the adversary is shadow tomography introduced by Aaronson et al. [Aar18]. At the core of our proof is the following theorem by Huang, Kueng, and Preskill [HKP20].

Theorem 4 (Theorem 1 and S16 [HKP20]). *Let O_1, \dots, O_M be M fixed observables and let ρ be an unknown n -qubit state. Given $T = O(\log(M/\delta)/\epsilon^2 \times \max_i \text{Tr}(O_i^2))$ copies of ρ , there exists a quantum algorithm that performs measurements in random Clifford basis on each copy and outputs $\tilde{p}_1, \dots, \tilde{p}_M$ such that, with probability at least $1 - \delta$*

$$\forall i, |\tilde{p}_i - \text{Tr}(O_i \rho)| \leq \epsilon$$

At a high level, the theorem states that outcomes of polynomially many random Clifford measurements on a state, i.e. a polynomial number of classical shadows, are enough to reconstruct an estimate of the statistics obtained by measuring an exponential number of observables. Note that, the post-processing required to reconstruct \tilde{p}_i values is often inefficient, however for our purpose, i.e. proving the impossibility of an information-theoretically secure quantum PKE the efficiency of the procedure is not of concern. Using Theorem 4 we are able to prove the impossibility statement.

Theorem 5. *For any security parameter λ and qPKE $\Gamma = (\text{Gen}, \text{QPKGen}, \text{Enc}, \text{Dec})$ there exists a polynomial m and a computationally unbounded adversary \mathcal{A} who can win the IND-CPA game with significant advantage only given $m(\lambda)$ copies of the public-key.*

Remark 4. Actually our attack allows us to recover the secret key with high probability, and thus the attack also breaks the one-wayness security of qPKE (which is a weaker security notion than IND-CPA). Thus, our theorem indeed shows a generic impossibility of unconditionally secure qPKE.

Proof. Let us describe the adversary given m copies of the public-key $|qp\kappa^*\rangle$ alongside a challenge ciphertext qc . We set the value of m later in the proof. For a value N , we define the following rank 1 projection ensemble $\{\Pi_{\text{dk}}^1 = |qp\kappa_{\text{dk}}\rangle\langle qp\kappa_{\text{dk}}|^{\otimes N}\}_{\text{dk} \leftarrow \text{Gen}(1^\lambda)}$. The adversary tries to find a decryption

key \mathbf{dk} such that $\text{Tr}(\Pi_{\mathbf{dk}}^1 |qp\kappa^*\rangle\langle qp\kappa^*|^{\otimes N})$ is relatively large. In order to do so the adversary computes $\text{Tr}(\Pi_{\mathbf{dk}}^1 |qp\kappa^*\rangle\langle qp\kappa^*|^{\otimes N})$ for all decryption keys \mathbf{dk} . By following the procedure from Theorem 4 on $\rho = |qp\kappa^*\rangle\langle qp\kappa^*|^{\otimes N}$, the adversary performs random Clifford measurements on

$$T = O\left(\log\left(\frac{\#\{\mathbf{dk}|\mathbf{dk} \leftarrow \text{Gen}(1^\lambda)\}}{\delta}\right)\right) \frac{1}{\epsilon^2} \text{Tr}(\Pi_{\mathbf{dk}}^{1^2})$$

copies of ρ to compute values $\tilde{p}_{\mathbf{dk}}$ such that with probability $1 - \delta$, for all \mathbf{dk}

$$\left|\tilde{p}_{\mathbf{dk}} - \text{Tr}(\Pi_{\mathbf{dk}}^1 |qp\kappa^*\rangle\langle qp\kappa^*|^{\otimes N})\right| \leq \epsilon.$$

Let us set $\epsilon < 1/6$ and δ to be a small constant, e.g. $1/100$. Immediately it can be noticed that as ϵ and δ are constants and $\text{Tr}(\Pi_{\mathbf{dk}}^{1^2}) = 1^{12}$, T is $O(\log(\#\{\mathbf{dk}|\mathbf{dk} \leftarrow \text{Gen}(1^\lambda)\}))$ which is $\text{poly}(\lambda)$ as the key-lengths should be polynomial in the security parameter.

