

Revocable Encryption, Programs, and More: The Case of Multi-Copy Security

Prabhanjan Ananth¹, Saachi Mutreja², and Alexander Poremba³

¹University of California, Santa Barbara

²Columbia University

³Massachusetts Institute of Technology

Abstract

Fundamental principles of quantum mechanics have inspired many new research directions, particularly in quantum cryptography. One such principle is *quantum no-cloning* which has led to the emerging field of revocable cryptography. Roughly speaking, in a revocable cryptographic primitive, a cryptographic object (such as a ciphertext or program) is represented as a quantum state in such a way that surrendering it effectively translates into losing the capability to use this cryptographic object. All of the revocable cryptographic systems studied so far have a major drawback: the recipient only receives one copy of the quantum state. Worse yet, the schemes become completely insecure if the recipient receives many identical copies of the *same* quantum state—a property that is clearly much more desirable in practice.

While multi-copy security has been extensively studied for a number of other quantum cryptographic primitives, it has so far received only little treatment in context of unclonable primitives. Our work, for the first time, shows the feasibility of revocable primitives, such as revocable encryption and revocable programs, which satisfy multi-copy security in oracle models. This suggests that the stronger notion of multi-copy security is within reach in unclonable cryptography more generally, and therefore could lead to a new research direction in the field.

1 Introduction

Designing mechanisms to provably revoke cryptographic capabilities is an age-old problem [Stu95, Riv98]. In the public-key infrastructure, certificate authorities have the ability to invalidate public-key certificates [Wik24], especially when the certificates have been compromised. Key rotation policies [Clo24] guarantee that outdated decryption keys become ineffective for future use. The existing approaches to tackle with this problem have their limitations owing to the fact that cryptographic secrets are represented as binary strings and hence, it is infeasible to provably ensure that the malicious attackers have erased information from their devices. In the context of key rotation, a compromised key can still be used to decrypt old ciphertexts. Another issue with using classical information to represent cryptographic keys is that it is difficult to detect compromise: a hacker could steal the classical key from a device without leaving a trace.

Recently, a line of works [Unr13, AL21, AKN+23, APV23, CGJL23, MPY23, AHH24] have leveraged quantum information and proposed new approaches for provable revocation of cryptographic objects, such as ciphertexts, programs and keys. These works studied revocation in the context of many Crypto 101 primitives, including pseudorandom functions, private-key and public-key encryption and digital signatures. In a revocable cryptographic primitive, an object (such as a program or a decryption key, etc.) is associated with a quantum state in such a way that *only with access to the state* the functionality of the original cryptographic object is retained. The common template for defining revocable security is in the form of a cryptographic game: The adversary receives *one copy* of the quantum state that it can use for a limited period of time after which it is supposed to return back the state to the owner. The security guarantee stipulates

that after the state is returned, the adversary effectively loses the capability to use the cryptographic object. At this point, it should be clear to the reader the necessity that such cryptographic objects are represented as quantum states: indeed, if they were classical, the adversary could always maintain a secret copy, while pretending to have erased everything from its device. On the other hand, the no-cloning theorem [WZ82, Die82] of quantum mechanics suggests that the above security experiment could very well be achieved.

Multi-Copy Security. Let us now zoom in on the part of the security experiment, where the adversary receives *only one copy* of the quantum state. In all of the prior works in the literature so far [AKN⁺23, APV23, CGJL23, MPY23, AHH24], this limitation persists. One could consider a more general definition, where the adversary receives k *identical* copies of the quantum state and is later asked to return back all of the copies of the state. The security guarantee is similar to before: after returning all of the k copies, the adversary should effectively lose all access to the underlying cryptographic object. We term this general security experiment to be *multi-copy security*.

There are a couple of reasons to study multi-copy security for revocable primitives.

- **HISTORICAL CONTEXT:** Multi-copy security is not new and has been extensively studied in quantum cryptography, especially in the context of foundational primitives such as pseudorandom states [JLS18, BS19, AQY22, AGQY22, BEM⁺23, BHHP24] and one-way state generators [MY22b, MY22a]. Indeed, multi-copy security has been crucial in the design of many cryptographic constructions. The works of [AGQY22, ALY24, KT24] used tomography, which inherently requires multiple copies in order to dequantize the communication in some of the quantum cryptographic primitives. Specifically for revocable primitives, a conceptual reason to study multi-copy security is to understand whether having more copies necessarily makes it easier for the adversary to clone quantum states. Investigating multi-copy security for revocable primitives is a starting step towards understanding multi-copy security for more advanced primitives such as public-key quantum money [AC12, Zha21a]. The question of whether multi-copy security is possible in unclonable cryptography was also recently raised in [MPSY24].
- **NESTED LEASING:** Having access to many more copies of the quantum state would also give more power to the user; for example, using a permutation test [KNY08] (a generalized version of SWAP test) where one is given $|\phi\rangle$ and polynomially many copies of $|\psi\rangle$, one can approximately test the overlap between $|\phi\rangle$ and $|\psi\rangle$. This ability allows for nested leasing of cryptographic objects, such as programs or keys. Suppose a user **A** is leased a large number of, say k , copies of a quantum program $|\psi\rangle$. User **A** could further lease a number, say $k' \ll k$, of copies of $|\psi\rangle$ to user **B**. At a later point in time, when user **B** is asked to return back its copies to user **A**. User **A** can then use its $k - k'$ copies to approximately test whether the returned copies are correct. If the test succeeds, user **A** then is in a position to return back all of its k copies to the true owner¹. Such a nested leasing approach could be especially useful in organizations with hierarchical structure.

The notion of multi-copy security we consider in this work is closely related to *collusion-resistant security*, considered in the works of [LLQZ22, QG23]. The crucial difference is that in the prior works, the adversary receives *i.i.d* copies of the quantum key whereas in our case, the adversary receives *identical* copies of the quantum key. In the nested leasing application discussed above, it was crucial that the user received many *identical* copies of the same state.

Multi-copy security using commonly studied unclonable states: Challenges. The first step towards addressing multi-copy security is to identify quantum states that are unclonable. We discuss the commonly studied unclonable quantum states below:

- **BB84 STATES:** these states are of the form $H^\theta |x\rangle$, where $\theta \in \{0, 1\}^n$ and $x \in \{0, 1\}^n$. These states have been influential in the design of private-key quantum money [Wie83] and in the design of encryption

¹A drawback of this approach is that there is some room for user **B** to cheat with noticeable probability in this approach without user **A** noticing, which means that the owner would not always be to pinpoint whether user **A** or user **B** cheated. Still, this approach offers a non-trivial solution to this challenging problem.

schemes with unclonable ciphertexts [BL19, BI20]. Given many copies of $H^\theta |x\rangle$, one can learn x and θ and hence, recover a complete description of $H^\theta |x\rangle$.

- **SUBSPACE AND COSET STATES:** these states are of the form $(\sqrt{|A|})^{-1} \cdot \sum_{\mathbf{x} \in A} |\mathbf{x}\rangle$, where $A \subseteq \mathbb{F}_2^n$ is a sparse subspace of \mathbb{F}_2^n , for some $n \in \mathbb{N}$. These states have been crucial in the design of public-key quantum money [AC12, Zha21a], among other primitives. Again given many copies, one can learn the basis of the subspace and hence a complete description of the subspace state. Another related class of unclonable states are coset states which are superpositions over a coset (rather than a subspace) and moreover, each term in the superposition has a phase that depends on a dual coset. Coset states have been influential in constructions of quantum copy-protection [CLLZ21]. Similar to subspace states, coset states are also learnable.
- **SIS-BASED STATES:** these states are of the form $\sum_{\mathbf{x} \in \mathbb{Z}_q^m, \mathbf{A}\mathbf{x}=\mathbf{y}} \alpha_{\mathbf{x}} |\mathbf{x}\rangle$, where $q, m \in \mathbb{N}$, $|\alpha_{\mathbf{x}}|^2$ is a discrete Gaussian distribution such that most of the weight is on low norm vectors \mathbf{x} and finally, $\mathbf{A} \in \mathbb{Z}_q^{n \times m}, \mathbf{y} \in \mathbb{Z}_q^n$. They were useful in designing traditional and advanced encryption systems with unclonable quantum keys [Por22, APV23, MPY23, AHH24]. Given many copies of this state, one can recover a short basis of the kernel of \mathbf{A} which can then be used to recover the above state.

In other words, all the above types of states are learnable and hence, they cannot be the basis of any unclonable cryptographic scheme satisfying multi-copy security. This suggests that we need to look for new unclonable quantum states that are unlearnable even given many copies. In the past, discovering new unclonable quantum states has led to pushing the frontier of unclonable quantum cryptographic primitives and we believe our endeavour could reap similar results.

2 Our Results

We now give an overview of our results.

Our Approach: Pseudorandomness Meets Unclonable Cryptography. We use *subset states* to tackle multi-copy security. Subset states have been recently studied in the context of quantum pseudorandomness. A subset state is associated with an unstructured and random subset $S \subset \{0, 1\}^n$ of the form $|S\rangle = (\sqrt{|S|})^{-1} \sum_{x \in S} |x\rangle$. Several recent works [GB23, JMW23] showed that random subset states (of non-trivial size) are approximate state k -designs for any polynomial k as a function of $n \in \mathbb{N}$, which make random subset states a natural candidate for multi-copy security—particularly since Haar states are unlearnable given polynomially many identical copies. At first sight, it would seem as though that the fact that random subset states are close to Haar states should discourage us from using them for unclonable cryptography, especially since Haar states have virtually no structure. Indeed, the structure of BB84 states, subspace states and others has been crucially exploited in various applications. Our main contribution is in showing that subset states, in some sense, have the minimal amount of structure to enable a number of interesting cryptographic primitives in the context of multi-copy revocable cryptography.

Multi-Copy Revocable Encryption. We first study revocable encryption [Unr13, AKN+23, APV23] with multi-copy security. A revocable encryption scheme is a regular encryption scheme but where the ciphertexts are associated with quantum states. Additionally, a revocable encryption scheme comes with the following security notion called *multi-copy revocable security*: informally, it states that any adversary that successfully returns k valid copies of a quantum ciphertext which it was given by a challenger, where k is an arbitrary polynomial, necessarily loses the ability to decrypt the ciphertexts in the future—even if the secret key is revealed. In more detail, the security game is formulated as follows:

- The adversary selects a pair of messages (m_0, m_1) and sends them to the challenger.
- The challenger randomly selects one of the two messages, say m_b , and encrypts it k times using a secret key \mathbf{sk} , and sends the ciphertext copies $|\psi_b\rangle^{\otimes k}$ to the adversary \mathcal{A} .

- At a later point in time, \mathcal{A} returns back all the copies of $|\psi_b\rangle$, which are then verified by the challenger.
- After successfully returning back the states, \mathcal{A} receives the secret key sk in the clear. Finally, \mathcal{A} outputs a guess b' .

The scheme is said to be secure if the probability that $b' = b$ is close to $1/2$.

Prior works [Unr13, AKN⁺23, APV23, CGJL23, AHH24] only studied variants of the above security game in the setting where the adversary receives only one copy of the quantum ciphertext. In fact, their schemes are easily seen to not satisfy multi-copy security. We show the following.

Theorem 2.1. *There exists an encryption scheme with (an oracular notion of) multi-copy revocable security in the random permutation model.*

Some remarks are in order. Firstly, our security proof does not fully achieve the standard notion of revocable security guarantee we stated above; rather, we consider a slightly different variant of the experiment, where in the second part of the game (instead of revealing the secret key in the clear) we allow the adversary to query an oracle that is powerful enough to enable decryption during the first phase of the game. While this constitutes a weaker notion of security, it nevertheless results in a meaningful notion of revocable security: once the adversary has successfully returned all of the copies of the ciphertext, it can no longer decrypt the ciphertext in the future—even if it gets access to an oracle that would have previously allowed it to do so. Second, our construction of multi-copy revocable encryption makes use of (ideal) random permutations; for example, one could instantiate such permutations in practice by using a quantum-secure pseudorandom permutation [Zha16]. In the security experiment which underlies our construction, the aforementioned oracle for decryption (i.e., that which is handed to the adversary after revocation has taken place) is in the form of an oracle for the permutation itself. We remark that this model resembles the random permutation model behind the international hash function standard SHA-3 [BDPA11a, BDPA11b], except that the adversary only receives oracle access to the permutation during the second part of the revocable security experiment. While our construction only achieves an oracular notion of revocable security, it is nevertheless the very first construction of revocable encryption which satisfies multi-copy security *in any model*.

Multi-Copy Revocable Programs. Our previous discussion on revocable encryption illustrates that encryption and decryption functionalities can be protected even if many copies of the quantum ciphertext are made available to the recipient. We generalize this result further and study whether arbitrary functionalities can be protected. We define and study revocable programs with multi-copy security. In this notion, there is a functionality preserving compiler that takes a program and converts it into a quantum state. The security guarantee is defined similar to revocable encryption:

- The challenger compiles a program P , sampled from a distribution \mathcal{D} on a set of programs \mathcal{P} , into a state $|\psi_P\rangle$. It then sends k copies of the state $|\psi_P\rangle$ to the adversary \mathcal{A} .
- At a later point in time, \mathcal{A} returns back all of the copies of $|\psi_P\rangle$.
- After returning back the state, \mathcal{A} is given x , where x is sampled from the input distribution of P . It then outputs a guess y .

The scheme is said to be secure if the probability that $y = P(x)$ is roughly close to the trivial success probability. Here, the trivial success probability is defined as the optimal probability of guessing $P(x)$ given just x (and the knowledge of \mathcal{D} and the input distribution).

Prior works propose revocable programs for specific functionalities in the plain model [CLLZ21] or for general functionalities in oracle models [ALL⁺21]. However, these works guarantee security only if the adversary receives one copy of the state. And as before, these constructions provably do not satisfy multi-copy security. We show the following.

Theorem 2.2. *There exist revocable programs which satisfy (an oracular notion of) multi-copy security in a classical oracle model.*

Unlike [Theorem 2.1](#), the above theorem relies upon more structured classical oracles. Similar to [Theorem 2.1](#), our construction guarantees information-theoretic security.

Finally, for the special case of point functions, we show that we can again only rely upon the random permutation model together with the (standard) quantum random oracle model. We show the following.

Theorem 2.3. *There exist revocable multi-bit point functions which satisfy (an oracular notion of) multi-copy security in both the random permutation model, as well as the quantum random oracle model.*

Our results and techniques opens the door for building more advanced unclonable primitives that preserve their security even if the adversary receives many copies of the unclonable quantum state.

Are Results in Oracle Models Interesting? It is natural for a reader to be skeptical of our results given that they are based in the oracle models. However, we would like to emphasize that achieving results in the oracle models still requires non-trivial amount of effort. As history suggests, constructions in the oracle models have eventually been adopted to constructions in the plain model. A classic example is the construction of public-key quantum money, which was first proposed in the oracle models by Aaronson and Christiano [[AC12](#)] and later, being instantiated in the plain model by Zhandry [[Zha21a](#)]. In a similar vein, our techniques could be useful for future works on achieving multi-copy security in the plain model.

Finally, we remark that some of our results only make use of the *ideal random permutation model*, a much more modest assumption which also resembles the security model behind SHA-3.

Applications to Sponge Hashing. As a complementary contribution, we show that the techniques we developed in this paper are more broadly applicable and extend to other cryptographic settings as well. Here, we single out the so-called *sponge construction* used in SHA-3 [[BDPA11a](#), [BDPA11b](#)].

In [Section 11](#), we study a simple query problem: Suppose that an adversary receives as input a hash table for a set of random input keys, where each hash is computed using a *salted* (one-round) sponge hash function. How many quantum queries are necessary to find a new *valid* element in the range of the hash function? Our contribution is a space-time trade-off which precisely characterizes the hardness of finding hash table elements in the presence of oracles that depend non-trivially on the sponge hash function.

Acknowledgements. The authors would like to thank Henry Yuen and Tal Malkin for many useful discussions. PA is supported by the National Science Foundation under Grant No. 2329938 and Grant No. 2341004. SM is supported by AFOSR award FA9550-21-1-0040, NSF CAREER award CCF-2144219, and the Sloan Foundation. AP is supported by the National Science Foundation (NSF) under Grant No. CCF-1729369.

2.1 Related work

We now discuss related notions which are relevant to this work.

Copy-Protection. This notion was first introduced by Aaronson [[Aar09](#)]. Informally speaking, a copy-protection scheme is a compiler that transforms programs into quantum states in such a way that using the resulting states, one can run the original program. Yet, the security guarantee stipulates that any adversary given one copy of the state cannot produce a bipartite state wherein both parts compute the original program. Copy-protection schemes have since been constructed for various classes of programs and under various different models, for example as in [[Aar09](#), [ALL⁺21](#), [CMP24](#), [AK22](#), [AKL23](#), [LLQZ22](#), [QG23](#)]. We remark, however, that all of the aforementioned works are completely insecure if multiple identical copies of the program are made available. The notion of multi-copy security we consider in this work is closely related to *collusion-resistant security*, considered in the works of [[LLQZ22](#), [QG23](#)]. The crucial difference is that in the prior works, the adversary receives *i.i.d* copies of the quantum key whereas in our case, the adversary receives *identical* copies of the quantum key.

Secure Software Leasing. Another primitive relevant to revocable cryptography is secure software leasing [AL21]. The notion of secure software leasing states that any program can be compiled into a functionally equivalent program, represented as a quantum state, in such a way that once the compiled program is returned, the (honest) evaluation algorithm on the residual state cannot compute the original functionality. Secure leasing has been constructed for various functionalities [AL21, CMP24, B JL⁺21, KNY21]. Similar to copy-protection, none of the aforementioned works consider multi-copy security.

Encryption Schemes with Revocable Ciphertexts. Unruh [Unr13] proposed a (private-key) quantum timed-release encryption scheme that is *revocable*, i.e. it allows a user to *return* the ciphertext of a quantum timed-release encryption scheme, thereby losing all access to the data. Broadbent and Islam [BI20] introduced the notion of *certified deletion*, which is incomparable with the related notion of unclonable encryption. This has led to the development of other certified deletion protocols, for example as in Ref. [Por22, BK23, BKP23, BKM⁺23, BBSS23]. However, the notion of multi-copy security, such as in our work, has not been studied.

3 Technical Overview

We first discuss the main technical lemma that underlies the construction of all the revocable primitives.

$k \rightarrow k + 1$ Unforgeability Of Subset States. Given k copies of a random subset state $|S\rangle$, where $|S| = \omega(\log(n))$, it is easy to produce k distinct elements in the set $S \subseteq \{0, 1\}^n$. This can be seen by just measuring all the copies of $|S\rangle$ in the computational basis. Since $|S|$ is sufficiently large, the probability that we get all the k elements to be distinct is high. Our main lemma states that it is computationally infeasible to produce $k + 1$ distinct elements in S given k copies of $|S\rangle$ and access to an oracle that tests membership in S . Roughly speaking, this can be seen as a multi-copy subset state analogue of the complexity-theoretic no-cloning result for *subspace states*, which was shown by Aaronson and Christiano [AC12].

To prove this, we appeal to the fact that random subset states are indistinguishable from Haar states, as shown in a couple of recent works [JMW23, GB23]. That is, for any k polynomial in n

$$\mathbb{E}_{\substack{S \subseteq \{0,1\}^n \\ |S| = \omega(\log(n))}} [|S\rangle\langle S|^{\otimes k}] \approx \mathbb{E}_{|\psi\rangle \leftarrow \mathcal{H}_n} [|\psi\rangle\langle\psi|^{\otimes k}]$$

Here, \approx refers to (negligible in n) closeness in trace distance and \mathcal{H}_n refers to the Haar distribution on n -qubit quantum states.

Furthermore, using the characterization of the symmetric subspace [Har13], the following holds:

$$\mathbb{E}_{\substack{S \subseteq \{0,1\}^n \\ |S| = \omega(\log(n))}} [|S\rangle\langle S|^{\otimes k}] \approx \mathbb{E}_{\substack{\mathbf{x} = (x_1, \dots, x_k) \leftarrow \{0,1\}^{nk} \\ x_1, \dots, x_k \text{ are distinct}}} [|\sigma_{\mathbf{x}}\rangle\langle\sigma_{\mathbf{x}}|],$$

Here, $|\sigma_{\mathbf{x}}\rangle = \frac{1}{k!} \sum_{\sigma \in S_k} |x_{\sigma(1)}, \dots, x_{\sigma(k)}\rangle$, where S_k is a symmetric group on k elements.

Why does this characterization help? note that if one were given k copies of $|\sigma_{\mathbf{x}}\rangle$, it is impossible to produce $k + 1$ distinct elements in $\{x_1, \dots, x_k\}$ because there are only k of them! This observation was first made in [BBSS23] although they considered subset states *with phase* for applications that are unrelated to our work.

However this observation is alone not sufficient to complete the proof for two main reasons. Firstly, the adversary receives access to the subset membership oracle. Secondly, the winning conditions are different in both the cases; in one case, given k copies of $|S\rangle$, $k + 1$ elements in S had to be produced while in the other case, given $|\sigma_{\mathbf{x}}\rangle$, $k + 1$ elements in \mathbf{x} need to be produced.

INTERMEDIATE LEMMA. To make progress on this lemma, we first study the following intermediate lemma. Consider the following two distributions:

- \mathcal{D}_0 : Sample t copies of a random subset state $|S\rangle$, i.e. S is sampled uniformly at random from all sets of size $\omega(\log(n))$.
- \mathcal{D}_1 : Sample $|\sigma_{\mathbf{x}}\rangle$, where $\mathbf{x} = (x_1, \dots, x_k) \leftarrow \{0, 1\}^{nk}$.

Similarly, we can define two oracles. The first oracle is \mathcal{O}_0 and the second oracle is \mathcal{O}_1 . We define \mathcal{O}_0 to test membership in S and \mathcal{O}_1 to test membership in $\{x_1, \dots, x_k\}$. The intermediate lemma states that a query-bounded adversary cannot distinguish whether it receives a sample from \mathcal{D}_0 and oracle access to \mathcal{O}_0 or it receives a sample from \mathcal{D}_1 and oracle access to \mathcal{O}_1 .

We prove this lemma in a series of steps. First, the adversary receives a sample from \mathcal{D}_0 and has oracle access to \mathcal{O}_0 .