We claim that if the adversary picks any key such that $\tilde{p}_{\mathbf{dk}} > 1/2$, they have found a key that has a high chance of decrypting the challenge ciphertext correctly. Let us elaborate. First of all, note that the adversary finds at least one such \mathbf{dk} with probability at least $1 - \frac{1}{100}$, as for the correct decryption key \mathbf{dk}^* , $\text{Tr}(\Pi_{\mathbf{dk}^*}^1 |qp\kappa^*\rangle\langle qp\kappa^*|^{\otimes N}) = 1$ hence $\tilde{p}_{\mathbf{dk}^*} > 1 - 1/6$ with probability at least $1 - \frac{1}{100}$.

The next thing to show is that any \mathbf{dk} such that $\tilde{p}_{\mathbf{dk}} > 1/2$ is a *good* decryption key. Note that due to Lemma 1 we have,

$$\text{Tr}(\Pi_{\mathbf{dk}}^1 |qp\kappa^*\rangle\langle qp\kappa^*|^{\otimes N}) = |\langle qp\kappa_{\mathbf{dk}} | qp\kappa^* \rangle|^{2N} \quad (17)$$

We note that for all \mathbf{dk} such that $p_{\mathbf{dk}} \leq 1 - \sqrt{\frac{3}{\log(N)}}$ we have:

$$p_{\mathbf{dk}} \leq 1 - \sqrt{\frac{3}{\log(N)}} \Rightarrow \langle qp\kappa_{\mathbf{dk}} | qp\kappa^* \rangle \leq 1 - \frac{1}{\log(N)} \quad (18)$$

$$\Rightarrow \text{Tr}(\Pi_{\mathbf{dk}}^1 |qp\kappa^*\rangle\langle qp\kappa^*|^{\otimes N}) \leq \left(1 - \frac{1}{\log(N)}\right)^{2N} \quad (19)$$

$$\leq e^{-2N/\log(N)} \ll 1/3, \text{ for a large enough } N \quad (20)$$

Given our choice of δ, ϵ , this ensures that if the adversary picks any \mathbf{dk} such that $\tilde{p}_{\mathbf{dk}} > 1/2$, with probability at least $1 - \frac{1}{100}$ we have that, $\left|\tilde{p}_{\mathbf{dk}} - \text{Tr}(\Pi_{\mathbf{dk}}^1 |qp\kappa^*\rangle\langle qp\kappa^*|^{\otimes N})\right| \leq 1/6$, $\text{Tr}(\Pi_{\mathbf{dk}}^1 |qp\kappa^*\rangle\langle qp\kappa^*|^{\otimes N}) > 1/3$ hence, $p_{\mathbf{dk}} > 1 - \sqrt{\frac{3}{\log(N)}}$.

As the last step, the adversary uses the \mathbf{dk} they obtain from the previous procedure to decrypt the challenge ciphertext qc^* . By union bound and following the discussion above the adversary's advantage to decrypt the challenge ciphertext correctly is greater than $1 - \frac{1}{100} - \sqrt{\frac{3}{\log(N)}}$ which by setting N to be a large constant is significantly larger than $1/2$. Finally note that this adversary uses $m = NT$ copies of the public-key, where $T = \text{poly}(\lambda)$ and N is a constant, so the total number of public-key copies used are polynomial in λ . \square

¹² this is due to $\Pi_{\mathbf{dk}}^1$ operators being rank-1 projections

Acknowledgments

The authors wish to thank Prabhanjan Ananth and Umesh Vazirani for related discussions, and Takashi Yamakawa for pointing out a simple argument to rule out the existence of information-theoretically secure qPKE. The argument is replicated here with his permission.

ABG and QHV are supported by ANR JCJC TCS-NISQ ANR-22-CE47-0004, and by the PEPR integrated project EPiQ ANR-22-PETQ-0007 part of Plan France 2030. GM was partially funded by the German Federal Ministry of Education and Research (BMBF) in the course of the 6GEM research hub under grant number 16KISK038 and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2092 CASA – 390781972. OS was supported by the Israeli Science Foundation (ISF) grant No. 682/18 and 2137/19, and by the Cyber Security Research Center at Ben-Gurion University. KB and LH are supported by the Swiss National Science Foundation (SNSF) through the project grant 192364 on Post Quantum Cryptography.