- We then modify the oracle that the adversary has access to. Instead of having oracle access to S , it instead has oracle access to T , where $S \subseteq T$ and moreover, S has negligible size compared to T and T has negligible size compared to 2^n . The intuition is that the adversary cannot find an element in $T \setminus S$ and then, thanks to O2H lemma [Unr07], the indistinguishability of the membership oracles (for S and T) follows.
- We change the process of sampling S . Instead of first sampling S and then sampling T subject to $S \subseteq T$. We first sample T and then sample S to be a random subset of T of size $\omega(\log(n))$. This process is identical to the previous step.
- Then, we switch k copies of $|S\rangle$ with $|\sigma_{\mathbf{x}}\rangle$, for a random $\mathbf{x} = (x_1, \dots, x_k)$, where (x_1, \dots, x_k) are drawn from T . Here, we prove a stronger statement that *even if the adversary has a description of T* , this indistinguishability holds. In other words, the adversary receives a sample from \mathcal{D}_1 although it has access to an oracle that tests membership in T . This follows from some propositions in [JMW23].
- Finally, we modify the previous step as follows. The adversary receives a sample from \mathcal{D}_1 with oracle access to \mathcal{O}_1 . To prove that the adversary cannot notice the modification, we need to show that the adversary cannot distinguish whether it has oracle access to membership in T or membership in $\{x_1, \dots, x_k\}$, where $x_1, \dots, x_k \leftarrow T$. In fact, we need to show that the indistinguishability should hold *even if the adversary receives x_1, \dots, x_k in the clear*. This again follows from another invocation of the O2H lemma.

The above intermediate lemma immediately implies the $k \rightarrow k + 1$ unforgeability of subset states. Let us see why. Suppose there is an adversary \mathcal{A} that violates the $k \rightarrow k + 1$ unforgeability property with non-negligible probability p . Then we can come up with a reduction \mathcal{R} that contradicts the above intermediate lemma as follows. \mathcal{R} receives as input a state $|\phi\rangle$ and oracle access to \mathcal{O} . First, it runs \mathcal{A} on input $|\phi\rangle$ while giving it oracle access to \mathcal{O} . \mathcal{A} then outputs $k + 1$ elements. \mathcal{R} inputs each of the $k + 1$ elements to the oracle \mathcal{O} . If the output of \mathcal{O} on each of the $k + 1$ distinct elements is 1 (i.e. the membership test passed) then output 1, else output 0. There are two cases to consider here:

- In the first case, $|\phi\rangle$ is k copies of a random subset state and moreover, \mathcal{O} tests membership in S . In this case, the probability that \mathcal{R} outputs 1 is exactly p .
- In the second case, $|\phi\rangle$ is a state of the form $|\sigma_{\mathbf{x}}\rangle$ and moreover, \mathcal{O} tests membership in $\{x_1, \dots, x_k\}$, where $\mathbf{x} = (x_1, \dots, x_k)$. In this case, the probability that \mathcal{R} outputs 1 is 0. This is because, $\{x_1, \dots, x_k\}$ only has k distinct elements.

Thus, the success probability of \mathcal{R} in violating the above intermediate lemma is p , which is non-negligible. This in turn is a contradiction.

Revocable Encryption. Armed with the $k \rightarrow k + 1$ unforgeability property, we will now tackle the first application of revocable encryption. Let us first formally state the syntax of such a scheme:

- $\text{KeyGen}(1^\lambda)$: on input the security parameter 1^λ , output a secret key sk .

- $\text{Enc}(\text{sk}, m)$: on input the secret key sk and a message m , output a (pure) ciphertext state $|\psi\rangle$ and a (private) verification key vk .
- $\text{Dec}(\text{sk}, \rho)$: on input the secret key sk and a quantum state ρ , output a message m' .
- $\text{Revoke}(\text{sk}, \text{vk}, \sigma)$: on input the secret key sk , a verification key vk and a state σ , output \top or \perp .

For multi-copy revocable security, we will consider the following experiment between a QPT adversary and a challenger.

1. \mathcal{A} submits two messages m_0, m_1 and a polynomial $k = k(\lambda)$ to the challenger.
2. The challenger samples a key $\text{sk} \leftarrow \text{KeyGen}(1^\lambda)$ and produces $|\psi_b\rangle \leftarrow \text{Enc}(\text{sk}, m_b)$. Afterwards, the challenger sends the quantum state $|\psi_b\rangle^{\otimes k}$ to \mathcal{A} .
3. \mathcal{A} returns a quantum state ρ .
4. The challenger performs the measurement $\left\{ |\psi_b\rangle\langle\psi_b|^{\otimes k}, \mathbb{I} - |\psi_b\rangle\langle\psi_b|^{\otimes k} \right\}$ on the returned state ρ . If the measurement succeeds, the game continues; otherwise, the challenger aborts.
5. The challenger sends the secret key sk to \mathcal{A} .
6. \mathcal{A} outputs a bit b' .

We say that the revocable encryption scheme has multi-copy revocable security, if no QPT adversary can guess b' with probability that's more than negligibly better than $\frac{1}{2}$.

Now, we will describe how we get revocable encryption in the *random permutation oracle* model. But before we explain our construction, we will begin by highlighting some challenges in achieving revocable encryption from subset states in the quantum random oracle model.

WARM-UP: QUANTUM RANDOM ORACLE MODEL. We will design revocable encryption schemes in which part of the quantum ciphertext will consist of a random subset state. Since generating a random subset state involves first sampling a random exponential-sized set, such states cannot be efficiently generated. However, let us first see if random oracles can be exploited to generate random subset states more efficiently.

Let $\mathcal{H}_{2n+m} = \{H : \{0, 1\}^{2n+m} \rightarrow \{0, 1\}^m\}$ be a hash function family, where H is modeled as the random oracle. To encrypt a message $\mu \in \{0, 1\}^m$, first sample a secret string salt $\leftarrow \{0, 1\}^{n+m}$ and coherently compute the following quantum state:

$$|\psi\rangle = 2^{-n/2} \sum_{x \in \{0, 1\}^n} |x\rangle |H(\text{salt}, x)\rangle$$

Now, we can measure the second register and obtain a random string y , in which case the state collapses to

$$|\psi_y\rangle \propto \sum_{x \in \{0, 1\}^n : H(\text{salt}, x) = y} |x\rangle$$

The ciphertext now consists of $|\psi_y\rangle$ together with $y \oplus \mu$. While this approach clearly solves the problem of efficiently generating the state, the construction has a major drawback: it only allows us to generate *one* copy of the state and, in particular, *identical* copies of the subset state $|\psi_y\rangle$ are not efficiently preparable! Therefore, an entirely new approach is necessary in order to make multi-copy revocable encryption possible.

REVOCABLE ENCRYPTION IN THE RANDOM PERMUTATION MODEL. In order to resolve the aforementioned issue, we instantiate our revocable encryption scheme using permutations instead; specifically, we make use of quantum-secure pseudorandom permutations (QPRPs), and we prove security in the random permutation

model: this means that, rather than revealing the QPRP key in the final part of the revocable security experiment, we instead allow the adversary to query an oracle for a (now random) permutation instead.²

Let $\Phi = \{\Phi_\lambda\}_{\lambda \in \mathbb{N}}$ be an ensemble of QPRPs $\Phi_\lambda = \{\varphi_\kappa : \{0, 1\}^{n+m} \rightarrow \{0, 1\}^{n+m}\}_{\kappa \in \mathcal{K}_\lambda}$, for some set \mathcal{K}_λ . Consider the scheme $\Sigma^\Phi = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$ which consists of the following QPT algorithms:

- $\text{KeyGen}(1^\lambda)$: sample a uniformly random key $\kappa \in \mathcal{K}_\lambda$ and let $\text{sk} = \kappa$.
- $\text{Enc}(\text{sk}, \mu)$: on input the secret key $\text{sk} = \kappa$ and message $\mu \in \{0, 1\}^m$, sample $y \sim \{0, 1\}^n$ and prepare the subset state given by

$$|S_y\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0, 1\}^n} |\varphi_\kappa(x||y)\rangle.$$

Output the ciphertext state $(|S_y\rangle^{\otimes k}, y \oplus \mu)$ and (private) verification key $\text{vk} = y$.

- $\text{Dec}(\text{sk}, \text{ct})$: on input the decryption key κ and ciphertext state $(|S_y\rangle, z) \leftarrow \text{ct}$, coherently apply the in-place permutation ϕ_κ^{-1} , measure to obtain a string $(x'||y')$, and output $y' \oplus z$.
- $\text{Revoke}(\text{sk}, \text{vk}, \rho)$: on input sk , a state ρ and verification key vk , it parses $\kappa \leftarrow \text{sk}$, $y \leftarrow \text{vk}$ and applies the measurement $\{|S_y\rangle\langle S_y|, \mathbb{I} - |S_y\rangle\langle S_y|\}$ to ρ ; it outputs \top if it succeeds, and \perp otherwise.

Observe that the state $|S_y\rangle$ can be efficiently generated using a unitary (that is part of the quantum ciphertext). Formally, we perform the following steps:

- Prepare a uniform superposition $\frac{1}{\sqrt{2^n}} \sum_{x \in \{0, 1\}^n} |x\rangle$.
- Then, append $|y\rangle$ to the above state, where y is sampled uniformly at random.
- Finally, compute a unitary that applies the permutation φ_κ in-place on each term in the state.

This in turn means that many copies of $|S_y\rangle$ can be prepared efficiently. Moreover, using this unitary, the algorithm Revoke can also be efficiently implemented.

PROOF OUTLINE: Now, we will give a proof outline for proving multi-copy secure revocable encryption. We will base the security of our revocable encryption scheme on the $k \rightarrow k + 1$ unforgeability of random subset states.

Let us assume for contradiction that there exists an adversary \mathcal{A} such that the revocable encryption scheme presented above is not secure. The main idea behind converting \mathcal{A} into an adversary \mathcal{A}' against the $k \rightarrow k + 1$ unforgeability of $|S\rangle$ is to construct hybrids that can be distinguished with non negligible probability, and then use the one-way to hiding lemma [AHU18] in order to construct an *extractor* \mathcal{E} which queries the membership oracle for S on a $(k + 1)^{\text{st}}$ distinct element *after* \mathcal{A} passes the revocation phase. Then, \mathcal{A}' can extract $k + 1$ elements given only k copies of $|S\rangle$ as follows:

1. \mathcal{A}' runs \mathcal{A} on input $|S\rangle^{\otimes k}$ and measures the registers returned by \mathcal{A} that are supposed to contain k copies of the subset state. With high probability (provided that the registers pass revocation), these measurements will result in precisely k distinct elements in S .
2. Then, \mathcal{A}' can run the extractor \mathcal{E} to extract the $(k + 1)^{\text{st}}$ distinct element in S .

Note that extra care is needed to argue that the two steps above succeed simultaneously with sufficiently high probability whenever \mathcal{A} is a successful adversary for the revocable security experiment. For simplicity, we omit these details in this overview and refer the reader to [Section 7.1](#).

²This is a weaker statement than the regular notion of revocable security; however, we believe it nevertheless meaningfully captures the fact that the adversary loses access to the ciphertext.

Now, we will briefly describe the sequence of hybrids which helps us construct the extractor. Before we state the hybrids, it will be useful to briefly motivate and define the one way to hiding lemma.

ONE-WAY TO HIDING. Consider the setting in which an adversary has access to an oracle H that outputs a value $H(x)$, and suppose we want to show that $H(x)$ is indistinguishable from a uniformly random string y . In the classical setting, we can do the following: the probability that an adversary distinguishes $H(x)$ and y is upper bounded by the probability that the adversary queries the oracle on x . That is, we have reduced an indistinguishability to the hardness of guessing x . The one-way to hiding lemma [AHU18] helps achieve the same reduction even when the adversary has *quantum* access to the oracle H . A natural approach to be able to use the one-way to hiding lemma would be the following:

- In the post-revocation phase, introduce a hybrid with a reprogrammed permutation oracle ϕ'_κ on elements in the set S , and prove that, if there exists an adversary \mathcal{A} against the revocable encryption scheme, then it must be able to detect the reprogrammed oracle.
- Use the one-way to hiding lemma to argue that this implies that \mathcal{A} queried the oracle on an element in the set S after the revocation phase.

The barrier to the above outlined approach is that currently, it is unclear how to apply the one-way to hiding lemma in the context of random permutation oracles. Thus, in order to use it to construct the extractor \mathcal{E} , the oracle must be modified.

However, we can resolve this issue by simulating a random permutation oracle with a *random function oracle*. In fact, Zhandry [Zha13] shows that random functions are indistinguishable from random permutations, provided that the adversary cannot query the permutation oracle in both the forward and inverse direction (random functions are trivially distinguishable from random permutations in this case). However, recall that, in the post revocation phase, the adversary gets access to κ itself, and therefore has the entire description of the permutation oracle, which is enough to simulate ϕ_κ as well as ϕ_κ^{-1} . Therefore, in order to be able to use the indistinguishability between random functions and random permutations, we modify our scheme in the following way: instead of giving the post-revocation adversary the key κ , our the post revocation adversary gets *oracle access* to ϕ_κ^{-1} instead. In fact, we carry out the proof in the random permutation model entirely, whereby we replace φ_κ with a perfectly random permutation φ . This is similar in spirit to the oracle-based security proof of unclonable encryption security by Broadbent and Lord [BL19].

PUTTING IT ALL TOGETHER. We are now ready to describe the sequence of hybrids we consider in order to prove the security of our construction in an oracular security model. The experiment is as follows:

RevokeExpt $^{\mathcal{A}}(1^\lambda, b)$:

1. \mathcal{A} submits two m -bit messages (μ_0, μ_1) and a polynomial $k = k(\lambda)$ to the challenger.
2. The challenger samples $y \sim \{0, 1\}^m$ and produces a quantum state

$$|S_y\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0, 1\}^n} |\varphi(x||y)\rangle.$$

The challenger then sends $|S_y\rangle^{\otimes k}$ and $y \oplus \mu_b$ to \mathcal{A} .

3. \mathcal{A} prepares a bipartite state on registers R and AUX, and sends R to the challenger and AUX to \mathcal{A} .
4. The challenger performs the projective measurement $\left\{ |S_y\rangle\langle S_y|^{\otimes k}, \mathbb{I} - |S_y\rangle\langle S_y|^{\otimes k} \right\}$ on R. If the measurement succeeds, the challenger outputs \perp . Otherwise, the challenger continues.
5. The challenger grants \mathcal{A} quantum oracle access to φ^{-1} .
6. \mathcal{A} outputs a bit b' .

We show that the following via hybrid experiments:

1. $\text{RevokeExpt}^{\mathcal{A}}(1^\lambda, b)$ is indistinguishable from the experiment H_2 which is the same as $\text{RevokeExpt}^{\mathcal{A}}(1^\lambda, b)$ except,

- The challenger samples a random subset $S \subseteq \{0, 1\}^{n+m}$ of size $|S| = 2^n$.
- The challenger samples a random $y \sim \{0, 1\}^m$.

after passing revocation, \mathcal{A} receives access to the following oracle: $f : \{0, 1\}^{n+m} \rightarrow \{0, 1\}^{n+m}$ subject to the constraint that $f(s) = *||y$, for all $s \in S$.

2. H_2 is indistinguishable from the experiment H_3 which is the same as H_2 except, after passing revocation, \mathcal{A} receives access to the following oracle: $f : \{0, 1\}^{n+m} \rightarrow \{0, 1\}^{n+m}$ subject to the constraint that $f(s) = *||u$, for all $s \in S$. Here, u is a randomly sampled string.

With some simple hybrid arguments, we show that if $\text{RevokeExpt}^{\mathcal{A}}(1^\lambda, b)$ is not secure, then \mathcal{A} can distinguish H_2 and H_3 . However, since H_2 and H_3 are the same except the the function f differs *only* on inputs in the subset S , we can now use one-way to hiding to come up with an extractor \mathcal{E} extracts an element in the subset S using the queries \mathcal{A} makes to f *after revocation succeeds*. To complete the proof, we construct an adversary that breaks the $k \rightarrow k + 1$ unforgeability of subset states, in violation of our main lemma. The high level idea is that the k copies of the subset state which are meant to be returned easily allow the adversary to obtain k distinct subset elements; whereas the additional $(k + 1)$ -st element can be extracted from the adversary’s side information via the aforementioned extractor.³

Revocable Programs. In addition to revocable encryption, we also achieve multi-copy secure revocable programs in the *classical* oracle model from the $k \rightarrow k + 1$ unforgeability of subset states. The proof of security in this setting is similar, and the main idea again is to construct an extractor extracting $k + 1$ elements given only k copies of S , though extra care is required when defining the hybrids in order to prove security. For the formal proof, we refer the reader to [Section 8](#).

Revocable Point Functions. As a final application of $k \rightarrow k + 1$ unforgeability, we give a revocable multi-copy secure scheme for a particular program; namely, multi-bit point functions $P_{y,m}$ of the form

$$P_{y,m}(x) = \begin{cases} m & \text{if } x = y, \\ 0^\lambda & \text{if } x \neq y, \end{cases}$$

in the random permutation oracle model, as well as the (standard) quantum random oracle model. The main idea is to show that revocable encryption schemes can be generically converted to schemes which satisfy a property called *wrong-key detection* (WKD); broadly speaking, this is a mechanism that allows a user with access to ciphertext to check whether a particular of choice decryption key is incorrect, in which case decryption algorithm simply outputs \perp . Inspired by [\[CMP24\]](#), we give a transformation that converts any revocable encryption scheme to one which satisfies the WKD property. Intuitively, this achieves our goal: to evaluate $P_{y,m}$ at point x , attempt to decrypt using x ; if decryption succeeds output the decrypted message, if decryption fails, output 0^λ . See [Section 10](#) for detailed proof.

4 Preliminaries

Let $\lambda \in \mathbb{N}$ denote the security parameter throughout this work. We assume that the reader is familiar with the fundamental cryptographic concepts.

³Note that there are a few subtleties that we gloss over in this overview, such as the fact that we need to condition on the event that revocation succeeds. We refer the reader to [Section 6](#) for the full proof.

For $N \in \mathbb{N}$, we use $[N] = \{1, 2, \dots, N\}$ to denote the set of integers up to N . The symmetric group on $[N]$ is denoted by S_N . In slight abuse of notation, we oftentimes identify elements $x \in [N]$ with bit strings $x \in \{0, 1\}^n$ via their binary representation whenever $N = 2^n$ and $n \in \mathbb{N}$. Similarly, we identify permutations $\pi \in S_N$ with permutations $\pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ over bit strings of length n . For a bit string $x \in \{0, 1\}^n$, we frequently use the notation $(x||*)$, where $*$ serves as a placeholder to denote the set $\{(x||y) : y \in \{0, 1\}^m\}$, where $m \in \mathbb{N}$ is another integer which is typically clear in context.

We write $\text{negl}(\cdot)$ to denote any *negligible* function, which is a function f such that, for every constant $c \in \mathbb{N}$, there exists an integer N such that for all $n > N$, $f(n) < n^{-c}$.

Quantum Computing For a comprehensive background on quantum computation, we refer to [NC11]. We denote a finite-dimensional complex Hilbert space by \mathcal{H} , and we use subscripts to distinguish between different systems (or registers). For example, we let \mathcal{H}_A be the Hilbert space corresponding to a system A . The tensor product of two Hilbert spaces \mathcal{H}_A and \mathcal{H}_B is another Hilbert space denoted by $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$. The Euclidean norm of a vector $|\psi\rangle \in \mathcal{H}$ over the finite-dimensional complex Hilbert space \mathcal{H} is denoted as $\|\psi\| = \sqrt{\langle \psi | \psi \rangle}$. Let $\mathcal{L}(\mathcal{H})$ denote the set of linear operators over \mathcal{H} . A quantum system over the 2-dimensional Hilbert space $\mathcal{H} = \mathbb{C}^2$ is called a *qubit*. For $n \in \mathbb{N}$, we refer to quantum registers over the Hilbert space $\mathcal{H} = (\mathbb{C}^2)^{\otimes n}$ as n -qubit states. We use the word *quantum state* to refer to both pure states (unit vectors $|\psi\rangle \in \mathcal{H}$) and density matrices $\rho \in \mathcal{D}(\mathcal{H})$, where we use the notation $\mathcal{D}(\mathcal{H})$ to refer to the space of positive semidefinite matrices of unit trace acting on \mathcal{H} . The *trace distance* of two density matrices $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ is given by

$$\text{TD}(\rho, \sigma) = \frac{1}{2} \text{Tr} \left[\sqrt{(\rho - \sigma)^\dagger (\rho - \sigma)} \right].$$

A quantum channel $\Phi : \mathcal{L}(\mathcal{H}_A) \rightarrow \mathcal{L}(\mathcal{H}_B)$ is a linear map between linear operators over the Hilbert spaces \mathcal{H}_A and \mathcal{H}_B . Oftentimes, we use the compact notation $\Phi_{A \rightarrow B}$ to denote a quantum channel between $\mathcal{L}(\mathcal{H}_A)$ and $\mathcal{L}(\mathcal{H}_B)$. We say that a channel Φ is *completely positive* if, for a reference system R of arbitrary size, the induced map $I_R \otimes \Phi$ is positive, and we call it *trace-preserving* if $\text{Tr}[\Phi(X)] = \text{Tr}[X]$, for all $X \in \mathcal{L}(\mathcal{H})$. A quantum channel that is both completely positive and trace-preserving is called a quantum CPTP channel. A *unitary* $U : L(\mathcal{H}_A) \rightarrow L(\mathcal{H}_A)$ is a special case of a quantum channel that satisfies $U^\dagger U = U U^\dagger = I_A$. An *isometry* is a linear map $V : L(\mathcal{H}_A) \rightarrow L(\mathcal{H}_B)$ with $\dim(\mathcal{H}_B) \geq \dim(\mathcal{H}_A)$ and $V^\dagger V = I_A$. A *projector* Π is a Hermitian operator such that $\Pi^2 = \Pi$, and a *projective measurement* is a collection of projectors $\{\Pi_i\}_i$ such that $\sum_i \Pi_i = I$. A positive-operator valued measure (POVM) is a set of Hermitian positive semidefinite operators $\{M_i\}$ acting on a Hilbert space \mathcal{H} such that $\sum_i M_i = I$.

Quantum algorithms. A polynomial-time *uniform* quantum algorithm (or QPT algorithm) is a polynomial-time family of quantum circuits given by $\mathcal{C} = \{C_\lambda\}_{\lambda \in \mathbb{N}}$, where each circuit $C \in \mathcal{C}$ is described by a sequence of unitary gates and measurements; moreover, for each $\lambda \in \mathbb{N}$, there exists a deterministic polynomial-time Turing machine that, on input 1^λ , outputs a circuit description of C_λ . Similarly, we also define (classical) probabilistic polynomial-time (PPT) algorithms. A quantum algorithm may, in general, receive (mixed) quantum states as inputs and produce (mixed) quantum states as outputs. We frequently restrict QPT algorithms implicitly; for example, if we write $\Pr[\mathcal{A}(1^\lambda) = 1]$ for a QPT algorithm \mathcal{A} , it is implicit that \mathcal{A} is a QPT algorithm that outputs a single classical bit. We say that a quantum algorithm \mathcal{A} has oracle access to a classical function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$, denoted by \mathcal{A}^f , if \mathcal{A} is allowed to use a unitary gate \mathcal{O}_f at unit cost in time. The unitary \mathcal{O}_f acts as follows on the computational basis states of a Hilbert space $\mathcal{H}_X \otimes \mathcal{H}_Y$ of $n + m$ qubits:

$$\mathcal{O}_f : |x\rangle_X \otimes |y\rangle_Y \longrightarrow |x\rangle_X \otimes |y \oplus f(x)\rangle_Y,$$

where the operation \oplus denotes bit-wise addition modulo 2. Oracles with quantum query-access have been studied extensively, for example in the context of quantum complexity theory [BBV97], as well as in cryptography [BDF⁺11, AHU18, AJOP20].

One-Way-to-Hiding Lemma. We use the following lemma which combines [AHU18, Theorem 3] and [AHU18, Lemma 8], where the latter is rooted in Vazirani’s *Swapping Lemma* [Vaz98].

Lemma 4.1 (One-Way-to-Hiding Lemma, [AHU18]). *Let \mathcal{X}, \mathcal{Y} be arbitrary sets and let $\mathcal{S} \subseteq \mathcal{X}$ be a (possibly random) subset. Let $G, H : \mathcal{X} \rightarrow \mathcal{Y}$ be arbitrary (possibly random) functions such that $H(x) = G(x)$, for all $x \notin \mathcal{S}$. Let z be a classical bit string or a (possibly mixed) quantum state (Note that G, H, \mathcal{S}, z may have arbitrary joint distribution). Let \mathcal{A} be an oracle-aided quantum algorithm that makes at most q quantum queries. Let \mathcal{B} be an algorithm that on input z chooses a random query index $i \leftarrow [q]$, runs $\mathcal{A}^H(z)$, measures \mathcal{A} ’s i -th query and outputs the measurement outcome. Then, we have*

$$|\Pr[\mathcal{A}^G(z) = 1] - \Pr[\mathcal{A}^H(z) = 1]| \leq 2q\sqrt{\Pr[\mathcal{B}^H(z) \in \mathcal{S}]}.$$

Moreover, for any fixed choice of G, H, \mathcal{S} and z (when z is a classical string or a pure state), we get

$$\|\psi_q^H\rangle - |\psi_q^G\rangle\| \leq 2q\sqrt{\frac{1}{q} \sum_{i=0}^{q-1} \|\Pi_{\mathcal{S}} |\psi_i^H\rangle\|^2},$$

where $|\psi_i^H\rangle$ denotes the intermediate state of \mathcal{A} just before the $(i+1)$ -st query, where the initial state at $i = 0$ corresponds to z , and $\Pi_{\mathcal{S}}$ is a projector onto \mathcal{S} .

Pseudorandom Permutations. A quantum-secure pseudorandom permutation is a bijective function family which can be constructed from quantum-secure one-way functions [Zha16].

Definition 4.2 (QPRP). *Let $\lambda \in \mathbb{N}$ denote the security parameter. Let $P : \{0, 1\}^\lambda \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a function, where $n(\lambda) = \text{poly}(\lambda)$ is an integer, such that each function $P_k(x) = P(k, x)$ in the corresponding family $\{P_k\}_{k \in \{0, 1\}^\lambda}$ is bijective. We say P is a (strong) quantum-secure pseudorandom permutation (or QPRP) if, for every QPT \mathcal{A} with access to both the function and its inverse, it holds that*

$$\left| \Pr_{k \sim \{0, 1\}^\lambda} [\mathcal{A}^{P_k, P_k^{-1}}(1^\lambda) = 1] - \Pr_{\varphi \sim \mathcal{P}_n} [\mathcal{A}^{\varphi, \varphi^{-1}}(1^\lambda) = 1] \right| \leq \text{negl}(\lambda),$$

where \mathcal{P}_n denotes the set of permutations over n -bit strings.

We use the following result due to Zhandry which applies to random permutations.

Theorem 4.3 ([Zha13], Theorem 3.1). *Any q -query quantum algorithm can distinguish random functions from random permutations over $\{0, 1\}^n$ with advantage at most $O(q^3/2^n)$.*

Subset States. We consider the following notations.

- We denote the set of distinct k -tuples over a set S by $\text{dist}(S, k)$.
- Suppose S is a set. We denote $|S\rangle = \frac{1}{\sqrt{|S|}} \sum_{x \in S} |x\rangle$.
- Suppose $X = \{x_1, \dots, x_t\} \subseteq \{0, 1\}^n$. We denote $|\sigma_X\rangle = \frac{1}{t!} \sum_{\sigma \in S_t} |x_{\sigma(1)}, \dots, x_{\sigma(t)}\rangle$, where S_t denotes the symmetric group on $[t]$.

We use the following lemma which follows from Propositions 3.3 and 3.4 in [JMW23].

Lemma 4.4 ([JMW23]). *Let $n, k \in \mathbb{N}$. Let $T \subseteq \{0, 1\}^n$ be a subset of size $|T| = t$. Then, it holds that*

$$\text{TD} \left(\mathbb{E}_{\substack{S \subseteq T \\ |S|=s}} [|S\rangle\langle S|^{\otimes k}], \mathbb{E}_{\substack{X \subseteq T \\ |X|=k}} [|\sigma_X\rangle\langle \sigma_X|] \right) \leq O \left(\frac{k}{\sqrt{s}} + \frac{sk}{t} \right).$$

5 $k \mapsto k + 1$ Unforgeability of Subset States

We now prove the following theorem. Roughly speaking, our theorem says that any quantum algorithm which receives k copies of a random subset state $|S\rangle$ (and a membership oracle for S) cannot find $k + 1$ distinct elements in S with high probability unless it makes a large number of queries.

Theorem 5.1 ($k \mapsto k + 1$ Unforgeability of Subset States). *Let $n \in \mathbb{N}$ and $k \in \mathbb{N}$. Then, for any q -query quantum oracle algorithm \mathcal{A} , and any $1 \leq k < s < t \leq 2^n$, it holds that*

$$\begin{aligned} & \Pr_{\substack{S \subseteq \{0,1\}^n \\ |S|=s}} \left[(x_1, \dots, x_{k+1}) \in \text{dist}(S, k+1) : (x_1, \dots, x_{k+1}) \leftarrow \mathcal{A}^{\mathcal{O}_S}(|S\rangle^{\otimes k}) \right] \\ & \leq O \left(q \cdot \sqrt{\frac{t-s}{2^n}} + q \cdot \sqrt{\frac{t-k}{2^n}} + \frac{k}{\sqrt{s}} + \frac{sk}{t} \right) + \text{negl}(n). \end{aligned}$$

In particular, we can let $k = \text{poly}(n)$, $q = \text{poly}(n)$ and $s(n) = n^{\omega(1)}$ be superpolynomial. Then, for any $t(n) = n^{\omega(1)}$ with $s(n)/t(n) = 1/n^{\omega(1)}$ and $t(n)/2^n = 1/n^{\omega(1)}$, the probability is at most $\text{negl}(n)$.

Proof. Using Lemma 5.2, we can prove Theorem 5.1 as follows: Let's assume for contradiction that there exists a q query quantum oracle algorithm \mathcal{A} , and a $1 \leq k < s < t \leq 2^n$, such that

$$\begin{aligned} & \Pr_{\substack{S \subseteq \{0,1\}^n \\ |S|=s}} \left[(x_1, \dots, x_{k+1}) \in \text{dist}(S, k+1) : (x_1, \dots, x_{k+1}) \leftarrow \mathcal{A}^{\mathcal{O}_S}(|S\rangle^{\otimes k}) \right] \\ & = O \left(q \cdot \sqrt{\frac{t-s}{2^n}} + q \cdot \sqrt{\frac{t-k}{2^n}} + \frac{k}{\sqrt{s}} + \frac{sk}{t} \right) + \delta(n). \end{aligned}$$

where $\delta(\cdot)$ is some non negligible function. Then, from Lemma 5.2, this implies that,

$$\begin{aligned} & \Pr_{X \subseteq \{0,1\}^n, |X|=k} \left[(x_1, \dots, x_{k+1}) \in \text{dist}(X, k+1) : (x_1, \dots, x_{k+1}) \leftarrow \mathcal{A}^{\mathcal{O}_X}(|\sigma_X\rangle) \right] \\ & \geq \delta(n) \end{aligned}$$

However, the probability that \mathcal{A} can succeed in this experiment is 0. This is because, X only contains k elements, and therefore, \mathcal{A} can never produce $k + 1$ distinct elements from X . This must imply that $\delta(n)$ is negligible. □

Technical Lemma. We need to show the following lemma which made use of in Theorem 5.1.

Lemma 5.2. *Let $n, k, t \in \mathbb{N}$ be integers such that $1 \leq k < s < t \leq 2^n$. Then, for any q -query quantum oracle algorithm \mathcal{D} which outputs a single bit, it holds that*

$$\begin{aligned} & \left| \Pr_{\substack{S \subseteq \{0,1\}^n \\ |S|=s}} \left[\mathcal{D}^{\mathcal{O}_S}(|S\rangle^{\otimes k}) = 1 \right] - \Pr_{\substack{X \subseteq \{0,1\}^n \\ |X|=k}} \left[\mathcal{D}^{\mathcal{O}_X}(|\sigma_X\rangle) = 1 \right] \right| \\ & \leq O \left(q \cdot \sqrt{\frac{t-s}{2^n}} + q \cdot \sqrt{\frac{t-k}{2^n}} + \frac{k}{\sqrt{s}} + \frac{sk}{t} \right). \end{aligned}$$

Proof. Consider the following hybrid distributions.

H₁: Output $\mathcal{D}^{\mathcal{O}_S}(|S\rangle^{\otimes k})$, where $S \subseteq \{0, 1\}^n$ is a random subset of size $|S| = s$.

H₂: Output $\mathcal{D}^{\mathcal{O}_S}(|S\rangle^{\otimes k})$, where the subset S is sampled as follows: first, sample a random subset $T \subseteq \{0, 1\}^n$ of size $|T| = t$, and then let $S \subseteq T$ be a random subset of size $|S| = s$.

Let $p(\mathbf{H}_i)$ be the probability that \mathbf{H}_i outputs 1, for some i . We now show the following.

Claim 5.3. $p(\mathbf{H}_2) = p(\mathbf{H}_1)$.

Proof. The distribution of sampling $S \subseteq \{0, 1\}^n$ of size $|S| = s$ is identical to the distribution of first sampling a superset $T \subseteq \{0, 1\}^n$ of size $|T| = t$, and letting $S \subseteq T$ be a random subset of size $|S| = s$. \square

H₃: Output $\mathcal{D}^{\mathcal{O}_T}(|S\rangle^{\otimes k})$, where the subset S is sampled as follows: first, sample a random subset $T \subseteq \{0, 1\}^n$ of size $|T| = t$, and then let $S \subseteq T$ be a random subset of size $|S| = s$.

Claim 5.4.

$$|p(\mathbf{H}_3) - p(\mathbf{H}_2)| \leq O\left(q \cdot \sqrt{\frac{t-s}{2^n}}\right).$$

Proof. We can model the quantum oracle algorithm $\mathcal{D}^{\mathcal{O}_S}$ on input $|S\rangle^{\otimes t}$ as a sequence of oracle queries and unitary computations followed by a measurement. Thus, the final output state just before the measurement can be written as

$$|\Psi_q^S\rangle = U_q \mathcal{O}_S U_{q-1} \dots U_1 \mathcal{O}_S U_0 |\psi_0\rangle |S\rangle^{\otimes k},$$

where U_0, U_1, \dots, U_q are unitaries (possibly acting on additional workspace registers, which we omit above), and where $|\psi_0\rangle$ is some fixed initial state which is independent of S .

In the next step of the proof, we will use the ‘‘subset flooding’’ technique to drown S in a random superset. Let $T \subseteq \{0, 1\}^n$ be a random superset of S of size $t > s$. We now consider the state

$$|\Psi_q^T\rangle = U_q \mathcal{O}_T U_{q-1} \dots U_1 \mathcal{O}_T U_0 |\psi_0\rangle |S\rangle^{\otimes k}.$$

We now claim that the states $|\Psi_q^S\rangle$ and $|\Psi_q^T\rangle$ are sufficiently close. From the definition of \mathcal{O}_T and \mathcal{O}_S , we have that $\mathcal{O}_T(x) \neq \mathcal{O}_S(x)$ iff $x \in T \setminus S \subset \{0, 1\}^n$. By the O2H Lemma (Lemma 4.1),

$$\begin{aligned} \mathbb{E}_{\substack{T \subseteq \{0,1\}^n, |T|=t \\ S \subseteq T, |S|=s}} \left\| |\Psi_q^S\rangle - |\Psi_q^T\rangle \right\| &\leq 2q \mathbb{E}_{\substack{T \subseteq \{0,1\}^n, |T|=t \\ S \subseteq T, |S|=s}} \sqrt{\frac{1}{q} \sum_{i=0}^{q-1} \left\| \Pi_{T \setminus S} |\Psi_i^S\rangle \right\|^2} \\ &\leq 2q \sqrt{\frac{1}{q} \sum_{i=0}^{q-1} \mathbb{E}_{\substack{T \subseteq \{0,1\}^n, |T|=t \\ S \subseteq T, |S|=s}} \left\| \Pi_{T \setminus S} |\Psi_i^S\rangle \right\|^2} \quad (\text{Jensen's inequality}) \\ &= O\left(q \cdot \sqrt{\frac{t-s}{2^n}}\right). \end{aligned}$$

Therefore, the probability (over the choice of S and T) that $\mathcal{D}^{\mathcal{O}_S}(|S\rangle^{\otimes k})$ succeeds is at most the probability that $\mathcal{D}^{\mathcal{O}_T}(|S\rangle^{\otimes k})$ succeeds—up to an additive loss of $O(q \cdot \sqrt{\frac{t-s}{2^n}})$. \square

H₄: Output $\mathcal{D}^{\mathcal{O}_T}(|\sigma_X\rangle)$, where the subset X is sampled as follows: first, sample a random subset $T \subseteq \{0, 1\}^n$ of size $|T| = t$, and then let $X \subseteq T$ be a random subset of size $|X| = k$.

Claim 5.5.

$$|p(\mathbf{H}_4) - p(\mathbf{H}_3)| \leq O\left(\frac{k}{\sqrt{s}} + \frac{sk}{t}\right).$$

Proof. Here, we make use of [Lemma 4.4](#) which says that, for any superset $T \subseteq \{0, 1\}^n$ of size $|T| = t$,

$$\text{TD} \left(\mathbb{E}_{\substack{S \subseteq T \\ |S|=s}} [|S\rangle\langle S|^{\otimes k}], \mathbb{E}_{\substack{X \subseteq T \\ |X|=k}} [|\sigma_X\rangle\langle \sigma_X|] \right) \leq O \left(\frac{k}{\sqrt{s}} + \frac{sk}{t} \right).$$

□

H₅: Output $\mathcal{D}^{\mathcal{O}_X}(|\sigma_X\rangle)$, where the subset X is sampled as follows: first, sample a random subset $T \subseteq \{0, 1\}^n$ of size $|T| = t$, and then let $X \subseteq T$ be a random subset of size $|X| = k$.

Claim 5.6.

$$|p(\mathbf{H}_5) - p(\mathbf{H}_4)| \leq O \left(q \cdot \sqrt{\frac{t-k}{2^n}} \right).$$

Proof. Suppose that $X \subseteq T$ is a random subset of size $|X| = k$, and let $|\sigma_X\rangle = \frac{1}{k!} \sum_{\sigma \in S_k} |x_{\sigma(1)}, \dots, x_{\sigma(k)}\rangle$, where S_k denotes the symmetric group on $[k]$. Consider the state

$$|\Phi_q^T\rangle = U_q \mathcal{O}_T U_{q-1} \dots U_1 \mathcal{O}_T U_0 |\psi_0\rangle |\sigma_X\rangle.$$

prepared by $\mathcal{D}^{\mathcal{O}_T}(|\sigma_X\rangle)$ just before the measurement. Similarly, we let

$$|\Phi_q^X\rangle = U_q \mathcal{O}_X U_{q-1} \dots U_1 \mathcal{O}_X U_0 |\psi_0\rangle |\sigma_X\rangle.$$

be the state prepared by $\mathcal{A}^{\mathcal{O}_X}(|\sigma_X\rangle)$. Using [Lemma 4.1](#) as before, we get that

$$\begin{aligned} \mathbb{E}_{\substack{T \subseteq \{0,1\}^n, |T|=t \\ X \subseteq T, |X|=k}} \left\| |\Phi_q^T\rangle - |\Phi_q^X\rangle \right\| &\leq 2q \mathbb{E}_{\substack{T \subseteq \{0,1\}^n, |T|=t \\ X \subseteq T, |X|=k}} \sqrt{\frac{1}{q} \sum_{i=0}^{q-1} \left\| \Pi_{T \setminus X} |\Phi_i^X\rangle \right\|^2} \\ &\leq 2q \sqrt{\frac{1}{q} \sum_{i=0}^{q-1} \mathbb{E}_{\substack{T \subseteq \{0,1\}^n, |T|=t \\ X \subseteq T, |X|=k}} \left\| \Pi_{T \setminus X} |\Phi_i^X\rangle \right\|^2} \quad (\text{Jensen's inequality}) \\ &\leq O \left(q \cdot \sqrt{\frac{t-k}{2^n}} \right). \end{aligned}$$

Therefore, the probability (over the choice of X and T) that $\mathcal{D}^{\mathcal{O}_T}(|\sigma_X\rangle)$ succeeds is at most the probability that $\mathcal{D}^{\mathcal{O}_X}(|\sigma_X\rangle)$ succeeds (up to an additive error of $\frac{q}{\sqrt{t-k}}$). □

Therefore, by applying the triangle inequality, we get that

$$|p(\mathbf{H}_1) - p(\mathbf{H}_5)| \leq O \left(q \cdot \sqrt{\frac{t-s}{2^n}} + q \cdot \sqrt{\frac{t-k}{2^n}} + \frac{k}{\sqrt{s}} + \frac{sk}{t} \right).$$

This proves the claim. □

6 Multi-Copy Revocable Encryption: Definition

In this section we formally define and construct multi-copy secure revocable encryption schemes. These are regular encryption scheme but where the ciphertexts are associated with quantum states. Moreover, the security property guarantees that any adversary that successfully returns k valid copies of a quantum ciphertext (which it received from a trusted party), where k is an arbitrary polynomial, necessarily loses the ability to decrypt the ciphertexts in the future—even if the secret key is revealed.

Our definition of revocable encryption is as follows:

Definition 6.1 (Revocable Encryption). Let $\lambda \in \mathbb{N}$ denote the security parameter. A revocable encryption scheme $\Sigma = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$ with plaintext space \mathcal{M} consists of the following QPT algorithms:

- $\text{KeyGen}(1^\lambda)$: on input the security parameter 1^λ , output a secret key sk .
- $\text{Enc}(\text{sk}, m)$: on input the secret key sk and a message $m \in \mathcal{M}$, output a (pure) ciphertext state $|\psi\rangle$ and a (private) verification key vk .
- $\text{Dec}(\text{sk}, \rho)$: on input the secret key sk and a quantum state ρ , output a message m' .
- $\text{Revoke}(\text{sk}, \text{vk}, \sigma)$: on input the secret key sk , a verification key vk and a state σ , output \top or \perp .

In addition, we require that Σ satisfies the following two properties:

Correctness of decryption: for all plaintexts $m \in \mathcal{M}$, it holds that

$$\Pr \left[m \leftarrow \text{Dec}(\text{sk}, |\psi\rangle) : \begin{array}{l} \text{sk} \leftarrow \text{KeyGen}(1^\lambda) \\ (|\psi\rangle, \text{vk}) \leftarrow \text{Enc}(\text{sk}, m) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

Correctness of revocation: for all plaintexts $m \in \mathcal{M}$, it holds that

$$\Pr \left[\top \leftarrow \text{Revoke}(\text{sk}, \text{vk}, |\psi\rangle) : \begin{array}{l} \text{sk} \leftarrow \text{KeyGen}(1^\lambda) \\ (|\psi\rangle, \text{vk}) \leftarrow \text{Enc}(\text{sk}, m) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

There are two properties we require the above scheme to satisfy.

Firstly, we require the above scheme to be correct. That is, we require the following to hold for all $m \in \{0, 1\}^\ell$,

$$\Pr \left[m \leftarrow \text{Dec}(\text{sk}, |\psi\rangle) : \begin{array}{l} \text{sk} \leftarrow \text{KeyGen}(1^\lambda) \\ |\psi\rangle \leftarrow \text{Enc}(\text{sk}, m) \end{array} \right] \geq 1 - \epsilon(\lambda),$$

for some negligible function $\epsilon(\cdot)$.

Multi-Copy Revocable Security. We use the following notion of security.

Definition 6.2 (Multi-Copy Revocable Security). Let $\lambda \in \mathbb{N}$ denote the security parameter and let $\Sigma = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$ be a revocable encryption scheme with plaintext space \mathcal{M} . Consider the following experiment between a QPT adversary \mathcal{A} and a challenger.