OS has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 756482). MW acknowledges support by the the European Union (ERC, SYMOPTIC, 101040907), by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2092 CASA - 390781972, by the BMBF through project QuBRA, and by the Dutch Research Council (NWO grant OCENW.KLEIN.267). Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.



References

- Aar18. Scott Aaronson. Shadow tomography of quantum states. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *50th ACM STOC*, pages 325–338. ACM Press, June 2018.
- ACC⁺22. Per Austrin, Hao Chung, Kai-Min Chung, Shiuan Fu, Yao-Ting Lin, and Mohammad Mahmoody. On the impossibility of key agreements from quantum random oracles. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 165–194. Springer, Heidelberg, August 2022.
- AGM18. Gorjan Alagic, Tommaso Gagliardoni, and Christian Majenz. Unforgeable quantum encryption. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 489–519. Springer, Heidelberg, April / May 2018.
- AGQY22. Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. Pseudorandom (function-like) quantum state generators: New definitions and applications. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 237–265. Springer, Heidelberg, November 2022.
- AHU19. Andris Ambainis, Mike Hamburg, and Dominique Unruh. Quantum security proofs using semi-classical oracles. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 269–295. Springer, Heidelberg, August 2019.
- AQY22. Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 208–236. Springer, Heidelberg, August 2022.
- BB84. Charles H. Bennett and Gilles Brassard. An update on quantum cryptography (impromptu talk). In G. R. Blakley and David Chaum, editors, *CRYPTO’84*, volume 196 of *LNCS*, pages 475–480. Springer, Heidelberg, August 1984.
- BBSS23. Amit Behera, Zvika Brakerski, Or Sattath, and Omri Shmueli. Pseudorandomness with proof of destruction and applications. *Cryptology ePrint Archive*, Paper 2023/543, 2023. <https://eprint.iacr.org/2023/543>.