RevokeExpt $_{\lambda, \Sigma, \mathcal{A}}(b)$:

1. \mathcal{A} submits two messages $m_0, m_1 \in \mathcal{M}$ and a polynomial $k = k(\lambda)$ to the challenger.
2. The challenger samples a key $\text{sk} \leftarrow \text{KeyGen}(1^\lambda)$ and produces $|\psi_b\rangle \leftarrow \text{Enc}(\text{sk}, m_b)$. Afterwards, the challenger sends the quantum state $|\psi_b\rangle^{\otimes k}$ to \mathcal{A} .
3. \mathcal{A} returns a quantum state ρ .
4. The challenger performs the measurement $\left\{ |\psi_b\rangle\langle\psi_b|^{\otimes k}, \mathbb{I} - |\psi_b\rangle\langle\psi_b|^{\otimes k} \right\}$ on the returned state ρ . If the measurement succeeds, the game continues; otherwise, the challenger aborts.
5. The challenger sends the secret key sk to \mathcal{A} .
6. \mathcal{A} outputs a bit b' .

We say that the revocable encryption scheme $\Sigma = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$ has multi-copy revocable security, if the following holds for all $\mu \in \{0, 1, \perp\}$:

$$|\Pr[\mu \leftarrow \text{RevokeExpt}_{\lambda, \Sigma, \mathcal{A}}(0)] - \Pr[\mu \leftarrow \text{RevokeExpt}_{\lambda, \Sigma, \mathcal{A}}(1)]| \leq \text{negl}(\lambda),$$

Remark 6.3. In this work, we will work with a weaker variant of this security game, one where the post revocation adversary is given access to an oracle that depends on the key sk instead. The reason for this will become clear from context.

Remark 6.4 (Search variant). Occasionally, we also consider the search variant of multi-copy revocable encryption. Here, the experiment is similar, except that in Step 1, the adversary only submits k to the challenger. Then, in Step 2, the challenger chooses a message m uniformly at random from the plaintext space, encrypts it k times and sends all of the copies to the adversary. Finally, the adversary is said to win the game if it guesses m correctly.

7 Construction of Multi-Copy Secure Revocable Encryption

In this section, we instantiate our revocable encryption scheme using quantum-secure pseudorandom permutations (QPRPs), and we prove security in the random permutation model: this means that, rather than revealing the QPRP key in the final part of the revocable security experiment, we instead allow the adversary to query an oracle for a the permutation instead. Moreover, the permutation is assumed to be random.

Construction 1. Let $\lambda \in \mathbb{N}$ be the security parameter. Let $n, m \in \mathbb{N}$ be polynomial in λ . Let $\Phi = \{\Phi_\lambda\}_{\lambda \in \mathbb{N}}$ be an ensemble of permutations $\Phi_\lambda = \{\varphi_\kappa : \{0, 1\}^{n+m} \rightarrow \{0, 1\}^{n+m}\}_{\kappa \in \mathcal{K}_\lambda}$, for some set \mathcal{K}_λ . Consider the scheme $\Sigma^\Phi = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$ which consists of the following QPT algorithms:

- $\text{KeyGen}(1^\lambda)$: sample a uniformly random key $\kappa \in \mathcal{K}_\lambda$ and let $\text{sk} = \kappa$.
- $\text{Enc}(\text{sk}, \mu)$: on input the secret key $\text{sk} = \kappa$ and message $\mu \in \{0, 1\}^m$, sample $y \sim \{0, 1\}^m$ and prepare the subset state given by

$$|S_y\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0, 1\}^n} |\varphi_\kappa(x||y)\rangle.$$

Output the ciphertext state $(|S_y\rangle^{\otimes k}, y \oplus \mu)$ and (private) verification key $\text{vk} = y$.

- $\text{Dec}(\text{sk}, \text{ct})$: on input the decryption key κ and ciphertext state $(|S_y\rangle, z) \leftarrow \text{ct}$, do the following:
 - Coherently compute φ_κ^{-1} on $|S_y\rangle$ and store the answer in a separate output register.
 - Measure the output register to get $x' || y' \in \{0, 1\}^{n+m}$.
 - Output $y' \oplus z$.
- $\text{Revoke}(\text{sk}, \text{vk}, \rho)$: on input sk , a state ρ and verification key vk , it parses $\kappa \leftarrow \text{sk}$, $y \leftarrow \text{vk}$ and applies the measurement $\{|S_y\rangle\langle S_y|, \mathbb{I} - |S_y\rangle\langle S_y|\}$ to ρ ; it outputs \top if it succeeds, and \perp otherwise.

Proof of Multi-Copy Revocable Security.

Theorem 7.1. *Construction 1 satisfies multi-copy revocable security in the random permutation model.*

Proof. Because we are working in the random permutation model, we will henceforth assume that Σ in [Construction 1](#) is instantiated with a perfectly random permutation φ rather than a QPRP permutation φ_κ .

Suppose that our construction does not achieve multi-copy revocable security. Then, there exists an adversary \mathcal{A} such that

$$|\Pr[\mu \leftarrow \text{RevokeExpt}_{\lambda, \Sigma, \mathcal{A}}(0)] - \Pr[\mu \leftarrow \text{RevokeExpt}_{\lambda, \Sigma, \mathcal{A}}(1)]| = \epsilon(\lambda),$$

for some non-negligible function $\epsilon(\cdot)$. For convenience, we model \mathcal{A} as a pair of quantum algorithms $(\mathcal{A}_0, \mathcal{A}_1)$, where \mathcal{A}_0 corresponds to the pre-revocation adversary, and \mathcal{A}_1 corresponds to the post-revocation adversary. We consider the following sequence of hybrid distributions.

\mathbf{H}_1^b : This corresponds to $\text{RevokeExpt}^{\mathcal{A}}(1^\lambda, b)$.

1. \mathcal{A}_0 submits two m -bit messages (μ_0, μ_1) and a polynomial $k = k(\lambda)$ to the challenger.
2. The challenger samples $y \sim \{0, 1\}^m$ and produces a quantum state

$$|S_y\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0, 1\}^n} |\varphi(x||y)\rangle.$$

The challenger then sends $|S_y\rangle^{\otimes k}$ and $y \oplus \mu_b$ to \mathcal{A}_0 .

3. \mathcal{A}_0 prepares a bipartite state on registers R and AUX, and sends R to the challenger and AUX to \mathcal{A}_1 .
4. The challenger performs the projective measurement $\left\{ |S_y\rangle\langle S_y|^{\otimes k}, \mathbb{I} - |S_y\rangle\langle S_y|^{\otimes k} \right\}$ on R. If the measurement succeeds, the challenger outputs \perp . Otherwise, the challenger continues.
5. The challenger grants \mathcal{A}_1 quantum oracle access to φ^{-1} .
6. \mathcal{A}_1 outputs a bit b' .

\mathbf{H}_2^b : This is the same experiment as in \mathbf{H}_1^b , except that we change how $|S_y\rangle^{\otimes k}$ is generated before revocation:

- The challenger samples a random subset $S \subseteq \{0, 1\}^{n+m}$ of size $|S| = 2^n$.
- The challenger samples a random $y \sim \{0, 1\}^m$.
- The challenger sends $|S\rangle^{\otimes k}$ and $y \oplus \mu_b$ to \mathcal{A}_0 .

Later, the challenger samples a random permutation $\pi : \{0, 1\}^{n+m} \rightarrow \{0, 1\}^{n+m}$ subject to the constraint that $\pi(s) = *||y$, for all $s \in S$. In other words, $\pi(S) = T_y$, where

$$T_y := \{x \in \{0, 1\}^{n+m} : x = (*||y)\}$$

and $|S| = |T_y| = 2^n$. After revocation, \mathcal{A}_1 receives oracle access to π .

Claim 7.2. \mathbf{H}_1^b and \mathbf{H}_2^b are identically distributed.

Proof. This follows immediately. We just changed the order in which we sample things. □

\mathbf{H}_3^b : This is the same experiment as in \mathbf{H}_2^b , except that we change the second part of the experiment: after the challenger sends $|S\rangle^{\otimes k}$ and $y \oplus \mu_b$ to \mathcal{A}_0 , he does the following:

- The challenger samples a random function $g : S \rightarrow T_y$.⁴
- The challenger samples a random permutation $\omega : (\{0, 1\}^{n+m} \setminus S) \rightarrow (\{0, 1\}^{n+m} \setminus T_y)$.

⁴Note that we have $g(S) \subseteq T_y$ in general.

After revocation, \mathcal{A}_1 receives oracle access to $f : \{0, 1\}^{n+m} \rightarrow \{0, 1\}^{n+m}$, where

$$f(x) = \begin{cases} g(x), & \text{if } x \in S \\ \omega(x), & \text{if } x \notin S. \end{cases}$$

Claim 7.3. \mathbf{H}_2^b and \mathbf{H}_3^b are $O(q^3/2^n)$ -close whenever \mathcal{A} makes q queries.

Proof. Here, we apply Zhandry's result from [Theorem 4.3](#) which says that random functions are indistinguishable from random permutations. Consider the following algorithm \mathcal{B} which receives oracle access to O , which is either a random function $F : \{0, 1\}^n \rightarrow \{0, 1\}^n$ or a random permutation $P : \{0, 1\}^n \rightarrow \{0, 1\}^n$:

1. \mathcal{B} samples a random subset $S \subseteq \{0, 1\}^{n+m}$ of size $|S| = 2^n$ and a random $y \sim \{0, 1\}^m$.
2. \mathcal{B} sends $|S_y\rangle^{\otimes k}$ and $y \oplus \mu_b$ to \mathcal{A}_0 .
3. When \mathcal{A}_0 replies with a bipartite state on registers R and AUX, \mathcal{B} performs the projective measurement $\left\{ |S\rangle\langle S|^{\otimes k}, \mathbb{I} - |S\rangle\langle S|^{\otimes k} \right\}$ on R. If it succeeds, \mathcal{B} outputs \perp . Otherwise, \mathcal{B} continues.
4. \mathcal{B} runs the post-revocation adversary \mathcal{A}_1 on input AUX. Whenever \mathcal{A}_1 makes a query, \mathcal{B} answers using the function

$$f(x) = \begin{cases} (\tau \circ O \circ \sigma)(x), & \text{if } x \in S \\ \omega(x), & \text{if } x \notin S \end{cases}$$

where we let

- σ be some canonical mapping from S (with $|S| = 2^n$) onto $\{0, 1\}^n$, i.e., σ is a function which assigns each element of $S \subseteq \{0, 1\}^{n+m}$ a unique bit string in $\{0, 1\}^n$.
- τ be the function which maps each $x \in \{0, 1\}^n$ to $(x||y) \in \{0, 1\}^{n+m}$.

Note that whenever O is a random permutation, the view of \mathcal{A} is precisely \mathbf{H}_3^b ; whereas, if O is a random function, the view of \mathcal{A} is precisely \mathbf{H}_4^b . Therefore, the claim follows from Zhandry's result in [Theorem 4.3](#). \square

\mathbf{H}_4^b : This is the same experiment as in \mathbf{H}_3^b , except that we change the second part of the experiment once again: after the challenger sends $|S_y\rangle^{\otimes k}$ and $y \oplus \mu_b$ to \mathcal{A}_0 , he does the following:

- the challenger samples a random function $f : \{0, 1\}^{n+m} \rightarrow \{0, 1\}^{n+m}$ subject to the constraint that $f(s) = *||y$, for all $s \in S$.

After revocation, \mathcal{A}_1 receives oracle access to f .

Claim 7.4. \mathbf{H}_3^b and \mathbf{H}_4^b are $O(q^3/(2^{n+m} - 2^n))$ -close whenever \mathcal{A} makes q queries.

Proof. The proof again follows from Zhandry's result in [Theorem 4.3](#), since distinguishing \mathbf{H}_3^b and \mathbf{H}_4^b amounts to distinguishing a random function from a random permutation mapping a random permutation $\omega : (\{0, 1\}^{n+m} \setminus S) \rightarrow (\{0, 1\}^{n+m} \setminus T_y)$. Since the domain and co-domain are of equal size $2^{n+m} - 2^n$, the advantage is at most $O(q^3/(2^{n+m} - 2^n))$. \square

\mathbf{H}_5^b : This is the same experiment as in \mathbf{H}_4^b , but now we change what \mathcal{A}_0 receives in the pre-revocation phase:

- The challenger samples a random subset $S \subseteq \{0, 1\}^{n+m}$ of size $|S| = 2^n$.
- The challenger samples a random $y \sim \{0, 1\}^m$.

- The challenger sends $(|S_y\rangle^{\otimes k}, y)$ to \mathcal{A}_0 .

After revocation, the challenger samples a random function $f : \{0, 1\}^{n+m} \rightarrow \{0, 1\}^{n+m}$ subject to the constraint that $f(s) = *||y \oplus \mu_b$, for all $s \in S$. \mathcal{A} receives oracle access to f .

Claim 7.5. \mathbf{H}_4^b and \mathbf{H}_5^b are indistinguishable.

Proof. This follows from the fact that we have just re-labeled the variables. \square

\mathbf{H}_6^b : This is the same experiment as in \mathbf{H}_5^b , except that we once again change what \mathcal{A}_0 receives in the pre-revocation phase:

- The challenger samples a random subset $S \subseteq \{0, 1\}^{n+m}$ of size $|S| = 2^n$.
- The challenger samples a random $y \sim \{0, 1\}^m$.
- The challenger samples a random $u \sim \{0, 1\}^m$.
- The challenger sends $(|S_y\rangle^{\otimes k}, y)$ to \mathcal{A}_0 .

After revocation, the challenger samples a random function $f : \{0, 1\}^{n+m} \rightarrow \{0, 1\}^{n+m}$ subject to the constraint that $f(s) = *||u$, for all $s \in S$. \mathcal{A} receives oracle access to f .

Note that hybrids \mathbf{H}_5^0 and \mathbf{H}_6^1 are identically distributed. By assumption, we also know that \mathcal{A} distinguishes \mathbf{H}_1^0 and \mathbf{H}_1^1 with non-negligible advantage. Therefore, our previous hybrid argument has shown that \mathcal{A} must also distinguish \mathbf{H}_5^0 and \mathbf{H}_6^0 (respectively, hybrids \mathbf{H}_5^1 and \mathbf{H}_6^1) with advantage at least $\epsilon'(\lambda)$, for some non negligible function $\epsilon'(\lambda)$. To complete the proof, we show the following claim which yields the desired contradiction to the $k \mapsto k + 1$ unforgeability property of subset states from [Theorem 5.1](#).

Claim 7.6. *Forge in Algorithm 1 is a poly(λ)-query algorithm (which internally runs \mathcal{A}) such that*

$$\Pr_{\substack{S \subseteq \{0,1\}^{n+m} \\ |S|=2^n}} \left[(x_1, \dots, x_{k+1}) \in \text{dist}(S, k+1) : (x_1, \dots, x_{k+1}) \leftarrow \text{Forge}^{\mathcal{O}_S}(|S\rangle^{\otimes k}) \right] \geq 1/\text{poly}(\lambda).$$

Proof. Since \mathbf{H}_5^b and \mathbf{H}_6^b can be distinguished by the adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ with advantage $\epsilon'(\lambda)$, for some non-negligible function $\epsilon'(\lambda)$, this implies that the post-revocation adversary \mathcal{A}_1 is an algorithm for which the following property holds: given as input uniformly random string y and an auxiliary register (conditioned on revocation succeeding), \mathcal{A}_1 can distinguish whether it is given an oracle for a function $H : \{0, 1\}^{n+m} \rightarrow \{0, 1\}^{n+m}$ which is random subject to the constraint that $H(s) = *||y \oplus \mu_b$ for all $s \in S$, or whether it is given an oracle for a function $G : \{0, 1\}^{n+m} \rightarrow \{0, 1\}^{n+m}$ which is a random function subject to the constraint that $G(s) = *||u$ for all $s \in S$. Crucially, the two functions differ precisely on inputs belonging to S , and are otherwise identical.

Consider the quantum extractor $\text{Ext}^G(y, \text{AUX})$ which is defined as follows:

1. Sample $i \leftarrow [q]$, where q denotes the total number of queries made by \mathcal{A} .
2. Run $\mathcal{A}_1^G(y, \text{AUX})$ just before the $(i - 1)$ -st query to G .
3. Measure \mathcal{A}_1 's i -th query in the computational basis, and output the measurement outcome.

From the O2H Lemma 4.1, we get that

$$\left| \Pr \left[\mathcal{A}_1^H(y, \text{AUX}_{|\top}) = 1 : \begin{array}{l} S \subset \{0,1\}^{n+m}, |S|=2^n \\ y \sim \{0,1\}^m \\ (R, \text{AUX}) \leftarrow \mathcal{A}_0(|S|^{\otimes k}, y) \\ H \text{ s.t. } H(s) = * || y \oplus \mu_b, \forall s \in S \end{array} \right] - \Pr \left[\mathcal{A}_1^G(y, \text{AUX}_{|\top}) = 1 : \begin{array}{l} S \subset \{0,1\}^{n+m}, |S|=2^n \\ y \sim \{0,1\}^m, u \sim \{0,1\}^m \\ (R, \text{AUX}) \leftarrow \mathcal{A}_0(|S|^{\otimes k}, y) \\ G \text{ s.t. } G(s) = * || u, \forall s \in S \end{array} \right] \right| \\ \leq 2q \sqrt{\Pr \left[\text{Ext}^G(y, \text{AUX}_{|\top}) \in S : \begin{array}{l} S \subset \{0,1\}^{n+m}, |S|=2^n \\ y \sim \{0,1\}^m, u \sim \{0,1\}^m \\ (R, \text{AUX}) \leftarrow \mathcal{A}_0(|S|^{\otimes k}, y) \\ G \text{ s.t. } G(s) = * || u, \forall s \in S \end{array} \right]},$$

where $\text{AUX}_{|\top}$ corresponds to the register AUX conditioned on the event that the projective measurement

$$\left\{ |S\rangle\langle S|^{\otimes k}, \mathbb{I} - |S\rangle\langle S|^{\otimes k} \right\}$$

succeeds on register R . Because the distinguishing advantage of the adversary \mathcal{A}_1 is non-negligible (conditioned on the event that revocation succeeds on register R) and $q = \text{poly}(\lambda)$, we get

$$\Pr \left[\text{Ext}^G(y, \text{AUX}_{|\top}) \in S : \begin{array}{l} S \subset \{0,1\}^{n+m}, |S|=2^n \\ y \sim \{0,1\}^m, u \sim \{0,1\}^m \\ (R, \text{AUX}) \leftarrow \mathcal{A}_0(|S|^{\otimes k}, y) \\ G \text{ s.t. } G(s) = * || u, \forall s \in S \end{array} \right] \geq 1/\text{poly}(\lambda).$$

To complete the proof, we now show that $\text{Forge}^{\mathcal{O}_S}(|S|^{\otimes k})$ in Algorithm 1 is a successful algorithm against the $k \rightarrow k+1$ unforgeability of subset states.

Algorithm 1: $\text{Forge}^{\mathcal{O}_S}(|S|^{\otimes k})$

Input: $|S|^{\otimes k}$ and a membership oracle \mathcal{O}_S , where $S \subset \{0,1\}^{n+m}$ is a subset.

Output: $x_1, \dots, x_{k+1} \in \{0,1\}^{n+m}$.

- 1 Sample uniformly random strings $y, u \sim \{0,1\}^m$;
- 2 Run $(R, \text{AUX}) \leftarrow \mathcal{A}_0(|S|^{\otimes k}, y)$;
- 3 Measure R in the computational basis to obtain x_1, \dots, x_k ;
- 4 Run the quantum extractor $\text{Ext}^G(y, \text{AUX})$ to obtain an element x_{k+1} , where the oracle G can be simulated via \mathcal{O}_S as follows: on input $x \in \{0,1\}^{n+m}$, we let

$$G(x) = \begin{cases} g_1(x) || u, & \text{if } \mathcal{O}_S(x) = 1 \\ g_2(x), & \text{otherwise} \end{cases}$$

where $g_1 : \{0,1\}^{n+m} \rightarrow \{0,1\}^n$ and $g_2 : \{0,1\}^{n+m} \rightarrow \{0,1\}^{n+m}$ are uniformly random functions.

- 5 Output (x_1, \dots, x_{k+1}) .
-

Let $\text{Revoke}(S, k, R)$ denote the projective measurement $\{|S\rangle\langle S|^{\otimes k}, \mathbb{I} - |S\rangle\langle S|^{\otimes k}\}$ of register R . Using the Simultaneous Distinct Extraction Lemma (Lemma 7.7), we get that

$$\Pr_{\substack{S \subset \{0,1\}^{n+m} \\ |S|=2^n}} \left[(x_1, \dots, x_{t+1}) \in \text{dist}(S, k+1) : (x_1, \dots, x_{k+1}) \leftarrow \text{Forge}^{\mathcal{O}_S}(|S|^{\otimes k}) \right] \\ \geq \left(1 - O\left(\frac{k^2}{2^n}\right) \right) \cdot \Pr \left[\text{Revoke}(S, k, R) = \top : \begin{array}{l} S \subset \{0,1\}^{n+m}, |S|=2^n \\ y \sim \{0,1\}^m \\ (R, \text{AUX}) \leftarrow \mathcal{A}_0(|S|^{\otimes k}, y) \end{array} \right] \\ \cdot \Pr \left[\text{Ext}^G(y, \text{AUX}_{|\top}) \in S : \begin{array}{l} S \subset \{0,1\}^{n+m}, |S|=2^n \\ y \sim \{0,1\}^m, u \sim \{0,1\}^m \\ (R, \text{AUX}) \leftarrow \mathcal{A}_0(|S|^{\otimes k}, y) \\ G \text{ s.t. } G(s) = * || u, \forall s \in S \end{array} \right]$$

which is at least inverse polynomial in λ . This proves the claim. \square

Thus, we obtain the desired contradiction to the $k \mapsto k + 1$ unforgeability property of subset states from [Theorem 5.1](#). This completes the proof of multi-copy revocable encryption security. \square

7.1 Simultaneous Distinct Extraction Lemma

The following lemma allows us to analyze the probability of simultaneously extracting $k + 1$ distinct subset elements in some subset $S \subseteq \{0, 1\}^n$ in terms of the success probability of revocation (i.e., the projection onto $|S\rangle\langle S|^{\otimes k}$) and the success probability of extracting another subset element from the adversary's state.

Lemma 7.7 (Simultaneous Distinct Extraction). *Let $n, k \in \mathbb{N}$ and let $\rho \in \mathcal{D}(\mathcal{H}_X \otimes \mathcal{H}_Y)$ be an any density matrix, where \mathcal{H}_X is an $n \cdot k$ -qubit Hilbert space and where \mathcal{H}_Y is an arbitrary Hilbert space. Let $|S\rangle$ denote a subset state, for some subset $S \subseteq \{0, 1\}^n$, and let $\mathcal{E} : \mathcal{L}(\mathcal{H}_Y) \rightarrow \mathcal{L}(\mathcal{H}_{X'})$ be any CPTP map of the form*

$$\mathcal{E}_{Y \rightarrow X'}(\sigma) = \text{Tr}_E \left[V_{Y \rightarrow X'E} \sigma V_{Y \rightarrow X'E}^\dagger \right], \quad \forall \sigma \in \mathcal{D}(\mathcal{H}_Y),$$

for some isometry $V_{Y \rightarrow X'E}$ and n -qubit Hilbert space $\mathcal{H}_{X'}$. Consider the POVM element

$$\Lambda = \sum_{\substack{s_1, \dots, s_{k+1} \in S \\ (s_1, \dots, s_{k+1}) \in \text{dist}(S, k+1)}} |s_1, \dots, s_{k+1}\rangle\langle s_1, \dots, s_{k+1}|_X \otimes V_{Y \rightarrow X'E}^\dagger (|s_{k+1}\rangle\langle s_{k+1}|_{X'} \otimes I_E) V_{Y \rightarrow X'E}.$$

Let $\rho_X = \text{Tr}_Y[\rho_{XY}]$ denote the reduced state. Then, it holds that

$$\text{Tr}[\Lambda \rho] \geq \left(1 - O\left(\frac{k^2}{|S|}\right) \right) \cdot \text{Tr}\left[|S\rangle\langle S|^{\otimes k} \rho_X\right] \cdot \text{Tr}\left[|S\rangle\langle S| \mathcal{E}_{Y \rightarrow X'}(\sigma)\right],$$

where $\sigma = \text{Tr}\left[\left(|S\rangle\langle S|^{\otimes k} \otimes I\right) \rho\right]^{-1} \cdot \text{Tr}_X\left[\left(|S\rangle\langle S|^{\otimes k} \otimes I\right) \rho\right]$ is a reduced state in system Y .

Proof. Because the order in which we apply Λ and $(|S\rangle\langle S|^{\otimes k} \otimes I)$ does not matter, we have the inequality

$$\begin{aligned} \text{Tr}[\Lambda \rho] &\geq \text{Tr}\left[\left(|S\rangle\langle S|^{\otimes k} \otimes I\right) \Lambda \rho\right] \\ &= \text{Tr}\left[\left(|S\rangle\langle S|^{\otimes k} \otimes I\right) \Lambda \rho \left(|S\rangle\langle S|^{\otimes k} \otimes I\right)\right] \\ &= \text{Tr}\left[\Lambda \left(|S\rangle\langle S|^{\otimes k} \otimes I\right) \rho \left(|S\rangle\langle S|^{\otimes k} \otimes I\right)\right]. \end{aligned} \tag{1}$$

Notice also that $(|S\rangle\langle S|^{\otimes k} \otimes I) \rho (|S\rangle\langle S|^{\otimes k} \otimes I)$ lies in the image of $(|S\rangle\langle S|^{\otimes k} \otimes I)$, and thus

$$\left(|S\rangle\langle S|^{\otimes k} \otimes I\right) \rho \left(|S\rangle\langle S|^{\otimes k} \otimes I\right) = \text{Tr}\left[\left(|S\rangle\langle S|^{\otimes k} \otimes I\right) \rho\right] \cdot \left(|S\rangle\langle S|^{\otimes k} \otimes \sigma\right), \tag{2}$$

for some $\sigma \in \mathcal{D}(\mathcal{H}_Y)$. Using [\(1\)](#) and [\(2\)](#), we get the lower bound,

$$\begin{aligned} \text{Tr}[\Lambda \rho] &\geq \text{Tr}\left[\Lambda \left(|S\rangle\langle S|^{\otimes k} \otimes I\right) \rho \left(|S\rangle\langle S|^{\otimes k} \otimes I\right)\right] \\ &= \text{Tr}\left[\left(|S\rangle\langle S|^{\otimes k} \otimes I\right) \rho\right] \cdot \text{Tr}\left[\Lambda \left(|S\rangle\langle S|^{\otimes k} \otimes \sigma\right)\right]. \end{aligned}$$

Next, we analyze the final right-most quantity separately. Specifically, we find

$$\begin{aligned}
& \text{Tr} \left[\Lambda \left(|S\rangle\langle S|^{\otimes k} \otimes \sigma \right) \right] \\
&= \sum_{\substack{s_1, \dots, s_{k+1} \in S: \\ (s_1, \dots, s_{k+1}) \\ \in \text{dist}(S, k+1)}} \text{Tr} \left[|s_1, \dots, s_k\rangle\langle s_1, \dots, s_k|_X \otimes V_{Y \rightarrow X'E}^\dagger (|s_{k+1}\rangle\langle s_{k+1}|_{X'} \otimes I_E) V_{Y \rightarrow X'E} \left(|S\rangle\langle S|^{\otimes k} \otimes \sigma \right) \right] \\
&= \sum_{s_{k+1} \in S} \left(\sum_{\substack{s_1, \dots, s_k \in S: \\ (s_1, \dots, s_{k+1}) \\ \in \text{dist}(S, k+1)}} |S|^{-k} \right) \text{Tr} \left[V_{Y \rightarrow X'E}^\dagger (|s_{k+1}\rangle\langle s_{k+1}|_{X'} \otimes I_E) V_{Y \rightarrow X'E} \sigma \right] \\
&= \sum_{s_{k+1} \in S} \left(1 - \sum_{\substack{s_1, \dots, s_k \in S: \\ (s_1, \dots, s_k) \notin \text{dist}(S, k) \\ \text{or } \exists i \in [k]: s_i = s_{k+1}}} |S|^{-k} \right) \text{Tr} \left[(|s_{k+1}\rangle\langle s_{k+1}|_{X'} \otimes I_E) V_{Y \rightarrow X'E} \sigma V_{Y \rightarrow X'E}^\dagger \right] \\
&\geq \left(1 - O\left(\frac{k^2}{|S|}\right) \right) \cdot \sum_{s_{k+1} \in S} \text{Tr} \left[(|s_{k+1}\rangle\langle s_{k+1}|_{X'} \otimes I_E) V_{Y \rightarrow X'E} \sigma V_{Y \rightarrow X'E}^\dagger \right] \\
&= \left(1 - O\left(\frac{k^2}{|S|}\right) \right) \cdot \sum_{s_{k+1} \in S} \text{Tr} \left[|s_{k+1}\rangle\langle s_{k+1}|_{X'} \cdot \text{Tr}_E \left[V_{Y \rightarrow X'E} \sigma V_{Y \rightarrow X'E}^\dagger \right] \right] \\
&= \left(1 - O\left(\frac{k^2}{|S|}\right) \right) \cdot \text{Tr} \left[|S\rangle\langle S| \mathcal{E}_{Y \rightarrow X'}(\sigma) \right].
\end{aligned}$$

Putting everything together, this gives the desired inequality

$$\text{Tr}[\Lambda\rho] \geq \left(1 - O\left(\frac{k^2}{|S|}\right) \right) \cdot \text{Tr} \left[|S\rangle\langle S|^{\otimes k} \rho_X \right] \cdot \text{Tr} \left[|S\rangle\langle S| \mathcal{E}_{Y \rightarrow X'}(\sigma) \right],$$

where $\sigma = \text{Tr} \left[(|S\rangle\langle S|^{\otimes k} \otimes I) \rho \right]^{-1} \cdot \text{Tr}_X \left[(|S\rangle\langle S|^{\otimes k} \otimes I) \rho \right]$ is a reduced state in system Y . \square

8 Multi-Copy Revocable Programs: Definition

In this section, we study whether arbitrary functionalities can be revoked. We define and study revocable programs with multi-copy security. In this notion, there is a functionality preserving compiler that takes a program and converts it into a quantum state, which can later be certifiably revoked.

We now give a formal definition of revocable programs.

Definition 8.1 (Revocable Program). *Let $\mathcal{P} = \bigcup_{\lambda \in \mathbb{N}} \mathcal{P}_\lambda$ be a class of efficiently computable program families $\mathcal{P}_\lambda = \{P : \mathcal{X}_\lambda \rightarrow \mathcal{Y}_\lambda\}$ with domain \mathcal{X}_λ and range \mathcal{Y}_λ . A revocable program compiler for the class \mathcal{P} is a tuple $\Sigma = (\text{Compile}, \text{Eval}, \text{Revoke})$ consisting of the following QPT algorithms:*

- **Compile**($1^\lambda, P$): *on input the security parameter 1^λ and a program $P \in \mathcal{P}_\lambda$ with $P : \mathcal{X}_\lambda \rightarrow \mathcal{Y}_\lambda$, output a quantum state $|\Psi_P\rangle$ and a (private) verification key vk .*
- **Eval**($|\Psi_P\rangle, x$): *on input a quantum state $|\Psi_P\rangle$ and input $x \in \mathcal{X}_\lambda$, output $P(x)$.*
- **Revoke**(vk, σ): *on input the verification key vk and a state σ , output \top or \perp .*

In addition, we require that Σ satisfies the following two properties for all $\lambda \in \mathbb{N}$:

Correctness of evaluation: *for all programs $P \in \mathcal{P}_\lambda$ and inputs $x \in \mathcal{X}_\lambda$, it holds that*

$$\Pr \left[P(x) \leftarrow \text{Eval}(|\Psi_P\rangle, x) : (|\Psi_P\rangle, \text{vk}) \leftarrow \text{Compile}(1^\lambda, P) \right] \geq 1 - \text{negl}(\lambda).$$

Correctness of revocation: for all programs $P \in \mathcal{P}_\lambda$, it holds that

$$\Pr [\top \leftarrow \text{Revoke}(\text{vk}, (|\Psi_P\rangle)) : (|\Psi_P\rangle, \text{vk}) \leftarrow \text{Compile}(1^\lambda, P)] \geq 1 - \text{negl}(\lambda).$$

Multi-Copy Revocable Security. We use the following notion of security.

Definition 8.2 (Multi-Copy Revocable Security for Programs). *Let $\mathcal{P} = \bigcup_{\lambda \in \mathbb{N}} \mathcal{P}_\lambda$ be a class of efficiently computable program families $\mathcal{P}_\lambda = \{P : \mathcal{X}_\lambda \rightarrow \mathcal{Y}_\lambda\}$ and let $\mathcal{D}_\mathcal{P} = \bigcup_{\lambda \in \mathbb{N}} \mathcal{D}_{\mathcal{P}_\lambda}$ be an ensemble of program distributions. Let $\mathcal{D}_\mathcal{X} = \bigcup_{\lambda \in \mathbb{N}} \mathcal{D}_{\mathcal{X}_\lambda}$ be an ensemble of challenge distribution families with $\mathcal{D}_{\mathcal{X}_\lambda} = \{\mathcal{D}_{\mathcal{X}_\lambda}(P)\}_{P \in \mathcal{P}_\lambda}$. Consider the following experiment between a QPT adversary \mathcal{A} and a challenger.*

RevokeExpt $_{\lambda, \Sigma, \mathcal{A}}^{\mathcal{D}_\mathcal{P}, \mathcal{D}_\mathcal{X}}$:

1. \mathcal{A} submits a polynomial $k = k(\lambda)$ to the challenger.
2. The challenger samples a program $P \sim \mathcal{D}_{\mathcal{P}_\lambda}$ with domain \mathcal{X}_λ and range \mathcal{Y}_λ , and runs $\text{Compile}(1^\lambda, P)$ to generate a pair $(|\Psi_P\rangle, \text{vk})$. Afterwards, the challenger sends the quantum state $|\Psi_P\rangle^{\otimes k}$ to \mathcal{A} .
3. \mathcal{A} returns a quantum state ρ .
4. The challenger performs the measurement $\left\{ |\Psi_P\rangle\langle\Psi_P|^{\otimes k}, \mathbb{I} - |\Psi_P\rangle\langle\Psi_P|^{\otimes k} \right\}$ on the returned state ρ . If the measurement succeeds, the game continues; otherwise, the challenger aborts.
5. The challenger samples a challenge input $x \sim \mathcal{D}_{\mathcal{X}_\lambda}(P)$ sends x to \mathcal{A} .
6. The adversary outputs $y \in \mathcal{Y}_\lambda$.
7. The challenger outputs 1 if and only if $P(x) = y$.

We say that a revocable program compiler $\Sigma = (\text{Compile}, \text{Eval}, \text{Revoke})$ has multi-copy revocable security for the ensembles $\mathcal{D}_\mathcal{P}$ and $\mathcal{D}_\mathcal{X}$, if the following holds for any QPT adversary \mathcal{A} :

$$\Pr [1 \leftarrow \text{RevokeExpt}_{\lambda, \Sigma, \mathcal{A}}^{\mathcal{D}_\mathcal{P}, \mathcal{D}_\mathcal{X}}] \leq p_{\text{triv}}^{\mathcal{D}_\mathcal{P}, \mathcal{D}_\mathcal{X}}(\lambda) + \text{negl}(\lambda),$$

where $p_{\text{triv}}^{\mathcal{D}_\mathcal{P}, \mathcal{D}_\mathcal{X}}(\lambda) = \sup_{\mathcal{A}} \{\Pr [P(x) \leftarrow \mathcal{A}(x) : x \leftarrow \mathcal{D}_{\mathcal{X}_\lambda}(P)]\}$ is the trivial guessing probability.

9 Construction of Multi-Copy Secure Revocable Programs in a Classical Oracle Model

In this section, we give a construction of multi-copy secure revocable programs; specifically, we work with a classical oracle model. Our construction is as follows:

Construction 2. *Let $\lambda \in \mathbb{N}$ be the security parameter. Let $n, m \in \mathbb{N}$ be polynomial in λ . Let $\Phi = \{\Phi_\lambda\}_{\lambda \in \mathbb{N}}$ be an ensemble of permutations $\Phi_\lambda = \{\varphi_\kappa : \{0, 1\}^{n+m} \rightarrow \{0, 1\}^{n+m}\}_{\kappa \in \mathcal{K}_\lambda}$, for some set \mathcal{K}_λ .*

- **Setup**(1^λ): sample a uniformly random key $\kappa \in \mathcal{K}_\lambda$ and let $\text{vk} = \kappa$.
- **Compile**($1^\lambda, P$): on input 1^λ and a program $P \in \mathcal{P}_\lambda$ with $P : \mathcal{X}_\lambda \rightarrow \mathcal{Y}_\lambda$, do the following:
 - Sample $y \sim \{0, 1\}^m$ and prepare the subset state given by

$$|S_y\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0, 1\}^n} |\varphi_\kappa(x|y)\rangle.$$

- Let $|\Psi_P\rangle = |S_y\rangle$ and, for brevity, let $S = \{\varphi_\kappa(x||y) : x \in \{0,1\}^n\}$ be the corresponding subset.
- Let $O = O_{P,S}$ denote an a (public) classical oracle, which is defined as follows:

$$O_{P,S}(x, s) = \begin{cases} P(x), & \text{if } s \in S \\ 0, & \text{otherwise} \end{cases}$$

- $\text{Eval}^O(|\Psi_P\rangle, x)$: on input $|\Psi_P\rangle$, $x \in \mathcal{X}_\lambda$, do the following:
 - Coherently evaluate $O_{P,S}$ on input $|x\rangle \otimes |\Psi_P\rangle$, and compute its output into an ancillary register.
 - Measure the ancillary register and then output the measurement outcome.
- $\text{Revoke}(\text{vk}, \rho)$: on input vk , a state ρ and verification key vk , it parses $y \leftarrow \text{vk}$ and applies the measurement $\{|S_y\rangle\langle S_y|, \mathbb{I} - |S_y\rangle\langle S_y|\}$ to ρ ; it outputs \top if it succeeds, and \perp otherwise.

The above scheme is easily seen to satisfy correctness.

Proof of Multi-Copy Revocable Security. Before we analyze the security of [Construction 2](#), let us first remark that we can instantiate the scheme using a QPRP family $\Phi_\lambda = \{\varphi_\kappa : \{0,1\}^{n+m} \rightarrow \{0,1\}^{n+m}\}_{\kappa \in \mathcal{K}_\lambda}$, for some key space \mathcal{K}_λ . In the security proof, however, we will work with the random permutation model instead. This means that we will consider random permutations throughout the security game.

Theorem 9.1. *Construction 2 satisfies multi-copy revocable security for any pair of distributions $\mathcal{D}_\mathcal{P}, \mathcal{D}_\mathcal{X}$ in the random permutation model, as well as in the classical oracle model, where the recipient receives an (ideal classical) oracle for the purpose of evaluation.*

Proof. Because we are working in the random permutation model, we will henceforth assume that Σ in [Construction 1](#) is instantiated with a perfectly random permutation φ rather than a QPRP permutation φ_κ . Suppose not. Then there exists an adversary \mathcal{A} such that,

$$\Pr \left[1 \leftarrow \text{RevokeExpt}_{\lambda, \Sigma, \mathcal{A}}^{\mathcal{D}_\mathcal{P}, \mathcal{D}_\mathcal{X}} \right] = \epsilon(\lambda) + p_{\text{triv}},$$

for some non negligible function $\epsilon(\lambda)$. For convenience, we model \mathcal{A} as a pair of QPT algorithms $(\mathcal{A}_0, \mathcal{A}_1)$, where \mathcal{A}_0 corresponds to the pre-revocation adversary, and \mathcal{A}_1 corresponds to the post-revocation adversary. We consider the following sequence of hybrid distributions.

$\mathbf{H}_{1,x}$: this corresponds to $\text{RevokeExpt}_{\lambda, \Sigma, \mathcal{A}}^{\mathcal{D}_\mathcal{P}, \mathcal{D}_\mathcal{X}}$.

$\text{RevokeExpt}_{\lambda, \Sigma, \mathcal{A}}^{\mathcal{D}_\mathcal{P}, \mathcal{D}_\mathcal{X}}$:

1. \mathcal{A} submits a polynomial $k = k(\lambda)$ to the challenger.
2. The challenger samples a program $P \sim \mathcal{D}_{\mathcal{P}_\lambda}$ with domain \mathcal{X}_λ and range \mathcal{Y}_λ , and a random $y \leftarrow \{0,1\}^m$, and prepares k copies of the subset state over $S = \{\varphi(x||y) : x \in \{0,1\}^n\}$ given by

$$|S\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |\varphi(x||y)\rangle.$$

3. The adversary sends $|S\rangle^{\otimes k}$ to \mathcal{A} , and grants \mathcal{A} access to the oracle $O = O_{P,S}$, where

$$O_{P,S}(x, s) = \begin{cases} P(x), & \text{if } s \in S \\ 0, & \text{otherwise} \end{cases}$$

4. The challenger performs the projective measurement $\{|S\rangle\langle S|^{\otimes k}, \mathbb{I} - |S\rangle\langle S|^{\otimes k}\}$ on R. If the measurement succeeds, the challenger outputs \perp . Otherwise, the challenger continues.
5. The challenger samples a challenge input $x \sim \mathcal{D}_{\mathcal{X}_\lambda}(P)$ sends x to \mathcal{A} .
6. The adversary outputs $y \in \mathcal{Y}_\lambda$.

$\mathbf{H}_{2,x}$: This is the same as hybrid 1 except, after the revocation phase, \mathcal{A} has access to the oracle $\tilde{O}_{P,S}$ which always outputs 0.

Claim 9.2. $\mathbf{H}_{2,x}$ and $\mathbf{H}_{1,x}$ are computationally indistinguishable.

Proof. Let's assume for contradiction that $\mathbf{H}_{2,x}$ and $\mathbf{H}_{1,x}$ can be distinguished by the adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ with advantage $\epsilon'(\lambda)$, for some non-negligible function $\epsilon'(\lambda)$. This implies that the post-revocation adversary \mathcal{A}_1 is an algorithm for which the following property holds: given as input a string $x \leftarrow \mathcal{D}_{\mathcal{X}}(P)$, and an auxiliary register $\text{AUX}_{|\top}$ (conditioned on revocation succeeding), \mathcal{A}_1 can distinguish whether it is given an oracle for a function $O_{P,S}$, which is defined as follows:

$$O_{P,S}(x, s) = \begin{cases} P(x), & \text{if } s \in S \\ 0, & \text{otherwise} \end{cases}$$

or whether it is given an oracle for a function $\tilde{O}_{P,S}$, which always outputs 0. Crucially, the two functions differ precisely on inputs belonging to S , and are otherwise identical.

Consider the quantum extractor $\text{Ext}^{O_{P,S}}(x, \text{AUX})$ which is defined as follows:

1. Sample $i \leftarrow [q]$, where q denotes the total number of queries made by \mathcal{A} .
2. Run $\mathcal{A}_1^{O_{P,S}}(x, \text{AUX})$ just before the $(i - 1)$ -st query to $O_{P,S}$.
3. Measure \mathcal{A}_1 's i -th query in the computational basis, and output the measurement outcome.