- BCKM21. James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. One-way functions imply secure computation in a quantum world. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part I*, volume 12825 of *LNCS*, pages 467–496, Virtual Event, August 2021. Springer, Heidelberg.
- BMW23. Khashayar Barooti, Giulio Malavolta, and Michael Walter. A simple construction of quantum public-key encryption from quantum-secure one-way functions. Cryptology ePrint Archive, Paper 2023/306, 2023. <https://eprint.iacr.org/2023/306>.
- BS23. Mohammed Barhoush and Louis Salvail. How to sign quantum messages. *arXiv preprint arXiv:2304.06325*, 2023.
- BZ13a. Dan Boneh and Mark Zhandry. Quantum-secure message authentication codes. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 592–608. Springer, Heidelberg, May 2013.
- BZ13b. Dan Boneh and Mark Zhandry. Secure signatures and chosen ciphertext security in a quantum computing world. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 361–379. Springer, Heidelberg, August 2013.
- CDMW18. Seung Geol Choi, Dana Dachman-Soled, Tal Malkin, and Hoeteck Wee. A black-box construction of non-malleable encryption from semantically secure encryption. *Journal of Cryptology*, 31(1):172–201, January 2018.
- CEV22. Céline Chevalier, Ehsan Ebrahimi, and Quoc Huy Vu. On security notions for encryption in a quantum world. In Takanori Ito and Santanu Sarkar, editors, *Progress in Cryptology - INDOCRYPT 2022 - 23rd International Conference on Cryptology in India, Kolkata, India, December 11-14, 2022, Proceedings*, volume 13774 of *Lecture Notes in Computer Science*, pages 592–613. Springer, 2022.
- Col23. Andrea Coladangelo. Quantum trapdoor functions from classical one-way functions. Cryptology ePrint Archive, Paper 2023/282, 2023. <https://eprint.iacr.org/2023/282>.
- Dol20. Javad Doliskani. Efficient quantum public-key encryption from learning with errors. Cryptology ePrint Archive, Paper 2020/1557, 2020. <https://eprint.iacr.org/2020/1557>.
- GGM86. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM (JACM)*, 33(4):792–807, 1986.
- GLSV21. Alex B. Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan. Oblivious transfer is in MiniQCrypt. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 531–561. Springer, Heidelberg, October 2021.
- Got05. Daniel Gottesman. Quantum public key cryptography with information-theoretic security. <https://www2.perimeterinstitute.ca/personal/dgottesman/Public-key.ppt>, 2005.
- GSV23. Alex B. Grilo, Or Sattath, and Quoc-Huy Vu. Encryption with quantum public keys. Cryptology ePrint Archive, Paper 2023/345, 2023. <https://eprint.iacr.org/2023/345>.
- HKP20. Hsin-Yuan Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10):1050–1057, 2020.
- IR90. Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In Shafi Goldwasser, editor, *CRYPTO’88*, volume 403 of *LNCS*, pages 8–26. Springer, Heidelberg, August 1990.
- JLS18. Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 126–152. Springer, Heidelberg, August 2018.
- KKNY05. Akinori Kawachi, Takeshi Koshihara, Harumichi Nishimura, and Tomoyuki Yamakami. Computational indistinguishability between quantum states and its cryptographic application. In Ronald Cramer, editor, *EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 268–284. Springer, Heidelberg, May 2005.
- KMNY23. Fuyuki Kitagawa, Tomoyuki Morimae, Ryo Nishimaki, and Takashi Yamakawa. Quantum public-key encryption with tamper-resilient public keys from one-way functions. Cryptology ePrint Archive, Paper 2023/490, 2023. <https://eprint.iacr.org/2023/490>.
- KQST22. William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. *arXiv preprint arXiv:2212.00879*, 2022.
- Kre21. William Kretschmer. Quantum pseudorandomness and classical complexity. In Min-Hsiu Hsieh, editor, *16th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2021, July 5-8, 2021, Virtual Conference*, volume 197 of *LIPICs*, pages 2:1–2:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- Ms09. Steven Myers and abhi shelat. Bit encryption is complete. In *50th FOCS*, pages 607–616. IEEE Computer Society Press, October 2009.

- MW23. Giulio Malavolta and Michael Walter. Non-interactive quantum key distribution. Cryptology ePrint Archive, Paper 2023/500, 2023. <https://eprint.iacr.org/2023/500>.
- MY22a. Tomoyuki Morimae and Takashi Yamakawa. One-wayness in quantum cryptography. Cryptology ePrint Archive, Paper 2022/1336, 2022. <https://eprint.iacr.org/2022/1336>.
- MY22b. Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 269–295. Springer, Heidelberg, August 2022.
- NC16. Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information (10th Anniversary edition)*. Cambridge University Press, 2016.
- NI09. Georgios M. Nikolopoulos and Lawrence M. Ioannou. Deterministic quantum-public-key encryption: Forward search attack and randomization. *Phys. Rev. A*, 79:042327, Apr 2009.
- OTU00. Tatsuaki Okamoto, Keisuke Tanaka, and Shigenori Uchiyama. Quantum public-key cryptosystems. In Mihir Bellare, editor, *CRYPTO 2000*, volume 1880 of *LNCS*, pages 147–165. Springer, Heidelberg, August 2000.
- Wie83. Stephen Wiesner. Conjugate coding. *SIGACT News*, 15(1):78–88, jan 1983.
- Zha12. Mark Zhandry. How to construct quantum random functions. In *53rd FOCS*, pages 679–687. IEEE Computer Society Press, October 2012.