Note that, since we are working with the random permutation model, the subset S is effectively a uniformly random subset of size 2^n . From the O2H Lemma 4.1, we get that

$$\begin{aligned} & \left| \Pr \left[\mathcal{A}_1^{O_{P,S}}(x, \text{AUX}_{|\top}) = 1 : \begin{array}{c} S \subset \{0,1\}^{n+m}, |S|=2^n \\ x \sim \{0,1\}^m \\ (\text{R}, \text{AUX}) \leftarrow \mathcal{A}_0(|S|^{\otimes k}) \\ O_{P,S} \end{array} \right] - \Pr \left[\mathcal{A}_1^{\tilde{O}_{P,S}}(x, \text{AUX}_{|\top}) = 1 : \begin{array}{c} S \subset \{0,1\}^{n+m}, |S|=2^n \\ x \sim \{0,1\}^m \\ (\text{R}, \text{AUX}) \leftarrow \mathcal{A}_0(|S|^{\otimes k}) \\ \tilde{O}_{P,S} \end{array} \right] \right| \\ & \leq 2q \sqrt{\Pr \left[\text{Ext}^{O_{P,S}}(\text{AUX}_{|\top}) \in S : \begin{array}{c} S \subset \{0,1\}^{n+m}, |S|=2^n \\ x \sim \{0,1\}^m \\ (\text{R}, \text{AUX}) \leftarrow \mathcal{A}_0(|S|^{\otimes k}) \\ O_{P,S} \end{array} \right]}, \end{aligned}$$

where $\text{AUX}_{|\top}$ corresponds to the register AUX conditioned on the event that the projective measurement

$$\{|S\rangle\langle S|^{\otimes k}, \mathbb{I} - |S\rangle\langle S|^{\otimes k}\}$$

succeeds on register R. Because the distinguishing advantage of the adversary \mathcal{A}_1 is non-negligible (conditioned on the event that revocation succeeds on register R) and $q = \text{poly}(\lambda)$, we get

$$\Pr \left[\text{Ext}^{O_{P,S}}(x, \text{AUX}_{|\top}) \in S : \begin{array}{c} S \subset \{0,1\}^{n+m}, |S|=2^n \\ x \sim \{0,1\}^m \\ (\text{R}, \text{AUX}) \leftarrow \mathcal{A}_0(|S|^{\otimes k}) \\ O_{P,S} \end{array} \right] \geq 1/\text{poly}(\lambda).$$

To complete the proof, we now show that $\text{Forge}^{O_S}(|S|^{\otimes k})$ in Algorithm 2 is a successful algorithm against the $k \rightarrow k + 1$ unforgeability of subset states. Let $\text{Revoke}(S, k, \text{R})$ denote the projective measurement

Algorithm 2: $\text{Forge}^{\mathcal{O}_S}(|S\rangle^{\otimes k})$

Input: $|S\rangle^{\otimes k}$ and a membership oracle \mathcal{O}_S , where $S \subset \{0, 1\}^{n+m}$ is a subset.

Output: $x_1, \dots, x_{k+1} \in \{0, 1\}^{n+m}$.

- 1 Run $(R, \text{AUX}) \leftarrow \mathcal{A}_0(|S\rangle^{\otimes k})$;
- 2 Measure R in the computational basis to obtain x_1, \dots, x_k ;
- 3 Sample a string $x \leftarrow \mathcal{D}_X(P)$;
- 4 Run the quantum extractor $\text{Ext}^{O_{P,S}}(x, \text{AUX})$ to obtain an element x_{k+1} , where the oracle $O_{P,S}$ can be simulated via \mathcal{O}_S as follows: on input $(x, s) \in \{0, 1\}^{2n+m}$, we let

$$O_{P,S}(x) = \begin{cases} P(x), & \text{if } \mathcal{O}_S(s) = 1 \\ 0, & \text{otherwise} \end{cases}$$

- 5 Output (x_1, \dots, x_{k+1}) .
-

$\{|S\rangle\langle S|^{\otimes k}, \mathbb{I} - |S\rangle\langle S|^{\otimes k}\}$ of register R . Using the Simultaneous Distinct Extraction Lemma (Lemma 7.7), we get that

$$\begin{aligned} & \Pr_{\substack{S \subseteq \{0,1\}^{n+m} \\ |S|=2^n}} \left[(x_1, \dots, x_{k+1}) \in \text{dist}(S, k+1) : (x_1, \dots, x_{k+1}) \leftarrow \text{Forge}^{\mathcal{O}_S}(|S\rangle^{\otimes k}) \right] \\ & \geq \left(1 - O\left(\frac{k^2}{2^n}\right) \right) \cdot \Pr \left[\text{Revoke}(S, k, R) = \top : \begin{array}{l} S \subseteq \{0,1\}^{n+m}, |S|=2^n \\ x \sim \{0,1\}^m \\ (R, \text{AUX}) \leftarrow \mathcal{A}_0(|S\rangle^{\otimes k}) \end{array} \right] \\ & \quad \cdot \Pr \left[\text{Ext}^{O_{P,S}}(x, \text{AUX}_{|\top}) \in S : \begin{array}{l} S \subseteq \{0,1\}^{n+m}, |S|=2^n \\ x \sim \{0,1\}^m \\ (R, \text{AUX}) \leftarrow \mathcal{A}_0(|S\rangle^{\otimes k}) \\ O_{P,S} \end{array} \right] \end{aligned}$$

which is at least inverse polynomial in λ . This proves the claim. \square

Notice that in $\mathbf{H}_{2,x}$, the oracle $\tilde{O}_{P,S}$ always outputs 0, thus, if $\mathbf{H}_{1,x}$ and $\mathbf{H}_{2,x}$ are indistinguishable, $(\mathcal{A}_0^H, \mathcal{A}_1^H)$ succeeds with probability $\epsilon^*(\lambda) + p_{\text{triv}}$, for some non negligible ϵ . However, this implies that $(\mathcal{A}_0^H, \mathcal{A}_1^H)$ breaks the revocable security of the program P . Therefore, it suffices to argue that $\mathbf{H}_{1,x}$ and $\mathbf{H}_{2,x}$ are close in order to complete the proof. \square

10 Construction of Multi-Copy Secure Revocable Point Functions in the QROM

In this section, we construct multi-copy secure revocable point functions. Our main ingredient is an underlying revocable encryption scheme with a so-called *wrong-key detection* mechanism. But first, we introduce some relevant tools, such as hybrid encryption schemes, which will be essential in our analysis.

10.1 Hybrid Encryption

We introduce a simple *hybrid encryption* scheme which we will make use of in order to construct revocable multi-copy secure point functions. Our hybrid encryption scheme combines a revocable encryption scheme with the classical *one-time pad* $\text{OTP} = (\text{KeyGen}, \text{Enc}, \text{Dec})$, where an encryption of a plaintext $m \in \{0, 1\}^\lambda$ via a key $k \in \{0, 1\}^\lambda$ is generated as $\text{OTP.Enc}(k, m) = k \oplus m$.

Construction 3 (Hybrid encryption scheme). *Let $\lambda \in \mathbb{N}$ denote the security parameter and let the scheme $\Sigma = (\text{KeyGen}, \text{Enc}, \text{Dec})$ be a revocable encryption scheme. Let $\text{OTP} = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$ be the classical one-time pad. We define the hybrid encryption scheme $\text{HE} = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$ as follows:*

- $\text{HE.KeyGen}(1^\lambda)$: This is the same procedure as $\Sigma.\text{KeyGen}(1^\lambda)$.
- $\text{HE.Enc}(\text{sk}, m)$: Given as input a key sk and a plaintext $m \in \{0, 1\}^\lambda$, encrypt as follows:
 1. Sample a random string $r \sim \{0, 1\}^\lambda$.
 2. Generate the pair $(|\text{ct}\rangle, \text{vk}) \leftarrow \Sigma.\text{Enc}(\text{sk}, r)$.
 3. Output the ciphertext $(|\text{ct}\rangle, \text{OTP.Enc}(r, m))$ and (private) verification key vk .
- $\text{HE.Dec}(\text{sk}, c)$: given as input a key sk and ciphertext $c = (c_0, c_1)$, decrypt as follows:
 1. Compute $r' = \Sigma.\text{Dec}(\text{sk}, c_0)$
 2. Output the plaintext $\text{OTP.Dec}(r', c_1)$.
- $\text{HE.Revoke}(\text{sk}, \text{vk}, \sigma)$: On input the keys sk and vk , as well as the state σ , output $\Sigma.\text{Revoke}(\text{sk}, \text{vk}, \sigma)$

We now prove the following theorem. Note that in the theorem below, we make use of the search variant of multi-copy revocable security which we mentioned in [Remark 6.4](#).

Theorem 10.1. *Let $\lambda \in \mathbb{N}$ denote the security parameter and let $\Sigma = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$ be a revocable multi-copy (search) secure encryption scheme. Then, the resulting hybrid encryption scheme given by $\text{HE} = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$ in [Construction 3](#) which is instantiated with Σ is also a revocable multi-copy (search) secure encryption scheme.*

Proof. Suppose that \mathcal{A} is a successful adversary against the (search) revocable security game with respect to HE , i.e., there exists a non-negligible function $\epsilon(\lambda)$ such that

$$\Pr [1 \leftarrow \text{RevokeEpt}_{\lambda, \text{HE}, \mathcal{A}}] = \epsilon(\lambda).$$

For convenience, we model \mathcal{A} as a pair of quantum algorithms $(\mathcal{A}_0, \mathcal{A}_1)$, where \mathcal{A}_0 corresponds to the pre-revocation adversary, and \mathcal{A}_1 corresponds to the post-revocation adversary.

Consider the following reduction $\mathcal{B} = (\mathcal{B}_0, \mathcal{B}_1)$ which breaks the security of Σ .

1. \mathcal{B}_0 runs \mathcal{A}_0 to obtain $k = k(\lambda)$, where $k(\lambda)$ is a polynomial, and sends it to the challenger.
2. Upon receiving a ciphertext of the form $\text{Enc}(\text{sk}, m)^{\otimes k}$, for a random $m \sim \{0, 1\}^\lambda$, \mathcal{B}_1 samples a random $r \sim \{0, 1\}^\lambda$, and runs $(R, \text{AUX}) \leftarrow \mathcal{A}_0(\text{Enc}(\text{sk}, m)^{\otimes k}, r)$.
3. \mathcal{B}_0 sends R to the challenger, and sends AUX to the post-revocation adversary \mathcal{B}_1 .
4. Upon receiving sk from the challenger, \mathcal{B}_1 runs $m' \leftarrow \mathcal{A}_1(\text{sk}, \text{AUX})$ and outputs $r \oplus m'$.

Therefore, \mathcal{B} succeeds with non-negligible probability $\epsilon(\lambda)$, which yields the desired contradiction. \square

10.2 Quantum Encryption with Wrong-Key Detection

We use the following variant⁵ of *wrong-key detection* for quantum encryption schemes of classical messages (QECM) schemes [[CMP24](#)].

Definition 10.2 (Wrong-key detection for QECM schemes). *A QECM scheme $\Sigma = (\text{KeyGen}, \text{Enc}, \text{Dec})$ satisfies the wrong-key detection (WKD) property if, for every $k' \neq k \leftarrow \text{KeyGen}(1^\lambda)$ and plaintext m ,*

$$\text{Tr}[(I - |\perp\rangle\langle\perp|)\text{Dec}_{k'} \circ \text{Enc}_k(m)] \leq \text{negl}(\lambda).$$

Next, we give a simple transformation that converts a multi-copy-search-secure revocable encryption scheme to one with the WKD property in the quantum random oracle model (QROM).

⁵[[CMP24](#)] consider the more general notion of quantum secret-key encryption schemes.

Construction 4 (Generic Transformation for WKD in the QROM). Let $\Sigma = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$ be a revocable encryption scheme and let λ be the security parameter. Fix a function $H : \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^\ell$. Then, the scheme $\Sigma^H = (\text{KeyGen}^H, \text{Enc}^H, \text{Dec}^H, \text{Revoke}^H)$ is defined by the following QPT algorithms:

- $\text{KeyGen}^H(1^\lambda)$: on input 1^λ , run $\text{KeyGen}(1^\lambda)$ to output a key $\text{sk} \in \{0, 1\}^\lambda$.
- $\text{Enc}^H(\text{sk}, m)$: on input $m \in \{0, 1\}^\lambda$, sample $x \sim \{0, 1\}^\lambda$, and output $(\text{Enc}(\text{sk}, x), H(\text{sk}||x), x \oplus m)$.
- $\text{Dec}^H(\text{sk}, \text{ct})$: on input ct , first parse $(\rho, c, y) \leftarrow \text{ct}$. Then, run $\text{Dec}(\text{sk}, \rho)$ and check (using an auxiliary register) whether $(\text{sk}||\cdot)$ (i.e., the outcome with pre-fix sk) maps to c under H . If yes, measure the register with outcome x' , and output $x' \oplus y$. Otherwise, output $|\perp\rangle\langle\perp|$.
- $\text{Revoke}^H(\text{sk}, \text{vk}, \sigma)$: On input the keys sk and vk , as well as the state σ , output the result of running the procedure $\text{Revoke}(\text{sk}, \text{vk}, \sigma)$ on the corresponding quantum part of the ciphertext.

Clearly, correctness is preserved. We now show the following lemma. For simplicity, we consider the standard notion of multi-copy revocable encryption security and remark that the oracular variant of security is analogous.

Lemma 10.3. Let Π be any multi-copy (search) secure revocable encryption scheme and let $H : \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^\ell$ be a hash function, for $\ell = 4\lambda$. Then, the scheme Σ^H in Construction 4 yields a multi-copy (search) secure revocable encryption scheme with WKD in the QROM.

Proof. Correctness is clearly preserved. Let us first verify the WKD property of the construction $\Sigma^H = (\text{KeyGen}^H, \text{Enc}^H, \text{Dec}^H, \text{Revoke}^H)$ in the QROM. We will analyze the probability that decryption with a wrong key k' passes the key detection test. Formally, for every $\kappa' \neq \kappa \leftarrow \text{KeyGen}^H(1^\lambda)$ and plaintext $m \in \{0, 1\}^\lambda$, we bound the probability as follows:

$$\begin{aligned} \Pr_H[\text{Dec}_{\kappa'}^H(\text{Enc}_{\kappa}^H(m)) \neq \perp] &= \Pr_{H,x}[H(\kappa' || \text{Dec}_{\kappa'}(\text{Enc}_{\kappa}(x))) = H(\kappa || x)] \\ &\leq \Pr_H[\exists w' \neq w \in \{0, 1\}^{2\lambda} : H(w') = H(w)] \\ &\leq \frac{2^{2\lambda} - 1}{2^{4\lambda}} = \text{negl}(\lambda). \end{aligned}$$

For security, suppose that \mathcal{A} is a successful adversary against the search variant of the revocable security game with respect to Σ^H , i.e., there exists a non-negligible function $\epsilon(\lambda)$ such that

$$\Pr[1 \leftarrow \text{RevokeExpt}_{\lambda, \Sigma^H, \mathcal{A}}] = \epsilon(\lambda).$$

For convenience, we model \mathcal{A} as a pair of quantum algorithms $(\mathcal{A}_0^H, \mathcal{A}_1^H)$, where \mathcal{A}_0^H corresponds to the pre-revocation adversary, and \mathcal{A}_1^H corresponds to the post-revocation adversary. Suppose also that \mathcal{A}^H makes at most $q = \text{poly}(\lambda)$ many queries to the random oracle H .

Now, consider the following hybrids:

H₁: This corresponds to $\text{RevokeExpt}_{\lambda, \Sigma^H, \mathcal{A}}$:

1. \mathcal{A}_0 submits $k = k(\lambda)$, where $k(\lambda)$ is a polynomial.
2. The challenger samples a message $m \sim \{0, 1\}^\lambda$ and $x \sim \{0, 1\}^\lambda$, lets $\text{sk} \leftarrow \text{KeyGen}(1^\lambda)$ and produces k copies of $|\psi\rangle \leftarrow \text{Enc}(\text{sk}, x)$. The challenger then sends $(|\psi\rangle^{\otimes k}, H(\text{sk}||x), x \oplus m)$ to \mathcal{A}_0^H .
3. \mathcal{A}_0^H prepares a bipartite state on registers R and AUX, and sends R to the challenger and AUX to the post-revocation adversary \mathcal{A}_1^H .
4. The challenger performs the projective measurement $\left\{ |\psi\rangle\langle\psi|^{\otimes k}, \mathbb{I} - |\psi\rangle\langle\psi|^{\otimes k} \right\}$ on register R. If the measurement succeeds, the game continues. Otherwise, the challenger outputs \perp .

5. The challenger sends sk to \mathcal{A}_1^H .
6. \mathcal{A}_1^H outputs a message m' .
7. The outcome of the experiment is 1, if and only if $m' = m$.

\mathbf{H}_2 : This is the same experiment as \mathbf{H}_2 , except that in Step 2 the challenger sends $(|\psi\rangle^{\otimes k}, z, x \oplus m)$ instead of $(|\psi\rangle^{\otimes k}, H(\text{sk}|x), x \oplus m)$ to \mathcal{A}_0 , where $z \sim \{0, 1\}^\ell$ is uniformly random.

Notice that in \mathbf{H}_2 , the oracle H is a completely random function and z is also completely random, i.e., both are independent of the rest of the scheme. Thus, if \mathbf{H}_1 and \mathbf{H}_2 are indistinguishable, $(\mathcal{A}_0^H, \mathcal{A}_1^H)$ breaks the revocable security of the hybrid encryption scheme in [Theorem 10.1](#), since both H and z can be simulated. Therefore, it suffices to argue that \mathbf{H}_1 and \mathbf{H}_2 are close in order to complete the proof.

Claim 10.4. \mathbf{H}_1 and \mathbf{H}_2 are negligibly close.

Proof. Suppose there exists an adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$ that can distinguish the two hybrids with non-negligible advantage $\epsilon(\lambda)$. In other words, \mathcal{A} satisfies

$$\left| \Pr \left[\begin{array}{l} \mathcal{A}_1^H(\text{sk}, \text{AUX}_{|\top}) = 1 : \\ \text{sk} \leftarrow \text{KeyGen}(1^\lambda) \\ m \sim \{0, 1\}^\lambda, x \sim \{0, 1\}^\lambda \\ |\psi\rangle \leftarrow \text{Enc}(\text{sk}, x) \\ H \sim \{H: \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^\ell\} \\ (\text{R}, \text{AUX}) \leftarrow \mathcal{A}_0^H(|\psi\rangle^{\otimes k}, H(\text{sk}|x), x \oplus m) \end{array} \right] \right. \\ \left. - \Pr \left[\begin{array}{l} \mathcal{A}_1^H(\text{sk}, \text{AUX}_{|\top}) = 1 : \\ \text{sk} \leftarrow \text{KeyGen}^H(1^\lambda) \\ m \sim \{0, 1\}^\lambda, x \sim \{0, 1\}^\lambda \\ |\psi\rangle \leftarrow \text{Enc}(\text{sk}, x) \\ H \sim \{H: \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^\ell\} \\ z \sim \{0, 1\}^\ell \\ (\text{R}, \text{AUX}) \leftarrow \mathcal{A}_0^H(|\psi\rangle^{\otimes k}, z, x \oplus m) \end{array} \right] \right| \geq \epsilon(\lambda). \quad (3)$$

Our goal is to use the O2H Lemma ([Lemma 4.1](#)) in order to argue that we can use $\mathcal{A}^H = (\mathcal{A}_0^H, \mathcal{A}_1^H)$ to extract $(\text{sk}|x)$. To this end, we introduce the following notation. For $H: \{0, 1\}^{2\lambda} \rightarrow \{0, 1\}^\ell$, a pre-image $x \in \{0, 1\}^{2\lambda}$ and image $z \in \{0, 1\}^\ell$, we use $H_{x,z}$ to denote the re-programmed function

$$H_{x,z}(w) := \begin{cases} H(x), & \text{if } w \neq x \\ z, & \text{if } w = x. \end{cases}$$

Consider the following oracle-aided quantum algorithm $\tilde{\mathcal{A}}$ which receives access to an oracle \mathcal{O} and an input of the form $(\text{sk}, x, |\psi\rangle^{\otimes k}, z, x \oplus m)$:

1. $\tilde{\mathcal{A}}$ runs $(\text{R}, \text{AUX}) \leftarrow \mathcal{A}_0^{\mathcal{O}}(|\psi\rangle^{\otimes k}, z, x \oplus m)$.
2. $\tilde{\mathcal{A}}$ performs the projective measurement $\{|\text{Enc}(\text{sk}, x)\rangle\langle \text{Enc}(\text{sk}, x)|^{\otimes k}, \mathbb{I} - |\text{Enc}(\text{sk}, x)\rangle\langle \text{Enc}(\text{sk}, x)|^{\otimes k}\}$ on register R. If the measurement fails, $\tilde{\mathcal{A}}$ aborts.
3. $\tilde{\mathcal{A}}$ runs $\mathcal{A}_1^{\mathcal{O}}(\text{sk}, \text{AUX}_{|\top})$ and outputs whatever it outputs.

Notice that $\tilde{\mathcal{A}}$ has precisely the same distinguishing advantage $\epsilon(\lambda)$ in [Equation \(3\)](#). Next, we also define the quantum extractor $\text{Ext}^{\mathcal{O}}(\text{sk}, x, |\psi\rangle^{\otimes k}, z, x \oplus m)$ acting as follows:

1. Sample $i \leftarrow [q]$, where q denotes the total number of queries made by $\tilde{\mathcal{A}}$.
2. Run $\tilde{\mathcal{A}}^{\mathcal{O}}(\text{sk}, |\psi\rangle^{\otimes k}, z, x \oplus m)$ just before the $(i-1)$ -st query to \mathcal{O} .
3. Measure $\tilde{\mathcal{A}}$'s i -th query in the computational basis, and output the measurement outcome.

Applying the O2H Lemma (Lemma 4.1) to the *particular* extractor $\text{Ext}^{\mathcal{O}}$ we defined above, we get

$$\begin{aligned}
\epsilon(\lambda) &\leq \left| \Pr \left[\tilde{\mathcal{A}}^H(x, \text{sk}, |\psi\rangle^{\otimes k}, z, x \oplus m) = 1 : \begin{array}{l} \text{sk} \leftarrow \text{KeyGen}(1^\lambda) \\ m \sim \{0,1\}^\lambda, x \sim \{0,1\}^\lambda \\ |\psi\rangle \leftarrow \text{Enc}(\text{sk}, x) \\ H \sim \{H: \{0,1\}^{2\lambda} \rightarrow \{0,1\}^\ell\} \\ z \sim \{0,1\}^\ell, H \leftarrow H_{\text{sk}||x,z} \end{array} \right] \right. \\
&\quad \left. - \Pr \left[\tilde{\mathcal{A}}^G(x, \text{sk}, |\psi\rangle^{\otimes k}, z, x \oplus m) = 1 : \begin{array}{l} \text{sk} \leftarrow \text{KeyGen}(1^\lambda) \\ m \sim \{0,1\}^\lambda, x \sim \{0,1\}^\lambda \\ |\psi\rangle \leftarrow \text{Enc}(\text{sk}, x) \\ G \sim \{G: \{0,1\}^{2\lambda} \rightarrow \{0,1\}^\ell\} \\ z \sim \{0,1\}^\ell \end{array} \right] \right| \\
&\leq 2q \sqrt{\Pr \left[(\text{sk}||x) \leftarrow \text{Ext}^G(\text{sk}, x, |\psi\rangle^{\otimes k}, z, x \oplus m) : \begin{array}{l} \text{sk} \leftarrow \text{KeyGen}(1^\lambda) \\ m \sim \{0,1\}^\lambda, x \sim \{0,1\}^\lambda \\ |\psi\rangle \leftarrow \text{Enc}(\text{sk}, x) \\ G \sim \{G: \{0,1\}^{2\lambda} \rightarrow \{0,1\}^\ell\} \\ z \sim \{0,1\}^\ell \end{array} \right]}.
\end{aligned}$$

In other words, the quantum extractor succeeds with probability at least

$$\Pr \left[(\text{sk}||x) \leftarrow \text{Ext}^G(\text{sk}, x, |\psi\rangle^{\otimes k}, z, x \oplus m) : \begin{array}{l} \text{sk} \leftarrow \text{KeyGen}(1^\lambda) \\ m \sim \{0,1\}^\lambda, x \sim \{0,1\}^\lambda \\ |\psi\rangle \leftarrow \text{Enc}(\text{sk}, x) \\ G \sim \{G: \{0,1\}^{2\lambda} \rightarrow \{0,1\}^\ell\} \\ z \sim \{0,1\}^\ell \end{array} \right] \geq \frac{\epsilon^2}{4q}.$$

But this implies that we can find $(\text{sk}||x)$ with $1/\text{poly}(\lambda)$ probability by measuring a random query which occurs either as a result of running the pre-revocation adversary $\tilde{\mathcal{A}}_0^G(|\psi\rangle^{\otimes k}, z, x \oplus m)$ or of the post-revocation adversary $\tilde{\mathcal{A}}_1^G(\text{sk}, \text{AUX}_{|\top})$. In either case, we can use x to immediately find m . Because G is a completely random function and z is also completely random (i.e., both can be simulated) this breaks the revocable security of the hybrid encryption scheme in [Theorem 10.1](#). \square

10.3 Revocable Multi-Copy-Secure Point Functions

We are now ready to define our revocable multi-copy-secure scheme for multi-bit point functions, which we obtain from any revocable encryption scheme with the aforementioned “wrong-key detection mechanism”. Here, we consider multi-bit point functions $P_{y,m}$ of the form

$$P_{y,m}(x) = \begin{cases} m & \text{if } x = y, \\ 0^\lambda & \text{if } x \neq y, \end{cases}$$

where $y, m \in \{0,1\}^\lambda$. Point functions are simply a special case of programs, so for a formal definition of syntax and revocable multi-copy security, see [Definition 8.1](#) and [Definition 8.2](#) respectively.

Our construction is the following

Construction 5 (Construction for revocable multi-copy secure point functions).

Let $\lambda \in \mathbb{N}$ be the security parameter. To construct a revocable multi-copy secure scheme for multi-bit point functions with input and output sizes λ , respectively, let $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$ be a revocable encryption scheme with WKD, with security parameter and message length equal to λ . We define the scheme $\Sigma = (\text{Compile}, \text{Eval}, \text{Revoke})$ as follows:

- **Compile** $(1^\lambda, P_{y,m})$: on input a security parameter λ and a multi-bit point function $P_{y,m}$, succinctly specified by the marked input y (of size λ) and message m (of size λ), it
 - generates a pair $(|\Psi_P\rangle, \Pi.\text{vk}) \leftarrow \Pi.\text{Enc}(y, m)$, and
 - outputs the quantum state $|\Psi_P\rangle$ and (private) verification key $\text{vk} = (y, \Pi.\text{vk})$.

- $\text{Eval}(|\Psi_P\rangle, x)$: On input $|\Psi_P\rangle$ and a string $x \in \{0, 1\}^\lambda$, do the following: append an ancillary qubit in the $|0\rangle$ state. Then, coherently perform a two-outcome measurement to check whether $\Pi.\text{Dec}(x, |\Psi_P\rangle)$ is in the state $|\perp\rangle \langle \perp|$, or not, and store the resulting bit in the ancilla. If true, output 0^λ . Otherwise, rewind the procedure and measure in the standard basis to obtain a message m' .
- $\text{Revoke}(\text{vk}, \sigma)$: on input the key vk and a state σ , parse $(y, \Pi.\text{vk}) \leftarrow \text{vk}$ and then run $\Pi.\text{Revoke}(y, \Pi.\text{vk}, \sigma)$.

We now prove the security of Construction 5.

Theorem 10.5. *Let $\Pi = (\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Revoke})$ be any multi-copy revocable secure encryption scheme which has the WKD property. Then, Construction 5 yields a revocable multi-copy secure scheme for multi-bit point functions with respect to the pair of ensembles $\mathcal{D}_\varphi, \mathcal{D}_\chi$ in the QROM, where \mathcal{D}_φ is the (natural) uniform distribution over multi-bit point functions and \mathcal{D}_χ is arbitrary.*

Proof. The correctness of the scheme follows directly from the WKD property of Π . Let \mathcal{A} denote the adversary for $\text{RevokeExpt}_{\lambda, \Sigma, \mathcal{A}}^{\mathcal{D}_\varphi, \mathcal{D}_\chi}$. We consider two cases, namely when $p_{\text{triv}}^{\mathcal{D}_\varphi, \mathcal{D}_\chi}(\lambda) = 1$ and when $p_{\text{triv}}^{\mathcal{D}_\varphi, \mathcal{D}_\chi}(\lambda) < 1$. In the former case, the scheme is trivially secure by definition and we are done. Hence, we will assume that $p_{\text{triv}}^{\mathcal{D}_\varphi, \mathcal{D}_\chi}(\lambda) < 1$ for the remainder of the proof. Note that, in this case, $\mathcal{D}_{\mathcal{X}_\lambda}(P_{y,m})$ has non-zero weight on the marked input x , for a random multi-bit point function $P_{y,m}$.

Let $x \leftarrow \mathcal{D}_{\mathcal{X}_\lambda}(P_{y,m})$ denote the input received by \mathcal{A} during the post revocation phase. We can express the probability that \mathcal{A} succeeds at $\text{RevokeExpt}_{\lambda, \Sigma, \mathcal{A}}^{\mathcal{D}_\varphi, \mathcal{D}_\chi}$ as follows:

$$\begin{aligned} \Pr[\mathcal{A} \text{ wins}] &= \Pr[\mathcal{A} \text{ wins} \mid y \neq x] \cdot \Pr[y \neq x] + \Pr[\mathcal{A} \text{ wins} \mid y = x] \cdot \Pr[y = x]. \end{aligned} \quad (4)$$

$$\leq \Pr[\mathcal{A} \text{ wins} \mid y \neq x] + \Pr[\mathcal{A} \text{ wins} \mid y = x] \cdot \Pr[y = x]. \quad (5)$$

$$= p_{\text{triv}}^{\mathcal{D}_\varphi, \mathcal{D}_\chi}(\lambda) + \Pr[\mathcal{A} \text{ wins} \mid y = x] \cdot \Pr[y = x]. \quad (6)$$

where the second equality follows because when $x \neq y$, then the post-revocation adversary \mathcal{A}_1 does not have access to the secret key y , and the winning probability is simply the optimal guessing probability in this case.

We complete the proof by showing that $\Pr[\mathcal{A} \text{ wins} \mid y = x] \leq \text{negl}(\lambda)$. This would imply that

$$\Pr[\mathcal{A} \text{ wins}] \leq p_{\text{triv}}^{\mathcal{D}_\varphi, \mathcal{D}_\chi}(\lambda) + \text{negl}(\lambda). \quad (7)$$

Suppose that \mathcal{A} succeeds with non negligible probability on the challenge x such that $x = y$. We will use \mathcal{A} to construct an adversary against the multi-copy (search) revocable security of the scheme Π . For convenience, we model \mathcal{A} as a pair of QPT algorithms $(\mathcal{A}_0, \mathcal{A}_1)$, where \mathcal{A}_0 corresponds to the pre-revocation adversary, and \mathcal{A}_1 corresponds to the post-revocation adversary. Consider the QPT adversary $\mathcal{B} = (\mathcal{B}_0, \mathcal{B}_1)$ against Π , which we define as follows:

1. \mathcal{B}_0 runs \mathcal{A}_0 to obtain $k = k(\lambda)$, where $k(\lambda)$ is a polynomial, and sends it to the challenger.
2. Upon receiving a multi-copy ciphertext of the form $|\psi\rangle^{\otimes k}$, where $|\psi\rangle = \text{Enc}(y, m)$, for a random $m \sim \{0, 1\}^\lambda$, \mathcal{B}_0 sends $|\psi\rangle^{\otimes k}$ to \mathcal{A}_0 , and runs $(R, \text{AUX}) \leftarrow \mathcal{A}_0(|\psi\rangle^{\otimes k})$.
3. \mathcal{B}_0 sends R to the challenger.
4. The challenger performs the projective measurement $\left\{ |\psi\rangle\langle\psi|^{\otimes k}, \mathbb{I} - |\psi\rangle\langle\psi|^{\otimes k} \right\}$ on R . If the measurement outcome is 1 then the challenger outputs \perp . Otherwise, the challenger continues.
5. \mathcal{B}_0 sends AUX to the post-revocation adversary \mathcal{B}_1 .
6. \mathcal{B}_1 receives the secret key $\text{sk} = y$ from the challenger, and it forwards y to \mathcal{A}_1 .
7. Finally, \mathcal{B}_1 outputs the output obtained by running \mathcal{A}_1 .

$\Pr[B \text{ wins}] = \Pr[A \text{ wins} \mid x = y] = \text{non-negl}(\lambda)$, which is a contradiction. \square

Finally, we end this section with the following corollary, which is immediate from the previous results in this section.

Corollary 10.6. *There exist revocable multi-bit point functions which satisfy (an oracular notion of) multi-copy security in both the random permutation model, as well as the quantum random oracle model.*

\square

11 Applications to Sponge Hashing

In this section, we show that the techniques we developed in this paper, i.e., our analysis of $k \mapsto k + 1$ unforgeability of random subsets via query lower bounds for permutations, are more broadly applicable and extend to other cryptographic settings as well. Here, we single out the so-called *sponge construction* and show how to prove new space-time trade-offs for finding elements in sponge hash tables. Our key insight is that the aforementioned query problem can be viewed as a particular instance of $k \mapsto k + 1$ unforgeability of random subsets in the presence of membership oracles.

11.1 Single-Round Sponge Hashing

The National Institute of Standards and Technology (NIST) recently announced a new international hash function standard known as SHA-3. This is a new family of cryptographic functions based on the idea of *sponge hashing* [BDPA11a, BDPA11b]. This particular approach allows for both variable input length and variable output length, which makes it particularly attractive towards the design of cryptographic hash functions. The internal state of a sponge function gets updated through successive applications of a so-called *block function* (modeled as a random permutation) $\varphi : \{0, 1\}^{r+c} \rightarrow \{0, 1\}^{r+c}$, where we call the parameters $r \in \mathbb{N}$ the *rate* and $c \in \mathbb{N}$ the *capacity* of the sponge.

The basic (single-round) sponge hash function $\text{Sp}_{\text{IV}}^\varphi : \{0, 1\}^r \rightarrow \{0, 1\}^r$ on input $x \in \{0, 1\}^r$ is defined by $h = \text{Sp}_{\text{IV}}^\varphi(x)$, where h corresponds to the first r bits of $\varphi(x \parallel \text{IV}) := (h \parallel w)$, for some fixed *initialization vector* $\text{IV} \in \{0, 1\}^c$. Here, the post-fix $w \in \{0, 1\}^c$ is simply discarded during the evaluation of the hash function, and we can think of $\text{IV} \in \{0, 1\}^c$ as a random *salt*. We give an illustration of the sponge hash function below.




Figure 1: The single-round sponge hash function.

While much is known about the security of the sponge in a classical world [BDPV08], the post-quantum security of the sponge construction in the case when the block function is a random (but invertible) permutation remains a major open problem [CBH⁺17, CMSZ19, Zha21b, CP24, MMW24]. The inherent difficulty in proving the post-quantum security of the sponge construction lies in the fact that an adversary needs to be modeled as a quantum query algorithm which has access to φ and its inverse φ^{-1} . This is where existing techniques [CBH⁺17, CMSZ19] seem to break down and require an entirely new approach, e.g., as the recent works of [CP24, MMW24] suggest.

11.2 Space-Time Trade-Offs for Finding Elements in Sponge Hash Tables

We now consider a simple query problem which seems to lie outside of the scope of existing techniques in [CBH⁺17, CMSZ19, CP24, MMW24]. Specifically, we consider the following task of finding elements in hash tables, which is also a potentially relevant problem in practice.

A simple query problem: Suppose that $\varphi : \{0, 1\}^{r+c} \rightarrow \{0, 1\}^{r+c}$ is a random permutation, for some $r, c \in \mathbb{N}$. Let $\text{IV} \in \{0, 1\}^c$ be a random salt. Given as input a hash table of size S which consists of

$$(h_1 = \text{Sp}_{\text{IV}}^\varphi(x_1), \dots, h_S = \text{Sp}_{\text{IV}}^\varphi(x_S))$$

where $x_1, \dots, x_S \leftarrow \{0, 1\}^r$ are randomly chosen distinct inputs, how many queries to φ^{-1} and a checking oracle $\text{Valid}_{\varphi, \text{IV}}$ (which outputs 1, if an input hash is in the range of $\text{Sp}_{\text{IV}}^\varphi$, and 0 otherwise) does a quantum algorithm need to find a valid new element in the range of $\text{Sp}_{\text{IV}}^\varphi$? Suppose that

$$\Pr \left[\text{Valid}_{\varphi, \text{IV}}(h') = 1 \wedge h' \neq h_i, \forall i \in [S] : h' \leftarrow \mathcal{A}^{\varphi^{-1}, \text{Valid}_{\varphi, \text{IV}}}(h_1, \dots, h_S) \right] = \epsilon > 0.$$

Can we prove that ϵ must satisfy a space-time trade-off, i.e., in terms of the size S of the hash table, and the number of queries T ? In other words, can we show that the only way an adversary can add a new element to the hash table is by making a large number of queries? Clearly, there is a subtle trade-off that appears in this scenario, since a sufficiently large hash table could potentially make it harder to find valid new elements.

Analyzing such a query problem is non-trivial for a number of reasons. The first challenge lies in the fact the oracle-aided algorithm \mathcal{A} receives as input list of hashes which depend on the salt itself (i.e., the algorithm can act *adaptively*). Typically, in traditional function inversion tasks [HXY19, ABPS24], the adversary only receives advice on the hash function itself, i.e., it does not get to receive information which depends on the actual salt as part of the pre-processing phase. The second challenge lies in the fact that \mathcal{A} receives not just one, but two oracles φ^{-1} and $\text{Valid}_{\varphi, \text{IV}}$ which depend on φ , and also the salt IV .

Surprisingly, our techniques for proving the unforgeability of random subsets via query lower bounds for permutations are amendable in this context and allow us to show the following theorem.

Theorem 11.1. *Let $r, c \in \mathbb{N}$ be integers and let $\lambda = \min(r, c)$. Let $\varphi : \{0, 1\}^{r+c} \rightarrow \{0, 1\}^{r+c}$ be a uniformly random permutation and $\text{IV} \leftarrow \{0, 1\}^c$ be a random salt. Suppose that $\mathcal{A}^{\varphi^{-1}, \text{Valid}_{\varphi, \text{IV}}}$ is a T -query quantum algorithm which receives as input a hash table $(h_1 = \text{Sp}_{\text{IV}}^\varphi(x_1), \dots, h_S = \text{Sp}_{\text{IV}}^\varphi(x_S))$ of size S , where $x_1, \dots, x_S \leftarrow \{0, 1\}^r$ are randomly chosen distinct inputs. Suppose that \mathcal{A} succeeds at finding a new element in the hash table with probability*

$$\Pr \left[\text{Valid}_{\varphi, \text{IV}}(h') = 1 \wedge h' \neq h_i, \forall i \in [S] : h' \leftarrow \mathcal{A}^{\varphi^{-1}, \text{Valid}_{\varphi, \text{IV}}}(h_1, \dots, h_S) \right] = \epsilon > 0.$$

Then, \mathcal{A} satisfies the space-time trade-off given by

$$\epsilon \leq O \left(T \cdot \sqrt{\frac{2^r - S}{2^{r+c}}} + \sqrt{\frac{2^r - S}{2^r}} + \frac{T^3}{2^{r+c} - 2^r} + \frac{T^3}{2^r} \right).$$

Proof. We give a direct reduction from the $S \mapsto S + 1$ unforgeability experiment in Theorem A.1. Suppose we are given as input (a_1, \dots, a_S) as well as a membership oracle \mathcal{O}_A for a random subset $A \subseteq \{0, 1\}^{r+c}$ of size $|A| = 2^r$. Consider the reduction $\mathcal{B}^{\mathcal{O}_A}(a_1, \dots, a_S)$ which does the following:

1. \mathcal{B} samples a random salt $\text{IV} \leftarrow \{0, 1\}^c$.
2. \mathcal{B} samples a random function $f : \{0, 1\}^{r+c} \rightarrow \{0, 1\}^{r+c}$ subject to the constraint that $f(a) = *||\text{IV}$, whenever $a \in A$.
3. \mathcal{B} runs $h' \leftarrow \mathcal{A}^{f, \mathcal{V}_{f, \text{IV}}}(h_1, \dots, h_S)$ to obtain a string $h' \in \{0, 1\}^r$, where we let h_1, \dots, h_S denote the first r bits of the strings a_1, \dots, a_S , respectively, and where we let

$$\mathcal{V}_{f, \text{IV}}(x) = \begin{cases} 1, & \text{if } f(x) = *||\text{IV} \\ 0, & \text{otherwise.} \end{cases}$$

4. \mathcal{B} does a standard Grover search to find a string $g \in \{0, 1\}^c$ such that $a'_{S+1} \in A$ and $a'_{S+1} \neq a_i$, for all $i \in [k]$, with $a'_{S+1} := h' || g$. If no such string g is found, \mathcal{B} outputs \perp . Note that this step requires at most $O(\sqrt{2^c})$ many additional queries in the worst case.

We now analyze the probability that \mathcal{B} solves the $S \mapsto S + 1$ unforgeability problem by finding $S + 1$ distinct elements in A . Using an identical analysis as in the proof of [Theorem 7.1](#), we can show that the probability that $\mathcal{A}^f, \mathcal{V}_{f, \mathcal{IV}}(h_1, \dots, h_S)$ finds a valid element such that $\mathcal{V}_{f, \mathcal{IV}}(h') = 1$ drops by at most

$$O\left(\frac{T^3}{2^{r+c} - 2^r} + \frac{T^3}{2^r}\right)$$

compared to ϵ . Moreover, \mathcal{B} solves the $S \mapsto S + 1$ unforgeability problem with precisely the same loss in probability. We can now invoke [Theorem A.1](#) to argue that

$$\epsilon - O\left(\frac{T^3}{2^{r+c} - 2^r}\right) - O\left(\frac{T^3}{2^r}\right) \leq O\left((T + \sqrt{2^c}) \cdot \sqrt{\frac{2^r - S}{2^{r+c}}} + \frac{2^r - S}{2^{r+c}}\right).$$

Re-arranging for ϵ and simplifying the bound, we get that

$$\epsilon \leq O\left(T \cdot \sqrt{\frac{2^r - S}{2^{r+c}}} + \sqrt{\frac{2^r - S}{2^r}} + \frac{T^3}{2^{r+c} - 2^r} + \frac{T^3}{2^r}\right).$$

Open questions. We now state a number of interesting open questions which are relevant to this section. While our analysis of finding elements in sponge hash tables does result in an interesting space-time trade-off, we do not believe that our bounds are potentially not tight. We conjecture that, if r is significantly larger than c , then \mathcal{A} requires at least

$$\Omega\left(\sqrt{\frac{2^{r+c}}{2^r - S}}\right)$$

many queries to succeed with constant success probability. We leave this as an interesting open problem. Finally, we believe that our analysis could also be extended in the case when \mathcal{A} receives access to the permutation φ —in addition to the oracles φ^{-1} and $\text{Valid}_{\varphi, \mathcal{IV}}$. Because our techniques do not seem apply in this setting, this suggest that an entirely new approach is necessary. □

References

- [Aar09] Scott Aaronson. “Quantum copy-protection and quantum money”. In: *2009 24th Annual IEEE Conference on Computational Complexity*. IEEE. 2009, pp. 229–242 (cit. on p. 5).
- [ABPS24] Gorjan Alagic, Chen Bai, Alexander Poremba, and Kaiyan Shi. “On the Two-sided Permutation Inversion Problem”. In: *IACR Communications in Cryptology 1.1* (Apr. 9, 2024). ISSN: 3006-5496. DOI: [10.62056/a0qj89n4e](https://doi.org/10.62056/a0qj89n4e) (cit. on p. 35).
- [AC12] Scott Aaronson and Paul Christiano. “Quantum money from hidden subspaces”. In: *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. 2012, pp. 41–60 (cit. on pp. 2, 3, 5, 6).
- [AGQY22] Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. “Pseudorandom (function-like) quantum state generators: New definitions and applications”. In: *Theory of Cryptography Conference*. Springer. 2022, pp. 237–265 (cit. on p. 2).
- [AHH24] Prabhanjan Ananth, Zihan Hu, and Zikuan Huang. “Quantum Key-Revocable Dual-Regev Encryption, Revisited”. In: *Cryptology ePrint Archive* (2024) (cit. on pp. 1–4).