A CCA-Secure Bit-Encryption from OWF

In this appendix, we describe a simple quantum public key bit encryption scheme that satisfies the strong notion of CCA security. The construction relies on a quantum-secure pseudorandom function

$$\text{PRF}: \{0, 1\}^\lambda \times \{0, 1\}^\lambda \rightarrow \{0, 1\}^{3\lambda}$$

which, as mentioned earlier in Section 2, can be constructed from any quantum-secure one-way function. Then our quantum PKE scheme $\Pi = (\text{Gen}, \text{QPKEGen}, \text{Enc}, \text{Dec})$ is defined as follows:

- The key generation algorithm $\text{Gen}(1^\lambda)$ samples two keys $\text{dk}_0 \xleftarrow{\$} \{0, 1\}^\lambda$ and $\text{dk}_1 \xleftarrow{\$} \{0, 1\}^\lambda$ and sets $\text{dk} = (\text{dk}_0, \text{dk}_1)$. The public-key generation $\text{QPKEGen}(\text{dk})$ prepares the states

$$|\text{qp}\kappa_0\rangle = \sum_{x \in \{0, 1\}^\lambda} |x, f_{\text{dk}_0}(x)\rangle \quad \text{and} \quad |\text{qp}\kappa_1\rangle = \sum_{x \in \{0, 1\}^\lambda} |x, f_{\text{dk}_1}(x)\rangle.$$

Where $\{f_{\text{dk}}\}_{\text{dk}}$ is a PRF. Note that both states are efficiently computable since the PRF can be efficiently evaluated in superposition in view of Fact 1. The quantum public key is then given by the pure state $|\text{qp}\kappa\rangle = |\text{qp}\kappa_0\rangle \otimes |\text{qp}\kappa_1\rangle$, whereas the classical secret key consists of the pair $\text{dk} = (\text{dk}_0, \text{dk}_1)$.

- Given a message $\text{pt} \in \{0, 1\}$, the encryption algorithm $\text{Enc}(|\text{qp}\kappa\rangle, \text{pt})$ simply measures $|\text{qp}\kappa_{\text{pt}}\rangle$ in the computational basis, and outputs the measurement outcome as the *classical* ciphertext $qc = (x, y)$ and the post measurement state $|x\rangle|y\rangle$.
- Given the ciphertext $qc = (x, y)$, the decryption algorithm $\text{Dec}(\text{dk}, qc)$ first checks whether $f_{\text{dk}_0}(x) = y$ and returns 0 if this is the case. Next, it checks whether $f_{\text{dk}_1}(x) = y$ and returns 1 in this case. Finally, if neither is the case, the decryption algorithm returns \perp .

Next, we establish correctness of this scheme.

Theorem 6. *If PRF is a quantum-secure pseudorandom function, then the quantum PKE scheme Π is correct.*

Proof. Observe that the scheme is perfectly correct if the ranges of f_{dk_0} and f_{dk_1} are disjoint. By a standard argument, we can instead analyze the case of two truly random functions f_0 and f_1 , and the same will hold for f_{dk_0} and f_{dk_1} , except on a negligible fraction of the inputs. Fix the range of f_0 , which is of size at most 2^λ . Then the probability that any given element of f_1 falls into the same set is at most $2^{-2\lambda}$, and the desired statement follows by a union bound. \square

Finally, we show that the scheme is CCA-secure. The main tool used in the proof is the one-way to hiding lemma [AHU19].

Lemma 2 (One-way to hiding). *Let $G, H : X \rightarrow Y$ be random functions and $S \subset X$ an arbitrary set with the condition that $\forall x \notin S, G(x) = H(x)$, and let z be a random bitstring. Further, let $\mathcal{A}^H(z)$ be a quantum oracle algorithm that queries H with depth at most d . Define $\mathcal{B}^H(z)$ to be an algorithm that picks $i \in [d]$ uniformly, runs $\mathcal{A}^H(z)$ until just before its i^{th} round of queries to H and measures all query input registers in the computational basis and collects them in a set T . Let*

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

$$P_{\text{guess}} = \Pr[S \cap T \neq \emptyset | T \leftarrow \mathcal{B}^H(z)]$$

Then we have that

$$|P_{\text{left}} - P_{\text{right}}| \leq 2d\sqrt{P_{\text{guess}}} \quad \text{and} \quad |\sqrt{P_{\text{left}}} - \sqrt{P_{\text{right}}}| \leq 2d\sqrt{P_{\text{guess}}} \quad (21)$$

Theorem 7. *If $\{f_{dk}\}_{dk}$ is a quantum-secure pseudorandom function ensemble, then the quantum PKE scheme Π is CCA-secure.*

Proof. It suffices to show that the CCA experiment with the bit b fixed to 0 is indistinguishable from the same experiment but with b fixed to 1. To this end we consider a series of hybrids, starting with the former and ending with the latter:

- **Hybrid 0:** This is the original CCA experiment except that the bit b fixed to 0.
- **Hybrid 1:** In this (inefficient) hybrid, we modify hybrid 0 to instead compute $|qp\kappa_0\rangle$ as

$$|qp\kappa_0\rangle = \sum_{x \in \{0,1\}^\lambda} |x, f(x)\rangle,$$

where f is a truly uniformly random function.

The indistinguishability between these two hybrids follows by a standard reduction against the quantum security of PRF: To simulate the desired n copies of $|qp\kappa_0\rangle$, and to answer decryption queries (except the one that contains the challenge ciphertext), the reduction simply queries the oracle provided by the PRF security experiment (possibly in superposition). Note that whenever the oracle implements PRF, then the view of the distinguisher is identical to hybrid 0, whereas if the oracle implements a truly random function, then the view of the distinguisher is identical to hybrid 1.

- **Hybrid 2:** In this (inefficient) hybrid, we modify hybrid 1 such that the challenge ciphertext is sampled as

$$x \xleftarrow{\$} \{0,1\}^\lambda \quad \text{and} \quad y \xleftarrow{\$} \{0,1\}^{3\lambda}.$$

The indistinguishability of hybrids 1 and 2 follows from the one-way to hiding lemma (Lemma 2). Let H be such that $H(x) = y$ and for all $x' \neq x$ we set $H(x') = f(x')$, and let $S = \{x\}$. Let \mathcal{A} be the adversary playing the security experiment. We claim that \mathcal{A}^f is the adversary playing in hybrid 1 whereas \mathcal{A}^H corresponds to the adversary playing hybrid 2: Observe that the public keys can be simulated with oracle access to f (H , respectively) by simply querying on a uniform superposition of the input domain, whereas the decryption queries can be simulated by query basis states. Importantly, for all queries after the challenge phase, the adversary is not allowed to query x to Dec^* . Hence the set T , collected by \mathcal{B} is a set of at most n uniform elements from the domain of f , along with Q basis states, where Q denotes the number of queries made by the adversary to the decryption oracle *before* the challenge ciphertext is issued. By a union bound

$$P_{\text{guess}} = \Pr[T \cap \{x\} \neq \emptyset] \leq \frac{(n + Q)}{2^\lambda} = \text{negl}(\lambda)$$

since x is uniformly sampled. Applying Lemma 2, we deduce that $|P_{\text{left}} - P_{\text{right}}|$ is also negligible, i.e., which bounds the distance between the two hybrids.

- **Hybrid 3:** In this (efficient) hybrid, we modify hybrid 2 to compute $|qp\kappa_0\rangle$ by using the pseudorandom function f_{dk_0} instead of the truly random function f . That is, we revert the change done in hybrid 1.

Indistinguishability follows from the same argument as above.

- **Hybrid 4:** In this (inefficient) hybrid, we modify hybrid 3 to compute $|qp\kappa_1\rangle$ as

$$|qp\kappa_1\rangle = \sum_{x \in \{0,1\}^\lambda} |x, f(x)\rangle$$

where f is a truly uniformly random function.

Indistinguishability follows from the same argument as above.

- **Hybrid 5:** In this (inefficient) hybrid, we modify hybrid 4 by fixing the bit b to 1 and computing the challenge ciphertext honestly, i.e., as

$$x \xleftarrow{\$} \{0,1\}^\lambda \quad \text{and} \quad y = f(x).$$

Indistinguishability follows from the same argument as above.

- **Hybrid 6:** In this (efficient) hybrid, we modify hybrid 5 to compute $|qp\kappa_1\rangle$ by using the pseudorandom function f_{dk_1} instead of the truly random function f . That is, we revert the change done in hybrid 4.

Indistinguishability follows from the same argument as above. The proof is concluded by observing that the last hybrid is identical to the CCA experiment with the bit b fixed to 1. \square