- [AHU18] Andris Ambainis, Mike Hamburg, and Dominique Unruh. *Quantum security proofs using semi-classical oracles*. Cryptology ePrint Archive, Paper 2018/904. <https://eprint.iacr.org/2018/904>. 2018. URL: <https://eprint.iacr.org/2018/904> (cit. on pp. 9, 10, 12, 13).
- [AJOP20] Gorjan Alagic, Stacey Jeffery, Maris Ozols, and Alexander Poremba. “On Quantum Chosen-Ciphertext Attacks and Learning with Errors”. In: *Cryptography* 4.1 (2020). ISSN: 2410-387X. DOI: [10.3390/cryptography4010010](https://doi.org/10.3390/cryptography4010010). URL: <https://www.mdpi.com/2410-387X/4/1/10> (cit. on p. 12).
- [AK22] Prabhanjan Ananth and Fatih Kaleoglu. “A Note on Copy-Protection from Random Oracles”. In: *arXiv preprint arXiv:2208.12884* (2022) (cit. on p. 5).
- [AKL23] Prabhanjan Ananth, Fatih Kaleoglu, and Qipeng Liu. “Cloning Games: A General Framework for Unclonable Primitives”. In: *arXiv preprint arXiv:2302.01874* (2023) (cit. on p. 5).
- [AKN⁺23] Shweta Agrawal, Fuyuki Kitagawa, Ryo Nishimaki, Shota Yamada, and Takashi Yamakawa. “Public key encryption with secure key leasing”. In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2023, pp. 581–610 (cit. on pp. 1–4).
- [AL21] Prabhanjan Ananth and Rolando L La Placa. “Secure software leasing”. In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer. 2021, pp. 501–530 (cit. on pp. 1, 6).
- [ALL⁺21] Scott Aaronson, Jiahui Liu, Qipeng Liu, Mark Zhandry, and Ruizhe Zhang. “New approaches for quantum copy-protection”. In: *Advances in Cryptology—CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part I 41*. Springer. 2021, pp. 526–555 (cit. on pp. 4, 5).
- [ALY24] Prabhanjan Ananth, Yao-Ting Lin, and Henry Yuen. “Pseudorandom strings from pseudorandom quantum states”. In: *ITCS*. 2024 (cit. on p. 2).
- [APV23] Prabhanjan Ananth, Alexander Poremba, and Vinod Vaikuntanathan. “Revocable cryptography from learning with errors”. In: *Theory of Cryptography Conference*. Springer. 2023, pp. 93–122 (cit. on pp. 1–4).
- [AQY22] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. “Cryptography from pseudorandom quantum states”. In: *Annual International Cryptology Conference*. Springer. 2022, pp. 208–236 (cit. on p. 2).
- [BBBV97] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. “Strengths and Weaknesses of Quantum Computing”. In: *SIAM Journal on Computing* 26.5 (Oct. 1997), pp. 1510–1523. ISSN: 1095-7111. DOI: [10.1137/S0097539796300933](https://doi.org/10.1137/S0097539796300933). URL: <http://dx.doi.org/10.1137/S0097539796300933> (cit. on p. 12).
- [BBSS23] Amit Behera, Zvika Brakerski, Or Sattath, and Omri Shmueli. “Pseudorandomness with proof of destruction and applications”. In: *Theory of Cryptography Conference*. Springer. 2023, pp. 125–154 (cit. on p. 6).
- [BDF⁺11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. “Random Oracles in a Quantum World”. In: *Advances in Cryptology – ASIACRYPT 2011*. Ed. by Dong Hoon Lee and Xiaoyun Wang. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 41–69. ISBN: 978-3-642-25385-0 (cit. on p. 12).
- [BDPA11a] G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche. *Cryptographic sponge functions*. Submission to NIST (Round 3). 2011. URL: <http://sponge.noekeon.org/CSF-0.1.pdf> (cit. on pp. 4, 5, 34).
- [BDPA11b] G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche. *The Keccak SHA-3 submission*. Submission to NIST (Round 3). 2011. URL: <http://keccak.noekeon.org/Keccak-submission-3.pdf> (cit. on pp. 4, 5, 34).

- [BDPV08] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. “On the Indifferentiability of the Sponge Construction”. In: *Advances in Cryptology – EUROCRYPT 2008*. Ed. by Nigel Smart. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 181–197. ISBN: 978-3-540-78967-3 (cit. on p. 34).
- [BEM⁺23] John Bostanci, Yuval Efron, Tony Metger, Alexander Poremba, Luowen Qian, and Henry Yuen. *Unitary Complexity and the Uhlmann Transformation Problem*. 2023. arXiv: 2306.13073 [quant-ph]. URL: <https://arxiv.org/abs/2306.13073> (cit. on p. 2).
- [BHHP24] John Bostanci, Jonas Haferkamp, Dominik Hangleiter, and Alexander Poremba. *Efficient Quantum Pseudorandomness from Hamiltonian Phase States*. 2024. arXiv: 2410.08073 [quant-ph]. URL: <https://arxiv.org/abs/2410.08073> (cit. on p. 2).
- [BI20] Anne Broadbent and Rabib Islam. “Quantum encryption with certified deletion”. In: *Theory of Cryptography: 18th International Conference, TCC 2020, Durham, NC, USA, November 16–19, 2020, Proceedings, Part III* 18. Springer. 2020, pp. 92–122 (cit. on pp. 3, 6).
- [BJL⁺21] Anne Broadbent, Stacey Jeffery, Sébastien Lord, Supartha Podder, and Aarthi Sundaram. “Secure software leasing without assumptions”. In: *Theory of Cryptography Conference*. Springer. 2021, pp. 90–120 (cit. on p. 6).
- [BK23] James Bartusek and Dakshita Khurana. “Cryptography with Certified Deletion”. In: *Advances in Cryptology – CRYPTO 2023: 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20–24, 2023, Proceedings, Part V*. Santa Barbara, CA, USA: Springer-Verlag, 2023, pp. 192–223. ISBN: 978-3-031-38553-7. DOI: 10.1007/978-3-031-38554-4_7. URL: https://doi.org/10.1007/978-3-031-38554-4_7 (cit. on p. 6).
- [BKM⁺23] James Bartusek, Dakshita Khurana, Giulio Malavolta, Alexander Poremba, and Michael Walter. *Weakening Assumptions for Publicly-Verifiable Deletion*. Cryptology ePrint Archive, Paper 2023/559. 2023. URL: <https://eprint.iacr.org/2023/559> (cit. on p. 6).
- [BKP23] James Bartusek, Dakshita Khurana, and Alexander Poremba. “Publicly-Verifiable Deletion via Target-Collapsing Functions”. In: *Advances in Cryptology – CRYPTO 2023: 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20–24, 2023, Proceedings, Part V*. Santa Barbara, CA, USA: Springer-Verlag, 2023, pp. 99–128. ISBN: 978-3-031-38553-7. DOI: 10.1007/978-3-031-38554-4_4. URL: https://doi.org/10.1007/978-3-031-38554-4_4 (cit. on p. 6).
- [BL19] Anne Broadbent and Sébastien Lord. “Uncloneable quantum encryption via oracles”. In: *arXiv preprint arXiv:1903.00130* (2019) (cit. on pp. 3, 10).
- [BS19] Zvika Brakerski and Omri Shmueli. “(Pseudo) random quantum states with binary phase”. In: *Theory of Cryptography Conference*. Springer. 2019, pp. 229–250 (cit. on p. 2).
- [CBH⁺17] Jan Czajkowski, Leon Groot Bruinderink, Andreas Hülsing, Christian Schaffner, and Dominique Unruh. *Post-quantum security of the sponge construction*. Cryptology ePrint Archive, Paper 2017/771. <https://eprint.iacr.org/2017/771>. 2017. URL: <https://eprint.iacr.org/2017/771> (cit. on pp. 34, 35).
- [ÇG23] Alper Çakan and Vipul Goyal. “Unclonable cryptography with unbounded collusions”. In: *Cryptology ePrint Archive* (2023) (cit. on pp. 2, 5).
- [CGJL23] Orestis Chardouvelis, Vipul Goyal, Aayush Jain, and Jiahui Liu. “Quantum key leasing for pke and fhe with a classical lessor”. In: *arXiv preprint arXiv:2310.14328* (2023) (cit. on pp. 1, 2, 4).

- [CLLZ21] Andrea Coladangelo, Jiahui Liu, Qipeng Liu, and Mark Zhandry. “Hidden Cosets and Applications to Unclonable Cryptography”. In: *Advances in Cryptology – CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part I*. Berlin, Heidelberg: Springer-Verlag, 2021, pp. 556–584. ISBN: 978-3-030-84241-3. DOI: [10.1007/978-3-030-84242-0_20](https://doi.org/10.1007/978-3-030-84242-0_20). URL: https://doi.org/10.1007/978-3-030-84242-0_20 (cit. on pp. 3, 4).
- [Clo24] Google Cloud. *Key Rotation*. <https://cloud.google.com/kms/docs/key-rotation>. 2024 (cit. on p. 1).
- [CMP24] Andrea Coladangelo, Christian Majenz, and Alexander Poremba. “Quantum copy-protection of compute-and-compare programs in the quantum random oracle model”. In: *Quantum* 8 (May 2024), p. 1330. ISSN: 2521-327X. DOI: [10.22331/q-2024-05-02-1330](https://doi.org/10.22331/q-2024-05-02-1330). URL: <https://doi.org/10.22331/q-2024-05-02-1330> (cit. on pp. 5, 6, 11, 29).
- [CMSZ19] Jan Czapkowski, Christian Majenz, Christian Schaffner, and Sebastian Zur. *Quantum Lazy Sampling and Game-Playing Proofs for Quantum Indifferentiability*. Cryptology ePrint Archive, Paper 2019/428. <https://eprint.iacr.org/2019/428>. 2019. URL: <https://eprint.iacr.org/2019/428> (cit. on pp. 34, 35).
- [CP24] Joseph Carolan and Alexander Poremba. “Quantum One-Wayness of the Single-Round Sponge with Invertible Permutations”. In: *Advances in Cryptology – CRYPTO 2024*. Ed. by Leonid Reyzin and Douglas Stebila. Cham: Springer Nature Switzerland, 2024, pp. 218–252. ISBN: 978-3-031-68391-6 (cit. on pp. 34, 35).
- [Die82] DGBJ Dieks. “Communication by EPR devices”. In: *Physics Letters A* 92.6 (1982), pp. 271–272 (cit. on p. 2).
- [GB23] Tudor Giurgica-Tiron and Adam Bouland. “Pseudorandomness from subset states”. In: *arXiv preprint arXiv:2312.09206* (2023) (cit. on pp. 3, 6).
- [Har13] Aram W. Harrow. *The Church of the Symmetric Subspace*. 2013. arXiv: [1308.6595 \[quant-ph\]](https://arxiv.org/abs/1308.6595). URL: <https://arxiv.org/abs/1308.6595> (cit. on p. 6).
- [HXY19] Minki Hhan, Keita Xagawa, and Takashi Yamakawa. *Quantum Random Oracle Model with Auxiliary Input*. Cryptology ePrint Archive, Paper 2019/1093. 2019. URL: <https://eprint.iacr.org/2019/1093> (cit. on p. 35).
- [JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. “Pseudorandom quantum states”. In: *Advances in Cryptology – CRYPTO 2018: 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19–23, 2018, Proceedings, Part III 38*. Springer. 2018, pp. 126–152 (cit. on p. 2).
- [JMW23] Fernando Granha Jeronimo, Nir Magrafta, and Pei Wu. “Subset States and Pseudorandom States”. In: *arXiv preprint arXiv:2312.15285* (2023) (cit. on pp. 3, 6, 7, 13).
- [KNY08] Masaru Kada, Harumichi Nishimura, and Tomoyuki Yamakami. “The efficiency of quantum identity testing of multiple states”. In: *Journal of Physics A: Mathematical and Theoretical* 41.39 (2008), p. 395309 (cit. on p. 2).
- [KNY21] Fuyuki Kitagawa, Ryo Nishimaki, and Takashi Yamakawa. “Secure software leasing from standard assumptions”. In: *Theory of Cryptography Conference*. Springer. 2021, pp. 31–61 (cit. on p. 6).
- [KT24] Dakshita Khurana and Kabir Tomer. “Commitments from quantum one-wayness”. In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. 2024, pp. 968–978 (cit. on p. 2).
- [LLQZ22] Jiahui Liu, Qipeng Liu, Luowen Qian, and Mark Zhandry. “Collusion resistant copy-protection for watermarkable functionalities”. In: *Theory of Cryptography Conference*. Springer. 2022, pp. 294–323 (cit. on pp. 2, 5).

- [MMW24] Christian Majenz, Giulio Malavolta, and Michael Walter. *Permutation Superposition Oracles for Quantum Query Lower Bounds*. 2024. arXiv: [2407.09655](https://arxiv.org/abs/2407.09655) [quant-ph]. URL: <https://arxiv.org/abs/2407.09655> (cit. on pp. 34, 35).
- [MPSY24] Tony Metger, Alexander Poremba, Makrand Sinha, and Henry Yuen. *Simple constructions of linear-depth t -designs and pseudorandom unitaries*. 2024. arXiv: [2404.12647](https://arxiv.org/abs/2404.12647) [quant-ph]. URL: <https://arxiv.org/abs/2404.12647> (cit. on p. 2).
- [MPY23] Tomoyuki Morimae, Alexander Poremba, and Takashi Yamakawa. “Revocable Quantum Digital Signatures”. In: *arXiv preprint arXiv:2312.13561* (2023) (cit. on pp. 1–3).
- [MY22a] Tomoyuki Morimae and Takashi Yamakawa. “One-wayness in quantum cryptography”. In: *arXiv preprint arXiv:2210.03394* (2022) (cit. on p. 2).
- [MY22b] Tomoyuki Morimae and Takashi Yamakawa. “Quantum commitments and signatures without one-way functions”. In: *Annual International Cryptology Conference*. Springer. 2022, pp. 269–295 (cit. on p. 2).
- [NC11] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. 10th. USA: Cambridge University Press, 2011. ISBN: 1107002176 (cit. on p. 12).
- [Por22] Alexander Poremba. “Quantum proofs of deletion for learning with errors”. In: *arXiv preprint arXiv:2203.01610* (2022) (cit. on pp. 3, 6).
- [Riv98] Ronald L Rivest. “Can we eliminate certificate revocation lists?”. In: *Financial Cryptography: Second International Conference, FC’98 Anguilla, British West Indies February 23–25, 1998 Proceedings 2*. Springer. 1998, pp. 178–183 (cit. on p. 1).
- [Stu95] Stuart G Stubblebine. “Recent-secure authentication: Enforcing revocation in distributed systems”. In: *Proceedings 1995 IEEE Symposium on Security and Privacy*. IEEE. 1995, pp. 224–235 (cit. on p. 1).
- [Unr07] Dominique Unruh. “Random oracles and auxiliary input”. In: *Annual International Cryptology Conference*. Springer. 2007, pp. 205–223 (cit. on p. 7).
- [Unr13] Dominique Unruh. *Revocable quantum timed-release encryption*. Cryptology ePrint Archive, Paper 2013/606. 2013. DOI: [10.1145/2817206](https://doi.org/10.1145/2817206). URL: <https://eprint.iacr.org/2013/606> (cit. on pp. 1, 3, 4, 6).
- [Vaz98] Umesh Vazirani. “On the power of quantum computation”. In: *Philosophical Transactions of the Royal Society of London A 365: 1759-1768* (1998) (cit. on p. 13).
- [Wie83] Stephen Wiesner. “Conjugate Coding”. In: *SIGACT News* 15.1 (Jan. 1983), pp. 78–88. ISSN: 0163-5700. DOI: [10.1145/1008908.1008920](https://doi.org/10.1145/1008908.1008920). URL: <https://doi.org/10.1145/1008908.1008920> (cit. on p. 2).
- [Wik24] Wikipedia. *Certificate revocation*. https://en.wikipedia.org/wiki/Certificate_revocation. 2024 (cit. on p. 1).
- [WZ82] William K Wootters and Wojciech H Zurek. “A single quantum cannot be cloned”. In: *Nature* 299.5886 (1982), pp. 802–803 (cit. on p. 2).
- [Zha13] Mark Zhandry. *A Note on the Quantum Collision and Set Equality Problems*. 2013. arXiv: [1312.1027](https://arxiv.org/abs/1312.1027) [cs.CC]. URL: <https://arxiv.org/abs/1312.1027> (cit. on pp. 10, 13).
- [Zha16] Mark Zhandry. *A Note on Quantum-Secure PRPs*. 2016. arXiv: [1611.05564](https://arxiv.org/abs/1611.05564) [cs.CR]. URL: <https://arxiv.org/abs/1611.05564> (cit. on pp. 4, 13).
- [Zha21a] Mark Zhandry. “Quantum lightning never strikes the same state twice. or: quantum money from cryptographic assumptions”. In: *Journal of Cryptology* 34 (2021), pp. 1–56 (cit. on pp. 2, 3, 5).

[Zha21b] Mark Zhandry. “Redeeming Reset Indifferentiability and Applications to Post-quantum Security”. In: *Advances in Cryptology – ASIACRYPT 2021: 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6–10, 2021, Proceedings, Part I*. Singapore, Singapore: Springer-Verlag, 2021, pp. 518–548. ISBN: 978-3-030-92061-6. DOI: [10.1007/978-3-030-92062-3_18](https://doi.org/10.1007/978-3-030-92062-3_18). URL: https://doi.org/10.1007/978-3-030-92062-3_18 (cit. on p. 34).

A Appendix

Tighter bound in the classical case. Here, we give a slightly tighter variant of [Theorem 5.1](#), which we apply in the context of sponge hashing.

Theorem A.1 (Classical $k \mapsto k + 1$ Unforgeability). *Let $n \in \mathbb{N}$ and $k \in \mathbb{N}$. Then, for any q -query quantum oracle algorithm \mathcal{A} , and any $1 \leq k < s \leq 2^n$, it holds that*

$$\begin{aligned} & \Pr_{\substack{x_1, \dots, x_k \leftarrow S \\ (x_1, \dots, x_k) \in \text{dist}(S, k)}}} \Pr_{\substack{S \subseteq \{0, 1\}^n \\ |S| = s}} \left[x' \in S \wedge x' \neq x_i, \forall i \in [k] : x' \leftarrow \mathcal{A}^{\mathcal{O}_S}(x_1, \dots, x_k) \right] \\ & \leq O \left(q \cdot \sqrt{\frac{s-k}{2^n}} + \frac{s-k}{2^n} \right). \end{aligned}$$

In particular, for $k = \text{poly}(n)$, $q = \text{poly}(n)$ and $s(n) = n^{\omega(1)}$, the probability is at most $\text{negl}(n)$.

Proof. Suppose that $S \subseteq \{0, 1\}^n$ is a random subset of size $|S| = s$. We can model the quantum oracle algorithm $\mathcal{A}^{\mathcal{O}_S}$ on input (x_1, \dots, x_k) as a sequence of oracle queries and unitary computations followed by a measurement. Thus, the final output state just before the measurement can be written as

$$|\Psi_q^S\rangle = U_q \mathcal{O}_S U_{q-1} \dots U_1 \mathcal{O}_S U_0 |\psi_0\rangle |x_1, \dots, x_k\rangle,$$

where U_0, U_1, \dots, U_q are unitaries (possibly acting on additional workspace registers, which we omit above), and where $|\psi_0\rangle$ is some fixed initial state which is independent of S .

Now let $X = \{x_1, \dots, x_k\} \subset S$. We now consider the state

$$|\Psi_q^X\rangle = U_q \mathcal{O}_X U_{q-1} \dots U_1 \mathcal{O}_X U_0 |\psi_0\rangle |x_1, \dots, x_k\rangle.$$

We now claim that the states $|\Psi_q^S\rangle$ and $|\Psi_q^X\rangle$ are sufficiently close. From the definition of \mathcal{O}_X and \mathcal{O}_S , we have that $\mathcal{O}_X(x) \neq \mathcal{O}_S(x)$ iff $x \in S \setminus X \subset \{0, 1\}^n$. By the O2H Lemma ([Lemma 4.1](#)), we get

$$\begin{aligned} \mathbb{E}_{\substack{S \subseteq \{0, 1\}^n, |S| = s \\ x_1, \dots, x_k \leftarrow S \\ (x_1, \dots, x_k) \in \text{dist}(S, k)}}} \left\| |\Psi_q^S\rangle - |\Psi_q^X\rangle \right\| & \leq 2q \mathbb{E}_{\substack{S \subseteq \{0, 1\}^n, |S| = s \\ x_1, \dots, x_k \leftarrow S \\ (x_1, \dots, x_k) \in \text{dist}(S, k)}}} \sqrt{\frac{1}{q} \sum_{i=0}^{q-1} \left\| \Pi_{S \setminus X} |\Psi_i^X\rangle \right\|^2} \\ & \leq 2q \sqrt{\frac{1}{q} \sum_{i=0}^{q-1} \mathbb{E}_{\substack{S \subseteq \{0, 1\}^n, |S| = s \\ x_1, \dots, x_k \leftarrow S \\ (x_1, \dots, x_k) \in \text{dist}(S, k)}}} \left\| \Pi_{S \setminus X} |\Psi_i^X\rangle \right\|^2} \quad (\text{Jensen's inequality}) \\ & = O \left(q \cdot \sqrt{\frac{s-k}{2^n}} \right). \end{aligned}$$

Therefore, the probability (over the choice of S and X) that $\mathcal{A}^{\mathcal{O}_X}(x_1, \dots, x_k)$ succeeds is at most the

probability that $\mathcal{A}^{\mathcal{O}_S}(x_1, \dots, x_k)$ succeeds—up to an additive loss of $O(q \cdot \sqrt{\frac{s-k}{2^n}})$. Thus,

$$\begin{aligned}
& \Pr_{\substack{x_1, \dots, x_k \leftarrow S \\ (x_1, \dots, x_k) \in \text{dist}(S, k)}}} \Pr_{\substack{S \subseteq \{0,1\}^n \\ |S|=s}} \left[x' \in S \wedge x' \neq x_i, \forall i \in [k] : x' \leftarrow \mathcal{A}^{\mathcal{O}_S}(x_1, \dots, x_k) \right] \\
& \leq \Pr_{\substack{x_1, \dots, x_k \leftarrow S \\ (x_1, \dots, x_k) \in \text{dist}(S, k)}}} \Pr_{\substack{S \subseteq \{0,1\}^n \\ |S|=s}} \left[x' \in S \wedge x' \neq x_i, \forall i \in [k] : x' \leftarrow \mathcal{A}^{\mathcal{O}_x}(x_1, \dots, x_k) \right] + O\left(q \cdot \sqrt{\frac{s-k}{2^n}}\right) \\
& \leq O\left(q \cdot \sqrt{\frac{s-k}{2^n}} + \frac{s-k}{2^n}\right).
\end{aligned}$$

where the last equality follows from the fact that

$$\Pr_{\substack{x_1, \dots, x_k \leftarrow S \\ (x_1, \dots, x_k) \in \text{dist}(S, k)}}} \Pr_{\substack{S \subseteq \{0,1\}^n \\ |S|=s}} \left[x' \in S \wedge x' \neq x_i, \forall i \in [k] : x' \leftarrow \mathcal{A}^{\mathcal{O}_x}(x_1, \dots, x_k) \right] \leq \frac{s-k}{2^n}.$$

This proves the claim. □